

# Co-evolutionary dynamics on a deformable landscape

**Marc Ebner**

Universität Würzburg  
Lehrstuhl für Informatik II  
Am Hubland  
97074 Würzburg, Germany  
ebner@informatik.uni-  
wuerzburg.de

**Richard A. Watson**

Brandeis University  
Volen Center for Complex  
Systems  
Mail Stop 18, Waltham,  
MA 02454-9110, USA  
richardw@cs.brandeis.edu

**Jason Alexander**

University of California, Irvine  
Logic & Philosophy of Science  
School of Social Science  
Irvine, CA 92697, USA  
jalex@uci.edu

**Abstract-** In order to use co-evolutionary techniques successfully one needs to investigate the dynamics of co-evolution. Assuming an open-ended evolutionary process, it would be desirable to establish the necessary and sufficient conditions which lead to a kind of arms race where species continually adapt in response to one another. In this paper we present a model of competitive co-evolution which is intended to investigate these conditions. In our model, co-evolving species are modeled as points which are placed randomly on a uniform landscape which is deformed by the species. The impact a species induces on its surrounding is not immediate. Instead, the deformation follows the species after some latency period. Evolution is modeled as a simple hill climbing process of the species. We investigate different conditions and their impact on the evolutionary dynamics. Some lead to stasis, some lead to cyclic behavior and others lead to an arms race.

## 1 Introduction

Most research in evolutionary algorithms has been done for a population of individuals moving on a fixed landscape described by a fitness function. Recent work also considers the use of dynamically changing fitness functions to evaluate the quality of an evolutionary algorithm to adapt to a new environment [11, 7]. In our research we are investigating evolutionary dynamics of co-evolving species.

Fitness landscapes of co-evolving species are coupled in that any move made by one species can have a significant impact on the fitness of another species [9]. For our work we have placed the individuals on a single fitness landscape as shown in Figure 1. It is possible that, using co-evolution, both sides may be able to reach a higher optimum than either side alone would be able to reach.

Co-evolution has been used successfully in sev-

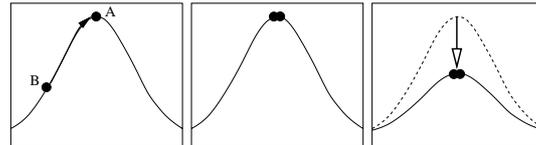


Figure 1: Two species coupled by a shared fitness landscape. In the example depicted above species A has reached a local fitness peak. Species B is adapting to the same niche. After species B has arrived at the optimum its presence begins to depress the fitness of both species.

eral different areas. Hillis [8] used co-evolution to evolve sorting networks. Reynolds [12], Miller and Cliff [10], and Cliff and Miller [1, 2] evolved predator-prey behaviors using computer simulations. Floreano and Nolfi [4] and Floreano et al. [5] evolved predator-prey behaviors for two Khepera robots. Funes et al. [6] used co-evolutionary techniques to evolve artificial game players. Kauffman [9] has developed a model for co-evolution and investigated the effect that the number of epistatic links between genes has on the smoothness or ruggedness of the landscape. Kauffman used coupled but different fitness landscapes for each species. In contrast to Kauffman's work we are using a single fitness landscape for all species.

## 2 Co-evolution, arms races and the Red Queen hypothesis

Co-evolution may lead to a kind of arms race where both sides try to out-compete their opponent [3]. Species engaged in an arms race have to do all they can just to retain the status quo. The effect where the fitness of co-evolving species remains at the same level over evolutionary time is called the Red Queen hypothesis after a figure from Lewis Carroll's novel *Through the Looking Glass* [14, 13]. In Figure 2 the effect is visualized for a simple land-

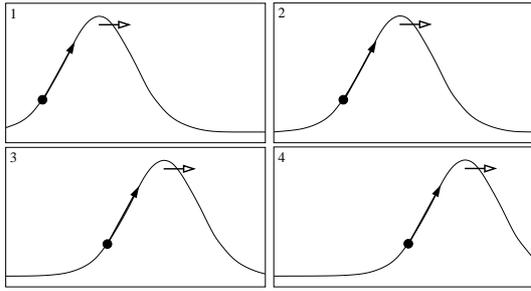


Figure 2: The Red Queen hypothesis: The fitness of co-evolving species may remain at the same level over evolutionary time. No matter how fast the species tries to climb the top, it never reaches it. As the species climbs towards the top, the landscape underneath the species deforms and it looks as if the species has remained exactly on the same spot.

scape. A species is located on one side of a hill and tries to reach the top of the hill. However, as the species climbs towards the top, the landscape underneath it deforms and it appears as if the species has not moved.

