


The role of Speciation in Spatial Coevolutionary Function Approximation

Folkert de Boer



Doctorate Thesis
July 2007

Cognitive Artificial Intelligence
University of Utrecht

1st Supervisor: dr. Marco Wiering
2nd Supervisor: Prof. dr. Paulien Hogeweg
3rd Supervisor: dr. ir. Edwin de Jong

Abstract

Evolution by reproduction, selection and mutation has given rise to complexity at many different levels, e.g. at the level of the many different species in the biosphere and at the level of complex information processing within individuals. We investigate how these two levels interact in an artificial setting. We investigate how diversification through speciation influences the evolution of complex information at the level of the individual. We use a coevolutionary computation model for function approximation. Two populations (one consisting of hosts, the other of parasites) co-evolve, their fitness defined on the interaction with the other population. The model is in a 2D space, and per generation a host encounters only a small subset of 'evolutionary problems', i.e. those in its direct neighborhood. This simple artificial setting allows us to study evolution of information processing, and integration of information over many generations. We study 2 evolutionary scenarios, one in which between generations all hosts and parasites are randomly distributed over space, and one in which the self induced spatial patterns of the system are maintained. We show that only in the latter case the hosts are in the end able to process all parasites correctly: i.e. the evolutionary target is reached. We analyze the evolutionary process by tracing back the lineages leading up to the final result. We show that in the undisturbed spatial system the host and parasite populations speciate in different subspecies which are adapted to each other and remain in the population until the target function is reached. In contrast in the well mixed case, although the diversity of parasites seen during the evolution is even higher than in the spatial case, we observe typical red queen evolution, in which the hosts continue to chase the parasites, unable to integrate enough information to reach the evolutionary target while only one host lineage survives the competition. We conclude that multiple levels of complexity reinforce each other in this coevolutionary setting: in our system spatial pattern formation and speciation appear to be prerequisites for the evolution of complex information processing.

Acknowledgements

First of all I would like to thank all the people from the theoretical biology & Bioinformatics group, where I wrote this thesis. Although in my opinion, the selfmade coffee is often much too strong, I really enjoy the many conversations in the coffeeroom, the lunches and the (sometimes) fruitful discussions. Also the suggestions and received help in a broad range of problems is really appreciated. Also thanks to my supervisors for reading and commenting the many draft versions of this thesis; Paulien, for the daily supervising, the helping and the guidance into the right (biological) direction, Marco for being a critical reader and for helping to make this thesis more readable and Edwin for the critical reflections and coevolutionary discussions. I thank Rients van Wijngaarden for initiating this project and the preceding work, but especially for the good time and discussions we had in London. I also would like to thank Daan for the 'stimulating' discussions about a lot of things and Marleen for her love and pointers when I got stuck again. Finally I would of course like to thank all my friends and family for support and understanding.

Folkert 2007

Contents

1	Introduction	1
1.1	Biological Evolution	1
1.1.1	Research Question I	2
1.2	Spatial Coevolution	2
1.2.1	Research Question II	4
1.3	Information Integration	4
1.3.1	Research Question III	5
1.4	CKI Relevance	5
1.5	GECCO '07	6
2	Evolutionary Computation	7
2.1	Inspiration from Biology	7
2.1.1	Adaptation and Speciation	7
2.1.2	Fitness	9
2.1.3	Fitness Landscapes	10
2.1.4	Survival of the Fittest	11
2.2	Artificial Evolution	12
2.2.1	Computational Application of Evolution	12
2.2.2	Problem Solving and Optimization	13
2.2.3	Direction	13
2.2.4	General Scheme of Evolutionary Computation	14
2.2.5	Exploration vs. Exploitation	15
2.3	Genetic Algorithms and Genetic Programming	15
3	Eco-evolutionary Dynamics	18
3.1	Evolutionary and Ecological Dynamics	18
3.2	Population Based vs. Individual Based Diversity	19
3.3	Sparse Fitness Evaluation	20
3.3.1	Biological Ground	20

3.3.2	Geological Grounding	21
3.3.3	Sparse Fitness in an Experimental Setup	21
3.4	Coevolution	22
3.4.1	Mutual Selective Pressure	22
3.4.2	A Coevolutionary Relationship	24
3.5	Arms Races	24
3.5.1	Red Queen Evolution	25
3.5.2	Life Dinner Principle	26
3.6	Coevolutionary Computation	27
3.6.1	Hosts and Parasites	28
3.6.2	The Ideal Trainer	28
3.7	Spatial Coevolution	29
3.7.1	Cellular Automata	30
4	Experimental Setup	32
4.1	The Model	32
4.1.1	Technical Aspects	33
4.2	Global Evolutionary Targets	33
4.3	Fitness Evaluation	34
4.4	Genetic Representation	36
4.5	Observables	37
5	Experimental Results	38
5.1	Efficacy of the Optimization	38
5.1.1	Variation	39
5.2	Coevolutionary Dynamics Depend on Fitness Evaluation	40
5.3	Parasite Speciation	41
5.4	Host Speciation	42
6	Conclusion	47
6.1	Summary and Conclusion	47
6.2	Discussion and Future Work	49

Chapter 1

Introduction

1.1 Biological Evolution

As Newton's notion of universal gravitation and his laws of motion brought both structure and innovation for the field of physics, Darwinian selection brought the biological field in a rapid. Darwin's concept of natural selection [Darwin, 1859] and his explanation for the *Origin of Species* have changed our view on nature irreversibly. When we look at earth, it is hard to believe that all present life has originated only from physical interactions and processes, operating on and within different populations of species. In its basic form, i.e. a process of *reproduction*, *mutation* and *selection*, evolution is the only scientific theory that provides a potential basis for the overwhelming diversity and complexity that we can observe in nature everywhere. The chance that someday another theory successfully will refute Darwinian selection as the driving force of evolution, like quantum mechanics did for Newtonian physics, is highly unlikely. It is the simplicity of the combination of the basic mechanisms which provide an enormous explanatory power, making Darwinian selection an extremely appealing and popular theory, subject to many studies. Although considerable work has been done, a lot of key features about the nature of evolutionary dynamics are still not properly understood, giving form to the field of evolutionary biology. The ecologist E.O. Wilson stated on this matter:

"What we understand best about evolution is mostly genetic, and what we understand least is mostly ecological. I will go further and suggest that the major remaining questions of evolutionary biology are ecological rather than genetic in content. They have to do with selection pressures from the environment as

revealed by the histories of particular lineages, not with genetic mechanisms of the most general nature.” [Wilson, 1992]

Maybe somewhat biased being an ecologist and not yet knowing of the deciphering of genomes, Wilson points out a very important notion in investigating evolutionary dynamics: evolutionary - and ecological timescales are heavily intertwined. Therefore our research will concern *eco-evolutionary dynamics*, considering the multiple levels of complexity which can reinforce and influence each other.

1.1.1 Research Question I

In explaining natural history, evolution can be described as a process of merely *adaptation* and *speciation* (see section 2.1.1). Adaptation and speciation are mainly driven by interactions and dynamics on the ecological timescale, and they are expressed in the evolutionary timescale, providing diversification through speciation. The role of adaptation in the process of speciation is a relatively clear one. When individuals within a species adapt to different circumstances, specialization of subpopulations can occur. This specialization can lead to divergence of genetic content within a population, which can give rise to the origin of different species. The opposite relation between speciation and adaptation however is still unclear. Therefore the main covering question of this thesis can be formulated as follows:

How does speciation of a population affect the adaptation of individuals?

1.2 Spatial Coevolution

The renowned phrase *survival of the fittest* [Spencer, 1864] seems to regard evolution merely as an optimization process¹. This notion of optimizing gave rise to a whole new field of research: the application of artificial evolution to the field of computational methods [Holland, 1975]. In nature the fitness of an individual depends in many ways on non-static features; organisms live in different environments, interacting and (co)evolving with other organisms. In order to use evolution for optimization, most artificial evolutionary models include a static fitness evaluation function, which clearly does not resemble the natural evolutionary process. When using the

¹for further discussion see subsection 2.1.4

terms and definitions of evolutionary computation strictly, probably all natural evolution has to be considered a *coevolutionary* process, because of the strong interconnection between the evolution of species and their environment. Evolutionary computation disconnects these aspects to be able to use and control the evolutionary process in a convenient way. This also points out why it is obvious that probably none of the artificial evolutionary settings resemble nature; it is clearly not the objective to give a realistic view on natural evolution. Only in the field of *coevolutionary computation*, these coevolutionary dynamics are explicitly taken into account and implemented in the evolutionary method.

The introduction of coevolutionary methods in the field of evolutionary computation is also not meant to achieve a 'realistic' model of natural evolution, it was purely introduced for computational optimization motives. Coevolution samples the problem domain more efficiently without having to evaluate all problems in the domain or defining very specific problem sets. As a consequence, we have to accept that the arising evolutionary dynamics in a coevolutionary computation system are much more complex and less transparent to understand.

Another inspiration from biology is to implement models in an explicit spatial environment. Evolutionary computation methods tend to neglect the role of a local environment where individuals live. In the perspective of biology, a world without *space* is a huge oversimplification. All organisms are constrained to the environment where they live. The interactions with other organisms and with their environment are purely local. In the nineties Hillis was the first to note that using the concept of space could also bring great advantage for artificial evolution. His computational coevolutionary process improved considerably when he placed the coevolving parties in a spatial environment, forcing the interactions to be local [Hillis, 1990]. How space really contributed to this success was not investigated until 1997, when Pagie and Hogeweg used a spatial coevolutionary parasite-host setting. They compared fitness evaluation based on a large static set of problems and fitness evaluation based on small coevolving sets of problems. With a spatial environment in the latter case, very little information is presented to the evolving hosts regarding the evolutionary target per evolutionary time step. In other words the fitness evaluation is very sparse. They showed surprisingly that the case with full fitness evaluation not only produced a better success rate, the produced solutions were also better generalizable and differed with respect to mutational stability; in the sparse fitness case, the sparse evaluated programs are less mutational stable than complete evaluated programs,

which preserves evolutionary freedom and can prevent overspecialization. Also it was suggested that spatiality causes an advantageous persistence in the population diversity over the generations and an even better targeting of weak points in the host-population by means of the coevolving parasites. These very promising results can probably be explained in terms of the emergence of spatial patterns. Spatial pattern formation creates the opportunity to form extra levels of selection. This multi-level character [Hogeweg, 1994] and the self-structuring in biological systems are already considered very important for natural evolution [Boerlijst, 1991, Pagie and Hogeweg, 1997, Savill et al., 1997]. Because space is neglected in most traditional studies, the consequences of these aspects in artificial evolving systems are mostly omitted.

1.2.1 Research Question II

Despite recent studies [Pagie and Mitchell, 2002, Williams and Mitchell, 2005, Mitchell et al., 2006] emphasize the importance and improved success of spatial embedding [Pagie and Hogeweg, 1997], the exact role of space and spatial pattern formation is still not fully understood. Therefore we formulate the next research question as:

What is the role of spatial pattern formation in spatial coevolutionary function approximation?

1.3 Information Integration

Evolution can be seen as a multi-level informatic process [Hogeweg, 1994]. Information presented to the system has to be integrated in order to be able to improve performance over the evolutionary timescale. The arising dynamics of the system are responsible for the way in which the provided information will be integrated. Information is presented and has to be integrated in the ecological timescale in order to speak of successful information integration on the evolutionary timescale. On the ecological timescale the dictating dynamics are determined by the interactions between and within populations. The interplay of these eco-evolutionary dynamics affect the information integration considerably.

Information present in a (eco)system can be distributed in roughly two ways. Individuals can adapt by regulation to 'all' possible circumstances (or 'problems') or the population can split into subpopulations, each specialized to a subset of circumstances. In terms of information integration, the variation

of all the information on the problemspace is reflected either in the variety of species (*population based diversity*), or present in all individuals within a species (*individual based diversity*). These different modes of information storage influence in their turn again several aspects of the ecological dynamics.

