

Coevolutionary dynamics of interacting species

Marc Ebner¹, Richard A. Watson², and Jason Alexander³

¹ Eberhard Karls Universität Tübingen, WSI für Informatik, Abt. RA, Sand 1,
72076 Tübingen, Germany, marc.ebner@wsii.uni-tuebingen.de

² University of Southampton, School of Electronics and Computer Science, Highfield,
Southampton, SO17 1BJ, UK, raw@ecs.soton.ac.uk

³ London School of Economics, Department of Philosophy, Logic and Scientific
Method, Houghton Street, London WC2A 2AE, UK, J.Alexander@lse.ac.uk

Abstract. One of the open questions in evolutionary computation is how an arms race may be initiated between coevolving entities such that the entities acquire new behaviors and increase in complexity over time. It would be highly desirable to establish the necessary and sufficient conditions which lead to such an arms race. We investigate what these conditions may be using a model of competitive coevolution. Coevolving species are modeled as points which are placed randomly on a two-dimensional fitness landscape. The position of the species has an impact on the fitness landscape surrounding the phenotype of the species. Each species deforms the fitness landscape locally. This deformation, however, is not immediate. It follows the species after some latency period. We model evolution as a simple hill climbing process. Selection causes the species to climb towards the nearest optimum. We investigate the impact different conditions have on the evolutionary dynamics of this process. We will see that some conditions lead to cyclic or stationary behavior while others lead to an arms race. We will also see spontaneous occurrence of speciation on the two-dimensional landscape.

1 Introduction

Most research in evolutionary algorithms assumes that a population of individuals is adapting to a fixed fitness landscape as defined by the fitness function. However, some researchers also have looked at the problem of a population of individuals adapting to a fitness function which changes over time, e.g [5,9,16]. In nature, the environment to which species adapt over the course of time is never static. It is constantly changing because it mostly consists of other individuals which themselves undergo evolution. In our research, we are focusing on coevolving species and their evolutionary dynamics. We try to isolate important factors which have an impact on the coevolutionary dynamics of different species.

Let us assume that we have two species adapting to their local environment. The fitness landscapes of coevolving species are coupled. If one species achieves a higher fitness, i.e. higher reproductive rate, this also has a significant impact on the fitness of the second species [7]. Two coevolving species may well reach a higher optimum than either one alone. This may be the result of an arms




Fig. 1. Two species are coupled by a shared fitness landscape.

race which is induced between the two species [1]. An important question in this respect is: what are the necessary and sufficient conditions for an arms race between two or more coevolving species. For this work, we will speak of an arms race when two or more species try to out-compete each other.

In our simulations, the species are placed on a single landscape to which they have to adapt. The behavior of two species A and B is illustrated in Figure 1. Each species is driven towards a local optimum due to the selective pressure. Species B is moving towards the same optimum, i.e. is located inside the same niche. It is assumed that if the niche is occupied by two or more species then the fitness of these species is reduced.

Kauffman [7] suggested an abstract model for coevolution in order to investigate coevolutionary dynamics. In his model, epistatic links between genes have an influence on the smoothness or ruggedness of the landscape. Each species had a separate landscape which was coupled to other species. According to Kauffman, one of the conditions necessary for his coevolving system to accumulate beneficial traits and to maintain fitness at a high level is that the couplings within each species is roughly equal to the product of the coupling between species and the number of other species each species interacts with. In contrast to Kauffman's work, our model uses a single shared fitness landscape for all species.

Other authors have also looked at evolutionary dynamics in several different problem domains. For instance, Spector and Klein [12] visualized evolutionary dynamics of a set of agents each representing a particular species inside a three-dimensional environment. Observed behaviors included genetic drift, as well as emergence of collective and multicellularity. In their work, interaction does not occur through an explicit fitness landscape. Evolutionary dynamics in the context of biological games are described by Nowak and Sigmund [10]. Game dynamics of finite populations are considered by Taylor et al. [14]. Evolutionary dynamics on graphs are explored by Lieberman et al. [8].

