

Information integration and red queen dynamics in coevolutionary optimization

Ludo Pagie*

Bioinformatics group
Utrecht University
Padualaan 8, 3584 CH, Utrecht
The Netherlands
l.pagie@bio.uu.nl

Paulien Hogeweg

Bioinformatics group
Utrecht University
Padualaan 8, 3584 CH, Utrecht
The Netherlands
p.hogeweg@bio.uu.nl

Abstract- Coevolution has been used as optimization technique both successfully and unsuccessfully. Successful optimization shows integration of information at the individual level over many fitness evaluation events and over many generations. Alternative outcomes of the evolutionary process, e.g. red queen dynamics or speciation, prevent such integration. Why coevolution leads to integration of information or to alternative evolutionary outcomes is generally unclear.

We study coevolutionary optimization of the density classification task in cellular automata in a spatially explicit, two-species model. We find optimization at the individual level, i.e. evolution of cellular automata that are good density classifiers. However, when we globally mix the populations, which prevents the formation of spatial patterns, we find typical red queen dynamics in which cellular automata classify all cases to a single density class regardless their actual density. Thus, we get different outcomes of the evolutionary process dependent on a small change in the model. We compare the two processes leading to the different outcomes in terms of the diversity of the two populations at the level of the genotype and at the level of the phenotype.

1 Introduction

In the context of evolutionary optimization techniques some studies show that coevolution leads to an increase in the performance or efficiency of the optimization process (Paredis, 1995; Husbans, 1994; Rosin & Belew, 1997). In these models coevolution is often compared to predator-prey or host-parasite interactions, i.e. a reciprocal antagonistic interaction (Bullock, 1995). The prey, or hosts, implement the potential solutions to the optimization problem. The predators, or parasites, implement individual ‘fitness-cases’. Coevolution does not always lead to general solutions of the optimization problem; red queen dynamics may hinder the optimization process (Paredis, 1997), the coevolving species may speciate (Hillis, 1990), or settle into “mediocre stable states” (Ficici & Pollack, 1998).

Hillis (1990) studied a coevolutionary optimization model in which sorter algorithms coevolved with sorter problems.

He found that coevolution of algorithms and problems resulted in a much more efficient optimization process that led to faster sorter algorithms than algorithms found in traditional evolutionary optimization processes. In addition to the coevolutionary, antagonistic relation between algorithms and problems Hillis embedded the evolutionary process in a spatial setting; algorithms and problems were situated on a 2-dimensional grid and interacted only locally. Similar, spatially embedded coevolutionary optimization models were studied by Husbans (1994) and Pagie & Hogeweg (1997). In all cases an improvement of the optimization process was reported.

In other non-spatial coevolutionary optimization models often techniques are used to ensure diversity of both populations and longevity of ‘good’ individuals (Paredis, 1995; Rosin & Belew, 1997; Juillé & Pollack, 1998). The increased longevity of solutions and the ensuring of diversity of both species help to evolve general behavior, i.e. in the optimization process (see also Paredis (1997)). The effects of such techniques, however, are automatic side-effects of local dynamics such as occur in spatial evolutionary systems (Collins & Jefferson, 1991; Husbans, 1994; Mahfoud, 1995; Pagie & Hogeweg, 1997; Rosin & Belew, 1997).

We present results of a study of a spatially explicit coevolutionary model in which two species have an antagonistic interaction. We compare two cases. The first case depicts coevolution in a spatial environment in which individuals interact and compete locally with each other so that spatial pattern formation occurs. The second case depicts coevolution in the same model except that the individuals of both populations are globally mixed every time step. In this case spatial pattern formation does not occur. In the first model the evolutionary process leads to individuals that have integrated adaptations to separate selection events into a general solution. In the second model, in which the individuals are mixed, we see typical cyclic red queen dynamics. Here, we present some results on the evolution of diversity in the two models. Elsewhere we will present additional results and focus on the presence of optimization, speciation, and red queen dynamics in evolutionary processes (Pagie & Hogeweg, 2000).

2 The model

We study the coevolutionary process in the context of the optimization of a computational task. Although the task is cho-

*Current address:
Santa Fe Institute, 1399 Hyde Park Road, NM 87501, ludo@santafe.edu

sen rather arbitrarily it lends itself easily for embedding in a two-species system with antagonistic interactions. The genetic encoding of the task is characterized by a non-linear genotype-phenotype mapping with strong epistatic interactions. We use a individual-based, discrete space, discrete time model with synchronous updating. The general structure of the model is very similar to the structure of the models that were studied by Hillis (1990) and Pagie & Hogeweg (1997). The two species present in the model are called CAs and ICs.

