

The Role of Speciation in Spatial Coevolutionary Function Approximation

Folkert de Boer

theoretical biology&Bioinformatics group
Padualaan 8
3584 CH Utrecht, The Netherlands
+3130 2533688

fkdeboer@gmail.com

Paulien Hogeweg

theoretical biology&Bioinformatics group
Padualaan 8
3584 CH Utrecht, The Netherlands
+3130 2533692

p.hogeweg@bio.uu.nl

ABSTRACT

The role of space is more and more accepted as a way to dramatically improve the success of coevolutionary function approximation. The process behind this success however is not yet fully understood. It is suggested that spatiality causes a persistence in the population diversity over generations and a better targeting of weak points in the host-population by means of the parasite.

In this paper we will discuss the role of spatial pattern formation and speciation in coevolutionary function approximation and the influence on the success rate of coevolution.

We observe specific patterns of speciation in the problems as well in the problem solving-population(LISP functions). These patterns depend on a combination of the functions and the fitness criteria. The success of the spatial coevolutionary process can be understood from the speciation patterns: only if the problems speciate such that 'easy ones' are first evaluated, the coevolutionary process is successful.

Categories and Subject Descriptors

F.1.0 [General]

General Terms

Theory

Keywords

Coevolution, spatial pattern formation, speciation, function optimization, CA

1. INTRODUCTION

More and more coevolution is regarded as a very successful evolutionary computational method [1,2,3]. Without having to evaluate all problems in the domain or defining very specific problem sets, coevolution with sparse fitness evaluation is successful because it samples the problem domain more efficiently.

Hillis[1] already noted that the coevolutionary process improved in a spatial environment. Pagie and Hogeweg[3] used a spatial model for function approximation. Their results are repeated several times and the role of space is more and more accepted as an effective method to improve coevolutionary computation[4,5].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'07, July 7-11, 2007, London, England, United Kingdom.
Copyright 2007 ACM 978-1-59593-698-1/07/0007...\$5.00.

How spatiality exactly contributes to this success is unknown. In contrast of the believe that the persisting diversity and the targeting of weak spots in the "host" population are the main reason for superior performance[6], we will conclude that the success crucially depends on the selforganizing dynamics of spatial patterns.

Coevolution is not always successful. Because of local competition and sparse fitness evaluation the 'arms races' can lead to red queen evolution when the host continuously adapts to the presented problem set and is unable to maintain the information.

We will investigate the role of spatial pattern formation in keeping the appropriate timescales and on the order of presenting the problems.

2. METHODS

2.1 The model

We will describe the role of spatial pattern formation, comparing it with a well mixed system.

Both systems are modeled in Cellular Automata and in the same general setup used by Pagie and Hogeweg[3]. This 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 x 50 cells, defining the population size at 2500. Competition for growth is local in space. Each individual has to compete within his 3 x 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 test without spatial pattern formation. Note that this well mixed model differs only in spatial pattern formation. Competition is still local, evaluation is asymmetric and sparse.

2.2 Target

We use 2 different functions as evolutionary target:

function 1:

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

function 2:

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

First we describe our results with a system approximating function 1 (further referred to as model #1). In the next section we will extend our research with the second function, first proposed by Pagie and Hogeweg[3].




Figure 1: function 1 landscape




Figure 2: function 2 landscape

3.2.2.3 Fitness

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 1 the domain is $X = 0.2, 5.0$ and $Y = 0.2, 5.0$ with an interval of 0.2. For function 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 on maximizing the distance between target function and solution of just one host. This asymmetry is found to give better results with respect to optimization than a symmetric evaluation[3]. The host fitness is less affected by changes in a parasite, the parasite however can react directly to fluctuations in the host. 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 proportional with the fitness and the selection pressure can be conveniently adjusted by using an exponential factor.