1.3.1 Research Question III

In a spatial model with only local interactions and sparse fitness evaluation, the assumption is made that per generation only a small subset of circumstances is encountered. Pagie and Hogeweg(1997) already showed that the information integration in the evolutionary timescale does work. A spatial coevolutionary system is able to evolve individual based diversity on the evolutionary timescale. However, dynamics on the ecological timescale are mainly caused by a population based diversity. Therefore the last research question is:

What is the role of population based diversity for the individual information integration?

1.4 CKI Relevance

This thesis is written as a conclusion for Cognitive Artificial Intelligence (in Dutch: Cognitieve Kunstmatige intelligentie, CKI). CKI is a combination of philosophical, psychological, linguistic and computational views on artificial intelligence. A quote which I met in a third year course, reflects in a good way my view on how to achieve this *Artificial Intelligence*:

Nature is far more ingenious than we are. The point is, evolution has already done it, so why not learn how that stupendous machine, our brain, actually works? [Churchland and Sejnowski, 1990]

In order to achieve AI we have to look at nature how it has to be done, so biology is probably the most appealing starting point for developing self-learning systems. The shape of (human) intelligence in present time has originated from millions of years of evolution. In my opinion, it is therefore presumptuous to think that we can design intelligence by ourselves. Our thinking is too much restricted to design a process which we do not yet fully understand. History has proven that evolutionary dynamics do not have such restrictions.

1.5 GECCO '07

Preliminary to writing this thesis I got the chance of submitting a late breaking paper for GECCO '07, a conference on evolutionary computation. Writing an article before the actual thesis may appear to be a strange order. In terms of optimal information integration it probably even is in principle the wrong order. However the article had to be submitted in the experimental stage of my research, and therefore I nevertheless adopted this order, thinking it would be a nice way to summarize part of my research in a scientific way. In consequence I chose to write my thesis with this article as a basis. Our article describes the implications of spatial modelling for the field of co-evolutionary function approximation with a biological character. The information density as well as the assumed preknowledge in the article is high and maybe hard to understand without the necessary knowledge in the specific field of evolutionary algorithms and genetic programming. Therefore the composition of my thesis is constructed in such a way that the rest of the thesis is written as a complement for the article "The Role of speciation in Spatial Coevolutionary Function Approximation" [De Boer and Hogeweg, 2007] written for GECCO '07. The content of this article is roughly the same as the chapters 4 and 5 of this thesis.

Chapter 2

Evolutionary Computation

2.1 Inspiration from Biology

Evolutionary Computation is the collective name for a range of computational techniques, inspired by the biological mechanisms of natural evolution. Darwin's concept of natural selection [Darwin, 1859] and his view on evolution are the main elements used in creating an artificial evolutionary process used for a variety of problems, ranging from practical applications in industry to a broad area of scientific research.

The vast majority of the history of life can be accounted for by physical processes operating on and within populations and species [Hoffman, 1989]. These physical processes are: reproduction, mutation, competition and selection. These processes are strongly intertwined and difficult to consider separately. It is the combination of these simple mechanisms what makes evolution such a strong process. This combination results again in *adaptation* and *speciation*, probably the most important notions of natural evolution.

2.1.1 Adaptation and Speciation

Adaptation in the biological sense, is a physiological process, a physical structure or a behavioral trait evolved by an organism as a result of natural selection such that it increases the expected long-term reproductive success of the organism. In this way adaptations enable living organisms to cope with environmental stresses and pressures.

Speciation is the evolutionary process by which new biological species arise. There are generally three modes of natural speciation considered, based on

the extent to which speciating populations are geographically isolated from one another: allopatric, parapatric and sympatric (see figure 2.1).

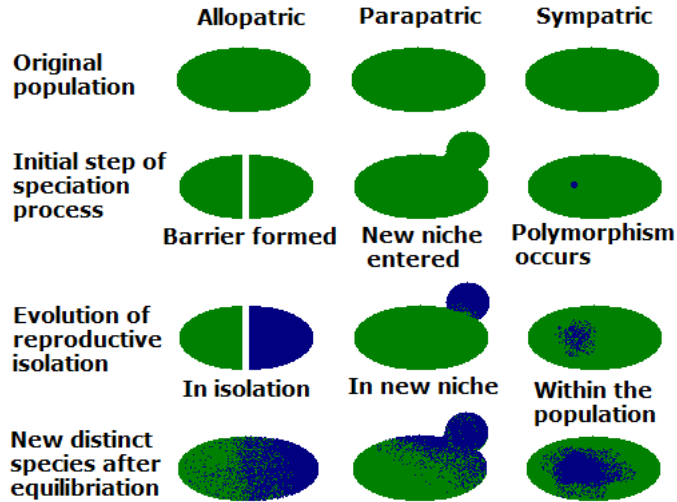


Figure 2.1: Different modes of speciation. Illustration by Dr. Dana Krempels 2006 ¹

Probably all these forms of speciation have taken place during the course of natural evolution, though it still remains a subject of debate as to the relative importance of each mechanism in driving biodiversity. The emphasis in the process of speciation is on biological differences between species. We speak of a species when a group of actually or potentially interbreeding individuals are reproductively isolated from other such groups. Thus biological speciation requires that the two species are unable to produce viable offspring together or that they simply avoid mating with members of the other group. However, regarding cloning or asexual reproduction reproductive isolation is undefined. In computational experiments the same problem arises. But for example in the case of bacteria, which reproduce mainly without obligate sexual reproduction, we can truly speak of different bacteria species. Because our model uses cloning for reproduction, the observed speciation in our model can be considered as an asexual form of sympatric speciation. Sympatric speciation refers to the formation of two or more new species from a single ancestral species, all present at the same geographic location, not separated by an environmental barrier.

¹This Illustration is licensed under the Creative Commons Attribution ShareAlike license versions.

The interplay of adaptation and speciation originates in the dynamical relation between mutation and selection. Where mutation (and/or recombination) provides a wider range of genetic variety in a population, natural selection reduces this variety again, selecting for useful traits, resulting in an adaptation to the environment. Sympatric speciation occurs as a consequence of adaptation. When individuals adapt to different circumstances, the specialization will cause new species to originate, bringing variety in the system. Allopatric speciation on the contrary, is considered to be caused mainly by genetic drift due to the isolation of subpopulations after an environmental separation, resulting in different adaptations.

The obtained variety of species has in its turn a positive effect through the change it brings in the evolutionary pressure. Each species has to maintain its niche by adapting even more in order to survive. Competition is now not only between individuals in a population, there is also evolutionary pressure through interactions between different populations with different evolutionary directions. However, the specific feedback of speciation on the individual adaptation is still unknown. In our experiments we will try to clarify how speciation of a population affects the adaptation of individuals.

2.1.2 Fitness

Fitness is a central concept in the process of natural (Darwinian) selection. Modern evolutionary theory defines fitness in terms of reproduction rather than survival. If an individual lives half as long as others of its species, but has twice as many surviving offspring, its genes are represented far more common in the population of the next generation. This so called *differential reproduction* expresses the difference in the rates of reproduction of differently adapted individuals. Therefore in biology, the fitness of an individual with a certain genotype, is generally considered an estimation of the capability of this individual to reproduce. Selection ensures that genotypic traits which provide a higher fitness will become more common in a population. So, if mutational differences in the individual genotype are said to affect the fitness, then it will directly affect the chance of reproduction.

Natural selection acts mostly on the differences in physical characteristics of individuals, the so called phenotypic differences. However, the genetic representation of these characteristics, the genotype, will also have a better chance for reproduction when associated with a favorable phenotype. Secondly fitness depends also on the degree of adaptation to the environment in which the individuals live. Thus an important observation is that the fitness of different individuals with the same genes, is not necessarily equal. For ex-

ample, residing in different ecological niches causes differences in the rate of reproduction (different fitness). Within the niche the (sub)species each increase their fitness again by adaptation, reinforcing the process of speciation. In short, fitness is a quantitative measurement of the success of a species, concerning all different important aspects of evolution and thus plays an important role in explaining the emergence of new species.

In evolutionary computation the concept of fitness is mostly implemented as a fitness function, where we can distinguish roughly 2 different classes: one where the fitness function does not change, as in optimizing a fixed function or testing on a static set of test cases; and one where the fitness function is variable, as in coevolving the set of problems as well.

2.1.3 Fitness Landscapes

To get an intuitive idea of the interplay between fitness (functions) and evolutionary dynamics, the relationship between genotypes (or phenotypes) and their according fitness can be visualized in a *fitness landscape*. This idea of studying evolution by using a landscape metaphor for the distribution of fitness, was first introduced in 1932 by Sewall Wright [Wright, 1932].

In a fitness landscape, or adaptive landscape, the fitness defines the height of the landscape. Genotypes which are very similar are said to be close to each other, while those that are very different are far away from each other. The two concepts of *height* and *distance* are sufficient to form the concept of a landscape. The set of all possible genotypes, their degree of similarity, and their related fitness values is then called a fitness landscape. When a genotype mutates, it moves through the landscape. By performing subsequent rounds of selection, growth, and mutation to a population of genotypes, the genotype space is searched for genotypes with high fitness values. This can get misleading when considering multiple mutation operators, because distance is then hard to define. For instance, when using crossover mutation, closeness, or the distance between different genotypes loses its meaning. The landscape metaphor also breaks down without a fixed fitness criterion. When using a dynamic fitness criterion instead, like in coevolution, the fitness landscape keeps changing during the course of evolution.

Fitness landscapes are often conceived of as a mountain range (figure 2.2). The highest peak represents the genotype of a global optimum. Local peaks are points in the landscape where all directions lead downhill resulting in a lower fitness. There can also exist valleys, regions from which most




Figure 2.2: Sketch of a 2D fitness landscape. The arrows indicate the preferred flow of a population on the landscape, and the points A, B, and C are local optima. The red ball indicates a population that moves from a very low fitness value to the top of a peak. Illustration by C.O. Wilke, 2001.

paths lead uphill to a higher fitness. A fitness landscape with many local peaks surrounded by deep valleys is called *rugged*. If genotypes in the same region have the same fitness, on the other hand, a fitness landscape is said to be *flat*.

However, a fitness landscape is generally multidimensional and therefore difficult to visualize. Reducing the landscape to a 3D-space can therefore be deceptive. But the key characteristics of a landscape can help in explaining evolutionary dynamics. A rugged landscape is regarded as a difficult landscape for most optimization algorithms, the chance of ending with a suboptimal solution is large and the course for an algorithm to keep improving its fitness is not clear. On the other hand, when alterations of the genotype keep you on a flat plateau, evolution is considered to be *neutral*, having no direct influence on the fitness of an individual.

2.1.4 Survival of the Fittest

The British economist Herbert Spencer was the first to use the popular term *survival of the fittest*, often used when referring to the process of *natural selection*.

”This survival of the fittest, which I have here sought to express in mechanical terms, is that which Mr. Darwin has called ’natural selection’, or the preservation of favoured races in the struggle for life.” [Spencer, 1864]

"Survival of the fittest" was introduced to clarify the principles of natural selection and the *struggle for life*. However this is a somewhat misleading description for the moving forces of evolution. It can be interpreted that only the individuals of a population that are *fittest* (eg. best in surviving) survive to generate offspring, enabling their genetic content to be passed on into the next generation. Because of this reasoning it is sometimes claimed that the phrase is a tautology. However when the full picture is considered, no tautology exists: the fittest individuals are the ones that fit the best in their environment, the best adapted. So, the complete mechanism leading from genetic differences, through differential reproductive success ("survival of the fittest"), to the actual evolutionary adaptation is a valid, informative reasoning, already confirmed by the actual existence of genetic variations. Individuals in a population surviving and generating offspring are by definition *the fittest*, which indeed is stated. However, *"survival of the fittest"* stays ambiguous and maybe confusing. Therefore it is not generally used by biologists, who prefer to just use *natural selection*.