The CAs are 1-dimensional, binary state cellular automata next-state rule-tables with a neighborhood size 3 (Wolfram, 1984; Toffoli & Margolus, 1987), the ICs are initial conditions of the cellular automata and are of length 149. Both CAs and ICs are represented as bit strings. The interaction between a CA and an IC, and therewith the basis on which their fitness is calculated, is based on the density-classification task of cellular automata (Mitchell *et al.*, 1994). In the density classification task the CAs must classify ICs on the basis of the number of 0s and 1s in the bit string of the IC. If the IC has a majority of zeros in its bit string it belongs to class 0, otherwise it is class 1[†]. The CA is allowed to iterate for maximally 320 time steps, starting with the IC as initial condition. If the CA settles into a homogeneous state of all zeros it classifies the IC as being of class 0. If the CA settles into a homogeneous state of all ones it classifies the IC as being of class 1. If the CA does not settle into a homogeneous state it answers "don't-know", and does not receive a fitness reward. Only if the CA classifies an IC correctly does it receive a fitness reward of 1. In all other cases the IC receives a fitness reward of f (see below).

This particular task for cellular automata and its evolutionary optimization is studied extensively by the EvCA-group in the Santa Fe Institute (see Mitchell *et al.* (1996) for a review). Coevolutionary models using this task were previously studied by Paredis (1997) and Juillé & Pollack (1998). The latter, however, used an intricate coevolutionary scheme incorporating global feedback strategies to prevent the occurrence of red queen dynamics (see also (Werfel *et al.*, 1999) for additional studies in that context). Here, we use the task of density classification primarily to study the process of coevolution between two antagonistic species. The (evolution of the) task itself is of little importance for this study although we are interested in its properties as evolutionary 'goal'. Below we will discuss some of these properties.

2.1 Spatial embedding and local dynamics

Individuals of both species are distributed in space which is a 2-dimensional regular grid of 30 by 30 cells with periodic boundary conditions. Each cell contains one CA and one IC, giving population sizes of 900 individuals. The CAs and ICs are evaluated with respect to each other locally in this space. The fitness of a CA is based on the ICs in its Moore ad-

joining, i.e. the eight cells directly neighboring the middle cell plus the middle cell itself. The fitness of an IC is based only on the CA in the same cell. This asymmetric fitness evaluation procedure was found to improve the evolutionary optimization process Pagie & Hogeweg (1997). The fitness evaluation scheme is characterized by a very sparse evaluation of the objective function, i.e. a general IC classification algorithm. Sparse evaluation is in fact unavoidable because the total number of ICs is 2^{149} and the total number of CAs is 2^{128} . Moreover, in (Pagie & Hogeweg, 1997) we showed that sparse fitness evaluation can help the evolutionary process rather than hinder it (see also Hillis (1990)). We call the fitness of CAs and ICs that they receive during fitness evaluation *local fitness*: it is the fitness value which is used in the selection process. In order to compare CAs objectively we calculate a general fitness measure (see below) which we call *performance fitness* (Mitchell *et al.*, 1994).

After fitness evaluation in each cell of the grid a selection procedure is performed between locally present CAs and between locally present ICs, and growth of the selected CA and IC in the cell. Selection is based, probabilistically, on the rank order of the nine individuals in the Moore neighborhood. The probability for an individual to be selected is 0.5^{rank} , where $rank = 1..8$. The last ranked individual (i.e. $rank = 9$) also has a probability 0.5^8 for being selected. Note that we have constant population sizes. Although this is usual in evolutionary optimization models it is of course less realistic from a biological point of view.

After selection and growth we apply mutations to the CAs and the ICs. We only use bit-flip mutations with rate 0.2 per CA and rate 0.5 per IC. The use of the bit-flip operator introduces a strong mutational bias, in terms of the density of bit strings, towards density values of 0.5. The presence of this bias appears to have a large influence on the evolutionary dynamics in the context of the task that we study here (see also (Mitchell *et al.*, 1994; Paredis, 1997)). For the initial conditions this bias pushes them directly towards the phenotype phase-transition in genotype space where it is easy to be difficult.

The two models that we study in this paper are as described above except that in the second model, i.e. the *mixed model*, we globally mix the individuals of both populations every time step. In the first model, i.e. the *base model*, spatial patterns can form and influence the evolutionary process (e.g. see Boerlijst & Hogeweg (1991); Savill & Hogeweg (1997)).

The *performance fitness* of a cellular automaton is defined as the number of correct classifications out of 10,000 randomly created initial conditions that have an unbiased density distribution (i.e. a binomial distribution around 0.5). We use this fitness measure, or *performance fitness*, when we compare CAs of different populations. Initial conditions with a density of approximately 0.5 are the most difficult to classify because bit string that are almost equal (e.g. differ on only one bit position) can belong to different density classes. In fact, the performance of a good cellular automaton, like

[†]The bit string of the ICs have an odd length, so the majority is always defined

for instance the GKL rule, decreases rapidly if it is evaluated on the basis of initial conditions whose density approaches 0.5 (Mitchell *et al.*, 1994; Juillé & Pollack, 1998). A ‘good’ cellular automaton has a fitness value of about 0.8 (e.g. the GKL rule; 0.81), although cellular automata have been found recently with fitness values of up to 0.86 (Juillé & Pollack, 1998).

