Learning in an Evolutionary Environment*

Wolfgang Ketter
Department of Computer Science
and Engineering
University of Minnesota
ketter@cs.umn.edu

January 31, 2003

Abstract

With the help of evolutionary models we are able to explain phenomena, which cannot be explained by other classical learning models. The intent of this paper is to look at different learning methodologies in an evolutionary environment, compare them with each other, and finally discuss how all this is applicable to the MAGNET system.

*This work was conducted in and as a part of the MAGNET framework. The MAGNET (Multi-AGent NEgotiation Testbed) research group is headed by Professor Maria Gini in the Artificial Intelligence, Robotics and Vision Lab (AIRVL), at the University of Minnesota. Many thanks for invaluable and stimulating conversations and support to Alexander Babanov.

1 Introduction

The learning of a society and of an individual inside an evolving society allow us to address many interesting problems that cannot be studied in the framework of individual or team learning. The following table summarizes the insight on the differences between the two frameworks. In the subsequent sections I examine each concept in this table in greater detail.

Individual	Society
Behavior	
Selfishness is often individually rational.	People as well as other life beings tend to co-operate.
Propagation of Knowledge	
Knowledge does not span beyond life-time of individual.	Society accumulate knowledge over generations of individuals.
Usability of Result	
Individual is too “erratic” to study on non-statistical basis.	A society is a relatively ”well-behaved” object to study.
Complexity and Precision	
The individual behavior depends on the behavior of other individuals.	Other individuals are “internalized” and studied simultaneously.

1.1 Behavior

The first part of this section is based on Robert Axelrod’s studies of the evolution of cooperation. In one of his earlier works on this subject [2] he poses the following question: “Under what conditions will cooperation emerge in a world of egotists without central authority?” One reason this question attracts our attention is because the selfish behavior often seems to be the best way of securing the maximum amount of resources for personal consumption. Still we observe patterns of cooperative behavior in almost every society from insects to higher animals to human beings. Robert Axelrod refers among many others to the following illustrative examples of emergence of cooperation:

- To protect the hive, honey bee thrones sting opponents when they come near the nest, but after that they die. So since they do not promote their genes, there should be no reason for them to do that, but they do it for the welfare of the society.
- Another example for cooperation is based on the location. On the high sea where a big fish is a predator for a small fish is no cooperation between those two. On the contrary near a cost line a small fish or crustacean is cleaning a big fish in which it removes and eats parasites from the body of the larger fish.
- Another case of the emergence of cooperation is the development of patterns of behavior in a legislative body such as the United State Senate. Mutually rewarding actions have led to the creation of an elaborate set of norms. Among the most important of these is the norm of reciprocity - a folk-way which involves helping out a colleague and getting repaid in kind.

If one has only a single interaction then cheating might be a viable solution. Whenever there are repeated interactions between individuals, we observe very complex patterns of strategic interaction. In case an individual can recognize a previous interactant and remember some aspects of the prior outcomes, the strategic situation can be modeled by the iterated Prisoner's Dilemma. A strategy could use the history of previous interactions to decide to cooperate or defect in the next move.

Axelrod has formalized Darwin's emphasis on individual advantages in terms of game theory. An important conclusion of Axelrod's studies is that foresight is not necessary for the evolution of cooperation.

1.2 Propagation of Knowledge

The knowledge of an individual and the modifications of behavior due learning does not span beyond lifetime. On the contrary the society accumulate knowledge over generations of individuals. Partial modifications persist through parent-child relationship and through promotion of genes that are more likely to learn the same knowledge. The parents are not the only medium which can transfer knowledge to their children. Throughout supervised learning, children learn in schools from teachers, what is right or wrong.

The scientist Lamarck proposed in the late nineteenth century, that evolution over many generations was directly influcend by individuals during their lifetime. In particular he proposed that experiences made by an individual can be directly passed on to their offspring [10]. Our current understanding of biological parent-child relationship objects this view point. Even that the strict view point of Lamarck is objected nowadays, there is a mechanism that suggest by which an individual which learns can affect the evolution process of the species which the individual is living in. This is called the Baldwin effect [4], which is named after J. M. Baldwin who was first writing of this in the year 1896.

The Baldwin effect states, that the ability of individuals to learn can have an indirect accelerating effect on the rate of evolutionary adaptation for the entire population. This accelerated adaptation may influence the evolution process to speed up. Some of the learned techniques may evolve into genetic traits and can so be inherited by their offspring, this means that such techniques do not have to be learned anymore. This means that the Baldwin effect provides an indirect approach that individual learning is able to influence the rate of evolutionary progress.