Another focus is on the dynamics of coevolution. Can we provide an answer to the question what conditions lead to an arms race between two coevolving species and which don't? What conditions produce a stable attractor where no further progress is possible once the attractor has been reached? If we work with open ended evolutionary systems, e.g. self-reproducing programs [11], we don't want evolution to halt. Hence, it is important to understand which set of conditions is necessary and sufficient in order to establish an arms race where the complexity of the coevolving species continuously improves.

2 Coevolutionary Dynamics on a Deformable Landscape

In order to visualize the dynamics of coevolution, we need an appropriate model. The model has to be simple enough such that computations can be carried out quickly. Hence, we have modeled evolution as a hill climbing process on a fitness landscape. This concept was originally introduced by Wright [18]. Modeling an entire population of individuals with reproduction and selection as would be found in the usual evolutionary model would be much more costly. Instead, we model only the population average phenotype of each species. Gradient ascent is then used to update this average. Thus, the evolutionary process consisting of replication, variation and selection is modeled as a single step. Each population is represented by an n -dimensional vector which describes its position in phenotype space. In the simplest case, i.e. a one-dimensional fitness landscape, we just have a scalar value. This scalar value describes the position of the species inside a one-dimensional world. Assuming that our world consists of n species then we require n scalar values. Experiments on a one-dimensional fitness landscape have been carried out by Ebner et al. [3]. In contrast to these experiments, we will be using a two-dimensional landscape. We will see that an interesting phenomena emerges as we switch from a one to a two-dimensional fitness landscape.

We experimented with three different models on how the population moves through phenotype space. The first method computes the sign of the gradient of the fitness landscape where the species is located and then moves the species one step to the right or left depending on the sign of this gradient. Let $f(x, t)$ be the height of the fitness landscape at time t . Let a species be located at position $x(t)$ then we compute its the velocity $\dot{x}(t)$ using

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

The second update method computes the gradient of the fitness landscape and then sets the velocity of the species to the gradient multiplied by a constant factor α . Whereas in equation (1) the movement of the population through phenotype space has constant speed (and fitness information merely controls the direction of movement), equation (2) presents a model where rate of change in phenotype is proportional to fitness gradient. This is more in alignment with a classical model where rate of change of fitness is proportional to fitness variance assuming that variance in phenotype is constant.

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

The third update method also computes the gradient of the fitness landscape and integrates this gradient over time.

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

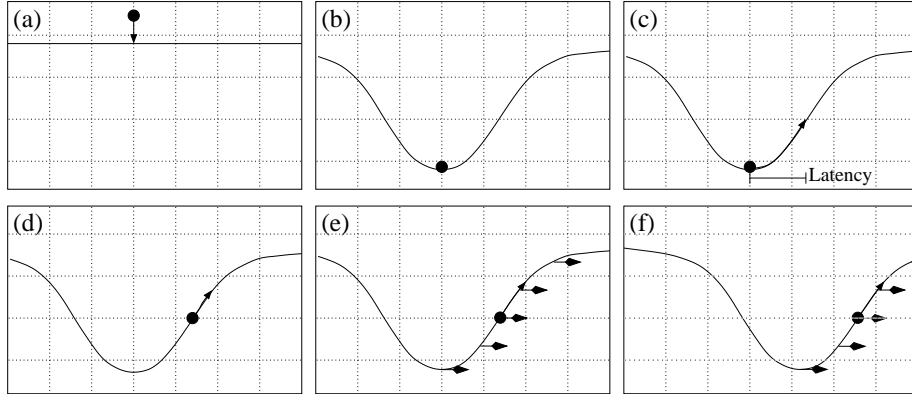


Fig. 2. Deformation of the fitness landscape.

This equation suggests that the mean of the population moves in response to the gradient at the current position inside the fitness landscape and also in response to the gradient at the previous time step. The effect of the previous location of the population mean can be interpreted as an effect arising from overlapping generations of individuals i.e. not all members of the population are replaced every generation, some are maintained from the previous generation. Such momentum terms are known to have interesting effects on many kinds of dynamical systems, and the experiments that follow indicate that if such an effect is active in biological populations then they may have significant effects on coevolutionary dynamics.

