From: AAAI Technical Report SS-03-02. Compilation copyright © 2003, AAAI (www.aaai.org). All rights reserved.

Modeling Artificial Multi-level Selection

Tom Lenaerts, Anne Defaweux, Piet van Remortel and Bernard Manderick
Computational Modeling Lab, Computer Science Department,
Vrije Universiteit Brussel,Belgium
{ tlenaert | adefaweu | pvremort | bmanderi } @vub.ac.be

Abstract

A major obstacle which limits the possibilities of a number of machine learning techniques like the genetic algorithm is the lack of mechanisms which allow for the dynamic construction of hierarchically organized solutions or artifacts. In the context of the GA, such mechanisms would require some form grouping or social behavior. Yet, this type of behavior is per definition maladaptive in the competitive context of the GA. Here, a biological model is analyzed and modeled which allows this social behavior to emerge. This model provides a first step for the computational synthesis of complex structures.

Motivation

The standard genetic algorithm (GA) is a learning/optimization algorithm based on the natural selection metaphor called 'survival of the fittest'. In the broader context of machine learning, the GA is considered to be a near-optimal, weak, search method. Near-optimal in the sense that, instead of providing an exact solution for a specific problem, it will try to approximate the solution as good as possible. Weak, in the sense that the algorithm does not use any information about for instance the structure of the domain. With respect to these two limitations, the GA searches the space of solutions to find the near-optimal one.

The search process is propelled by the competition between the solutions in the population. Good solutions for a particular problem will be able to increase their number more when compared to bad solutions. In more technical terms, solutions compete to survive and reproduce and the success in this competition depends on the relative adaptedness (expected fitness) of the solution in relation to the rest of the population.

Like any model, the GA is a simplified model of natural selection and is based on particular notions that were defined in that context. One way in which these models are simplifications is that in a simple GA the notion of an environment is static. By this we mean that there is a deterministic relation between solution representation (genotype) and expected fitness. This is opposed to biological reality where, on the one hand the genotype-environment interactions specify a norm of reaction, i.e. a genotype does not specify a unique outcome but a pattern of different developmental outcomes in function of the different environments (Lewontin 2000). On the other hand, phenotype-environment interactions determine the adaptedness of the phenotype and this outcome is of course not static due to any kind of environmental changes.

The environment can be defined in different ways. There are roughly two main categories which can be combined in different ways. On the one hand there is a category which is based on the 'organization' of the population or its structure, i.e. spatial or temporal organization. On the other hand, there is a category which is based on interactions within the population, i.e. within- or between-species frequency effects.

Over the years a number of GA researchers have recognized the importance of the environment on the outcome of the evolutionary process. Some have investigated the dynamics of evolution in a spatially organized population (Manderick & Spiessens 1993) or between spatially separated subpopulations (Cantú-Paz 1995). Others have focused on simulating co-evolutionary arms-races which explicitely model the interaction between the members in the population (Rosin 1997). Yet, in both cases the major effort has been put on improving the GA so that the competitive force can produce better solutions for problems which are hard to tackle by the standard GA.

It is remarkable how little effort is spend on incorporating positive interactions, like collaboration, coordination or cooperation, in a GA models (or extensions of it) to produce better results. First, these kind of behaviors might help the GA in traversing the search space towards high-fitness regions. Second, by explictly introducing these positive interactions, techniques like the GA might become *scalable* since this would allow the creation of complex solutions. As a result of this creative process, the solution would become a hierarchical construct.

This problem of scalability is not restricted to the context of this search method. In machine learning, different techniques like hierarchical learning and layered learning have been introduced to improve the generalization of a single learning task or to combine the effects of different sub-tasks (Stone & Veloso 2000). Even more, this issue of hierarchical

Copyright © 2003, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

organization has been a recurring theme in artificial intelligence (Simon 1996), Alife (Lenaerts, Groß, & Watson 2002) and Biology (Pattee 1973; Scott 2002).

In this paper we will not try to tackle the scalability issue in the overall machine learning domain. We start from the premiss that their is a competitive base algorithm which simulates the biological differential survival and reproduction metaphors. Competition between the elements in these algorithms, by itself, is not adequate to perform the synthesis of new constructs since their exists no incentive towards social behavior. Due to the very nature of the evolutionary process, individuals which try to behave socially will be punished for their efforts since it reduces their chances for survival and reproduction.