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

The fitness for the host is defined as:

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

where the distance d is defined as:

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

with $M = 1$ for absolute evaluation and $M = T(p_i)$ for relative evaluation and n normalizes the distance. In all our experiments $n = 3/4$.

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 nodes in a correct solution) an exponential penalty is subtracted from the fitness. This penalty has a maximum of 1/9th of the maximal fitness.

This is sufficient to highly improve calculation time and keep the solutions short.

2.4 Genetic representation

The genetic representation of the solutions is based on genetic programming and is essentially the same as proposed by Pagie and Hogeweg[3]. The genotype of the host is a (LISP) list-representation 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 $\{x, y, R\}$, where 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 below), the crossover is internal. 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 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 (i.e. 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. 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 a 1000 timesteps the simulation will be stopped, otherwise simulations are stopped after reaching the solution and staying in the population 50 time steps.

2.5 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.

4. RESULTS

3.1 Efficacy of the optimization




Figure 3: Spatial patterns in problem space. 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[3, 7, 8].

figure 3 visualizes the pattern formation from parasites in our model #1, showing waves of parasites with alternating high X,Y values. This pattern formation is the driving force in the dynamics.

Table 1. Success rates of both models

Model #1	Success rate
Spatial patterns	27/30(90%)
Well mixed	3/30(10%)

The results in table 1 show that the well-mixed system fails in almost all simulations in finding a correct solution. In our system with pattern formation however, 90 % of the runs gives a correct solution, almost always within the 300 time steps. This dramatic difference in performance can be totally attributed to the occurrence of pattern formation.

3.2 Coevolutionary dynamics depend on fitness evaluation

[9, 10, 11, 12] all describe a ideal trainer as an trainer which presents problems of increasing difficulty. This is often done with a learning gradient or a domain specific fitness evaluation. We will 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. Because of this particular landscape we can better follow the evolutionary process.

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 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. This is the case for relative evaluation in the polynomial function and absolute for function 2.

For similar reasons Pagie[10] found that for the density classification test fitness for solving easy problems(extreme densities) had to be larger than for harder problems.

3.3 Parasite speciation

To truly understand the role of pattern formation in our system we




Figure 4: parasite speciation in evaluating function 1

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 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 2 is the evolutionary target (and the evaluation is absolute), we also observe speciation in the parasite population(fig. 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: parasite speciation in function 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.




Figure 6: combined ancestor trace in well mixed system

3.3 host speciation

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, this results in wavelike patterns of subsequent host-subpopulations. In model #1 we observe subpopulations of hosts speciating into a different dimension of the target 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 case. These trees consist 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 outer most branches are the present population at the time indicated by the corresponding color. The hosts present at $t=1$ are indicated.

In the well mixed scheme (fig. 6) all the offspring comes from only one ancestor. If in contrast we look at (fig. 7) 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.

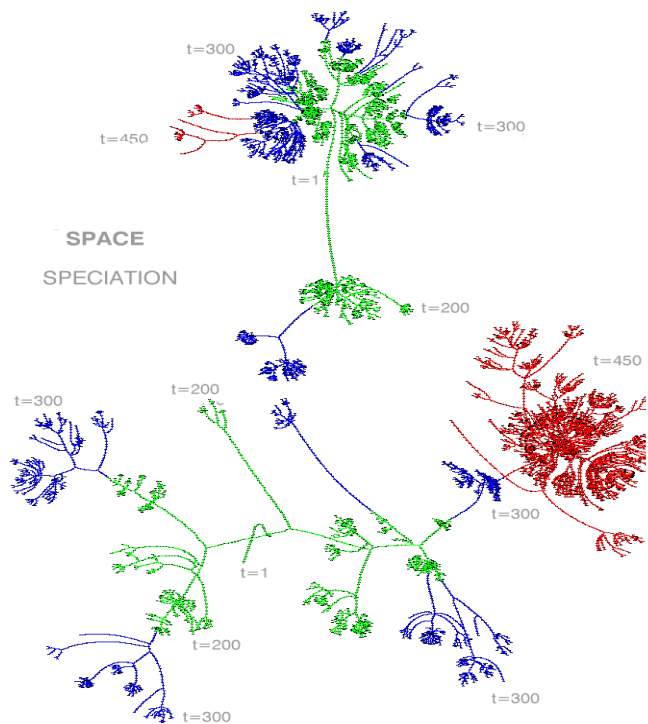


Figure 7: combined ancestor trace with pattern formation

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 subpopulation.