Due to the fact that the movement of one species can have a significant influence on the fitness of another species it is very difficult to measure the progress of co-evolution [1]. It could happen that the fitness values (as measured by reproductive success) of two co-evolving species remain at exactly the same value, nevertheless they both have become considerably more sophisticated in performing their task. So one of the questions currently under investigation in the field of co-evolutionary algorithms is how to measure evolutionary progress. Another area of particular importance is the dynamics of co-evolution. What are the conditions which lead to an arms race between two co-evolving species? What conditions lead to a stable attractor where all evolutionary progress comes to a halt?

We would like to establish conditions that give rise to an arms race where individuals continually adapt to one another. In the next section we discuss our model which was created to study co-evolutionary dynamics on a deformable landscape. It is hoped that the study of this model will eventually identify a set of necessary conditions for continuous evolution of increasingly complex individuals.

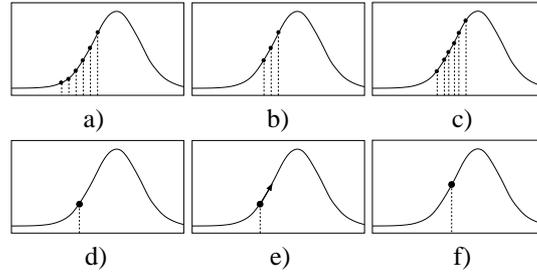


Figure 3: Modeling the movement of a population. The top row in the figure a), b), and c) shows a population climbing a fitness landscape. a) a population of individuals exist on the side of an incline. b) after selection, only highly fit individuals are retained, c) when these individuals reproduce their offspring cover a new area that is slightly higher up the fitness incline. The bottom row d), e), and f) shows the same population represented just as the population average as in our experiments. d) a point representing the population is placed at the average position of individuals in the initial population of figure a), e) the gradient of the landscape at this point is measured, f) the population is moved up the incline by a distance calculated by the measured gradient.

### 3 Co-evolutionary dynamics on a deformable landscape

In order to allow us to visualize the dynamics of co-evolution, the model has to be simple such that it can be calculated fast enough. Therefore we have modeled evolution as a hill climbing process (Figure 3). We do not use a population of individuals with reproduction and selection as would be found in the usual evolutionary model. Instead, we model only the population average of each species and this average is updated by gradient ascent. Thus, this process of replication, variation and selection is modeled as a single step of the species. In this way, each population is represented by a single scalar value - the position of the population in phenotype space - in this case, its horizontal position in our one dimensional world. If there are  $n$  species the whole model requires  $n$  scalar values. The equations below give various methods of update that we use to calculate the movement of the populations. Each provides a method to estimate the movement of a real population. We experimented with three different update rules.

The first update rule simply moves the species to the left or right depending on the sign of the land-

scape's gradient.

$$\dot{x}(t) = \begin{cases} -1 & \text{if } \frac{\partial}{\partial x}l(x, t) < 0, \\ 0 & \text{if } \frac{\partial}{\partial x}l(x, t) = 0, \\ 1 & \text{if } \frac{\partial}{\partial x}l(x, t) > 0 \end{cases} \quad (1)$$

The second update rule sets the velocity proportional to the gradient.

$$\dot{x}(t) = \alpha \frac{\partial}{\partial x}l(x, t) \quad (2)$$

The third update rule integrates the gradient with a damping factor. The following rule is used to update velocity in this case:

$$\dot{x}(t) = \alpha \left( \frac{\partial}{\partial x}l(x, t) \right) + \beta \dot{x}(t-1) \quad (3)$$

where  $l(x, t)$  specifies the height of the landscape at time  $t$ . In order for this to be biologically plausible one has to assume that although the external mutation rate is constant, the phenotypic effect the mutation has on the species is able to vary over time. This could be the case if the genetic code also evolves along with the species. That is, species that climb up faster (those with the better code) reproduce more because they reach higher fitness values faster.