This raises the question; How can an evolutionary algorithm like the GA promote its population members to work together? What is the minimal set of conditions which should be imposed on the structure of this artificial evolutionary system? A direct approach would be to just engineer it into the existing algorithm and tune it until it performs the job for a specific task. Yet, this approach can lead to different kinds of problems, like, for instance, the fact that the resulting algorithm is completely problem dependent. To create a generally applicable algorithm, a more scientific approach would require us to re-investigate the foundations of the GA, i.e return to Biology and try to come up with an improvement of the GA from there, to overcome the current limitations.

In the following two sections, the general biological context related to the increase in complexity is sketched. These sections will focus on the particular biological models we are working on; multi-level selection or intrademic group selection models. The reader will be made aware of the importance and the issues in these models. Afterwards, our analysis and simulation of the selection dynamics in standard and multi-level selection models is discussed. At the end the paper is concluded.

Complexity and Population Structure

Eörs Szathmáry lists two phenomena that, he claims, underly all increases in complexity and cooperation (Szathmáry 1999). These phenomena are identified as *synergism* and *limited dispersal*.

Synergism is defined as the working together of two things to produce an effect greater than the sum of the two individual effects. Hence it refers to a particular process which creates a state S that combines different primitive elements. Yet, the interaction between the elements in state S must produce certain effects which allow it to survive in the competitive environment. Survival is possible when the different population members collaborate and when this collaboration results in an increase in the expected fitness of each of them relative to the other members of the population. Hence, in this context, it might be useful to examine the different mechanisms which might produce this effect, e.g. symbiosis (Frank 1995).

The second phenomenon, limited dispersal, refers to the notion that phenotypic properties of population members are unable to spread rapidly (if at all) over the entire population. This can occur when the population is structured into different geographically separated modules (colonies) or cell-like groups (propagules). The distinction between these mechanisms is defined by the manner in which the isolated clusters of elements are created. In the first case, migration produces the colonies. In the second case, this migration process is not present. The propagules are physical entities which isolate certain elements from the other elements in the population. We refer to these two mechanisms respectively as *localization* and *encapsulation*.

Richard Michod argues about similar issues (Michod 1999). His work discusses the (newly) established perspective in Biology that individuals (organisms) can be considered social communities of parts that have become functionally integrated. Yet, social behavior has been defined as a critical problem in the context of evolution by natural selection since this kind of behavior derives its benefit from collaboration and shared benefit, which is vulnerable to cheating (Wilson 2003). How was this problem of cheating resolved? It seems that some form of control needs to be introduced to regulate the conflict. In other words, social communities of primitive entities become a new modular (higher-level) entity when certain functionalities emerge which keep the components in line and punish any attempt to cheat. In this way the complexity of the entities in the system can be increased.

Note, though that for the formation of social communities of primitive elements in an evolutionary system, certain conditions need to be imposed on structure of the system. As was described earlier, in standard models of differential survival and reproduction, social behavior is punished and will go extinct since it is considered to be maladaptive. By imposing a structure on the population, certain asocial behaviors will have difficulties in spreading in the population and as a result social communities might be able to evolve (Michod 1999).

Note that the problem described by Michod is not new. The question was originally defined by Edward Wilson as the central problem of Sociobiology (Wilson 1975):

> How can altruism, which by definition reduces personal fitness, possibly evolve by natural selection?

There have been different attempts and approaches to investigate the central problem of Sociobiology. These attempts resulted in different biological models which can be used to answer the question on how complexity can be promoted autonomously in computational artifacts or solutions for particular problems.

One avenue of research, multi-level selection, argues that cooperative behavior can evolve as a result of higher-level selective forces which counter the lower-level ones. Hence natural selection is not restricted to the level of the organism. For this event to occur certain specific conditions need to be imposed on the structure of the population. There are different examples of biological models which discuss this. Yet, we have limited our investigation to a particular type of models; intrademic group selection models (Wade 1978) .

More information on these models and related biological issues like selection and fitness can be found in (Brandon &

Burian 1984; Maynard Smith & Szathmáry 1995; Sober & Wilson 1998; Keller 1999; Michod 1999).

Promoting Social Behavior

We want to make it very clear from the start that we are aware of the biological controversy surrounding group selection or it's new "revival" called multi-level selection (MLS). It is not our intention to argue in favor of certain biological perspectives or put forward new biological claims.

The scientific goal, here, is to design and create algorithms which perform certain tasks. To pursue this goal, useful concepts from other scientific domains can be examined. It is with this goal in mind, that models of MLS are investigated. We try to evaluate whether these biological models can contribute anything to the computational synthesis of hierarchically structured artifacts or solutions.

Static vs. Dynamic Environments

The importance of the models in the context of artificial evolutionary systems like the GA, lies in the fact that MLS models incorporate environmental variation. It is this environment which is an influential parameter that can change the selection dynamics. As will be shown in a later example, it can promote social behavior in a system like the GA.