As an example of the Baldwin effect you can imagine some new change in the environment of some species, such as a new predator. This new event will trigger individuals in the species which are capable of learning how to deal with this new predator. Then these individuals are able to influence to rate of evolution in the society. One can see in this example that change in the environment promotes learning.

1.3 Usability of Result

An individual is too "erratic" to study on non-statistical basis. A society on the other hand is a relatively "well-behaved" object to study and becomes predictive in a certain way. In machine learning one has to assume a fixed state space, but when doing learning in an evolutionary environment the paradigm of individual learning can not be applied, it is impossible to describe the state space of a whole society while is is equally impossible to predict state transitions in the model that limits the number of states to individual ones and ignores the rest of a society.

1.4 Complexity and Precision

The individual behavior depends on the behavior of other individuals. This increases the complexity of the system drastically, since there are too many external variables which influence the behavior of an individual. However if we study a whole society many of otherwise external variables, e.g. other individuals, are no longer outside entities. Other individuals are "internalized" and studied simultaneously.

2 Evolution, Natural and Artificial

2.1 Natural Evolution

For getting a better understanding of the mentioned algorithms and techniques I will briefly speak about evolution in a natural environment. This allows us to make comparisons between a natural and an artificial evolutionary environments later on.

Darwin explains the natural evolution for human beings with slight variations that occur between individuals (from whatever cause). Those variations which are profitable to this species tend to be preserved, because the offspring who inherits variation of this sort has a better case of surviving compared to other offspring of the same generation. "Natural selection is a power incessantly ready for action." (Darwin) All this is what Darwin means when he speaks from the "struggle for life." Spencer calls it: "Survival of the fittest" [14].

Nowadays we have synthesis of genetics and Darwin's theory is updated as follows:

- The modern theory of evolution counts on natural selection and random genetic drift.
- Characteristics are inherited as discrete entities (genes).
- "Variations" (Darwin) are due to the presence of multiple alleles (parameter values) of a gene.

I think Dawkins theory [5] is very interesting and thought stimulating: "living organisms exist for the benefit of DNA rather than the other way round." So in principle everything is governed by genes and we are only the outer skins to help them to exist and execute their decisions.

Another observation about learning and evolution in general is, that both are adaptation processes, but learning is on the individual level (individual agents) and evolution is on the level of a whole population. This leads us to the research field of artificial life which is covered in the next sub section.

2.2 Artificial Evolution

Artificial life (AL) is defined as the study of man-made systems that behave in ways characteristic of natural living systems. Researches in AL are concerned about building complex systems out of simple entities. Much of the lifelike behavior in these kind of systems is emergent (global behavior) and is not explicitly programmed by the designer (local behavior is programmed).

People working in symbolic AI have been criticized quite often that their input is often preprocessed by humans and their output interpreted by humans. On the other hand creatures in AL sense their environments and act on them directly. Their actions affect the environment directly and so the creatures future's perceptions. There is no human anymore in the loop who could influence the result.

As an example for an artificial life creature we can take "Animat" [17] as a simple animal. Those simple animals can: a) Sense with their sensory signals, b) are capable of actions, c) certain signals have special status and d) they operate internally and externally to optimize the rate of occurrence of the special signals. The basic problem in AL system is for the creature: "What to do next?"

The Animat learns via classifier systems, an amalgam of production systems and genetic algorithms. Classifier systems use genetic algorithms to evolve rules. Animat's classifier rules consists of a taxon, an action, and a strength. A taxon is a template capable of matching any certain set of detector vectors. An action is one of eight possible moves, and the strength is a numerical measure of the classifier's value to Animat. Animat has a population of classifiers. Their number is a fixed parameter of the system. The population evolves as Animat lives. Animat adapts by learning. All the knowledge what Animat has emergent out of the system, it is neither programmed nor explicitly represented. Only reactions to local situations are programmed into or learned by Animat. The global behavior emerges.

In AL systems we have four major approaches: genetic algorithms, neural networks, cellular automata and finite automata. Also there are two major possibilities of structuring the society: spatial (limited neighborhood) and mean-field (every one can interact with everybody else).