The different species are placed on the same fitness landscape Figure 2(a). Each species deforms the landscape in its vicinity Figure 2(b). We use a Gaussian shaped deformation. The shape of the deformation can be considered to be due to the distribution of a population around the population average assuming Gaussian mutation. The species has a negative impact on fitness in its vicinity. Such a negative impact on fitness is also used in evolutionary algorithms (where it is called sharing [4]) when trying to populate different local optima. We use the same type of deformation for all species. The deformation of this landscape gets larger as more and more species occupy the same position.

The species are pushed towards local optima by the selective force of evolution Figure 2(c). The deformation caused by the species can either have an immediate or a latent effect on the local fitness. A latent effect only becomes apparent after some time steps Figure 2(d-f). We observe quite interesting behavior if the deformation caused by the species occurs after a latency period of several time steps. The analog with natural evolution would be the excessive use of resources which are depleted after some time. Thus, we have basically two different modes. In the non-latent model, the deformation of the landscape is positioned wherever the species is located inside the fitness landscape. In the la-

tent model, the deformation is positioned wherever the species was located some number of time steps in the past. In this model, the species may climb toward a local optimum. However, after some time, this local optimum is depressed by the presence of the species and the species then has to adapt to a new optimum. We set the latency to a constant value for all species. Once we use such a latency period, we will observe the Red Queen effect [13,17] as shown in Figure 2(e-f). The species tries to climb to the top of a local optima. However, the deformation follows the species. It thus seems as if no progress has been achieved. The species still has the same fitness even though it moved through phenotype space.

In summary, our model has the following parameters. Each species is described by its position on the landscape, an update rule, which describes how the position of the species for the next time step is derived, a function describing the deformation caused by the species, and an update method which describes how the deformation follows the species. The deformation, that is the impact a species has on the fitness landscape, could also be made dependent on additional external parameters. However, in our view, the model should be made as simple as possible.

3 Experimental Results

Experiments were performed on a two-dimensional fitness landscape. We have used circular boundary conditions. MPEG as well as AVI movies of these experiments are available for download from ⁴. The source code for these experiments is also available for download.

Exp.	Species	Update Rule	Latency	Hills	Observed Behavior	Fig.
1	40	Eqn (1)	0	0	spread out then shift	3
2	40	Eqn (1)	50	0	clumped shift (arms race)	4
3	40	Eqn (2)	0	0	spread out then tiny shift	3
4	40	Eqn (2)	50	0	cyclic	5
5	40	Eqn (3)	0	0	spread out then shift	3
6	40	Eqn (3)	2	0	cyclic	6
7	40	Eqn (3)	50	0	arms race	7 & 8
8	40	Eqn (3)	0	100	stasis	9
9	40	Eqn (3)	50	100	arms race	10

Table 1. Parameter settings which were used for the experiments on the two-dimensional landscape.

The different parameter settings which were used are shown in Table 1. We used 40 species due to the fact that the size of the landscape has increased compared to the experiments on the one-dimensional landscape. A flat fitness

⁴ <http://www.ra.cs.uni-tuebingen.de/mitarb/ebner/welcome.html>

landscape was used for experiments 1 through 7 while experiments 8 and 9 used a non-flat fitness landscape. For the latter two experiments, 100 Gaussian hills (with the same standard deviation that was used for the depression caused by the species) were distributed uniformly over the entire area.




Fig. 3. Experiment 1, 3 and 5: synchronous shift of the species. The movement is very small for experiment 3.

Figure 3 shows the results for experiments 1, 3 and 5. For all three of these experiments, the species spread out evenly over the environment and then move over the environment. For experiment 2, we observe a clumped shift behavior with continued motion. Figure 4 shows the results. The number in the upper left hand corner shows the simulation time step. Eventually, the species synchronize their motion and all move in the same direction. We consider this to be an arms race type of movement as illustrated in Figure 2. The results of experiments 4, and 6 are shown in Figure 5 and 6 respectively. For both of these experiments, we observe a mixture between cyclic behavior and continued motion. This is best observed in the movie or by using the program code.