The environment is a parameter which is often neglected in GAs since it is assumed that one individual codes for one solution which contains all the information necessary for evaluation, i.e. optimization in a static environment. Yet, in most interesting applications, there is no deterministic relation between representation and fitness, the environment is dynamic. MLS models become important when a trait is controlled not only by genetic informations but also by the interactions with the environment and the other individuals in that environment (Dicks 2000).

In this discussion we focus on selection at two levels; the individual and the group. Selection among groups requires that the groups differ in actualized fitness and that this difference is due to the adaptedness of these groups to a common environment. What is meant by differences in adaptedness at the group level? How is a fitness assigned and what does it mean? In the next section a particular type of models, intrademic models, will be discussed. These models come in different shapes and make different assumptions. We will discuss only what is relevant to us.

Intrademic group Selection

Intrademic group selection models are a specific category of group selection models. They assume a single (panmictic) population whose members are distributed into isolated neighborhoods or "trait groups" (Wilson 1979). This formation is performed by a process called dispersion or colonization. Actions which affect the fitness of the individuals occur inside the groups and, as a consequence, produce differences in the viability between the members of the different groups. For instance, as can be seen in Figure 1, after dispersion five groups are formed which consist of different frequencies of individual types in the population. Since, it is assumed that

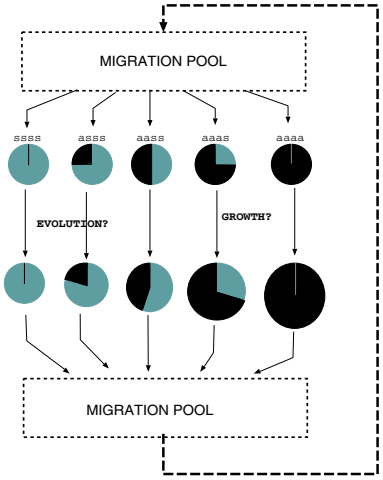


Figure 1: Illustration of intrademic group selection models (Sober & Wilson 1998).

interactions influence the selection dynamics, the actualized fitness of the members of each group will be different.

After the colonies of groups are constructed, the members of these groups reproduce and the population grows. At a certain point in time, this growth process stops and the groups contribute migrants to a common pool, i.e. the migrant pool from which the next generation will be constructed. The frequency of each type in this migrant pool represent the frequency of the surviving types.

If all populations are assumed to contribute the same amount of migrants to the pool, then the frequency of each type in this pool is the average frequency over all contributing groups. If this assumption is used, group selection becomes impossible. If however each group contributes proportionally to its size, the frequencies of each type become average *weighed* frequencies and as a result group selection becomes possible.

For simplicity, it is assumed that the initial group sizes are equal. This assumption simplifies the design and analysis of the algorithm. Yet, the intensity of group selection can increase if their is initially a positive correlation between the size of the group and the frequency of the types in the group.

The variation between the initial groups depends on the sampling strategy which creates the colonies from the migrant pool. It is the variation between the groups that will eventually determine the impact of group selection on the entire process. Wade suggests two extreme models for this sampling strategies (Wade 1978): migrant pool and propagule pool.

Variation Between the Groups?

Group selection models, like the one described in the previous section, depend completely on how the groups are constructed at each generation. Selection at this higher level is only possible when there is enough variation between the groups.

A first assumption made by all group selection models is

about the size of the initial groups, i.e. the groups should be small. When the groups are small samples of the pool, there will be a large between-group variance. Consequently, the selection pressure at the higher-level will be increased.

A second important facet of these group selection models is the way in which the colonists are chosen from the pool, i.e. the sampling strategy. There are different strategies to perform this task. Each strategy will influence, in a different way, the impact of group selection on the overall selection dynamics in the model. The migrant and propagule strategies can be considered as two extreme sampling strategies. This actual sampling strategy can be one of these two extremes or a intermediate combination of them.

The first strategy, the migrant pool, chooses colonists at random from the pool. Hence, the frequencies of the different types in the migrant pool will be uniformly distributed over the different groups. In the context of the small groups, there might still be some variation between the groups and, as a result, group selection can occur. Yet, due to the repetitive application of this strategy, the population is extensively mixed and in the long term, standard individual selection will determine the outcome of this process.

This form of group selection can not be very effective unless the initial between-group variance can be re-generated at each colonizing episode. An alternative would be that the sampling strategy produces a non-random assortment of individuals within each group. For instance, kin groups are examples of groups which consist of siblings or closely related relatives and which result in an increased between-group variance.