Cellular automata provide a way of viewing whole populations of interacting "cells", each of which is itself a computer (automaton). By building appropriate rules into a cellular automaton, we can simulate many kinds of complex behavior. The most famous example for cellular automata is probably “The Game of Live” [1], which was introduced by the Cambridge mathematician John Conway. This is an example of learning as social interaction leading to survival. There are many simulations as applets available on the web, I have myself since last fall a simulation running under my account and it is fun to watch to see the populations evolve.

There are two major differences in this kind of system compared with genetic algorithms. First, with evolutionary systems there is no explicit fitness function, the environment acts as an implicit fitness function, it decides which agent lives long enough and eats well to have the energy to reproduce (strength). Here, as opposed to genetic algorithms, the population size varies over time. Births can occur with no corresponding deaths, and conversely.

3 Evolutionary Methodologies

3.1 Genetic Algorithms

Any abstract task to be accomplished can be thought of solving a problem, and this itself can be thought of as a search through a space of possible solutions. Since we want to find “an optimal” solution, the problem can be viewed as an optimization process. If the search needs only to cover a small space then we can use classical exhaustive methods, but if the state space gets large we need special artificial intelligence techniques. Appropriate in such situations are stochastic algorithms whose search methods model some natural phenomena: genetic inheritance and as mentioned before the Darwinian strife for survival. Genetic algorithms fall into this class of algorithms, too.

I will explain what to understand when speaking of genetic algorithms. Genetic algorithms use a vocabulary borrowed from natural solution. A genetic algorithm is started with a set of solutions (represented by chromosomes) called population. Then the solutions from one population are taken and used to form a new population. This is motivated by a hope, that the new population will be better then the old one. After that solutions which are selected to form new solutions (offspring) are selected according their fitness — the more suitable they are the more chances they have to reproduce. This is done with the help of a fitness function. After that reproduction is accomplished by crossover and mutation. All the individuals which have been selected for reproduction are randomly paired. In the next step we select for each pair a crossover point. If we think of the parents as strings of the same length, then the crossover point is somewhere in the range of the length. The first offspring will inherited all bits from parent one until this point and the rest of the string will be inherited from parent two, and the opposite is true for the second offspring. However, each gene can be altered by random mutation to a different value (e.g. change from 0 to 1 or from 1 to 0, if the string has a binary encoding). The mutation rate is usually very low, otherwise there is too much randomization in the system. Mutation to a certain degree is good, since it gives the chance to find solutions, that would have not been found in the population without it. If one wants to apply a genetic algorithm to a problem, one needs to answer the following questions [13]:

- What is the fitness function?
- How is an individual represented?
- How are individuals selected?
- How do individuals reproduce?

This is the classical approach to genetic algorithms, nowadays there are many variations of them available. Let's have a closer look at the evolution process. An evolution process run on a population of chromosomes corresponds to a search through a search space of potential solutions. When doing such a search we have to find a balance between exploiting the best solutions and exploring the search space. A good strategy for exploiting the best solutions for possible improvement is Hill-climbing, but this algorithm has the disadvantage, that it neglects the exploration of the search space, and often gets stuck in a local optimum. On the other hand a random search, like simulated annealing, explores the search space very well, but ignores the exploitations of the promising regions of the state space. Genetic algorithms offer a good balance between exploitation and exploration of the search space and are domain independent.

Genetic algorithms belong to the class of probabilistic algorithms and because they combine elements of directed and stochastic search they are suitable for solving complex problems.

Genetic algorithms for example are used for solving the following problems [9]:

- Optimization of a function.
- Solving the prisoner's dilemma.
- Traveling salesman problem. For example, [6] approaches the problem with a random-key representation for ordering problems. The random-key encoding has the advantage that any bit string represents a legal tour, which eliminates the need for specialized crossover operators.

While referring to the help of genetic algorithms one has to play "God," because of the need for the fitness function. One has to know when applying genetic algorithms which parameter to optimize, but in nature we do not have a fitness function. What if we have a system with too many parameters and we do not know which parameter, or property to optimize? In this case the evolutionary environment is better, it does not require one to specify what parameter to optimize up front, those will be discovered in the process of evolution. In this process parameters may come up, which one has never thought of to optimize beforehand.

Another issue with genetic algorithms is that they ignore the experience gained during an individual's lifetime. If we want to simulate real evolution processes than this type of learning (reinforcement) is not enough. As mention earlier other methods such as supervised learning need to be integrated in the process, such like a teacher student relationship. That is why I like for certain problem categories evolutionary environments much better than genetic algorithms. They are more flexible than genetic algorithms in the sense that certain properties will emerge out of them. Evolution modifies the whole organism.