As an evolutionary optimization task evolving good cellular automata appears to be difficult; in only a small number of evolutionary runs are cellular automata found with fitness values in the same range as the fitness of the handwritten cellular automata (Mitchell *et al.*, 1996).




Figure 1: IC fitness function Φ . The fitness f an IC gets if it is not correctly classified depends on its density. As result we get stabilizing selection toward minimal or maximal density values, which are the ‘easy’ ICs.

Previous studies of coevolutionary optimization of this task showed a rapid evolution of ICs towards the most difficult part of their genotype space, i.e. where they have a density around 0.5. In order to circumvent this problem we used a density dependent fitness function to calculate the fitness of ICs (fig. 1). This fitness function implements stabilizing selection towards minimum (i.e. 0.0) and maximum (i.e. 1.0) density values. The actual values of f do not matter, only the symmetry of Φ around density = 0.5 and the fact that f increases monotonically when it approaches the minimum and maximum density values. In fact we simply used $\Phi(IC_i) = |(density(IC_i) - \frac{\max\ density}{2})|$. In Pagie & Hogeweg (2000) we compare and discuss the evolutionary dynamics in this model when we use the fitness function of fig. 1 and if we use a flat fitness function.

3 Results

In this section we will describe a simulation of the base model and a simulation of the mixed model. We found that the two simulations are typical for the possible outcomes of the evolutionary process in the model. The precise parameter settings do not influence the general results to a great extent. The values that we used in the simulations that we describe here were actually chosen rather arbitrarily, e.g. we did not optimize our results in any particular way. The two simulations discussed below are run with the same parameter settings. Simulations are started with randomly created CAs, i.e. CAs with a density around 0.5; the ICs have an initial density of 0.0.

3.1 Two typical simulations

From the point of view of optimization of density classification the most important variable is the performance fitness. In fig.2 we plot the evolution of the performance fitness of the best CA in the population in the base model (solid line) and in the mixed model (dashed line). The performance fitness of the best individual in the mixed model fluctuates between 0.50 and 0.55. Even the best CAs in this model do not classify random initial conditions much more accurately than random classification into class 0 or class 1.

The performance fitness of the best CA in the base model initially increases and then fluctuates between 0.70 and 0.75. These values for performance fitness of the CAs are in the same range as the performance fitness values for the best cellular automata found in the evolutionary optimization models studied by Mitchell *et al.* (1994), Crutchfield & Mitchell (1995), and Paredis (1997). Clearly, they are much more general than the CAs from the mixed model. Following the concepts of Crutchfield & Mitchell (1995), and Hordijk *et al.* (1998), the CAs use particle-based strategies in order to compute the density of ICs, as does, for instance, the rule GKL.

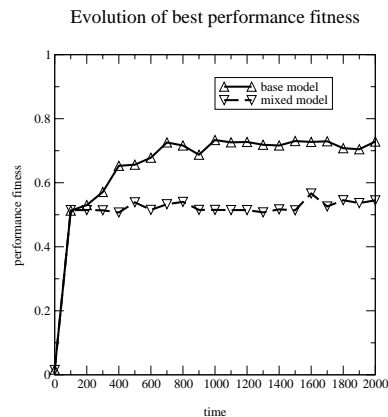


Figure 2: Evolution of absolute fitness of the best individuals in the population in the base model (solid line) and the model with global mixing (dashed line). Whereas the CAs in the base model show an increase in the fitness of the best individual the best CA in the globally mixed model remains around 0.55.

In the base model we see that the CAs evolve a generalized classification algorithm whereas in the mixed model the performance fitness ≈ 0.5 . However, the success of individual CAs in the models depend on their local fitness values. In fig.3 we show time-plots of the base model (A) and of the mixed model (B) again, but now we plot the average of the *local* fitness values of the CAs and the ICs, and their average density. All averages are normalized between 0.0 and 1.0, but the true ranges are given in the legend.

Although there is a great difference between the performance fitness values of CAs of the two models their local fitness is approximately equal when we average over time (≈ 0.9). The IC populations do better in the base model; they

have a time-averaged local fitness of 0.08 in the base model whereas in the mixed model they have a time-averaged local fitness of 0.04. The IC fitness value, however, also depends on the density values of the ICs in the population.

Thus, seen as a biological system the CAs do equally well in both models. Of course, in these models we do not take into account the population dynamics which may alter the results in this particular respect. Thus, although from an optimization point of view there is a clear difference in fitness values from a local point of view the CAs perform well in both cases. In (Pagie & Hogeweg, 2000) a similar equality of time-averaged local fitness values was found in the base model and in the mixed model, using a different IC fitness function.