2.2 Artificial Evolution

2.2.1 Computational Application of Evolution

In most scientific research, evolutionary computation is not meant to model evolution best, rather evolution is used as a metaphor to solve or optimize all sorts of problems. Although the basic evolutionary principles are easy to understand, much is unclear about evolutionary mechanisms and the specific role of selective forces by which natural species evolve. The controlled and well-understood conditions from evolutionary models can possibly help to bring insight into the details of how life and intelligence evolved in the natural world.

The three main fields of research in simulated evolution are genetic algorithms, evolution strategies and evolutionary programming. Each method emphasizes a different facet of natural evolution. Genetic algorithms stress chromosomal operators. Evolution strategies emphasize behavioral changes at the level of the individual. Evolutionary programming stresses behavioral change at the level of the species [Fogel, 1994]. Our methods cannot be classified solely under one of these methods, but uses techniques mostly associated with genetic algorithms and evolutionary (or genetic) programming(see section 2.3).

2.2.2 Problem Solving and Optimization

As mentioned, the phrase *survival of the fittest* already suggests that natural evolution can be seen as merely a process of optimization. The variation within a population in combination with selection pressure, leads to an 'improved' fitness in the subsequent populations; the individuals in a population are being 'optimized'. This notion of optimizing gave rise to a whole new field of research: the application of artificial evolution to the field of computational methods [Holland, 1975]. Evolutionary computation utilizes the biological idea that the level of adaptation determines the chance for reproduction of an individual. In artificial evolution only those individuals in a population that are adapted best, are kept in the population and given a chance to reproduce. The rest will be removed, leaving the population with an increased average level of adaptation. Optimization does not imply perfection, yet evolutionary computation can in this way discover highly precise functional solutions, with only the need to define the particular problems in an environment.

2.2.3 Direction

Because natural evolution is an apparently undirected mechanism, selection pressure originates solely from the interactions with other individuals and environmental influences. These influences and interactions change all the time, thus there is no predefined global evolutionary target. Individuals live up only to the objectives which are important for them: *surviving* and *passing on genetic material*. The course of evolution depends on the direction caused by pursuing those aims in life. By defining a fitness criterion one can force a system to evolve in a pre-determined direction, enabling evolutionary computation methods to solve or optimize specific problems. The use of natural selection already eliminates one of the greatest difficulties of problem solving: specifying in advance all the features of a problem and the actions a program should take to deal with them, is no longer necessary. Using evolution, it is possible to find solutions that solve problems even when no person can fully understand the structure of the solution. No knowledge of the problem has to be incorporated into the algorithm; the boundaries of problemspace and an evolutionary direction are enough to solve problems without full knowledge of the underlying mechanisms.

2.2.4 General Scheme of Evolutionary Computation

Usually, the initial population is formed by randomly generated candidate solutions. Koza compares this initial population to an army of parachutists dropping onto the landscape of a problem's search space, with each one being given orders to find the highest peak [Koza, 2003]. Small mutations enable each individual to explore its immediate neighborhood, while selection enforces progress, guiding the subsequent generations uphill to more promising parts of the solution space.

The solution for a problem or the performance to be optimized can act as the global evolutionary target. The "environmental pressure" forces evolution in the desired direction towards a better solution. Each generation all individuals are evaluated with a fitness function on how well they are 'adapted'. When fitness is assigned, the individuals are stochastically selected to provide the genetic material, forming the individuals in the next generation. Selection is stochastic in the sense that a higher fitness of an individual will bring a greater chance for this individual to reproduce. In this way individuals with a very low fitness still have a (small) possibility of reproducing in order to increase the exploration rate(see next subsection). Then variation operators are random applied on the genetic material of the selected individuals. Usually the variation operators are *mutation*, randomly modifications of the genotype, and *recombination* or *crossover*, whereby existing genetic material is exchanged, between or within a genotype. After applying these variation operators, the population is replaced and competition between individuals begins all over again. It differs per method if a generation as a whole is replaced with a total new generation or if only certain (outcompeted) individuals are replaced, giving subsequent generations the possibility to compete with each other.

<p>Begin INITIALIZE <i>first generation</i> Repeat until Termination condition is met EVALUATE <i>all individuals</i> (assign fitness) SELECT <i>individuals</i> APPLY <i>variation operators</i> (recombination and mutation) REPLACE <i>population</i> end</p>

Table 2.1: general scheme of evolutionary computation

2.2.5 Exploration vs. Exploitation

One of the appealing qualities of using artificial evolution for computational purposes is said to be the good tradeoff between *exploration* and *exploitation* [Eiben and Smith, 2003]. These concepts are best understood with the fitness landscape metaphore in mind. Exploration means that in every generation new individuals are localised in new untested regions of the search space. Exploitation means that individuals concentrate in the vicinity of already known good solutions. The trade-off between these concepts is important because too much exploration leads to an inefficient search and too much exploitation leads to premature convergence, possibly trapping the population on a local peak, representing a suboptimal solution. Evolutionary computational methods are able to prevent premature convergence because they can explore a far greater range of potential solutions to a problem than do conventional methods. This exploration of the landscape is done by random mutations providing a wider range of genetic variety in a population. Selecting for useful characteristics, stochastic selection reduces the variety again, resulting in the exploitation of potential good candidate solutions with a fitness advantage. Changing the ratio between selection pressure and random mutations, enables the adjustment of the trade-off between exploration and exploitation. Whereas classical algorithms like hill climbing often get stuck on a local optima, evolutionary algorithms can in this way be effective in escaping local optima and discovering the global optimum in even a very rugged and complex fitness landscape.

2.3 Genetic Algorithms and Genetic Programming

Just as a child creates magnificent fortresses through the arrangement of simple blocks of wood [building blocks], so does a Genetic Algorithm seek near optimal performance through the juxtaposition of short, low-order, high-performance schemata, or building blocks [Goldberg, 1989].

Although Genetic algorithms, evolutionary strategies and genetic programming are all different methods and used for various objectives, they essentially are all different variants of population-based generate-and-evaluate algorithms. Therefore they share more similarities than differences.

Genetic Algorithms (GA) originated mainly from the work of John Holland [Holland, 1975, Holland, 1992] and later David Goldberg [Goldberg, 1989].

GA's are said to be the most obvious mapping of natural evolutionary processes into a computer system. Genetic algorithms involve encoded strings (*chromosomes*) that represent particular candidate solutions to a problem. The evolutionary emphasis is on the role of recombination or crossover, trying to create a robust, adaptive system. GA's try to evolve a solution by the rearrangement of building blocks. The chromosomes are evaluated and the best strings are mixed to form a new generation.

The recent unraveling of whole genomes, shows that duplication is probably one of the most important mutational processes in nature. Despite genes have a meaning and some fitness consequences in GAs, this duplication of genes is not possible in a traditional GA. Therefore they are still quite constrained in their evolutionary potential and it seemed quite desirable to have a more flexible coding to create more evolutionary freedom. For this Koza developed genetic programming(GP) to indicate the application of GAs to the evolution of computer programs [Koza, 1992]. Where GA's can be seen merely as an evolutionary search for the best parameters, genetic programming strives towards a system which is implemented without too much predefined meaning, but with an interesting *genotype-phenotype mapping*. Instead of encoding a representation of a solution(just being a chromosome), GP incorporates genomes as computer code and allows mutational operators to expand and contract this code, thus leading to a great deal of flexibility. These computer programs are traditionally represented as tree structures (see figure 2.3).




Figure 2.3: Tree representation of $x - 1 - x^3$. LISP-representation would be `(- (- x 1) (* x (* x x)))`

Trees can be easily evaluated in a recursive manner. Every tree node has an operator function and every terminal node has an operand, making

mathematical expressions easy to evolve and evaluate. Therefore traditionally GP favors the use of programming languages that naturally embody tree structures; functional programming languages. Koza used the LISP programming language for this. LISP has the generality and power of a general program language, however there is no a priori distinction between the program itself and the data the program is running on. That brings a requirement: each list should have a meaning, so dividing by zero needs to be defined in order to get closure of the system. The primary (LISP) coding structure is expanded and contracted by adding genetic operators which can change and make modifications in the tree: point mutations, branch swapping, branch copying and deletion (see figure 2.4). Genetic operators act directly on the tree itself and, although at first sight this might appear advantageous, it also limits this technique because the genetic operators must be very carefully applied so that only valid structures are formed.




Figure 2.4: In crossover, the most used operator in GP, selected branches are exchanged between two parent trees to create offspring.

GP also uses populations of individuals, selects them according to fitness, and introduces genetic variation using genetic operators, just like GA's. Thus, the fundamental difference between GP and GAs lies in the nature and behavior of the individuals and, consequently, in the way they are reproduced and mutated to allow adaptation. In contrast to GAs, GP trees may assume different sizes and shapes in the process of evolution. As such, populations evolve and may discover solutions of greater complexity. The use of such a flexible coding instead of a predefined meaning, GP enables the altering of the fitness landscape (see subsection 2.1.3) by the genotype itself. The choice and structuring of the genes that form the genotype, determines the fitness landscape that has to be searched, and defines thereby the ease of exploring.

Chapter 3

Eco-evolutionary Dynamics

3.1 Evolutionary and Ecological Dynamics

Evolutionary computation is often based solely on evolutionary dynamics; ecological processes are excluded or neglected. How an evolutionary process acts relative to a fitness criterion can be evaluated because most methods use an artificial predefined fitness function. In this way evolution is forced in the desired direction (see subsection 2.2.3). However, this can also be considered as a constraint for the evolutionary process; the possibilities of the system are restricted and the 'freedom' of evolution is limited for computational purposes. Despite the focus is mainly on just solving the predefined problem, problems are still allowed to be solved in many different ways. Thus in the course of evolution there will be several partial solutions, coexisting in the population. This variety of different possible solutions can be considered as some sort of ecosystem diversity. As in nature, an important mechanism in maintaining this 'ecosystem diversity' can be the evolutionary process of continuous adaptation through mutations, which indeed is incorporated and evaluated in the artificial setting. However, the role of the dynamics within the ecosystem itself are excluded. And the various dynamics within a ecosystem influence not only the diversity and individual interactions, but also the stability of the system as a whole. Therefore it is difficult to see both processes separated, so maybe we should take both evolutionary and ecological dynamics in consideration.

The most striking difference between ecological and evolutionary models, is the total separation of timescales. Ecological models, mostly used for population dynamics, assume that evolution is slow enough that it will not influence ecological processes, and evolutionary models mostly assume that ecol-

ogy is fast enough for a population to be in an equilibrium or a Quasi Steady State. However, ecological processes are not necessarily fast and evolutionary processes are not by definition slow [van der Laan and Hogeweg, 1995]. The necessary interlocking of the processes occurring on ecological and evolutionary timescales, i.e. the interaction between ecology and evolution, can't be neglected.

Therefore in this chapter we will focus not only on the evolutionary dynamics within and between populations; also the important dynamics within the ecosystem formed by individuals of different (sub)populations are considered. Despite the fixed population size in our experiments make the 'ecological' dynamics somewhat misleading in a traditional sense, the dynamics between the subpopulations within our fixed population can be regarded as ecological dynamics. In short, we will discuss the eco-evolutionary dynamics which are accountable for the behaviour of our coevolutionary computational model.

3.2 Population Based vs. Individual Based Diversity

As discussed in subsection 2.1.1 the main concepts in evolution are *adaptation* and *speciation*. In an evolving population, individuals can adapt by regulation to all encountered circumstances or the population can split up into subpopulations, each specializing to a subset of the possible circumstances. This is called *individual based diversity* and *population based diversity*, respectively. Pagie and Hogeweg(2000a) studied under what conditions individual- and population based diversity evolve, using a model based on the restriction-modification(RM) system in bacteria.