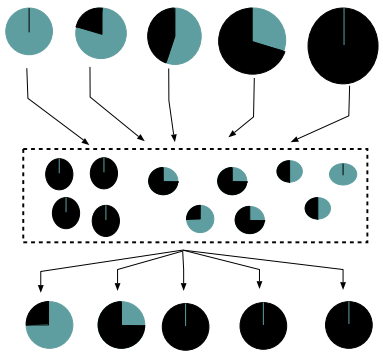


Figure 2: Illustration of the propagule pool in an intrademic group selection model. Note that the biggest groups contribute more propagules to the pool. This will influence the probability that a particular propagule will be selected to start a new group (Wade 1978).

The second strategy is one which was used by Wade for his experimental research and is visualized in Figure 2 (Wade 1978). The general idea is that, at the end of a generation, each group contributes a number of "propagules" proportional to its size. A propagule is defined as a representative sample of the group. Due to this proportional contribution, the pool will contain a weighed average of the different types of propagules. To establish new colonies, the propagules are selected from the propagule pool and no mixing of colonizing individuals is permitted. Of course, the effect of this scheme on the between-group variance is different from the uniform sampling strategy. Note that the propagules resemble samples from the groups at the previous generation. In order for these samples to be meaningful they can not be too small. Yet remember that small groups were considered better for the between-group variance at the beginning of a generation. Hence some balance needs to found for this process to work properly.

The following resemblances and differences between both approaches can be summarized.

- Both strategies contribute elements proportional to the size of the group. Yet, in the first case these elements are individuals and in the second case they are small groups of elements.
- The propagules persist as distinct units in the pool as opposed to the mixing of the elements in the migrant pool.
- The groups which are constructed out of the propagules selected from the pool resemble the composition of the groups at the previous generation. In the migrant strategy, in the worst case, the groups are created by a random uniform sampling process and resemble the composition of the migrant pool.

It is assumed that in biological circumstances, natural populations will not lie on either side of the spectrum defined by the two models. It will probably somewhere in between.

A further discussion could be initiated on the meaning of fitness and inheritance on the level of the group. Since it was not the goal of this paper to contribute to this debate, it will be ignored here. Yet a number of interesting articles exits which elaborate on these issues. We refer the interested reader to (Uyenoyama & Feldman 1980; Mayo & Gilinsky 1987; Brandon 1990)

Modeling and Analyzing Levels of Selection

In this section, our research is outlined. We will refer to the resulting computational model as a *multi-level evolutionary algorithm* (MLEA). The research approach consists of three phases:

1. a mathematical analysis of the selection dynamics in a standard EA and the MLEA model,
2. the development of a computational model and its application on a set of experiments related to the evolution of strategies and
3. the validation of the evolutionary capabilities of the MLEA model to combine lower-level entities to achieve higher-level functionality.

First, the selection dynamics of both the standard EA and the MLEA and their relation are modeled mathematically. This step provides us with a clear understanding of the selection dynamics in both models.

Second, an algorithm was developed and evaluated on a set of experiments related to the evolution of strategies

for symmetric two-player games (Axelrod 1997). The algorithm evaluates how this stability emerges from the interaction of the lower-level components. Note that we do not discuss the emergence of networks yet. The approach here is to verify the analytical results in a computational model. Important was that the infinite population model no longer holds and certain problems relating to this had to be resolved. The standard EA model was adapted to the mechanism visualized in Figure 1

Third, to validate the applicability, the model is incorporated in a set of existing EA: co-evolutionary algorithms, artificial immune systems and learning classifier systems. These models were chosen because they already incorporate the notion of interaction and relative fitness.

In the next subsection, we will present an example and discuss the selection dynamics in the intrademic group selection model. This account should provide a clear view on the significance of these models in the context of artificial evolutionary systems. Due to space limitations, we will not discuss the computational models. This choice was made to keep the content of the paper broad in order *not* to restrict the discussions at the symposium to some specific implementations of the model. More information on the first two phases can be found in (Lenaerts 2003). The third phase is reserved for future work.

Orgin of Altruism

Altruism is a behavioral property which has fascinated evolutionary biologists for a long time. This fascination resulted from the difficulty in explaining altruistic behavior in biological organisms from a Darwinian perspective. Generally, altruistic behavior causes the bearer of this property to pay a price for its behavior in terms of reproductive success. In other words, natural selection can not evolve this trait since it causes individuals to have *less* offspring than their conspecifics. Hence there is a selective advantage to being selfish.