The different species are placed on the same fitness landscape and they interact through deformations of this landscape. For the results which are presented here, we have used a one dimensional phenotype space where fitness is represented as the height of the landscape. However, the model can also be expanded easily to an N-dimensional phenotype space. The goal of each species is to climb to the highest peak it can find and remain there. However, each species deforms the landscape in its immediate vicinity. The deformation can be arbitrarily shaped. For our experiments we mostly used a Gaussian. The use of a Gaussian distribution of deformation models the distribution of a population around the population average assuming Gaussian mutation. The Gaussian deforms the landscape much like a heavy ball placed on a rubber landscape. The heavier the ball the larger the deformation. We used the same type of deformation for all species. Thus, each species corresponds to a point with unit mass. The deformation becomes larger as more and more species occupy the same position on the landscape.

Interesting behavior emerges if the deformation caused by the species is not immediate but occurs with a latency period. That is, whereas in the non-latent model, the deformations of the landscape are

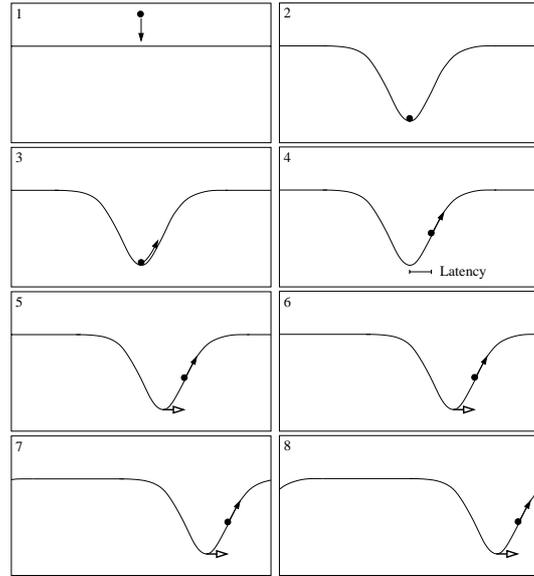


Figure 4: Landscape deformation. We start with a completely flat landscape. If a species is placed on this landscape it deforms the landscape in its vicinity. This is similar to the deformation one gets by placing a heavy ball on a rubber landscape. The heavier the ball the larger the deformation. In our model each species has a unit mass and the influence of the species on the environment is not immediate. Instead a latency of some time steps is used. That is, the deformation follows the current position of the species with some delay. Evolution is modeled as a step towards the top of the hill. Due to the latency, the species is able to begin its climb out of the local minimum. As the species tries to climb to the top the deformation follows the species which leads to the Red Queen effect.

calculated from the current positions of the species, in the latent model, the deformations are calculated from the position that the species had occupied some number of time steps in the past. In this way, a species has a chance to climb a peak, for example, *before* that same peak is depressed by its own presence. However, if the species remains still for a while, or moves slowly, then the deformation that is following behind the species will take effect on it. An analog in nature would be the use of resources which are available but are depleted after some time. For our experiments we have always used a constant latency which was the same for all species. The use of a latency period leads to the Red Queen effect as shown in Figure 4. As the species tries to climb the top of the hill the deformation follows the species and it seems as if no