Bacteria have RM-systems to protect themselves from foreign DNA, such as bacteriophages. A bacteria can have different plasmids coding for the RM-systems, however having a RM-system affects the growth of the bacteria in a negative way, restricting its use. Phages on the other hand can become immune to specific RM-systems, causing the defense system to fail. In their model, Pagie and Hogeweg(2000a) coevolved a population of bacteria, with a fixed number of possible RM-systems, and viruses which can infect the bacteria. The results of this model show that the system, more or less arbitrarily, goes to one of two attractors:

- An individual based diversity where every bacteria has as many plasmids as possible for maximum individual defense. In this system in-

vading virus populations can survive easily besides the bacteria population.

- A population based diversity where each individual has only one plasmid and therefore only a weak defense against the different virus populations. The different bacterial plasmids are spatially spread over different regions in the population. Population dynamics in this mode result in only a few viable viruses populations, outcompeted by the bacteria population.

Under some circumstances the system can switch from individual based to population based diversity in order to gain stability on the level of the population. This population based stability is due to the formation of a static pattern of subpopulations, each specialized with a different plasmid. The result is an invasion resistant, stable population, consisting of different subpopulations. In spite of the difference in how the information is integrated, both modes cause a viable stable population with almost the same amount of information integrated, i.e. almost the same amount of plasmids is present in both modes. In the one mode this information is divided over different individuals while in the other mode all plasmids are present in every individual.

In our study we look at the role of a population based diversity in the process of individual adaptations in a population. Can population dynamics lead to information integration on the individual level?

3.3 Sparse Fitness Evaluation

3.3.1 Biological Ground

The standard biological dogma on fitness is that it should bring immediate benefits. However there are several examples of evolving species which can cope with all sorts of problems, even if they haven't seen those problems every generation. The selection for those traits has to be on the longer, evolutionary timescale. The ability to cope with situations seen only once in a long time tends to negate the idea of immediate benefits. Two examples:

- Bacteria often use a flagellum to move. A flagellum is a long slender projection from the cell body which can propel the cell by beating in a whip-like motion. Bacteria spend several generations to develop such a flagellum. There is no immediate fitness advantage in doing this.

- *E. coli* is one of the main species of bacteria living in the lower intestines of mammals. *E. coli* bacteria spend many generations in the gut, however they can also live for many generations in an external environment outside of the gut. The regulation for survival in both environments remains through all generations. The long term benefits are preferred above the immediate benefits.

These two examples show that selection over the evolutionary time scale with no selection pressure for immediate benefits, cannot be neglected. How can sparse information, encountered in the short timescale, lead to integration of information on the long time scale?

For what we know the first organisms in early evolution had a pretty short life span. Therefore they probably encountered only a limited amount of circumstances to cope with per generation. This is called *sparse fitness evaluation*. Later in time the life span of organisms expanded and the number of circumstances encountered within a lifetime also increased. However the environmental variation on the longer time scale, which whole ecosystems have to handle, always persisted.

3.3.2 Geological Grounding

In geology the so called *Snowball Earth* hypothesis as it currently stands, proposes that the Earth was entirely covered by ice for several periods in earth history. The cold climate periods were cyclical, punctuated with well defined warm-humid interglacial periods [Rieu et al., 2007]. In spite of these harsh climate changes, fossil records reveal little change during the glacial interval and do not show the expected eco-evolutionary fluctuations [Corsetti et al., 2006], i.e. evolution allows ecosystems to handle not only seasonal variation, ecosystems are also robust against (extreme) climate variation.

3.3.3 Sparse Fitness in an Experimental Setup

In our experiments we make the assumption that per generation only a small subset of the whole problemspace is encountered. A host 'sees' only the problems present in his spatial neighborhood. By embedding the model in space(see section 3.7), sparse fitness evaluation is induced. Pagie and Hogeweg(1997) already studied the role of such sparse fitness evaluation on information integration in the evolutionary timescale. This was done in the same coevolutionary setting we use for our experiments. They compared sparse fitness evaluation with complete fitness evaluation. In the setting

with complete fitness evaluation each individual sees all the possible problems in problemspace every generation. In the setting with sparse fitness evaluation each individual sees only a few of the possible problems per generation. The fitness of each individual is then based on how well it solves the problems seen that generation.

After evolving the system, the sparse evaluated setting produces individuals with a simpler genotype, better coping with all circumstances. The evolved genotypes were better generalizable to circumstances never seen before and as most important result the success rate on finding an appropriate solution for the whole problem space was found with a much higher success rate than in the complete evaluation scheme. One important evolutionary mechanism in reaching these results, is a lower mutational robustness in the sparse fitness setting. When using a fixed target, evolution will select for robust (neutral) solutions, while in coevolutionary sparse setting, there is selection for solutions in which mutations have extra effect [Pagie and Hogeweg, 1997]. Selection of the system for this lower mutational robustness causes greater mutational effects when the phenotype is not yet generalizable. This allows fast and easy adaptation to new problems, forcing the final adaptation to the complete set of possible problems to be highly generalizable and more compact in comparison with the complete evaluation case. When all the possible problems are encountered every generation (a static evolutionary target), the system will try to adapt to all the circumstances simultaneously, trying to code for a very rugged fitness landscape. This brings a great risk for evolution to get stuck on the local (sub)optima. The pitfall is that evolution in this setting does not really 'adapt' to the circumstances, it tries only to extend to other local optima by adding more genetic material. This causes great mutational robustness with almost no generalizability. As a result the evolved genotypes are very complex and overfitted to the local optima. Adaptation to new problems can only be reached by adding even more genetic material, resulting in a population of overspecialized individuals.

3.4 Coevolution

3.4.1 Mutual Selective Pressure

In biology, *coevolution* is the mutual evolutionary influence between two species. Each party in a coevolutionary relationship exerts selective pressure on the other, thereby affecting each other's evolution. The exerting

of selective pressure on each other, forces both of the coevolving parties to keep evolving at higher rates than they would do solitary. This intrinsic property of influencing each other's evolution is also the base for the success in applying coevolution in the field of evolutionary computation.

When species coevolve, their respective levels of fitness depend not only on their own genetic structure and their level of adaptation, but on the development of another species as well. The gene pool of one species creates an extra selection pressure on the other species. Mutual evolutionary influence between populations does of course not necessarily imply mutual dependence. The host of a parasite, or the prey of a predator, does clearly not depend on its enemy for survival, however they truly influence each other's evolutionary direction. This exerted pressure leading to coevolutionary dynamics can either be described as positive or negative pressure.

Certain plants, which have evolved valuable food sources in order to attract insects are an example of exerting *positive coevolutionary pressure*. The attracted insects, such as bees, provide a good means to distribute pollen for the plant, far more efficient than the distribution of pollen by wind. The coevolutionary process has led to deep flowers which fit the tongue-length of specific species of bees. The flowers are sturdy and irregularly-shaped and can even have a specifically-designed landing platform. For example, snapdragons will only open for a bee of the right weight. The bees on the other hand have nectar as their main food source. Bees have evolved a long proboscis (a complex sort of "tongue") that enables them to obtain the nectar from flowers and they are totally specialized for feeding on nectar and pollen.

The danger of this positive kind of coevolution is that both species could evolve into a position of total dependency through increased specialization, increasing the risk of extinction if either species declines or 'cheats' in the coevolutionary relation by changing the mutualistic role into a parasitic one. In some cases mutual coevolution is quite more specific, for example between two cellular functions. Current day mitochondria and chloroplasts were once free-living unicellular individuals. The mitochondrial genomes coevolved with some of the nuclear genes that now function within the mitochondria. During this coevolutionary relationship, these cells entered the cytoplasm of other cells and melted together, with the combination of genes now regarded as eukaryote mitochondria as a result. This *endosymbiotic* transition from multiple cells, each representing a different function, to one cell with the combined functions can be considered as a case of population based diversity, which has led to individual based diversity (see section 3.2).

Strange as it sounds, *negative coevolutionary pressure* can be considered more constructive. In order to survive, prey species will continuously have to adapt defensive or evasive systems to avoid predation. Predators in their turn, respond by developing mechanisms to evade these defences. Both species are evolving and luckily in nature neither of the strategies will ever be perfect. Due to variation in the population some prey are always more vulnerable, and some predators in the population will be less efficient. The effect of selection pressure within the populations will therefore be high, forcing a higher rate of selection for preferable traits and thereby faster adaptation for both species. The likelihood of extinction in the negative case is limited, because this kind of evolution mostly does not lead to a total dependency on each other.

3.4.2 A Coevolutionary Relationship

Several factors influence the likelihood and strength of a coevolutionary relationship. Of course it is necessary to have a certain level of co-occurrence in order to coevolve; species have to reside at least partly in the same environment to interact. The type of the relationship between the two species before we can speak of coevolution, is also important. Will coevolution bring mutualistic advantage for both species, strengthen the existing predator-prey dynamics, or will it only push on the competition? Species that have intimate relationships, such as a specialistic predator and its prey or a host-specific parasite, already interact more actively and thus are more likely to influence each other's selection pressure. Species that do not directly encounter each other but interact through competition for resources, are less likely to coevolve, but if the strength of the competition is high enough, an indirect form of coevolution can occur.

Most species experience natural selection from many environmental aspects (see subsection 2.1.2). It seems therefore unlikely that in nature one species would act as the sole or even the primary selection pressure on another. Nonetheless, there are good examples of tightly coevolved relationships where two populations have a highly specialized interaction. In these cases, the selective advantages gained by responding to only one source of selection outweighs other environmental factors.

3.5 Arms Races

In a coevolutionary relationship, an adaptation in one species changes the selection pressure on the other species, giving rise to a counter-adaptation.

If this keeps occurring reciprocally, these refined mutual counter-adaptations may result in an escalation of adaptations from both species, called an '*arms race*' [Dawkins and Krebs, 1979]. An arms race is mostly considered to be a negative unstable runaway of the driving force which is generally responsible for the improvement of species by subsequent adaptations. A nice example of a (chemical) arms race is the coevolutionary case of the rough-skinned newt and the common garter snake. The rough-skinned newts produce a potent nerve toxin that concentrates in their skin, meant for deterring predators. Garter snakes have evolved resistance to this toxin through a set of genetic mutations, and prey upon the newts where no other animal is capable of preying on this newt without adverse effects. The relationship between the rough-skinned newt and the common garter snake has resulted in an evolutionary arms race that has driven toxin levels in the newt to extreme levels.

Dawkins(1979) described the occurring dynamics on the different timescales of this phenomenon, in terms of a metaphor. He compares a fox chasing a rabbit with a particular submarine and the ship it is trying to sink. This race of individuals is on the behavioural time scale. But on the evolutionary time scale another kind of race takes place. Submarine designers learn from earlier failures. As technology progresses, later submarines are better equipped to detect and sink ships, and later-designed ships are better equipped to resist. This '*arms race*' takes place in the longer, historical time scale and leads to improvement of the design of both submarines and ships. Similarly, over the evolutionary time scale the fox population may evolve improved adaptations for catching rabbits, and the rabbit population improved adaptations for escaping.

3.5.1 Red Queen Evolution

Evolutionary arms races can end in several ways. The most drastic way is when one species drives the other to extinction. In case of overspecialization or advanced adaptation, this can mean the end for both species. Another possibility, exemplified by the still existing flower-bee coevolution, is that both sides reach a mutual local optimum. Lastly, arms races may have no stable end at all, when counter-adaptations keep cycling continuously without integrating any new information. This is also known as Red Queen Evolution [van Valen, 1973].

The fitness of individuals is highly dependent on the environmental influences (see subsection 2.1.2). So when two (coevolving) species keep adapting

to each other, the effective environmental selection pressure on the other changes repeatedly. When species are unable to maintain the integrated information in the evolutionary time, we speak of a Red Queen evolution: The populations must continuously adapt to maintain the same level of fitness. Preventing a Red Queen evolution is one of the major difficulties in coevolutionary computation.