When we look in more detail at the dynamics of fitness and density in fig. 3 we can see that the initial transients are

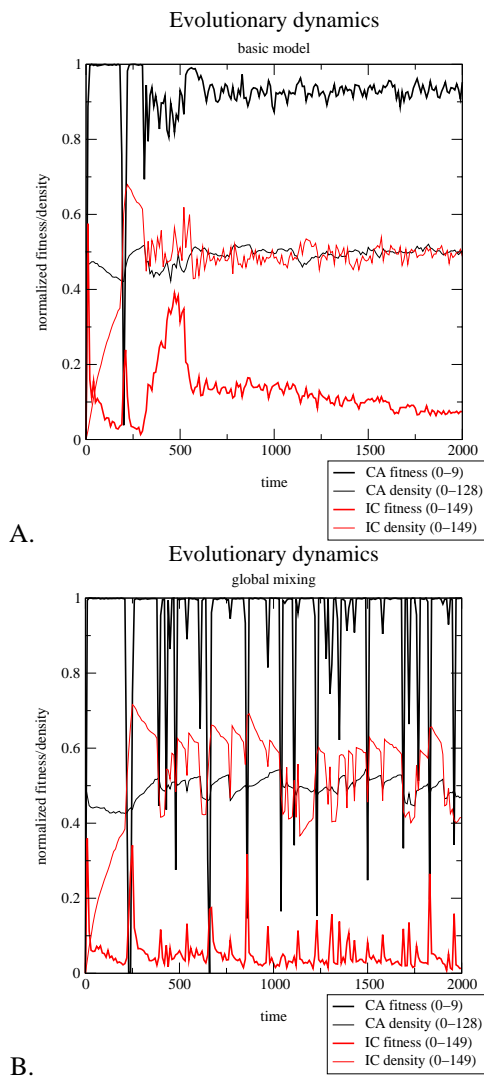


Figure 3: Evolutionary dynamics of basic model (A) and the model with global mixing (B). The simulations start with the same parameter values and with the same initial state.

roughly the same; large fluctuations of the average fitness values of the CAs and the ICs together with large fluctuations of the average densities of the CAs and the ICs. The simulations start with ICs that have a density of 0.0 which are very easy to classify correctly. Indeed, the average local fitness of the CA population quickly increases toward maximum values. As a result of mutations, ICs will arise with density values higher than 0.0. But initially these ICs are still very easy to classify correctly and the CAs maintain the high local fitness values.

The subsequent evolution of the IC population towards ICs with still higher densities increases the difficulty of the ICs. However, even when the density of the ICs approaches 0.5 the CAs in the population still classify them correctly and maintain high local fitness values. This is because the CAs simply settle into a homogeneous state of zeros independent of the state of the IC. Up to this time this strategy of the CAs in fact performs perfectly and this behavior is easily evolved and easily maintained.

At $t \approx 200$ ICs arise that have a density larger than 0.5. Now the CAs have a problem; settling into a homogeneous state of zeros is no longer the correct behavior. Indeed, the average local fitness of the CAs drops to very low values. During this stage the IC population experiences strong selection pressure towards ICs with very large density values as a result of the IC fitness function Φ (fig.1). Soon after the switch in the average density of the ICs, however, we see that the average local fitness of the CA population rises again to very high values in both simulations. At this point the same general behavior can be seen as at the beginning of the runs, except that the density of the ICs is now larger than 0.5.

From this point, the dynamics of the two simulations diverge. The mixed model continues to show fluctuations in the average density of the ICs and sharp drops in the average local fitness of the CAs for short periods of time. In the base model a different evolutionary phase unrolls. The fluctuations in the average IC density value become smaller, as do the fluctuations in the average local fitness of the CA population. The CAs, however, no longer attain maximum local fitness although they did initially, and continue to do so in the mixed model.

3.2 Dynamics of IC densities

In order to understand the different evolutionary dynamics in the base model and in the mixed model we look at the distribution of ICs in the population. Here, we already see that in terms of their phenotype the IC population is homogeneous in the mixed model but heterogeneous in the base model.