3.2 Core War and Evolution of Computer Programs

Core War[1] is a game created by A. K. Dewdney in which rival programs battle to the death. First written about in Scientific American in the 1980's. In the following I refer to a research project “EVOLVING WARRIORS” [16] by Thorsell, in which he shows results of an implementation of Core War. The core war project uses the tools of genetic programming. From [16]:”A Core Warrior is a program whose ultimate goal is to take total control over a computer by throwing out all other processes, similar to what a malevolent virus would do to your personal computer.

The instruction set for Core Wars, Redcode, is a limited assembly language, but it is computationally (Turing) complete. Thus anything, which can be done by a computer program, can be done in Core Wars, given a large enough memory space.

The Core War system can be compared to a biological system, where the production of synthetic organisms based on a computer metaphor of organic life in which CPU time is the ”energy” resource and memory is the ”material”' resource. Memory is organized into informational patterns that exploit CPU time for self-replication. Mutations generate new forms, and evolution proceeds by natural selection as different genotypes compete for CPU time and memory space.”

In Core War there are the warriors (computer programs) and they try to take over as much system resources as possible. Two or more warriors compete in a simulated computer with a simulated instruction set, the warriors are positioned at random in the (simulated) memory and each warrior is given an execution thread each. The one warrior wins who has the most resources at the end. Warriors may also die, during fights.

[1] “Core” refers to the core (magnetic rings) computer memory.

Core War is an example of an approach for learning a specific goal. Here a strategy evolves explicitly and is not a gene sequence. This is a good application for genetic programming.

3.3 Game-Theoretic Models

There are many more evolutionary models out there then mentioned so far. I have found a very good article called "Evolutionary Dynamics in Game-Theoretic Models" [8] by Kristian Lindgren. In this article there is a good chart which shows different mechanisms and dimensionality of different evolutionary models. Biologist have found other examples for building of cooperation in the animal world. Hamilton [7] gave a famous example about a fish school. Many small fishes gather together when they are swimming out in the ocean. This seems in the first place not as a very good idea, since an individual small fish would have a higher chance of survival against a predator than in a big school. When alone it is quick and easily over seen by a big predator, but when in a big group it can be an easy victim, since the predator can eat many of the small fishes at once. So why do that then at all? Because it increase the survival of the swarm rather than an individual and raises the chances for reproduction, but also in a big school the small fishes appear from the distance as a big fish compared to other smaller predators and this formation then saves their life, since the other predator is not able to distinguish between the shape of a bigger predator or a school of many small fishes. In this case the school acts a defense against other predators. What is also very interesting, that a fish school is an example of a truly equalitarian, decentralized society in which a leader often trades places with those behind it, this happens especially when the fish school changes abruptly directors when chased by a predator. When we compare this to our human military organization, we see that those are not so flexible, they have often a centralized command structure.

Behaviors like in a fish school can for instance only be explained with an evolutionary system. It can explain observations from the real world which can not be explained considering individual agents (from the point of view of individual rationality).

4 An Evolutionary Framework for the Large-scale Testing of the MAGNET System[2]

After finalizing the theory behind our Expected Utility (EU)-based customer agent, as well as developing its supplier counterpart, we will devote efforts to testing them against various criteria. In particular, we are interested in testing how well individual agents interact in a populated market. This will help us understand the nuances of the application of EU to (Request for Quotes) RFQ generation. The major goals of this part of the study are:

- provide the statistical data necessary for the evaluation of the theoretical assumptions and derivations;
- facilitate the understanding of the nuances of the EU-based RFQ generation and to drive improvements to the theory and implementation;
- study the relative performance of agents in the simulated market, developing an understanding of the properties of automated and mixed-initiative combinatorial auction-based trading societies.

We will design our large-scale test suite atop an evolutionary approach to economic simulation. The structure of the simulation will be defined like this:

- The society will initially consist of one customer agent and many heterogeneous supplier agents. The choice of this setup is due to the assumption that there is no competition on the customer side, so customers can be replaced by one representative agent who issues RFQs

[2] This section is an extention of the correspondent section in the NSF02 proposal.

with high frequency. Each supplier will initially be provided with a "factory" that produces one type of good or service and maintains the schedule of production.