Figure 3.1: Red Queen: *"In this place, it takes all the running you can do to stay in the same place."*

3.5.2 Life Dinner Principle

Between coevolving populations, Dawkins(1979) argues that as a predator develops offensive adaptations, the prey will have to counter with defensive adaptations. This results in a selection pressure that is unequal, because as Dawkins puts it "if the predator loses the race, he simply loses a meal. If the prey loses the race, he loses his life". This is called the *life-dinner principle*. The selection pressure on the prey to evolve new adaptations is higher. Thus in predator-prey coevolution there is a built-in imbalance between predator and prey with respect to the penalty of failure. "Mutations that make foxes lose races against rabbits might therefore survive in the fox gene pool longer than mutations that cause rabbits to lose races" [Dawkins and Krebs, 1979].

3.6 Coevolutionary Computation

Being inspired by biology, it was a logical next step for evolutionary computation to incorporate coevolutionary dynamics into computational methods. Nowadays, more and more coevolution is regarded as a very successful method to strongly improve evolutionary computation [Hillis, 1990, Juillé and Pollack, 1996, Pagie and Hogeweg, 1997], but it is still a great minority in the field of evolutionary methods. Coevolution is considered to be a very efficient method to guide the development of a population of problems in the best way, enabling the population of candidate solutions to solve the predefined problem. The exerting of selective pressure on each other (see subsection 3.4.1) forces both of the coevolving parties to keep evolving and keeps the problems optimally challenging for evolving candidate solutions at different stages of evolution until the evolutionary target is reached.

Coevolution is implemented in many ways, however roughly three main forms of coevolutionary models can be distinguished [Pagie and Hogeweg, 1997]:

- Host-parasite models in which candidate solutions are evaluated on the basis of small subsets of the total problemset which defines the evolutionary target. The subsets of evaluated problems coevolve with the population of candidate solutions [Hillis, 1990].
- Competitive evolutionary models have candidate solutions competing with each other in game-like tournaments. The fitness of the solutions depends on the ratio of wins and losses in these tournaments. The coevolutionary dynamics operate solely within one population, the evolutionary target can be either predefined or dynamically defined in terms of the behavior of the opponent [Angeline and Pollack, 1993, Sims, 1994, Juillé and Pollack, 1996, Rosin and Belew, 1997].
- Cooperative evolutionary models use individuals of several different (coevolving) species. Individuals from the different populations are combined before they are evaluated with respect to an evolutionary target [Husbands, 1994, Potter and De Jong, 1994].

In our experiments we use the same general setup as [Pagie and Hogeweg, 1997], which uses the first option, the host-parasite implementation, in a model for coevolutionary function approximation. Therefore the term 'coevolution' is used in our experiments for the coevolutionary dynamics between two populations of different 'species'.

3.6.1 Hosts and Parasites

In order to get coevolutionary dynamics, the evolutionary goals of individuals in the one population are opposed to those in the other population (fig 3.2). Because these are nevertheless dependent on each other, Hillis called them in describing his coevolutionary model, *hosts* and *parasites* [Hillis, 1990]. These terms referring to a population of *solvers* or *candidate solutions* and a population of *problems* respectively, can give rise to confusions because of the counter intuitive meaning of parasites and hosts, when compared to the biological perspective. The sometimes used notions *predator* and *prey* would maybe already be a more fitting description in the biological sense of our model. The most unambiguous in the sense of evolutionary computation still keeps to be *problems* and *solvers*. However, because we use the same general setup as [Pagie and Hogeweg, 1997], we also adopted the parasite-host terminology.




Figure 3.2: the evolutionary goal of the host-population is opposed to the parasite-population

3.6.2 The Ideal Trainer

The adaptation of a host population in order to solve certain problems can be considered a process of *learning*. Optimal learning is a complex notion. In evolutionary computation a system evolves by learning to cope with different circumstances or by solving different problems. Solving a problem, individuals try to improve their 'knowledge' on the problem domain. In this way the system gathers information which finally has to lead to a solution, which can account for all the problems.

As in most learning processes, both human and machine, the level of training difficulty is highly important. When the problems are too easy or too difficult, there is generally no progress whatsoever. The nice aspect of using coevolution for computational learning is that not only the learners evolve to become better, the problems also evolve, trying to pinpoint the weak spots of the learner. It is, however, not enough to just point out the weaknesses of the system, for improving its overall performance. In order to gain the most progression, the learning individuals need problems which fit their current 'knowledge' the best. Pagie and Hogeweg(2000) already pointed out that in the case of the majority-problem the best way to evolve good solutions is to give the system such a fitness-advantage that easy problems were first evaluated and only when those 'easy' problems were solved, the system proceeded, trying to solve the more difficult cases. In [Epstein, 1994, Juillé and Pollack, 1998, Pagie and Hogeweg, 2000b] and [De Jong and Pollack, 2004] is an ideal trainer described. They all describe such as a trainer which presents problems of increasing difficulty. This is often done with such a learning gradient or a domain specific fitness evaluation.

We will demonstrate in our experiments that *spatial pattern formation* can force the system dynamics in such a way that it can be considered as a general guide to ideal learning.

3.7 Spatial Coevolution

In the perspective of biology, a world without *space* is not possible. All organisms are constrained to the environment where they live. The interactions with other organisms and with their environment are purely local. A species cannot have a direct influence on the evolution of another species, located somewhere else in space.

In evolutionary computation, spatial embedded models weren't considered useful until the nineties of the last century. The calculation time of all the necessary parallel computations are a great problem for serial computers. Hillis produced the first massive parrallel computer as a huge Cellular Automata, a so called connection machine, enabling a different use of computation. Using his connection machine, he was in 1990 the first to use space in a coevolutionary optimization proces, confining the interactions of coevolving individuals to their local neighborhood. Hillis already noted that the coevolutionary process improved when he placed the two coevolving populations in a spatial environment, forcing the interactions to be local. In

his model on the evolution of sorting algorithms, this spatial induced sparse fitness evaluation gave rise to very good results [Hillis, 1990].

Pagie and Hogeweg also used a spatial environment in their function approximation model [Pagie and Hogeweg, 1997] and already noted that spatial embedded models had some important features which in a non-spatial system had to be explicitly declared by the use of different heuristics. More genetic diversity is naturally preserved in the population and the weak spots in the population of candidate solutions are targeted more effectively. Their results are repeated several times and the role of space is more and more accepted as an effective method to improve coevolutionary computation [Pagie and Mitchell, 2002, Williams and Mitchell, 2005, Mitchell et al., 2006].

How spatiality exactly contributes to the success of coevolutionary function approximation is still unknown. The formation of spatial patterns is probably the most important feature brought in by a spatial environment [Boerlijst, 1991, Savill et al., 1997], also described in [Pagie and Hogeweg, 1997, Johnson and Boerlijst, 2002]. Spatial patterns give the coevolutionary system a better chance to organize itself in a convenient way, forcing the necessary information for evolutionary improvement to be presented at the appropriate timescales and in the right subsequent order (see subsection 3.6.2). These self-organised spatial structures can form a new level of selection [Boerlijst, 1991] acting on (sub)populations, increasing competitive pressure on the individuals. The spatial patterns also effectively constrain the evolutionary dynamics in such a way that the information can be integrated without sliding down into a Red Queen evolution as discussed in subsection 3.5.1 or other problems associated with coevolution [Cartlidge and Bullock, 2004, De Jong and Pollack, 2004, Paredis, 1997].

In contrast to the belief that the persisting diversity and the targeting of weak spots in the population of hosts are the main reason for superior performance of coevolutionary computation [Williams and Mitchell, 2005], we will conclude that the success crucially depends on the selforganizing dynamics of spatial patterns.

3.7.1 Cellular Automata

In order to investigate the role of space we implemented our model in a two dimensional Cellular Automata (CA). A CA consists of a regular grid of cells, each one with a finite number of states. Time is discrete, and the state of a cell at time t is based on the values in its neighbourhood at time $t - 1$. When the rules are applied to the whole grid a new generation is

created. Behavior of a subclass of Cellular Automata is considered to be 'on the edge of chaos' in the sense that the only way to predict the state of the whole population at a certain time is by means of going through all the subsequent preceding states. A well-known example of a two dimensional CA is Conway's game of life. In this game local interaction in combination with only 3 simple rules already leads to 'complex' macro-scale behavior. In this way a variety of spatial patterns can emerge on such a grid as meso-scale patterns.

Our CA consists of a 50×50 2-D toroidal square lattice, with one host and one parasite per grid cell.

Chapter 4

Experimental Setup

4.1 The Model

For our experiments we use the same general setup as [Pagie and Hogeweg, 1997]. This setup uses coevolutionary optimization in order to approximate a certain predefined function. This predefined function acts as the global evolutionary target. The system consists of two coevolving populations. One population of candidate solutions, the *hosts*, coevolves with a population of problems, the *parasites*. Population size is kept the same. Each host has a mathematical equation as a genotype, which codes for a corresponding function, the phenotype. The problems consist of an (x,y) -value as a genotype, which code for a phenotypic unique value, comprised to the global evolutionary target. The coevolutionary dynamics arise because the 'evolutionary goals' of the parasites are opposed to the global evolutionary target of the hosts, but fitness is assigned depending on each other (see subsection 3.6.1). The hosts try to evolve the phenotypic form of the global evolutionary target and the parasites evolve within this global evolutionary target trying not to be solved.

A good biological metaphor is to regard the (x,y) -values of the parasites as circumstances in which the host resides. The degree of adaptation to a particular circumstance is defined as the host's given approximation of the unique solution based on the (x,y) -value. When a host can solve all possible problems, the host is considered to be adapted to all possible circumstances, having a phenotypic structure identical to the predefined function. This will give maximum fitness for all the possible problems; the global evolutionary target is reached.

In our experiments we will describe the role of spatial pattern formation, comparing it with a well mixed system. Both the systems are modeled in Cellular Automata. Our CA consists of a 2-D toroidal square lattice, with one host and one parasite per grid cell. The size of the lattice is 50×50 cells, defining the population size at 2500. Competition for growth is local in space. Each individual has to compete within their 3×3 neighborhood for reproduction. A selected individual will grow into the central cell of the nine cells under consideration. In our well mixed counterpart the whole population of hosts and parasites is mixed after each timestep in order to exclude spatial pattern formation. Note that this well mixed model differs only in spatial pattern formation. Competition is still local, evaluation is asymmetric and sparse.

4.1.1 Technical Aspects

We used the CASH2-library for implementing a spatial environment in our model. This is all done with the C++ programming language. Like in genetic programming the solvers consist of a tree like representation of functions, i.e LISP-functions. We used DrScheme for the functional programming part of our model. With Guile, a Scheme-interpreter for C++, we could use both the power of CASH2 in simulating CA's and a practical functional representation of the genetic functions. Despite of the problems which arise when using different programming languages and interfaces simultaneously in one model, the overall usage of the model and the possibilities to substract data were sufficient.