Fig. 4. Experiment 2: clumped shift behavior after convergence.

A latency of 50 for experiment 7 resulted in an arms race similar to that of the one-dimensional case. The first few steps and the resulting arms race are shown in Figure 7. However, it also shows a very interesting difference. Even though we work with separate species in our model, converged points in the search space can be considered to be a single species. In the two-dimensional environment, the racing species do not all arrive at the same point as they climb out of the depression - rather they spread-out across the rim as can be seen in Figure 8. This slight stochastic separation is exaggerated by the competitive relationships of the species dynamics and occasionally produces a division




Fig. 5. Experiment 4: a mixture between cyclic behavior and continued motion.




Fig. 6. Experiment 6: A latency of 2 produces cyclic behavior.




Fig. 7. Experiment 7: arms race between different species.




Fig. 8. Spontaneous speciation occurs during experiment 7. The species spread along the rim of the depression. This causes the creation of several sub-sets of species engaged in an arms race.




Fig. 9. (a) Non-flat environment (used for experiments 8 and 9). (b) Experiment 8: The species spread over the landscape to avoid the negative influence of other species and to exploit and fitness advantages present in the landscape.




Fig. 10. Experiment 9: An arms race occurs. The species can eventually sweep over all local optima.

- separate depressions each with a separate sub-set of species engaged in an arms race. This spontaneous speciation does not involve island models or other mechanisms of segregation such as reproductive isolation from one another [6] or mate preference due to marker traits [2]. In our model, speciation occurs through the coupled dynamics of the coevolving species. Thus, in our model sympatric speciation [15], speciation without isolation, is also possible without modeling preferences for ecological traits or preferences for certain marker traits of other individuals.

For experiments 8 and 9 we have used the non-flat environment (shown in Figure 9(a)). Experiment 8, which used a latency of 0, resulted in a state of stasis with small local oscillations (shown in Figure 9(b)). This figure shows very clearly how competitive coevolution is able to automatically promote a niching behavior that distributes the species so as to equalize the landscape - the almost flat resulting landscape (as compared to the underlying landscape shown in Figure 9(a)) indicates that every part of the fitness landscape is being used near optimally. Experiment 9 with a latency of 50 resulted in an arms race. The result of this experiment is shown in 10.

We now address the question of why the use of a latency factor causes the species to explore the entire landscape. Essentially species adapt to the nearest local optimum. Consider species adapting to a non-flat fitness landscape with a latency factor of zero. The species immediately deform the landscape but each species still adapts towards a local optimum. The result is a stationary landscape where the species are distributed over the landscape in such a way that any move of the species would result in a lower fitness value for that species. Now suppose that a sufficiently large latency factor is involved. If this is the case, then the species adapt towards the nearest optimum. However, due to the latency factor, the deformation caused by the species follows the species after some time. This local optimum invites other species to climb towards it as long as the combined negative influence of other species is smaller than the height of this optimum. Once several species have reached the same optimum and the combined negative influence of all of these species is larger than the height of this optimum, then the deformation pushes the species away from its previously optimal position in phenotype space. The species now have to climb towards new optima.

When several species climb towards the same optimum, they initially clump together before they are pushed away from that optimum. The result is a Red Queen type of arms race where the species are distributed along the rim of the depression. This arms race lasts until the species are sufficiently distant from each other. Note that the direction the species approach the optimum and also the exact timing of this approach influences the way the species are pushed away from the optimum. The result is that over time the species may reach any point of the landscape provided that in the model, the combined depression caused by the species is sufficiently large to push the species away from the highest optimum of the landscape.