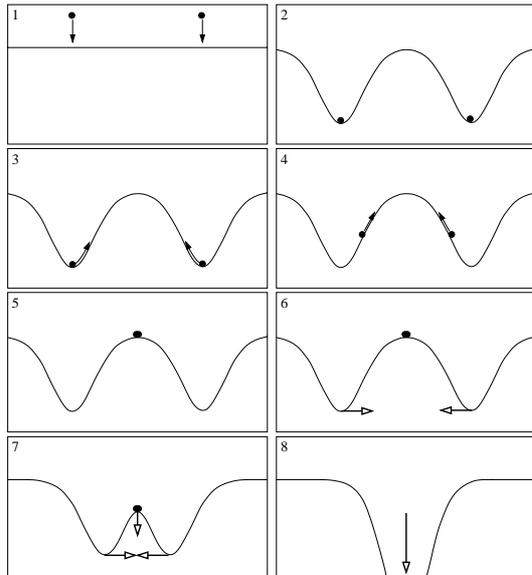


Figure 5: A local optimum is formed in between two species if they are placed next to each other. For the sequence shown here it is assumed that the latency is rather long. In this case both species are able to reach the local optimum if they happen to move towards the local optimum which lies in between the two species. As they are staying put on the local optimum the deformation of the landscape follows them which seriously reduces their fitness. The species have no choice but to wait until they have reached the bottom of the valley and are able to climb up again.

progress has been achieved. Although the species has moved through phenotype space it still has exactly the same fitness as before.

Figure 5 shows what can happen if two species happen to climb the same hill. If two species are placed next to each other and if they both happen to climb the hill in between them they will first reach a local optimum. After some time the deformation follows them which leads to a depression of the local optimum. They have no choice but to wait until they have reached the bottom of the valley before they are able to climb up again.

In summary, each species is described by the following parameters: its position on the landscape, an update rule for the next position, a function which specifies the deformation caused by the species, and an update rule which describes how the deformation follows the species. The deformation, that is the impact of one species on other species, could also be made dependent on additional external parameters.

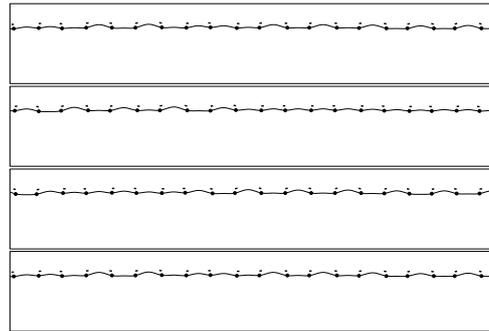


Figure 6: Experiment 1 uses a velocity update rule where the velocity is determined by the sign of the environment's gradient. The parameters of experiment 1 may lead to the small cyclic behavior shown here. Lines above the species represent the species' current velocity.

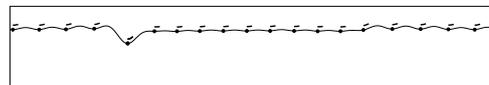


Figure 7: Experiment 1 may also lead to a synchronous shift if two species happen to end up close to each other.

## 4 Experimental results

A number of experiments were performed to investigate the dynamics of co-evolution. The different experiments are shown in Table 1. For all our experiments 20 species were used. Each species caused a Gaussian deformation. The update rule, latency and type of environment were varied and the resulting effects on co-evolution analyzed. As update rules we used the three different types of rules which were described above. We also experimented with different types of environments in particular a completely flat environment and an environment where some Gaussian hills were distributed at random over the landscape. The deformation caused by the hills is analogous to the deformation caused by the species themselves except that the deformation is caused upward and the hills are stationary.

The first update rule simply moves to the left or right depending on the sign of the gradient. For experiment 1 we used update rule (1) to update the velocity of the species. A latency of 0 with a completely flat environment was used. We either observed a small cyclic behavior as shown in Figure 6 or a synchronous shift as shown in Figure 7.

For experiment 2 we used a latency of 50 and

Experiment	Species	Update Rule	Deformation	Latency	Hills	Observed Behavior
1	20	(1)	Gaussian	0	0	cyclic or shift
2	20	(1)	Gaussian	50	0	clumped shift
3	20	(2)	Gaussian	0	0	stasis
4	20	(2)	Gaussian	50	0	cyclic or arms race
5	20	(3)	Gaussian	0	0	stasis
6	20	(3)	Gaussian	2	0	cyclic
7	20	(3)	Gaussian	50	0	arms race
8	20	(3)	Gaussian	0	50	stasis
9	20	(3)	Gaussian	50	50	arms race

Table 1: Experiments were made with different update rules, latency periods and type of environments.