The standard model of altruistic behavior assumes a population of size n which contains altruistic and non-altruistic (selfish) individuals at frequencies p and $(1-p)$ respectively. Hence the population contains $np = n_a$ altruists and $n(1-p) = n_s$ selfish organisms.

	C	D
C	b-c, b-c	-c, b
D	b, -c	0,0

Table 1: Fitness increment due to pairwise interaction

It is typically assumed that each organism, independent of its behavior, produces certain amount of offspring. The specific behavior of each organism increases or decreases this amount (see Table 1) . Since each altruist benefits other organisms at a personal cost, his number of offspring is reduced with the amount $c > 0$. A single other member of the population receives from this altruist b more offspring. In most discussions of altruism it is assumed that $b > c$.

As will be shown, the expected fitness of a selfish individual will always be higher than that of an altruistic individual and as a consequence these altruists will, over time, go extinct in a standard selection model.

To simplify the calculations, we assume an infinite population. Note that in any computational model, this assumption can not be made.

This problem will be used in the next section to demonstrate the influence of group selection in an intrademic model. An extensive account on the problem from a biological perspective can be found in (Sober & Wilson 1998; Fox Keller & Lloyd 1992) and from an economical perspective in (Bergstrom 1993; 2001).

Other examples were analyzed in a similar manner. These examples correspond to instances of the symmetric two player games defined in (Weibull 1996). Next to the altruism example this was the matching pennies game, the battle of the sexes and the hawk-dove game.

Selection Dynamics

In the previous sections it was argued that introducing different levels of selection adds extra complexity to the dynamics of an artificial evolutionary algorithm like the GA. In order to understand these dynamics and how they can be tuned, an analysis was made of the differential survival and asexual reproduction of simple individuals. These simple individuals have an implicitly encoded behavior and this behavior determines their evolutionary success (see the altruism example in the previous section).

Different analytical methods exist to analyze selection. We focused on two methods: replicator dynamics (RD) defined in evolutionary game theory (EGT) (Weibull 1996) and George Price's covariance equation (Price 1970; 1972). Price's covariance equation was selected because it is often cited in the context of multi-level selection. Here we will only show the derivation of the RD in the context of the altruism example. The derivation of the Price equation produces the same results.

RDs are dynamical systems which highlight the role of selection in evolutionary models. A RD is formalized as a system of discrete difference equations or continuous ordinary differential equations. Initially this system of equations does not incorporate any mutation mechanisms to introduce variation. This mutation mechanism has, in the context of EGT, more relation to the notion of evolutionary stability than with the selection dynamics. Yet, extensions of the RD exist which incorporate both.

The population is generally represented by a state $\mathbf{x}$ which contains the different frequencies of each type in the population. In the context of the example, the population state is $\mathbf{x} = (p, (1-p))$, where p represents the frequency of altruists and $(1-p)$ the frequency of the selfish individuals in the population. We want to derive a (discrete) replicator dynamic which calculates the change in the frequency of the altruists between two consecutive points in time t and $(t+1)$.

The number of altruists after one generation and the number of selfish individuals after one generation are

$$n_a(t+1) = n_a(1 + p(t)b - c)$$

$$n_s(t+1) = n_s(1 + p(t)b)$$

And the total number of new individuals after one generation is

$$n(t+1) = n(t)(1 + p(t)(b-c)$$

Using these values, the relative frequency of the altruists in the population can be determined easily:

$$p(t+1) = \frac{n_a(t+1)}{n(t+1)} = n(t)\frac{1+p(t)b-c}{1+p(t)(b-c)}$$

The change in the altruist frequency can again be determined easily:

$$\Delta p = p(t+1) - p(t) = \frac{-cp(t)(1-p(t))}{1+p(t)(b-c)} \qquad (1)$$

This equation represents the discrete replicator dynamic of the altruism example in a standard population and a cobweb diagram visualizes the dynamics of this system.

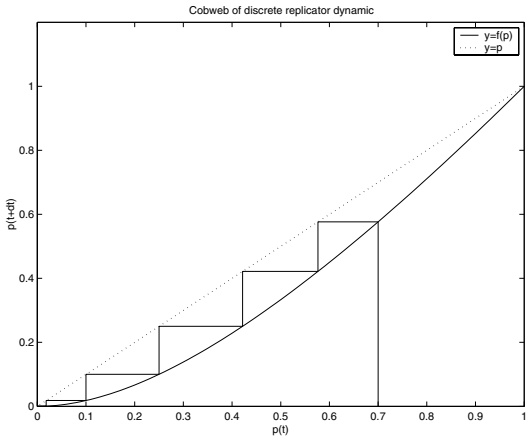


Figure 3: Cobweb diagram of the selection dynamics in the altruism example in a standard population model. This diagram also visualizes the multi-level selection dynamics in the altruism example when the variance between the groups $\sigma^2 = 0$

In this diagram, the frequency of the altruists in the population at time t is plotted against the number at the next generation $t+1$. The diagonal line represents the equilibrium case where $p(t) = p(t+1)$. On this diagram a "progression line" is plotted. This line shows that starting from even high frequencies of altruists in the population, the selfish will be the only that survive in the end.