4.2 Global Evolutionary Targets

As global evolutionary target we evolved our system towards two different functions:

$$Target(x, y) = x^3 + y^3 + 5x^2 + x \quad (4.1)$$

$$Target(x, y) = \frac{1}{1 + x^{-4}} + \frac{1}{1 + y^{-4}} \quad (4.2)$$

Model #1 refers to our model with function 4.1 as a global target. This function has a landscape which is monotonously rising through the domain, slightly skewed into one dimension. Because of this particular landscape we

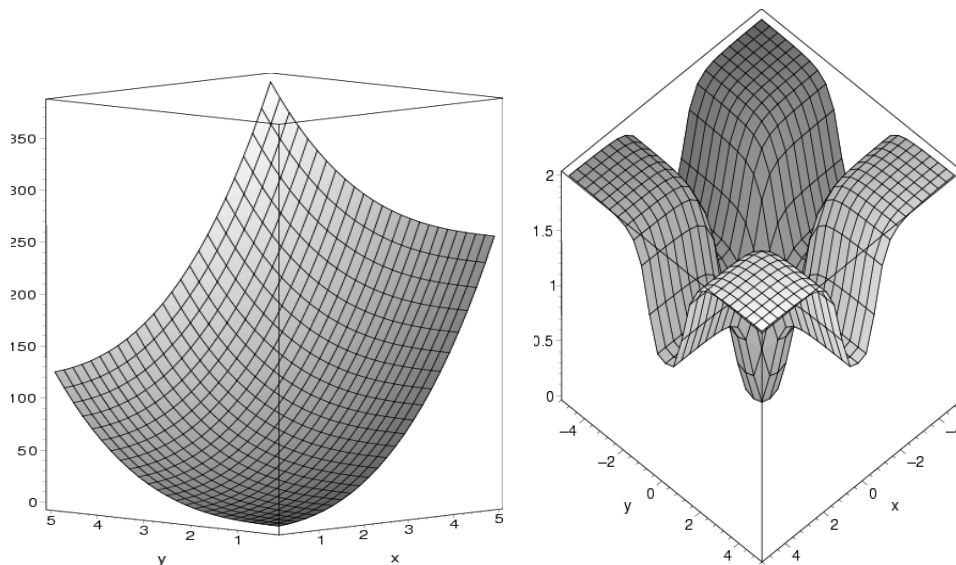


Figure 4.1: functionlandscape of functions 4.1 and 4.2

can better follow the evolutionary process and therefore study what role the appearance of a function has on the process.

Model #2 uses the same target function which [Pagie and Hogeweg, 1997] used, function 4.2. This function is considered to be a nice benchmark for coevolutionary function approximation and is thoroughly used investigating spatial coevolution. [Pagie and Mitchell, 2002, Williams and Mitchell, 2005, Mitchell et al., 2006]

4.3 Fitness Evaluation

The problems on which fitness of the host is evaluated, are based on X,Y-values. For both functions there are 26 x 26 problems regularly distributed over the corresponding domain. For function 4.1 the domain is $X = 0.2, 5.0$ and $Y = 0.2, 5.0$ with an interval of 0.2. For function 4.2 the distribution is in the domain $X = -5.0, 5.0$ and $Y = -5.0, 5.0$ with an interval of 0.4. The coevolving parasites can only adopt these values.

The fitness of a host is based on minimizing the distance between the function solution and the target of the nine parasites in its neighborhood. The parasite fitness is defined as maximizing the distance between target function and solution of just one host (further explained below). This asymmetry is found to give better results with respect to optimization than a

symmetric evaluation [Pagie and Hogeweg, 1997]. The host’s fitness is less affected by changes in a parasite, the parasite however can react directly to fluctuations in the host. These asymmetric dynamics can also be seen in nature as Dawkins life-dinner principle shows (see subsection 3.5.2). The built-in imbalance between host and parasite with respect to penalizing failure, is even strengthened in our model, using a worst case scenario for the parasites. Pagie and Hogeweg(1997) based the fitness of a parasite on the host in the same grid cell. In our model the fitness of parasites is based upon the best given approximation of a host in the 3×3 neighborhood. Because the evolutionary goals of both populations are opposed, this is the worst case scenario for the parasite. In this way the selection-pressure for parasites is increased even more.

Note also that the sparse fitness evaluation has in this way an important extra feature. Because each host has his own neighborhood with corresponding parasites, competition is between hosts which have a fitness assigned from a different subset of the problem space. This increases the effectiveness of sparse fitness evaluation by, in a way, comparing the different parasites on distinguishing capability.

The selection of hosts is proportional. Each fitness of all the hosts in a neighborhood are summed and a random value in between is generated, selecting the corresponding host for reproduction. With this method the chance of reproduction is increased proportionally with the fitness and the selection pressure can be conveniently adjusted by using an exponential factor. Selection is done simultaneously in order to prevent overlap in generations.

Parasites are ranked according to their fitness and the i th ranked parasite is selected¹ with probability $(\frac{1}{2})^i$.

The fitness for the host is defined as

$$Fitness(host) = e^{-n*d}$$

where distance d is defined as

$$d = \sum_{i=1}^9 \frac{|T(p_i) - h(p_i)|}{M}$$

and $M = 1$ for absolute and $M = T(p_i)$ for relative evaluation.

$T(p_i)$ is the target of parasite i subjected to the global evolutionary target.

¹to ensure $\sum_{i=1}^{i=max} probability = 1$ the two last ranked parasites use the penultimate rank.

$h(p_i)$ is the host's approximation of $T(p_i)$ based on the X,Y value of parasite i . The term n normalizes the distance. In all our experiments $n = \frac{3}{4}$.

To prevent the solutions of the hosts getting too large, we added a small size penalty. When the length of the solution passes 65 nodes (which is far above the necessary number of nodes in a correct solution) an exponential penalty is subtracted from the fitness. This penalty has a maximum of $\frac{1}{9}th$ of the maximal fitness. This is sufficient to highly improve calculation time and keep the solutions short. This penalty is biologically justified in terms of the error threshold. In nature there is also a 'penalty' on increasing a genome size, because the chance for a mutation is per position, not per genome. When increasing the genome size, the chance for a mutation somewhere in the genome also increases. In our model mutation is defined per genome instead of per position. In simulations without some sort of restriction on the genome size, it can be advantageous to evolve redundant code in order to reduce the change of getting a 'lethal' mutation in the important part of the genome. Without disadvantages an extreme long genome-part can be evolved and multiplied with zero. With our small size penalty, this phenomenon is excluded.

4.4 Genetic Representation

The genetic representation of the solutions is based on genetic programming and is essentially the same as proposed by [Pagie and Hogeweg, 1997]. The genotype of the host is a hierarchal list-representation(LISP) of a function tree. The function set consists of the operators $\{+, -, *, \%\}$ where we use the protected division operator $\%$, such that division by zero gives 1.0. The possible terminals are $\{x, y, \mathfrak{R}\}$, where \mathfrak{R} is a constant. This constant is defined at declaration as a random constant between -1.0 and 1.0. Note that almost in all cases the system prefers to make a constant by dividing one variable by itself. Constants are kept as a possibility to extend the possibilities of the system and maintain the 'freedom' of the evolutionary process. The genotype of the parasite represents one problem in problem space, specified by only one X,Y-value. Hosts selected after evaluation are subject to point mutations and crossover, with a 20% and 40% chance, respectively. In order to use an ancestor trace (explained in section 4.5), we use clonal reproduction with gross chromosomal rearrangements. This kind of internal crossover replaces a randomly chosen sub-tree in the selected host with a randomly chosen subtree of a copy of itself. In our experiments we observed that this internal crossover has the same success rate as external

crossover. Next to the fact that the fundamental mechanism is the same, deletion and duplication in combination with point mutation seem powerful enough for fully integrating the necessary information in an evolutionary run. In this way we can easily examine the information flow within the genomes of candidate solutions through evolution. Every time step also 10% of the parasite-population are mutated. This means that one of the variable values is changed into an adjacent value(.e.g plus or minus 0.4). The genotypic space is not toroidal. When a value of a parasite is on the border of the problem domain, it can only mutate in one direction. A solution is considered completely 'correct' if the sum of the absolute differences with the target in evaluating all 676 problems in the domain is less than 0.01. In practice all our evolved correct solutions exactly matched the evolutionary target, giving no difference in evaluation. A solution is only marked as 'correct' when the solution stays in the population for at least 50 time steps. In practice we observed a total domination of the solution, causing it to spread through the whole population in a couple of generations. Each run is started with small, randomly created functions for hosts of maximum depth 3. Because initialization with random values for the parasites gives considerably better results, this will be the case when not mentioned otherwise. When no solution is found within 1000 timesteps the simulation will be stopped, otherwise simulations are stopped after reaching the solution and staying in the population 50 time steps.

4.5 Observables

In order to observe the spatial pattern formation we output the different values present in each grid-cell of our CA with a color palette. We can plot the fitness of the parasite, the fitness of the host and the X,Y value of each parasite (divided in 2 separate plots) in space.

In our comparison between coevolution with and without pattern formation, we use an ancestor trace. All new individuals entering the population, originating by mutation, get an identifier and a list of its ancestors. In this way it is possible to trace back all differentiation and speciation leading to an individual in a time step. Every time step 60% (40% internal crossover and 20% point) of the whole population mutates, so constructing a tree with all the mutations in one time step gives a representative view on the mutational branches leading to all the hosts in that population.

Chapter 5

Experimental Results

5.1 Efficacy of the Optimization




Figure 5.1: Spatial patterns in problem space of model #1. Left for X, right for Y. Red denotes a high value, black denotes zero. Intermediate values are in between.

As in most spatial coevolutionary processes, pattern formation is an important feature [Boerlijst, 1991, Savill et al., 1997, Pagie and Hogeweg, 1997]. Figure 5.1 and 5.2 visualize the pattern formations from parasites in both our models, showing waves of parasites with alternating high X,Y values. This pattern formation is the driving force in the dynamics in both systems.

Considering the results for model #2, table 5.1 shows that the well-mixed system fails in almost all simulations in finding a correct solution. In




Figure 5.2: Spatial patterns in problem space of model #2. Left for X, right for Y. Red denotes a high value, black denotes zero. Intermediate values are in between.

our system with pattern formation however, 90% of the runs gives a correct solution, almost always within 300 time steps. This dramatic difference in performance can be totally attributed to the occurrence of pattern formation.

Method	Success rate
Spatial patterns	27/30(90%)
Well mixed	3/30(10%)

Table 5.1: Success rates of model #2

5.1.1 Variation

In contrast to expectation we see a higher parasite diversity in the well mixed model than in the spatial model (Shannon diversity is 15% higher at all levels of course graining). Also the distance between observed parasites at one location of the grid in subsequent time steps is higher in the mixed case. In the spatial model more than 50% are the same in subsequent time steps (see figure 5.3), which corresponds to the observed waves of similar parasites. The hosts observe even less variation because they travel along with the waves. Intuitively, one would expect this lack of information to harm the coevolutionary optimization. Our results show the opposite effect. The scale of the waves is a selforganizing property, which depends on the

rate of adaptation. In this way, in combination with the right fitness evaluation, the interconnecting subpopulations organize themselves into an 'ideal learner/trainer' system.




Figure 5.3: Distance between subsequent parasites at random position. Black denotes the spatial system, grey the well mixed.

5.2 Coevolutionary Dynamics Depend on Fitness Evaluation

As discussed in 3.6.2 an ideal trainer is described by many as a trainer which presents problems of increasing difficulty [Epstein, 1994, Juillé and Pollack, 1998, Pagie and Hogeweg, 2000b, De Jong and Pollack, 2004]. This is often done with a learning gradient or a domain specific fitness evaluation. Can we demonstrate that spatial pattern formation can force the system dynamics in such a way that it can be considered as a guide to ideal learning? In order to grasp the full dynamics of coevolutionary function approximation, we first test on model #1 in which the landscape is monotonously rising through the domain, slightly skewed into one dimension. At first sight it seems that this function has to be far more easy to solve for our coevolutionary system, but this is not the case. It even turned out that the parameters which give a 90% performance in model #2, cannot find any solution for this polynomial function. In order to get correct solutions in model #1 the fitness evaluation has to be relative instead of absolute. This can be understood if we examine the parasite behavior induced by the fitness evaluation. When the problem domain is monotonously rising and the evaluation is absolute, all the parasites clump in the highest part of the domain, $X=5.0, Y=5.0$. Here the absolute fitness advantage for parasites is highest. Because of this clumping, hosts can only evaluate one instance out of all problems. Variation is

minimized and the solution cannot be found through the coevolutionary process. Relative fitness evaluation for function 4.2, in contrast, will result in all the parasites clumping in the lowest part of the problem domain. Small inaccurate approximations of the hosts compared to the small target give a relatively high fitness advantage for the parasites. Because this is also the most difficult part of the landscape to solve, information integration is stopped. Only when the preferred problems are solved, the parasites will evolve to other regions of the problem space enabling the host to integrate enough information over the generations to solve the complete problem. In combination with the spatial properties of the system, this is the case for relative evaluation in the polynomial function and absolute for function 4.2. For similar reasons [Pagie and Hogeweg, 2000b] found that for the majority voting task, fitness for solving easy problems (the extreme densities) had to be larger than for harder problems.