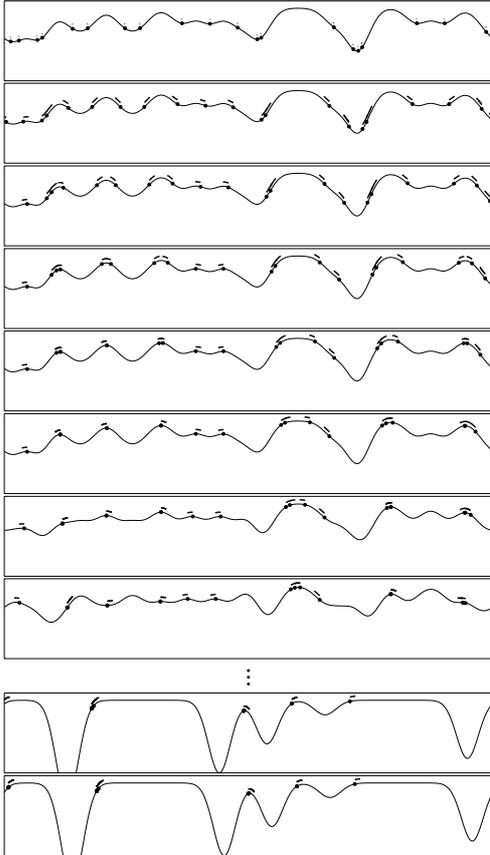


Figure 8: Experiment 2 uses a velocity update rule where velocity is set according to the sign of the landscape's gradient. The parameters of experiment 2 may lead to the clumped shifting behavior.

observed the clumped shifting behavior as shown in Figure 8. With a large latency we are more likely to get a clumping behavior because the species are able to climb the local hills before the deformation sets in.

For experiment 3 we used update rule (2) with  $\alpha = 10$  to update the velocity of the species. No la-

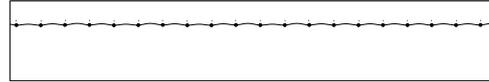


Figure 9: Experiment 3 uses a velocity update rule where velocity is always equal to the gradient of the environment. The small dots above the species indicate that the species are almost standing still. No further improvement is possible. Evolutionary space is only partially explored.

tency and a completely flat environment was used. The species spread over the landscape which leads to an almost stationary state with little movement. The resulting state of experiment 3 is shown in Figure 9. After the species have spread apart, which can be viewed as each species having found a niche where they do not interfere very much with the other species, no further improvement is possible. Exploration of evolutionary space is no longer possible in this case.

For experiment 4 we used update rule (2), a latency of 50, and a completely flat environment. We either observed a cyclic behavior as shown in Figure 10 or an arms race as shown in Figure 11. The movement of the species on the landscape was rather slow. That is why we used update rule (3) for the remainder of the experiments. This update rule produced much more interesting dynamics.

For experiment 5 we used update rule (3) with  $\alpha = 9$  and  $\beta = 0.9$ , no latency and a completely flat environment. Again we observed the same behavior as shown in Figure 9 for experiment 3.

For experiment 6 we increased the latency from 0 to 2. This may lead to the cyclic behavior shown in Figure 12. Neighboring species are trying to climb the same hill. The deformation follows them closely and new hills are created at the position where they came from. This leads to cyclic behavior. Again no further improvement is possible

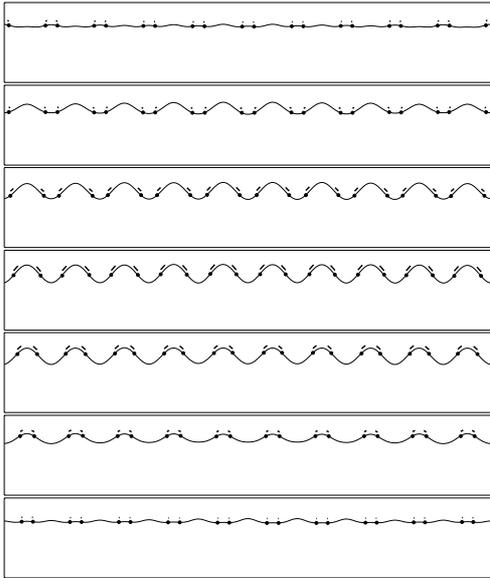


Figure 10: Experiment 4 uses a velocity update rule where velocity is always equal to the gradient of the environment. The parameters of experiment 4 may lead to the cyclic behavior shown here.