4 Conclusion

We have developed a simple evolutionary model which is used to investigate the dynamics of competitive coevolution. Due to the simplicity of the model, it is easy to implement and easy to understand. Most importantly, it can be calculated fast enough which allows us to visualize the behavior in real time. Adaptation of a species to its environment is modelled as a local hill-climbing process which moves the population mean towards the nearest local optimum. Each species has an influence on the environment where it is located. In our model, each species deforms the fitness landscape and thereby interacts with other species located nearby. This allows us to study the dynamics of coevolving species. We have observed several qualitatively different behaviors. Depending on the choice of parameters, the species distribute evenly over phenotype space or engage in an arms race. In case of an arms race, all of phenotype space is explored. The determining factor which causes the switch between the two regimes was the latency parameter. The latency parameter determines the time lapse for a species to have an impact on its current environment. This model has

shown that such delays in adaptive interactions can have a significant effect on the dynamics of coevolving species.

References

1. Dawkins, R., Krebs, J.R.: Arms races between and within species. *Proc. R. Soc. Lond. B* 205, 489–511 (1979)
2. Dieckmann, U., Doebeli, M.: On the origin of species by sympatric speciation. *Nature* 400, 354–357 (Jul 1999)
3. Ebner, M., Watson, R.A., Alexander, J.: Co-evolutionary dynamics on a deformable landscape. In: *Proceedings of the 2000 Congress on Evolutionary Computation*, San Diego, CA, Vol 2. pp. 1284–1291. IEEE Press (2000)
4. Goldberg, D.E.: *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Company, Reading, MA (1989)
5. Grefenstette, J.J.: Evolvability in dynamic fitness landscapes: A genetic algorithm approach. In: *Proc. of the 1999 Congress on Evolutionary Computation*. vol. 3, pp. 2031–2038. IEEE Press, Washington D.C. (1999)
6. Higgs, P.G., Derrida, B.: Genetic distance and species formation in evolving populations. *Journal of Molecular Evolution* 35, 454–465 (1992)
7. Kauffman, S.A.: *The Origins of Order. Self-Organization and Selection in Evolution*. Oxford University Press, Oxford (1993)
8. Lieberman, E., Hauert, C., Nowak, M.A.: Evolutionary dynamics on graphs. *Nature* 433, 312–316 (Jan 2005)
9. Morrison, R.W., Jong, K.A.D.: A test problem generator for non-stationary environments. In: *Proc. of the 1999 Congress on Evolutionary Computation*. vol. 3, pp. 2047–2053. IEEE Press, Washington D.C., (1999)
10. Nowak, M.A., Sigmund, K.: Evolutionary dynamics of biological games. *Science* 303, 793–799 (2004)
11. Ray, T.S.: Is it alive or is it GA? In: Belew, R.K., Booker, L.B. (eds.) *Proc. of the 4th Int. Conf. on Genetic Algorithms*, University of California, San Diego. pp. 527–534. Morgan Kaufmann Publishers, San Mateo, CA (1991)
12. Spector, L., Klein, J.: Evolutionary dynamics discovered via visualization in the breve simulation environment. In: *Workshop Proc. of the 8th Int. Conf. on the Simulation and Synthesis of Living Systems*, Sydney, Australia. pp. 163–170. University of New South Wales (2002)
13. Stenseth, N.C., Smith, J.M.: Coevolution in ecosystems: Red queen evolution or stasis? *Evolution* 38(4), 870–880 (1984)
14. Taylor, C., Fudenberg, D., Sasaki, A., Nowak, M.A.: Evolutionary game dynamics in finite populations. *Bulletin of Mathematical Biology* 66, 1621–1644 (2004)
15. Tregenza, T., Butlin, R.K.: Speciation without isolation. *Nature* 400, 311–312 (1999)
16. Trojanowski, K., Michalewicz, Z.: Searching for optima in non-stationary environments. In: *Proc. of the 1999 Congress on Evolutionary Computation*. vol. 3, pp. 1843–1850. IEEE Press, Washington D.C. (1999)
17. Valen, L.V.: A new evolutionary law. *Evolutionary Theory* 1, 1–30 (1973)
18. Wright, S.: The roles of mutation, inbreeding, crossbreeding and selection in evolution. In: *Proc. of the 6th Int. Congress on Genetics*. pp. 356–366. Ithaca, NY (1932)