5.3 Parasite Speciation




Figure 5.4: parasite speciation in model #1 (function 4.1).

To truly understand the role of pattern formation in our system we have to look at the parasite behavior. Within model #1 we observe parasite speciation, originating in the spatial pattern formation. A wave like pattern of subsequent alternating subpopulations travels through the whole space. These waves are separated subpopulations of parasites within the population. As can be seen in figure 5.4 we observe separate subpopulations of parasites in the population in the whole simulation until a correct solution is found. Because these subpopulations target different weaknesses in the host population, they form the wavelike patterns and can in this way coexist,

maintaining parasite-variation. The present parasite phenotypes plotted in an intermediate time step show that the parasites have speciated into very specific regions. These regions are the same for every simulation of model #1. When function 4.2 is the evolutionary target (and the evaluation is absolute), we also observe speciation in the parasite population (figure 5.5). Again specific regions in problem space are targeted by the speciated parasites. Because of the symmetry of the landscape, the regions differ per simulation, however always out of a small subset of the problem space.




Figure 5.5: parasite speciation in model #2 (function 4.2).

In both cases the parasites speciate within the first 25 steps of evolution into different subpopulations. The whole parasite population exists only of these subpopulations. Although the regions where the subpopulations linger differ per time and subpopulations do travel to other regions, great parts of the problem space will not ever be covered during the simulation. However, correct solutions are found.

5.4 Host Speciation

In the well mixed scheme (figure 5.4) all the offspring comes from only one ancestor. If in contrast we look at (fig. 5.4) the most striking difference is that in the spatial model, within all colored time steps there is offspring traceable back to two different ancestors present at initialization.

The above described parasite pattern formation is only one side of the coevolutionary coin. The wavelike patterns of the parasites force the hosts to adapt to different subpopulations, competing in space, which results in wavelike patterns of subsequent host-subpopulations. In model #1 we observe subpopulations of hosts speciating into a different dimension of the target

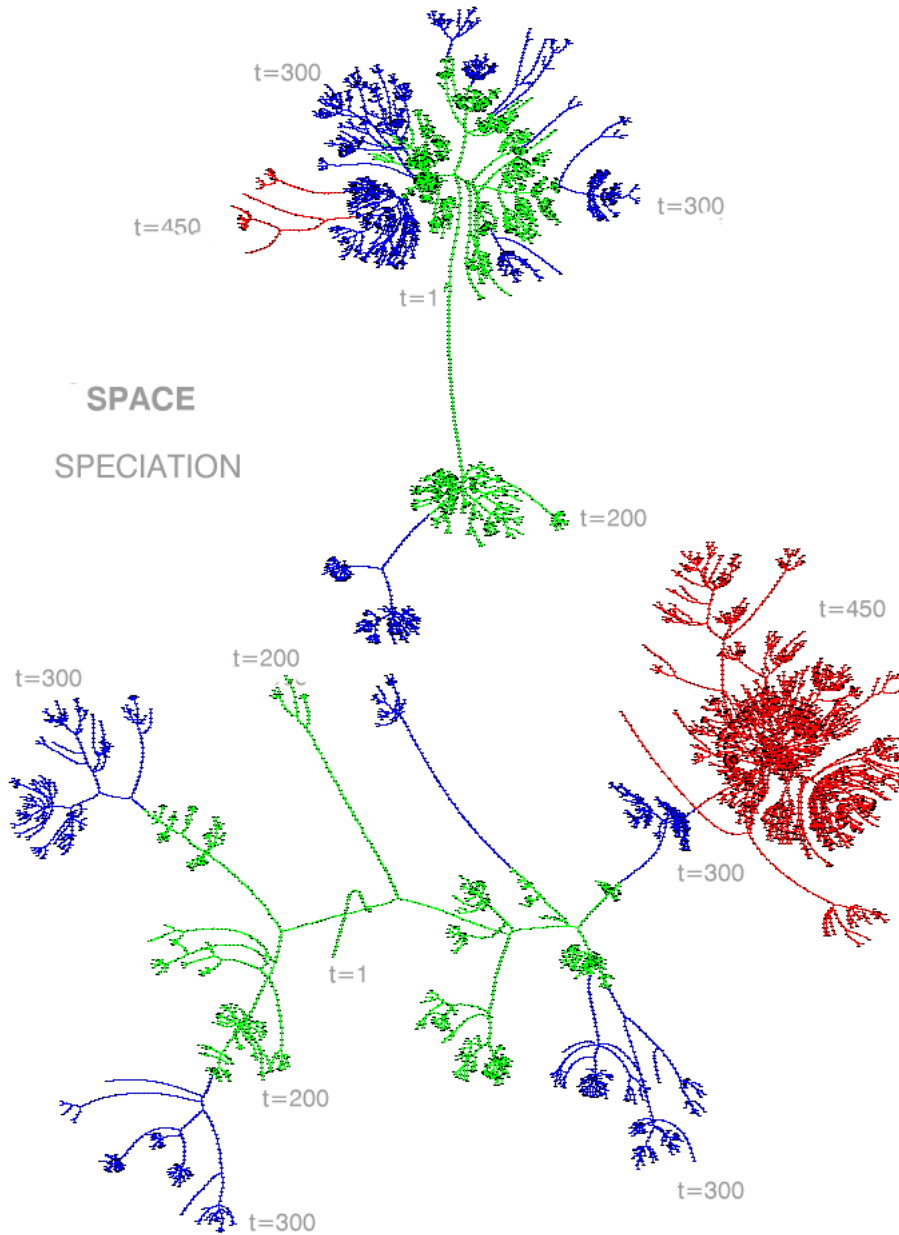


Figure 5.6: combined ancestor trace in system with spatial patterns shows Multiple sustaining lineages:genetic variation maintained

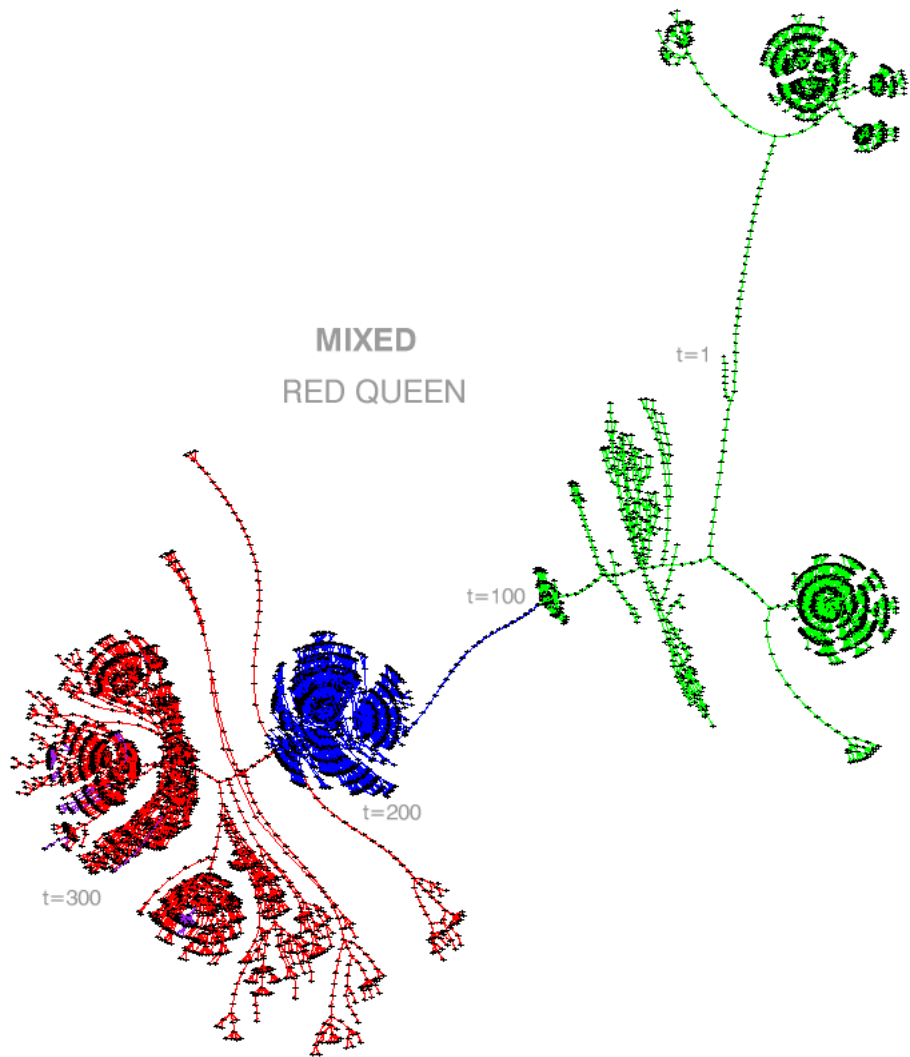


Figure 5.7: combined ancestor trace in well mixed system shows only one sustaining lineage

function, in order to specialize towards one of the parasite-subpopulations. In space we can again distinguish the subsequent waves of different subpopulations alternating each other. These patterns fit the parasite patterns, each subpopulation almost entirely on top of each other, following each other through the space. The spatial patterns we also observe in model #2, again indicating speciation. Only the waves are more difficult to distinguish because of the faster dynamics and different amount of subpopulations. In order to verify the host-speciation in model #2, we used the ancestor trace. In this way we can nicely see the evolutionary path and the mutational branches leading to the final population. To identify the role of space in this process, we constructed trees from the ancestor traces, both in the pattern formation case of model #2 as in the well mixed case¹. These trees consists of the whole mutated population at three different time steps, indicated by different colors. Evolution takes place in the direction of the arrows, so the outmost branches are the present population at the time indicated by the corresponding color. The hosts present at $t = 1$ are indicated.

A second observation here is that the different mutational branches keep expanding in the different time steps, indicating different genetic branches present in the same population. In the tree from the well mixed model we observe only one sustaining lineage leading to the final population. At the intermediate colored time steps the tree seems to expand its branches, however only one single branch takes care for the whole population in the next depicted time step. This difference in ancestor trees shows exactly the influence of spatial pattern formations. In the spatial model subpopulations adapt to different subpopulations of the parasite, so that genetic variation only diminishes when a genotype with the correct solution comes into existence and naturally dominates and takes over the population. The different branches are forced to compete with each other and force the different subpopulations to adjust to the parasite subpopulations, sustaining progressive evolution. In the well mixed model, all variation originates from the same ancestor. The differentiated branches at t , are gone at time step $t = t + 50$ and the whole population at $t = t + 50$ originated out of only one genotype present at t . This is the typical pattern, known as Red Queen evolution (subsection 3.5.1), which is observed in many host/virus interactions like Influenza in humans.

¹These two ancestor traces are randomly taken samples of both methods. However, we did observe ancestor traces with the expected corresponding clear cut patterns in almost all randomly examined simulations(5/6 in the spatial case, 4/6 in the wellmixed counterpart).