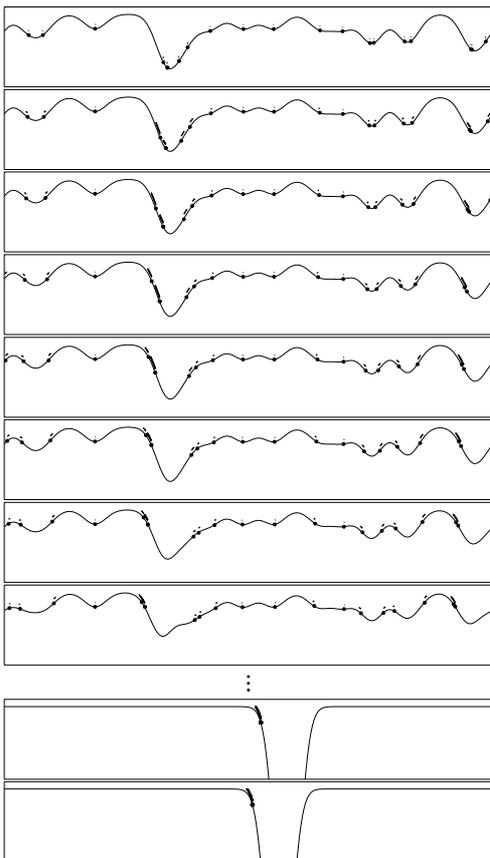


Figure 11: Experiment 4 may also lead to an arms race.

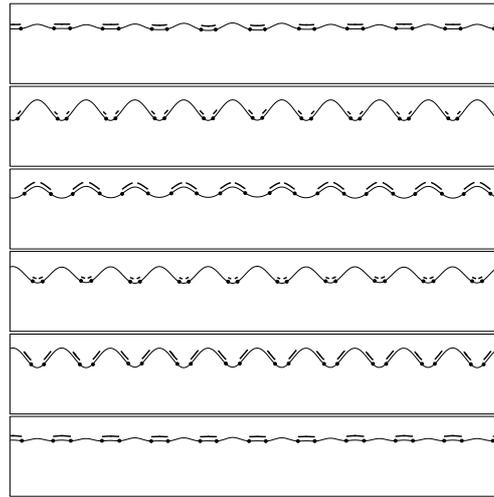


Figure 12: Attractor of experiment 6 . Groups of two species try to climb the same hill which leads to a cyclic behavior. Again no further improvement is possible. Evolutionary space is only partially explored.

and evolutionary space is only partially explored. Experiment 6 may also lead to different behaviors where the species are not evenly spread over the environment.

For experiment 7 we increased the latency from 2 to 50. A typical run of experiment 7 is shown in Figure 13. This time we again experience an arms race. Initially species are distributed at random over the landscape. However some will be closer than others. If two species happen to fall on the same side of a deformation they start running away from the valley but the valley follows them. After a while more and more species are caught by this deformation which leads to the arms race.

Next we experimented with a non-zero environment. 50 Gaussian hills are distributed additively over the entire landscape. The resulting landscape is shown in Figure 14. This landscape which could be deformed by the species, just as before, was used for experiments 8 and 9 .

For experiment 8 we used the same parameters as for experiment 5 except that we now used a non-zero environment. Again the species spread over the landscape and an almost stationary state results. The resulting state is shown in Figure 15. No further improvement is possible.