Figure 4 shows the distribution of the densities of all ICs in the mixed model (A) and the base model (B) between $t=2100$ and $t=2200$ and the average local fitness of the CAs. In fig. 4A the population of ICs switches back and forth between high and low density values. At $t=2100$ the average local fitness of the CAs is very low, and the ICs experience only a strong selection toward lower density values due to the fitness function Φ (fig. 1). As soon as CAs arise that classify

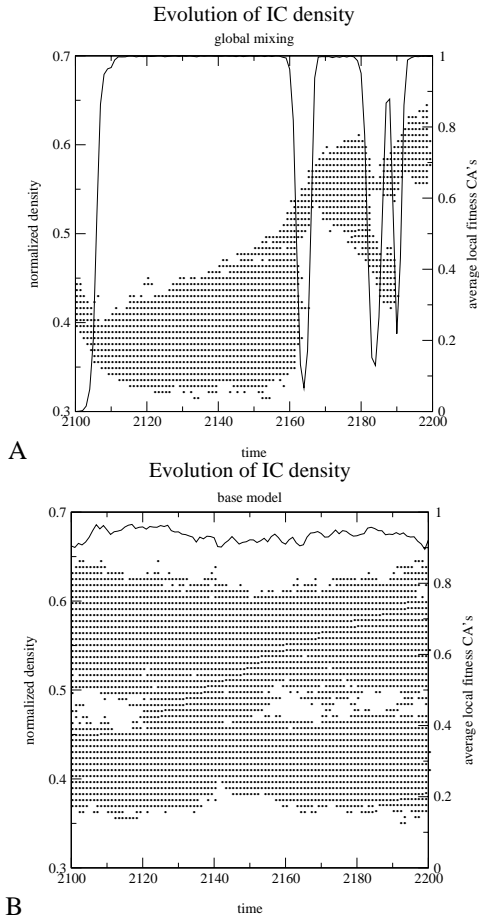


Figure 4: Evolution of IC density over 100 time steps in mixed model (A), and base model (B). The densities of all individuals are plotted as dots, the average local fitness values are plotted as thin lines. The population of ICs in the mixed model switches en masse from their density class, also characterized by short drops in average CA fitness. In the base model two subpopulations of ICs exist.

the ICs correctly, here $t \approx 2110$, the density distribution of the IC population starts to broaden considerably. When all CAs classify all ICs correctly the ICs experience neutral selection rather than selection toward low density values. The mutational bias brings about the drift of ICs toward density values of 0.5. As soon as ICs arise with a density larger than 0.5 the IC population jumps from class 0 to class 1 en masse and the same picture is seen again.

In fig.4B we see that in the base model the IC population has speciated into two distinct subpopulations of ICs, with densities around 0.4 and 0.6, which stably coexist. Thus, at the phenotypic level the diversity of ICs is much larger in the base model than in the mixed model. Nevertheless, the number of unique IC genotypes in the mixed model is actually larger than in the base model. For the simulations discussed above we counted the number of unique genotype in 15 generations between $t=600$ and $t=2000$, at every 100th time step.

In the mixed model on average 454.1 unique genotypes are present ($\sigma = 83.1$), in the base model on average 369.3 unique genotypes are present ($\sigma = 14.1$). For the CAs we find in the mixed model an average of 321.8 ($\sigma = 88.9$), and in the base model an average of 179.7 ($\sigma = 17.0$). Thus, in the mixed model we find much more unique genotypes than in the base model. However, this may largely result from the long periods of neutral selection in the mixed model.

3.3 Evolution and maintenance of diversity

Rather than looking at the number of unique genotypes we can also look at how genotypes are distributed over the genotype space. In the lower panel in fig.5 we have plotted the hamming distance of ICs in the base model within and between the subpopulations of the different density-classes. Whereas the hamming distance between ICs of different subpopulations is expected to be relatively large, the hamming distance of ICs of the same subpopulation of one generation is also very large. In fact, the distance between ICs of different subpopulations peaks near the distance which is expected between two random ICs, i.e. 75 bits; within a subpopulation the hamming distance distribution still peaks around 60.

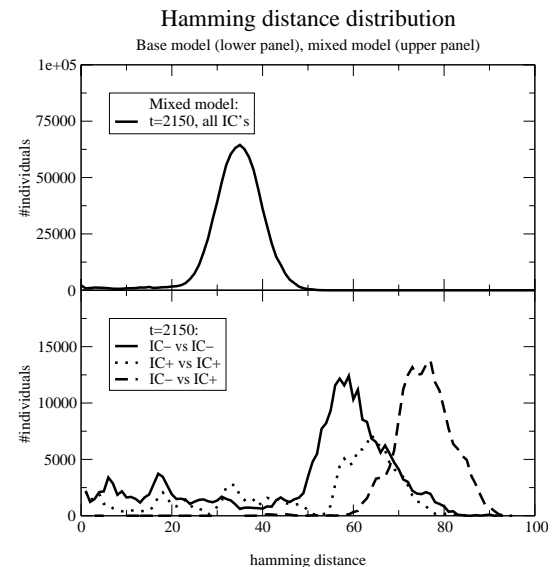


Figure 5: Distributions of the hamming distance between ICs of base model (lower panel) and between ICs of mixed model (upper panel), at generation $t=2000$. Lower panel: hamming distances are calculated between all ICs with a density lower than 0.5 (solid), all ICs with a density higher than 0.5 (dotted), and between ICs of different subpopulations (dashed). Upper panel: hamming distances are calculated between all ICs of the population.