Chapter 6

Conclusion

6.1 Summary and Conclusion

In this thesis three research questions were postulated. In answering the first question *How does speciation of a population affect the adaptation of individuals?* we compared a coevolutionary model including spatial pattern formation with a variant where all the individuals were spatially mixed after each timestep. This well mixed model differs only in the lack of the emergence of spatial patterns. The competition is still local and the evaluation is still asymmetric and sparse. The difference in behavior as well as the success rate can in this way be subscribed to the presence or absence of space. The results are clear on the matter of success rate: In 90% of the simulations, our spatial model produced a global solution against an unsatisfactory success rate of 10% in the well mixed case. Thus, the success of spatial coevolutionary function approximation can almost entirely be ascribed to the presence of spatial pattern formation. Where the spatial embedding enables speciation, the spatial patterns are formed by this speciation of the population, suggesting that the speciation of the population has a positive effect on the individual adaptation.

How this advantage is reached by the emergence of spatial patterns is addressed in the second research question: *What is the role of space in spatial coevolutionary function approximation?* The behavior and emergence of spatial patterns is best explained, beginning within the parasite population. In both models we see clear speciation in problemspace within the population of parasites, only the behavior of the subpopulations differs per model. In the spatial variant the different subpopulations are quite stable, quite smoothly travelling through problem space and only gradually changing in number

of individuals within the different subpopulation. On the contrary, the well mixed case displays major fluctuations in the number of individuals per subpopulation. As a result, in the spatial model it is possible for the spatial distributed host population to adapt to different subpopulations of parasites, each lingering in another part of problemspace. Due to this specialization, the host population also splits into subpopulations. The combination of an ancestor trace and the use of clonal reproduction with gross chromosomal rearrangements instead of traditional crossover, make it possible to visualize speciation within the population of hosts, showing the mutational branches leading to the final population. Where the ancestor tree of the well mixed model only shows one lineage sustaining through the course of evolution, in the spatial model multiple lineages can keep expanding. Through these multiple lineages the genetic variation is maintained in the population over the long timescale.

Concerning information integration we now can observe a population based diversity. The different subpopulations are specialized on different subsets of problemspace and hence have integrated different parts of the information. Due to this self structuring of the system, a new level of selection arises at the level of the subpopulations. This newly obtained multilevel character of our system is probably the key element in answering our last research question, namely *What is the role of population based diversity for the individual information integration?* The population based diversity results in competing subpopulations in space and the resulting spatial patterns are formed in such a way that the system has to compete within- and between emerged subpopulations. The patterns 'guide' the learning process in space and provide in combination with the right fitness evaluation a sufficient learning gradient, presenting problems of different difficulty in the *right* order, necessary for adaptation of the host population. Relevant information is also presented at *multiple* timescales, guiding the coevolving subpopulations. Therefore, due to the selfstructuring of the system, the selforganizing dynamics of spatial patterns and multiple levels of selection, population based diversity can lead to individual information integration. We conclude that the influence of space cannot be attributed to local competition, or higher diversity alone, but depends on the selforganizing dynamics of the spatial patterns formed by the speciation of the population, and the appropriate timescales for adaptation generated.

6.2 Discussion and Future Work

After the work on spatial coevolutionary computation from Pagie and Hogeweg(1997), the role of space is more and more accepted as an effective method to improve coevolutionary computation. However there still is a lack of basic research on the fundamental mechanisms of space in a coevolutionary setting. The results shown in this thesis are at least promising and hopefully will revive the interest in the spatial variant of coevolutionary computation. A lot of work still has to be done. Upcoming research will include a statistical analysis in order to reinforce our results. The exact role of space is still not entirely explained, so future research may focus on investigating as many facets of space as possible.

A different important aspect which is still relatively unknown is to which extent the spatial location of a mutational branch over time is important to evolve into a global solution. We have to investigate where the lineage leading to the global solution originates and lingers in space through the evolutionary course. Is it advantageous for a mutational lineage to dwell on the borders of subpopulations in order to fully experience all levels of selection or does the global solution originates in the 'safe' center of a competing subpopulation?

A very surprising result is the fact that the encountered diversity of problems at a random position in space seems to be higher in the well mixed system. We have to investigate if this is also the case for the encountered diversity of individuals, when individually followed in space and time. Related to this, it might also be interesting to study the role of sparse fitness evaluation in this context. However, because of the sparse fitness evaluation each host has its own neighborhood with corresponding parasites, therefore competition is between hosts which have a fitness assigned from a different subset of the problem space. This increases the effectiveness of evaluation by comparing the distinguishing capabilities of the different parasites, maybe compensating for a lower variety. We think that only if the problems speciate such that 'easy ones' are first evaluated, the coevolutionary process is successful. Can we consider space to be in some senses an ideal trainer?

Bibliography

- [Angeline and Pollack, 1993] Angeline, P. and Pollack, J. (1993). Competitive environments evolve better solutions for complex tasks. In Forrest, S., editor, *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 264–270.
- [Boerlijst, 1991] Boerlijst, M.C. and Hogeweg, P. (1991). Spiral wave structure in prebiotic evolution: hypercycles stable against parasites. *Physica D*, 48:17–28.
- [Cartlidge and Bullock, 2004] Cartlidge, J. and Bullock, S. (2004). Combating coevolutionary disengagement by reducing parasite virulence. *Evolutionary Computation*, 12(2):193–222.
- [Churchland and Sejnowski, 1990] Churchland, P. M. and Sejnowski, T. J. (1990). Neural representation and neural computation. *Philosophical Perspectives*, 4:343–382.
- [Corsetti et al., 2006] Corsetti, F., Olcott, A., and Bakermans, C. (2006). The biotic response to neoproterozoic snowball earth. *Palaeogeography, Palaeoclimatology, Palaeoecology*, 232(2-4):114–130.
- [Darwin, 1859] Darwin, C. (1859). *'The Origin of Species', the mechanism of natural selection*. Murray, , London, United Kingdom.
- [Dawkins and Krebs, 1979] Dawkins, R. and Krebs, J. (1979). Arms races between and within species. In *Proceedings of the Royal Society of London. Series B, Biological Sciences*, volume 205 of *The Evolution of Adaptation by Natural Selection*, pages 489–511.
- [De Boer and Hogeweg, 2007] De Boer, F. and Hogeweg, P. (2007). The role of speciation in spatial coevolutionary function approximation. In *Proceedings of the 2007 Genetic and Evolutionary Computation Conference*, pages 2437–2441. GECCO-2007.

- [De Jong and Pollack, 2004] De Jong, E. and Pollack, J. (2004). Ideal evaluation from coevolution. *Evolutionary Computation*, 12(2):159–192.
- [Eiben and Smith, 2003] Eiben, A. and Smith, J. (2003). *Introduction to Evolutionary Computing*. Natural Computing Series. Springer.
- [Epstein, 1994] Epstein, S. (1994). Toward an ideal trainer. *Machine learning*, 15(3):251–277.
- [Fogel, 1994] Fogel, D. B. (1994). An introduction to simulated evolutionary optimization. *IEEE Trans. on Neural Networks: Special Issue on Evolutionary Computation*, 5(1):3–14.
- [Goldberg, 1989] Goldberg, D. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Boston, MA, USA.
- [Hillis, 1990] Hillis, D. (1990). Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, 42:228–234.
- [Hoffman, 1989] Hoffman, A. (1989). *Arguments on Evolution: a Paleontologist's View*. Oxford University Press, New York.
- [Hogeweg, 1994] Hogeweg, P. (1994). Multilevel evolution: replicators and the evolution of diversity. *Physica D*, 75:275–291.
- [Holland, 1975] Holland, J. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor.
- [Holland, 1992] Holland, J. (1992). Genetic algorithms. *Scientific American*, pages 44–50.
- [Husbands, 1994] Husbands, P. (1994). Distributed coevolutionary genetic algorithms for multi-criteria and multi-constraint optimisation. In *Evolutionary Computing, AISB Workshop*, pages 150–165.
- [Johnson and Boerlijst, 2002] Johnson, C. and Boerlijst, M. (2002). Selection at the level of the community: the importance of spatial structure. *Trends in Ecology and Evolution*, 17(2):83–90.
- [Juillé and Pollack, 1998] Juillé, H. and Pollack, J. (1998). Coevolving the "ideal" trainer: Application to the discovery of cellular automata rules. In *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 519–527.

- [Juillé and Pollack, 1996] Juillé, H. and Pollack, J. B. (1996). Co-evolving intertwined spirals. In *Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming*, page 461–467.
- [Koza, 2003] Koza, J., editor (2003). *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*. Kluwer Academic Publishers.
- [Koza, 1992] Koza, J. R. (1992). *Genetic programming: on the programming of computers by means of natural selection*. MIT Press, Cambridge, MA, USA.
- [Mitchell et al., 2006] Mitchell, M., Thomure, M., and Williams, N. (2006). The role of space in the success of coevolutionary learning. In *Proceedings of ALIFE X, The Tenth International Conference on the Simulation and Synthesis of Living Systems*.
- [Pagie and Hogeweg, 1997] Pagie, L. and Hogeweg, P. (1997). Evolutionary consequences of coevolving targets. *Evolutionary Computation*, 5(4):401–418.
- [Pagie and Hogeweg, 2000a] Pagie, L. and Hogeweg, P. (2000a). Individual- and population-based diversity in restriction-modification systems. *Bulletin of Mathematical Biology*, 62(4):759–774.
- [Pagie and Hogeweg, 2000b] Pagie, L. and Hogeweg, P. (2000b). Information integration and red queen dynamics in coevolutionary optimization. In *The 2000 Congress on Evolutionary Computation*, pages 1260–1267.
- [Pagie and Mitchell, 2002] Pagie, L. and Mitchell, M. (2002). A comparison of evolutionary and coevolutionary search. *International Journal of Computational Intelligence and Applications*, 1(2):53–69.
- [Paredis, 1997] Paredis, J. (1997). Coevolving cellular automata: Be aware of the red queen! In *Proceedings of the Seventh International Conference on Genetic Algorithms*.
- [Potter and De Jong, 1994] Potter, M. and De Jong, K. (1994). A cooperative coevolutionary approach to function optimization. In Davidor, Y., Schwefel, H.-P., and Manner, R., editors, *Proceedings of the Third Conference on Parallel Problem Solving from Nature*, volume 866 of *Lecture notes in computer science*, pages 249–257.

- [Rieu et al., 2007] Rieu, R., Allen, P., Pltze, M., and Pettke, T. (2007). Climatic cycles during a neoproterozoic "snowball" glacial epoch. *Geology*, 35(4):299–302.
- [Rosin and Belew, 1997] Rosin, C. and Belew, R. (1997). New methods for competitive coevolution. *Evolutionary Computation*, 5(1):1–29.
- [Savill et al., 1997] Savill, N., Rohani, P., and Hogeweg, P. (1997). Reinforcing spatial patterns enslave evolution in a host-parasitoid system. *Journal of Theoretical Biology*, 188(1):11–20.
- [Sims, 1994] Sims, K. (1994). Evolving 3d morphology and behavior by competition. In *Highlights of the AlifeIV Conference*, volume 1 of *Artificial Life IV*, pages 353–372.
- [Spencer, 1864] Spencer, H. (1864). *The Principles of biology*, volume 1. Williams and Norgate, London.
- [van der Laan and Hogeweg, 1995] van der Laan, J. and Hogeweg, P. (1995). Predator-prey coevolution: Interactions across different timescales. In *Proceedings: Biological Sciences*, volume 259, pages 35–42.
- [van Valen, 1973] van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, 1:1–30.
- [Williams and Mitchell, 2005] Williams, N. L. and Mitchell, M. (2005). Investigating the success of spatial coevolutionary learning. In *Proceedings of the 2005 Genetic and Evolutionary Computation Conference*, pages 523–530. GECCO-2005.
- [Wilson, 1992] Wilson, E. (1992). *The diversity of life*. Penguin Books.
- [Wright, 1932] Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. In *Proceedings of the Sixth International Congress on Genetics*.