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, which is observed in many host/virus interactions like eg. Influenza in humans.

5. 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.(fig 8), 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 the interconnecting subpopulations organizes itself into an "ideal learner/trainer" system.

6. CONCLUSION

The success of spatial coevolutionary function approximation can almost entirely be ascribed to the presence of spatial pattern formation. Spatial patterns guide the learning process in space and time. Chaotic waves present relevant information at multiple timescales and guide the coevolving subpopulations, preventing a red queen race in which the sparse information is not integrated into a full solution.

Interestingly the encountered diversity of problems at a random position in space is higher in the well mixed system, both at a very short and at the long timescale. Nevertheless the success rate of the spatial system is much higher.

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 and the appropriate timescales for adaptation generated.

7. ACKNOWLEDGMENTS

We thank Rients van Wijngaarden for initiating this project and his dedication.

8. REFERENCES

- [1] Hillis, D. W. (1990); *Co-evolving parasites improves simulated evolution as an optimization procedure*. Physica D 42:228-234.
- [2] Juillé, H., & Pollack, J. B. (1996). *Co-evolving intertwined spirals*. In L. Fogel, P. Angeline, & T. Baeck (Eds.), *Evolutionary Programming V*:

Proceedings of the Fifth Annual Conference on Evolutionary Programming (pp. 461–467). Cambridge, MA: The MIT Press.

- [3] Pagie, L. & Hogeweg, P. (1997); *Evolutionary consequences of coevolving targets*. Evolutionary Computation 5(4):401-418.
- [4] Pagie, L. & Mitchell, M. (2002). A comparison of evolutionary and coevolutionary search. International Journal of Computational Intelligence and Applications, 2(1), 53–69.
- [5] Mitchell, M., Thomure, M.D. And Williams, N.L.(2006) *The role of space in the Success of Coevolutionary Learning*. Proceedings of ALIFE X—The Tenth International Conference on the Simulation and Synthesis of Living Systems, 2006.
- [6] Williams, N. L. and Mitchell, M. (2005). *Investigating the success of spatial coevolutionary learning*. In Proceedings of the 2005 Genetic and Evolutionary Computation Conference, GECCO-2005, pages 523–530, New York. ACM Press.
- [7] M.C. Boerlijst and P. Hogeweg, *Spiral wave structure in prebiotic evolution: hypercycles stable against parasites*, Physica D 48 (1991) 17-28.
- [8] N. J. Savill, P. Rohani, P. Hogeweg (1997); *Self-Reinforcing Spatial Patterns Enslave Evolution in a Host-Parasitoid System*, Journal of Theoretical Biology Vol. 188, no. 1, pp. 11-20.
- [9] Epstein, S. L. (1994). *Toward an ideal trainer*. Machine Learning , 15 (3), 251–277.
- [10] Pagie , L. (1999); *Coevolutionary dynamics: information integration, speciation, and red queen dynamics*. In: Information integration in evolutionary processes, chap. 5. pp. 67–93.
- [11] Juillé, H. and Pollack, J. B. (1998). *Coevolving the 'ideal' trainer: Application to the discovery of cellular automata rules*. In Genetic Programming 1998: Proceedings of the Third Annual Conference , San Francisco, CA. Morgan Kaufmann.
- [12] De Jong, E. D. and Pollack, J. B. (2002). *Ideal evaluation from coevolution*. Evolutionary Computation , 12(2):159–192.