In the top panel we have plotted the hamming distance of all individuals in a population of ICs in the mixed model at $t = 2150$. These ICs have a large range of density values and have undergone a long period of neutral selection (fig.4A).

Nevertheless, the hamming distance of the ICs peaks at much lower values than in the base model.

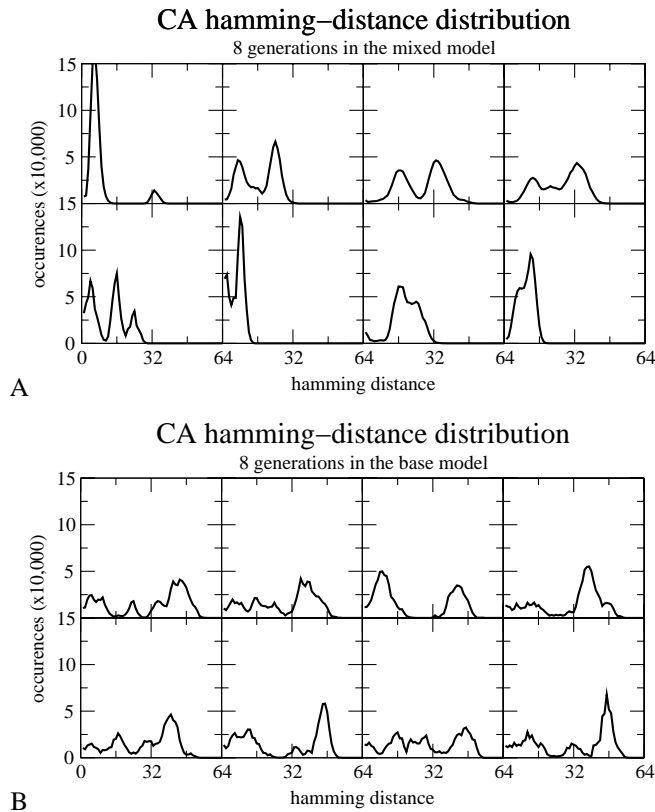


Figure 6: Hamming distance distributions in base model (A) and in mixed model (B). For each model the distribution of eight generations are plotted from $t=1300$ (lower left) to $t=2000$ (upper right) for every 100 time steps. CAs in the base model are always widely distributed with many CAs having hamming distances around 40 bits. In the mixed model the distribution varies per generation but peaks generally at lower values and sometimes at extremely low values.

Note that in the base model some ICs within subpopulations have small hamming distances. In the mixed model almost all ICs are at least 20 bits apart from all other ICs. This suggests that in the mixed model all ICs are distributed over a large area in the genotype space. In the base model, however, the ICs are distributed in a few distinct clusters.

In Fig. 6 we show hamming distance distributions of the CAs in the mixed model (6A) and in the base model (6B). For each model we plotted the hamming distance distribution of 8 generations, between $t=1300$ and $t=2000$, every 100 time steps. In the base model we see that the distribution always has a large peak for large hamming distances as well as appreciable numbers of small hamming distances. In the mixed model the distribution tends to center around one or two peaks, the median of which varies considerably. Again, in the base model CAs are distributed in a number of distinct clusters. In the mixed mode, on the other hand, the CAs

are sometimes focussed in a small region in genotype space, sometimes they form clusters as well (see also below). This results from the evolutionary dynamics that the CAs experience in the mixed model; long periods of neutral selection separated by evolutionary bottlenecks and short periods of strong selection.

When we perform a cluster analysis on the genotypes of CAs in the mixed and the base model we do not find clearcut differences between the two models, although further analysis is needed. When we perform a cluster analysis on the *phenotypes* of CAs we find a clear difference between the two models. For this purpose the phenotype of a CA is defined as the classifications it makes on a set of 500 initial conditions. The initial conditions of this set have density values that are uniformly distributed between 0.4 and 0.6. The phenotype of each CA is determined on the basis of the same set of initial conditions.

In the base model the cluster analysis shows that often a small number of distinct clusters exist, optimal numbers being mostly between four and nine. A cluster analysis on the phenotypes of CAs in the mixed model shows that indeed only one large cluster exists: all CAs behave equally “single-mindedly” in their classification behavior.

3.4 Spatial pattern formation




Figure 7: A: Snapshots of spatial distribution of the IC population at 5 consecutive time steps in the base model. ICs of density class 0 are depicted in black, ICs of density class 1 are depicted in grey. B: Space-time plot of IC population over 180 time steps, time going from left to right. ICs of different density classes are distributed in complex wave patterns which overtake each other continuously.