Now, the same process can be repeated for the intrademic group selection model. The general idea is that we calculate the change in the frequency from one migrant pool to the next (see Figure 1). It is assumed that the size of the initial groups is the same and that there are k groups after colonization. It is also assumed that in each group i, equation (1) describes the growth.

Remember that the frequency of the elements in the migrant pool (non-propagule model) was a weighed average of the frequencies in the different groups at the end of the generation. Hence the actual number of altruists in the migrant pool (n_a)is a of the number of altruists in each group (n_{a_i}):

$$\begin{aligned} n_a(t+1) &= n_{a_1}(t+1) \quad n_{a_2}(t+1) \quad \ldots + n_{a_k}(t+1 \\ &= \sum_{i=1}^{k} n_{a_i}(t)(1 + p_i(t)b - c) \\ &\cdots \\ &= n_a(t)(1 + \sum_{i=1}^{k} \frac{n_{a_i}(t)}{n_a(t)}(p_i(t)b - c) \\ &= n_a(t)(1 + \mathrm{E}_a[(p_i(t)b - c)]) \end{aligned}$$

In this derivation p_i represents the frequency of the altruists in group i . The original sum at the beginning of the derivation is (in the end) expressed in terms of the number of altruists in the previous time step. It is important to note that the actual number of altruists will depend on relative number of altruists in each group i. The weighed sum of the fitness of an altruist corresponds to an expected value. It was therefore that this sum is written as $\mathrm{E}_a[(p_i(t)b - c)]$.

The same derivations can be performed for the actual number of selfish individuals in the migrant pool ($n_s(t+1)$) and as a consequence, the total number of new individuals ($n(t+1)$). In turn the relative frequency ($p(t+1)$) and the change in relative frequency (Δp) can be determined. We refer for the complete derivation to (Lenaerts 2003). Here we just show the outcome:

$$\Delta p = \frac{b(\frac{1}{k}\sum_{i=1}^{k} p_i^2(t) - p^2(t)) - cp(t)(1-p(t))}{1+p(t)(b-c)}$$

If the previous equation is examined carefully, then it can be noticed that

$$\frac{1}{k}\sum_{i=1}^{k} p_i^2(t) - p^2(t) = \frac{1}{k}\sum_{i=1}^{k}(p_i(t) - p(t))^2 = \sigma^2(t)$$

or in other words, the between-group variance in altruist frequency ($\sigma^2(t)$) will determine the dynamics of these model. As a result we obtained:

$$\Delta p = \frac{b\sigma^2(t) - c(1-p(t))p(t)}{1+p(t)(b-c)} \qquad (2)$$

The right hand side of equation (2) consist of two parts. The second part is the individual selection which is exactly equal to the discrete dynamic we previously derived for the single-population model (see equation (1)). The first part represents the group selection parameter and this can alter the dynamics in the system. In Figure 4, Figure 5 and Figure 6 a cobweb diagram is shown of the MLS dynamics in this example when the variance $\sigma^2(t)$ is 0.05, 0.125 and 0.5 respectively. As can be deduced from these diagrams, if the between-group variance increases then altruistic behavior can prosper and even take over the population. Social behavior can evolve when $b\sigma^2(t) > c\ (t)(1-p(t))$. Note that when $\sigma^2(t) = 0$, the diagram for equation (2) corresponds to Figure 3.

Although this example is very simple it gives a first indication on the importance of group selection models. Furthermore, these analytical results confirm the observations

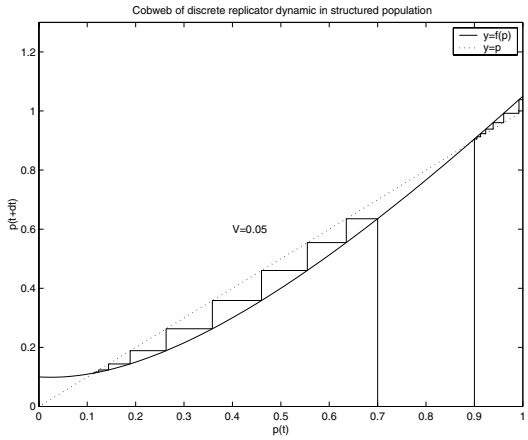


Figure 4: Cobweb diagram of the multi-level selection dynamics in the altruism example. In this diagram the $\sigma^2(t) = 0.05$