Finally, we increased the latency to 50. So experiment 9 uses the same parameters as experiment 7 except that a non-zero environment was used. The results are shown in Figure 16. Just as in ex-

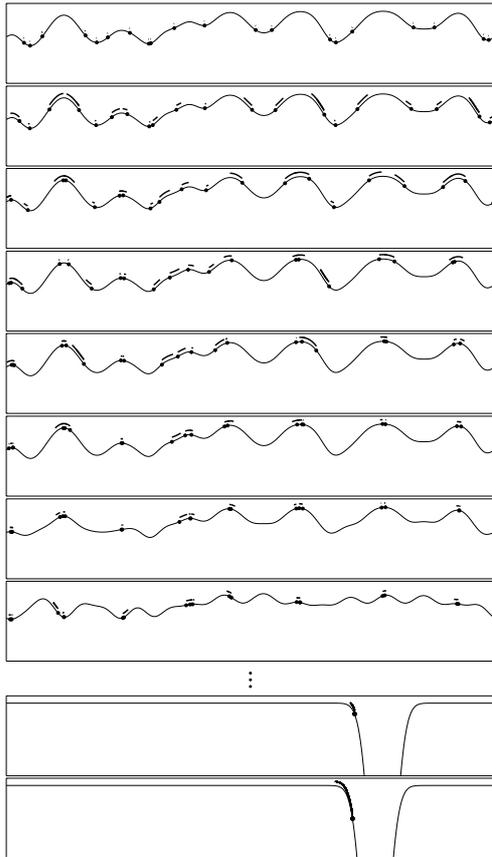


Figure 13: Attractor of experiment 7 . The parameters of the experiment have lead to an arms race.



Figure 14: Non-zero environment used for experiments 8 and 9 .

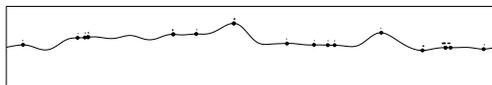


Figure 15: Experiment 8 is similar to experiment 5 except that a non-zero environment was used. The species spread over the landscape to avoid the negative influence of other species and to exploit and fitness advantages present in the landscape. As in experiment 5 species are almost stationary which allows no further improvement.

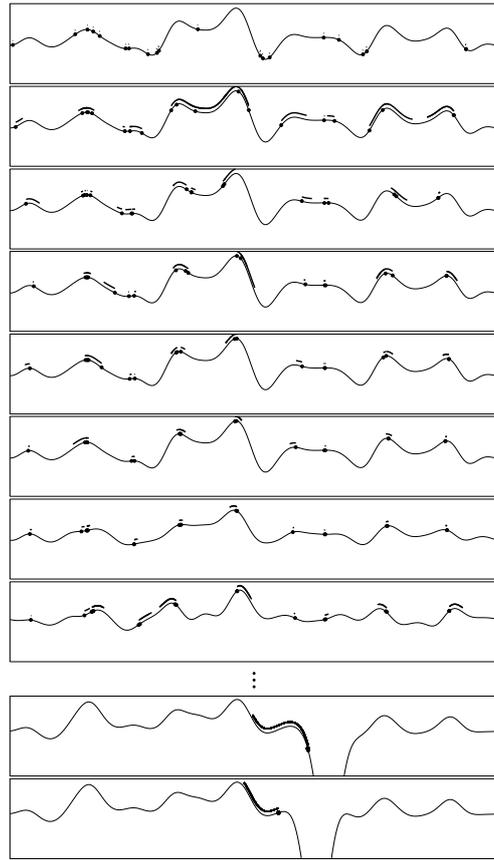


Figure 16: Experiment 9 is similar to experiment 7 except that a non-zero environment was used. Again an arms race results as was the case for experiment 7 .

periment 7 we observed an arms race. To get this type of behavior it is important that the deformation caused by the species is of comparable size as the hills. Only then is it possible for the species to move over the entire landscape.

## 5 Conclusion

We have developed a model to investigate the dynamics of competitive co-evolution. The model is simple to implement, easy to understand and can be calculated fast enough to allow real time visualization. Adaptation of a species is modeled as hill-climbing of the population mean. Species interact by deforming the landscape in their vicinity. The model allows us to study the dynamics of many co-evolving species. So far, we have experienced phenomena like even distribution of species over phenotype space as well as a kind of arms race where species clump together and race through all

of phenotype space. The determining factor in the switch between these two regimes was the latency with which a population affected the environment. The model is easily extensible to investigate other phenomena. In particular, it would be interesting to investigate the effects of different types of deformations (e.g. increase as well as decrease of the fitness landscape).