The question that arises is how the phenotypic diversity of the CAs and the ICs and the distinctiveness of the species into which they cluster is maintained in the base model. Figure 7A shows five snapshots of consecutive time steps of the spatial distribution of the ICs of different density classes. The ICs of class 0 are colored black and the ICs of class 1 are grey. The ICs are distributed in many small patches rather than in only a few large patches. In fig.7B we show a space-time plot of the ICs over a period of 180 time steps in which we plot a vertical cross-section of the grid at consecutive time steps. The space-time plot shows that complex wave patterns are

present; patches of black ICs grow into patches of grey ICs, and vice versa. As a consequence, at any one point in space ICs of the two density classes alternate frequently. This alternation of the two density classes is not primarily a result of mutation, which causes the global oscillations of the average IC density in the mixed model, but it is a result of spatial dynamics. In the base model, ICs ‘chase’ CAs not only in genotypes-space, as in the mixed model, in addition they ‘chase’ them in space-space.




Figure 8: Snapshots of the spatial distribution of the CAs in the base model. The color coding is based on cluster analysis of the CAs. In A) and B) the cluster analysis is based on the phenotype of the CAs (see text), in C) and D) the cluster analysis is based on the genotype of the CAs. A) and C) are snapshots at $t=400$, B) and D) are snapshots at $t=900$. CAs that belong to a single cluster are distributed in small patches in several locations in space.

We studied the spatial dynamics of the CAs in the base model using a color coding which shows the results of the cluster analysis discussed above (fig. 8). The CAs are clustered according to their phenotype (8A and 8B) or their genotype (8C and 8D), each at two time steps. Similar to the spatial distribution of the ICs, also in this case the CAs belonging to a single cluster occur in patches, and sometimes in more than one.

Given the results presented above the following picture arises of the eco-evolutionary process as it occurs in the base model. At the population level the CAs and ICs are split in several distinct clusters, or species, which come about by local specialization. A species is present in one or more patches in the field. At the boundaries of the patches species compete with other species, but the number of players in a single competitive interaction is small, typically two or three. Furthermore, each competition occurs in the context of the op-

ponent population. This opponent population is also heterogeneously distributed in space. Thus, competition between species within one population occurs, in parallel, in different contexts. Because species are distributed in a number of patches they will generally ‘see’, i.e. compete with, all other species simultaneously, as in the mixed model or other models with global interaction structures.

However, the main difference between the base model and the mixed model is that in the base model at the individual level the interactions are localized in space and time; during some time individuals interact only with individuals of one type, or of very few types. Also the short-term effect of the coevolutionary process remains local in space; individuals can have prolonged interactions with a small number of other individuals, thus giving the opportunity of local specialization. Despite (or thanks to; Pagie (1999)) the lack of local diversity we see evolution of general behavior, i.e. behavior that is successful under many different circumstances.

4 Conclusion

We have studied a coevolutionary model of two antagonistically interacting species. We compared the evolutionary dynamics that occur if individuals remain localized in space, i.e. when spatial pattern generation occurs, and the evolutionary dynamics that occur if individuals are globally mixed every time step. In the first case we find that individuals evolve a generalized response to environmental circumstances, whereas in the second case the systems exhibit evolutionary oscillatory dynamics. In that case we see the evolution of much simpler behavior, which is optimized with respect to one of the possible states of the other species. This strategy makes them easily exploitable, however. As a result we see red queen dynamics where both coevolving species oscillate between two states.

We found that the diversity of the populations of CAs and ICs differs greatly in the two models, albeit differently at different levels. The mixed model shows higher numbers of unique genotypes in the population, whereas the diversity on the level of phenotypes is almost completely absent. Also, the diversity that we find in the base model, both at the level of the genotypes as at the level of the phenotypes, is structured, i.e. it occurs in distinct clusters, or species. This variety in the base model, however, is distributed over time and space. At the level of the individuals the interactions and selection pressures may be of a single type for extended periods of time. Locally, individuals experience sparse fitness evaluation, i.e. they ‘see’ only a few fitness cases which are often of a single type, but overall we see the evolution of general behavior, i.e. information is integrated at the individual level (Pagie, 1999).

Other studies of coevolutionary optimization models that include techniques to maintain high diversity in populations also show evolution of general behavior (e.g. Paredis (1995); Rosin & Belew (1997); Juillé & Pollack (1998)). In these models individuals still interact at a global level, i.e. every

individual ‘sees’ all other individuals and diversity is not necessarily preserved in terms of distinct species. In our model speciation is the result of the *local* eco-evolutionary dynamics and diversity is therefor ‘cheap’. It is an open question to what extent the evolution of diversity and speciation depends on parameters of the evolutionary process, and how it influences the evolution of general behavior. It will be interesting to compare the evolution of diversity, speciation and general behavior under different parameter settings and also to models that use other techniques.