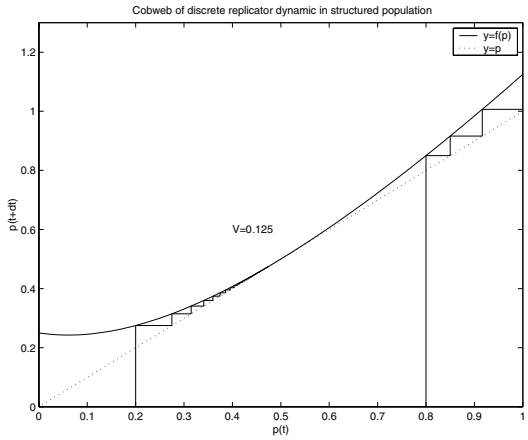


Figure 5: Cobweb diagram of the multi-level selection dynamics in the altruism example. In this diagram the $\sigma^2(t) = 0.125$

made in biological models. Further analysis of these models were made in the context other interaction games and N-population models where the interaction between different species was modeled.

Furthermore, the analytical model here assumes that the interactions are between two individuals, symmetric and occur within a single population. Symmetric refers to the fact that the role of the individuals is interchangeable. Many strategically relevant interactions as in Biology are not symmetric. These interactions occur between different individuals of different populations. In (Lenaerts 2003), the analysis described here was also extended to multi-population models.

Conclusions

As was argued in the beginning of this paper, this research was conducted in the context of the computational synthesis

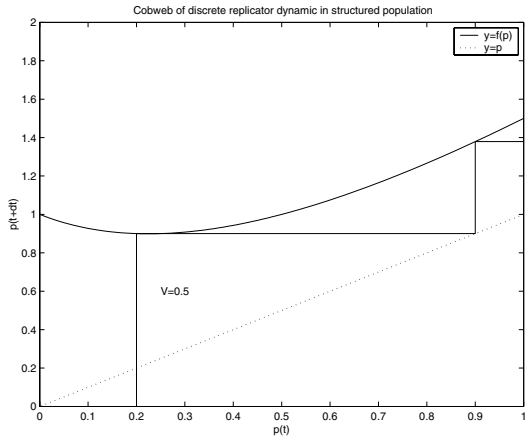


Figure 6: Cobweb diagram of the multi-level selection dynamics in the altruism example. In this diagram the $\sigma^2(t) = 0.5$

of complex artifacts or solutions. We discussed the importance of this issue in the context of evolutionary algorithms and more generally machine learning and argued in favor of returning to Biology for answers to how an evolutionary algorithm should be adapted to promote social behavior in the competitive population.

In Biology there is recently been an interest in similar issues. These issues are collectively termed evolutionary transitions and describe the evolution of instances of different levels of complexity in the biological realm. One of the mechanisms which is considered responsible for the increase in complexity and social behavior is called "limited dispersal". This mechanism refers to the fact that certain maladaptive behaviors were limited in spreading over the population which provided socially interacting individuals with a chance to survive. Hence this mechanism promotes social behavior in a competitive environment.

Here, we investigated a particular model of limited dispersal; intrademic group selection and showed that it is indeed possible for social behavior (altruism example) to survive when certain conditions are met.

Yet, this investigation was only a first step towards the real construction of higher-level computational constructs. As a second step, these models will be used to analyze the emergence of networks of primitive components. This second step could for instance be performed using learning classifier systems or models of artificial immune systems.

To conclude, the question on how to model computational synthesis was recently examined at the workshop on the modeling of dynamical hierarchies (DHs) at the Alife VIII conference in Sydney. At this workshop a working definition or DHs was proposed:

- a hierarchically structured collection of simple and complex entities which are linked through a part-whole relationship,
- where different types of interactions between the parts result in non-linear dynamics which produce and sustain the

whole and

- the whole supervenes the parts through encapsulation and/or the emergence of new functionalities at the level of the whole.

It is our opinion that any algorithm which tries to produce solutions or artifacts of higher complexity, will need to investigate how it corresponds to these three items in the DH definition.

References

Axelrod, R. 1997. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton University Press, Princeton, New Jersey.

Bergstrom, T. 1993. How altruism can prevail in an evolutionary environment. *American Economic Review* 83(2):149–155.