- The customer agent will issue an RFQ for one or several tasks according to a (stationary or not) Poisson process. Multiple RFQs will be open concurrently, so that each supplier must evaluate several RFQs at once. Upon receiving bids, the customer agent will find the winning bundle of bids and award bids. After that, it will monitor the execution of the plan and make appropriate payments to suppliers.
- The performance of supplier agents will be evaluated on the basis of profit averaged over a substantial period of time. The market will run in an evolutionary fashion, i.e. by removing suppliers with negative profits over periods of time, and introducing new suppliers with strategies from the pool of all available strategies whenever the average profit in the market exceeds some positive value.
- The information on successful bids and completed tasks will be collected, processed, and provided to the customer agent to be used in the RFQ generation and winner determination procedures.

The rationale behind our choice of an evolutionary framework is that it provides the necessary information without requiring any complex theory on agent motivation, optimization criteria, or strategic interaction. Unsuccessful species of supplier agents will be washed away from the market, creating places for the more fit. At the same time, the market will provide customer agents with dynamic information on supplier availability, market prices, and cumulative success probabilities.

Evolutionary frameworks have been used extensively in Economics [11, 12, 15]. The framework will allow us to tune the market by tweaking the frequency of issuing RFQs and will allow for the dynamic introduction of new supplier strategies, without imposing any assumptions on the nature of strategies. We will later extend the framework to support trade games to be played with human subjects. This will be a tool specially useful for teaching, as a tool to explore strategic behaviors and to study the emergence of cooperation [2, 3].

5 Conclusion and Future Work

If one knows which properties to optimize in a system than probabilistic algorithms, such as genetic algorithms with a fitness function are a good choice. On the contrary, if one has a very complex system with many parameters possible for optimization and one does not know which parameter to optimize, than building an evolutionary system is a good choice. In an evolutionary system one can observe phenomena which cannot be obsered in classical systems. In evolutionary systems there is no fitness function, but there is a function which removes unsuccessful members. The outcome of implementing an evolutionary system in MAGNET for example could be several different level of strategies, not only one optimal one. One could be a very fast one, but not so profitable for the customer, another could be more a deliberate one, but it would more profitable for the customer. Furthermore we could find strategies dependent on the size of the company. Intially we will distribute all the strategies of all different sizes of enterprises, but during the simulation (evolution) certain strategies will be apoted by a certain kind/size of the enterprise. We will observe that new behavior/strategies evolve over time.

References

[1] Scientific American. NOT JUST FUN AND GAMES. http://www.sciam.com/1999/0499issue/0499profile.html.

[2] R. M. Axelrod. *The evolution of cooperation.* Basic Books, 1984.

[3] Robert Axelrod. *The complexity of cooperation.* Princeton University Press, 1997.

[4] J. M. Baldwin. A new factor in evolution. http://www.santafe.edu/sfi/publications/Book-inforev/baldwin.html, 1896.

[5] Richard Dawkins. *The Blind Watchmaker.* Norton, 1987.

[6] Stephanie Forrest. Genetic algorithms: Principles of natural selection applied to computation. *Science*, 261:872–878, 1993.

[7] Janet T. Landa. Bioeconomics of some nonhuman and human societies: new institutional economics approach. In *Journal of Bioeconomics*, pages 95–113. Kluwer Academic Publishers, 1999.

[8] Kristian Lindgren. Evolutionary dynamics in game-theoretic models. In *The Economy as an Evolving Complex System II*, pages 337–367, 1997.

[9] Zbigniew Michalewicz. *Genetic Algorithms + Data Structures = Evolution Programs.* Springer, 1996.

[10] Tom M. Mitchell. *Mchine Learning.* McGRAW-HILL, 1997.

[11] Richard R. Nelson. Recent evolutionary theorizing about economic change. *Journal of Economic Literature*, 33(1):48–90, March 1995.

[12] David Rode. Market efficiency, decision processes, and evolutionary games. Department of Social and Decision Sciences, Carnegie Mellon University, March 1997.

[13] Stuart Russel and Peter Norvig. *Artificial Intelligence - A Modern Approach.* Prentice Hall, 1995. ISBN: 0131038052.

[14] Herbert Spencer. *The Principles of Sociology.* D. Appleton, 1897.

[15] Leigh Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. ISU Economics Working Paper No. 1, Department of Economics, Iowa State University, December 2001.

[16] Linus Thorsell. EVOLVING WARRIORS. http://www.dd.chalmers.se/~f95lith/CoreWar/Paper.html.

[17] S. W. Wilson. Knowledge growth in an artificial animal. In *In Procedings of the First International Conference on Genetic Algorithms and Their Applications*, pages 16–23, 1985.