In order to use co-evolution as a successful search strategy which leads to increasingly complex individuals in an environment that allows open-ended evolution one needs to determine the necessary conditions which lead to a kind of arms race. So far we have been able to produce stasis as well as arms races with our model and latency seems to be an important factor.

## Acknowledgment

Part of this work was supported by the Santa Fe Institute, Santa Fe, NM.

## References

- [1] D. Cliff and G. F. Miller. Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. In F. Morán, A. Moreno, J. J. Merelo, and P. Chacón (Eds.), *3rd Europ. Conf. on Artificial Life*, pages 200–218, Berlin, 1995. Springer-Verlag.
- [2] D. Cliff and G. F. Miller. Co-evolution of pursuit and evasion II: Simulation methods and results. In P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson (Eds.), *From Animals to Animats 4: Proc. of the 4th Int. Conf. on Simulation of Adaptive Behavior*, pages 506–515, Cambridge, MA, 1996. The MIT Press.
- [3] R. Dawkins and J. R. Krebs. Arms races between and within species. *Proc. R. Soc. Lond. B*, 205:489–511, 1979.
- [4] D. Floreano and S. Nolfi. God save the red queen! Competition in co-evolutionary robotics. In J. R. Koza, K. Deb, M. Dorigo, D. B. Fogel, M. Garzon, H. Iba, and R. L. Riolo (Eds.), *Genetic Programming 1997: Proc. of the 2nd Int. Conf. on Genetic Programming, July 13-16, 1997*, pages 398–406, San Francisco, CA, 1997. Morgan Kaufmann Publishers.
- [5] D. Floreano, S. Nolfi, and F. Mondada. Competitive co-evolutionary robotics: From theory to practice. In R. Pfeifer, B. Blumberg, J.-A. Meyer, and S. W. Wilson (Eds.), *From Animals to Animats 5: Proc. of the 5th Int. Conf. on Simulation of Adaptive Behavior*, pages 515–524. The MIT Press, 1998.
- [6] P. Funes, E. Sklar, H. Juillé, and J. Pollack. Animal-animat coevolution: Using the animal population as fitness function. In R. Pfeifer, B. Blumberg, J.-A. Meyer, and S. W. Wilson (Eds.), *From Animals to Animats 5: Proc. of the 5th Int. Conf. on Simulation of Adaptive Behavior*, pages 525–533, Cambridge, MA, 1998. The MIT Press.
- [7] J. J. Grefenstette. Evolvability in dynamic fitness landscapes: A genetic algorithm approach. In *Proc. of the 1999 Congress on Evolutionary Computation*, volume 3, pages 2031–2038, Mayflower Hotel, Washington D.C., 6-9 July 1999. IEEE Press.
- [8] W. D. Hillis. Co-evolving parasites improve simulated evolution as an optimization procedure. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen (Eds.), *Artificial Life II, SFI Studies in the Sciences of Complexity*, pages 313–324. Addison-Wesley, 1991.
- [9] S. A. Kauffman. *The Origins of Order. Self-Organization and Selection in Evolution*. Oxford University Press, Oxford, 1993.
- [10] G. F. Miller and D. Cliff. Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In D. Cliff, P. Husbands, J. Meyer, and S. W. Wilson (Eds.), *From Animals to Animats III: Proc. of the 3rd Int. Conf. on Simulation of Adaptive Behavior*, pages 411–420, Cambridge, MA, 1994. The MIT Press.
- [11] R. W. Morrison and K. A. De Jong. A test problem generator for non-stationary environments. In *Proc. of the 1999 Congress on Evolutionary Computation*, volume 3, pages 2047–2053, Mayflower Hotel, Washington D.C., 6-9 July 1999. IEEE Press.
- [12] C. W. Reynolds. Competition, coevolution and the game of tag. In R. Brooks and P. Maes (Eds.), *Artificial Life IV*, pages 59–69, Cambridge, MA, 1994. The MIT Press.
- [13] M. L. Rosenzweig, J. S. Brown, and T. L. Vincent. Red Queens and ESS: The coevolution of evolutionary rates. *Evolutionary Ecology*, 1:59–94, 1987.
- [14] L. Van Valen. A new evolutionary law. *Evolutionary Theory*, 1:1–30, July 1973.