Bergstrom, T. 2001. Evolution of social behaviour: Individual and group selection models. Paper provided by California Santa Barbara - Department of Economics in its series University of California Santa Barbara - Department of Economics with number 14-01.

Brandon, R., and Burian, R., eds. 1984. *Genes, Organisms, Populations; Controversies Over the Units of Selection*. Cambridge Massachusetts, London England: MIT Press.

Brandon, R. N. 1990. *Adaptation and Environment*. Princeton, New Jersey: Princeton University Press.

Cantú-Paz, E. 1995. A summary of research on parallel genetic algorithms. Technical Report Illegal Report No 95007, University of Illinois at Urbana-Campaign, IL.

Dicks, L. 2000. All for one! *New Scientist* 167(2246):30–35.

Fox Keller, E., and Lloyd, E., eds. 1992. *Keywords in Evolutionary Biology*. Harvard University Press, Cambridge Massachusetts, London England: Harvard University Press.

Frank, S. 1995. The origin of synergistic symbiosis. *Journal of Theoretical Biology* 176:403–410.

Keller, L., ed. 1999. *Levels of Selection in Evolution*. Monographs in Behaviour and Ecology. Princeton, New Jersey: Princeton University Press.

Lenaerts, T.; Groß, D.; and Watson, R. 2002. On the modelling of dynamical hierarchies. In *The Workshop Proceedings of the 8th Interantional Conference on the Simulation and Synthesis of Living Systems*, 37–44. University of New South Wales.

Lenaerts, T. 2003. *Different Levels of Selection in an Artificial Evoluitionary Algorithm; Analysis and simulations of the selection dynamics*. Ph.D. Dissertation, Vrije Universiteit Brussel, Computer Science Department, Pleinlaan 2, 1050 Brussels, Belgium (forthcoming).

Lewontin, R. 2000. *The Triple Helix: gene, organism and environment*. Cambridge, Massachusetts; London, England: Harvard University Press.

Manderick, B., and Spiessens, P. 1993. Fine-grained paralled genetic algorithms. In Schaffer, J., ed., *Proceedings of the Third International Conference on Genetic Algorithms*. San Mateo, California: Morgan Kaufmann Publishers, Inc.

Maynard Smith, J., and Szathmáry, E. 1995. *The major transitions in evolution*. Oxford University Press.

Mayo, D., and Gilinsky, N. 1987. Models of group selection. *Philosophy of Science* 54:515–538.

Michod, R. 1999. *Darwinian Dynamics; evolutionary transitions in fitness and individuality*. Princeton, New Jersey: Princeton University Press.

Pattee, H., ed. 1973. *Hierarchy Theory; The challenge of complex systems*. New York: George Braziller.

Price, G. 1970. Selection and covariance. *Nature* 227:520–521.

Price, G. 1972. Extension of covariance selection mathematics. *Annals of Human Genetics, London* 35:485–490.

Rosin, C. D. 1997. *Coevolutionary search among adversaries*. Ph.D. Dissertation, San Diego, CA.

Scott, A. 2002. Biological reductionism. In *The Workshop Proceedings of the Alife VIII Conference*. University of New-South Wales, Sydney, Australia: The 8th International Conference on th Simulation and Synthesis of Living Systems.

Simon, H. 1996. *The Sciences of the Artificial*. The MIT Press, Cambridge, Massachusetts, London, England., third edition.

Sober, E., and Wilson, D. 1998. *Unto Others, the Evolution and Psychology of Unselfish Behaviour.* Cambridge, MA:Harvard University Press.

Stone, P., and Veloso, M. M. 2000. Layered learning. In *Machine Learning: ECML 2000, 11th European Conference on Machine Learning, Barcelona, Catalonia, Spain, May 31 - June 2, 2000, Proceedings*, volume 1810, 369–381. Springer, Berlin.

Szathmáry, E. 1999. *The First Replicators.*, In Keller (1999). 31–52.

Uyenoyama, M., and Feldman, M. 1980. Theories of kin and group selection: A population genetics perspective. *Theoretical Population Biology* 17:380–414.

Wade, M. 1978. A critical review of the models of group selection. *The Quarterly Review of Biology* 53(2):101–114.

Weibull, J. 1996. *Evolutionary Game Theory*. Cambridge, MA: MIT Press.

Wilson, E. 1975. *Sociobiology: The New Synthesis*. Cambridge, Massachusetts and London, England: The Belknap Press of Harvard University Press, fourth edition.

Wilson, D. 1979. Structured demes and trait group variation. *American Naturalist* 113:606–610.

Wilson, D. 2003. Social groups all the way down. *Human Biology* (forthcoming).