Bibliography

- BOERLIJST, M. C. & HOGEWEG, P. (1991); Selfstructuring and selection: spiral waves as a substrate for evolution. In: Langton, C. G. (ed.), *Artificial Life*. Addison-Wesley, Redwood City, CA, vol. 2, pp. 255–276, pp. 255–276.
- BULLOCK, S. G. (1995); Co-evolutionary design: Implications for evolutionary robotics. Cognitive Science Research Paper 384, School of Cognitive and Computing Sciences, University of Sussex, Brighton.
- COLLINS, R. J. & JEFFERSON, D. R. (1991); Selection in massively parallel genetic algorithms. In: Belew, Richard K.; Booker, L. B. (ed.), *Proceedings of the 4th International Conference on Genetic Algorithms*. Morgan Kaufmann, San Diego, CA, pp. 249–256, pp. 249–256.
- CRUTCHFIELD, J. P. & MITCHELL, M. (1995); The evolution of emergent computation. *Proceedings of the National Academy of Sciences* **92**:10742–10746.
- FICICI, S. G. & POLLACK, J. B. (1998); Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In: Adami, C.; Belew, R. K.; Kitano, H. & Taylor, C. (eds.), *Proceedings of the 6th International Conference on Artificial Life (ALIFE-98)*. MIT Press, Cambridge, MA, USA, pp. 238–247, pp. 238–247.
- HILLIS, D. W. (1990); Co-evolving parasites improve simulated evolution in an optimization procedure. *Physica D* **42**:228–234.
- HORDIJK, W.; CRUTCHFIELD, J. P. & MITCHELL, M. (1998); Mechanisms of emergent computation in cellular automata. In: *PPSN-V*.
- HUSBANDS, P. (1994); Distributed coevolutionary genetic algorithms for multi-criteria and multi-constraint optimization. In: Fogarty, T. (ed.), *Evolutionary Computing, AISB Workshop Selected Papers*. Springer-Verlag, vol. 865 of *Lecture Notes in Computer Science*, pp. 150–165, pp. 150–165.
- JUILLÉ, H. & POLLACK, J. B. (1998); Coevolving the ideal trainer: Application to the discovery of cellular automata rules. In: Koza, J. R.; Banzhaf, W.; Chellapilla, K.; Deb, K.; Dorigo, M.; Fogel, D. B.; Garzon, M. H.; Goldberg, D. E.; Iba, H. & Riolo, R. (eds.), *Genetic Programming 1998: Proceedings of the Third Annual Conference*. Morgan Kaufmann, University of Wisconsin, Madison, Wisconsin, USA, pp. 519–527, pp. 519–527.
- MAHFOUD, S. W. (1995); A comparison of parallel and sequential niching methods. In: Eshelman, L. J. (ed.), *Proceedings of the 6th International Conference on Genetic Algorithms*. Morgan Kaufmann Publishers, San Francisco, pp. 136–143, pp. 136–143.
- MITCHELL, M.; CRUTCHFIELD, J. P. & DAS, R. (1996); Evolving cellular automata with genetic algorithms: A review of recent work. In: *Proceedings of the First International Conference on Evolutionary Computation and its Applications (EvCA'96)*.
- MITCHELL, M.; CRUTCHFIELD, J. P. & HRABER, P. T. (1994); Evolving cellular automata to perform computations: Mechanisms and impediments. *Physica D* **75**:361–391. SFI Working Paper 93-11-071.
- PAGIE, L. (1999); *Information integration in evolutionary processes*. Ph.D. thesis, Bioinformatics group, Utrecht University.
- PAGIE, L. & HOGEWEG, P. (1997); Evolutionary consequences of coevolving targets. *Evolutionary Computation* **5**(4):401–418.
- PAGIE, L. & HOGEWEG, P. (2000); Coevolutionary dynamics: information integration, speciation and red queen dynamics. in prep.
- PAREDIS, J. (1995); Coevolutionary computation. *Artif. Life* **2**:355–375.
- PAREDIS, J. (1997); Coevolving cellular automata: be aware of the red queen! In: Baeck, T. (ed.), *Proceedings of the 7th Int. Conference on Genetic Algorithms (ICGA 97)*, pp. 393–400, pp. 393–400.
- ROSIN, C. D. & BELEW, R. K. (1997); New methods for competitive coevolution. *Evolutionary Computation* **5**(1):1–29.
- SAVILL, N. J. & HOGEWEG, P. (1997); Evolutionary stagnation due to pattern-pattern interactions in a co-evolutionary predator-prey model. *Artificial Life* **3**:81–100.
- TOFFOLI, T. & MARGOLUS, N. (1987); *Cellular Automata Machines: A new environment for modeling*. Scientific Computation. MIT Press, Cambridge, Massachusetts.
- WERFEL, J.; MITCHELL, M. & CRUTCHFIELD, J. P. (1999); Resource sharing and coevolution in evolving cellular automata. Submitted to IEEE Trans. Evol. Comp.
- WOLFRAM, S. (1984); Cellular automata as models of complexity. *Nature* **311**:419–424.