

---

# Niching and Evolutionary Transitions in MAS

---

A. Defaweux,

T. Lenaerts,

S. Maes,

B. Manderick,

A. Nowé,

K. Tuyls,

P. van Remortel,

K. Verbeeck

Computational Modeling Lab

Vrije Universiteit Brussel

Belgium

## Abstract

In this paper we address two topics which are currently under investigation at our research lab. The first concerns the question of how cooperation can emerge in a system with antagonistic agents and how this can be modelled through a system of Reinforcement Learning(RL) agents. Current problems result from the fact that RL systems try to model all agents active in the environment. As a solution we are examining biological niching models and measures in order to reduce the complexity of the agent's learning model. The second topic is closely related to the first since it addresses the emergence of cooperating evolving groups: Evolutionary Transitions. We are convinced that the general mechanisms which can be found in the biological transition examples can be used to construct complex agents from simpler ones.

## 1 MAS and EC

When looking at current Multi-agent System(MAS) research there is a lot of discussion about how to define an agent (Wooldridge 1999), what the architecture of the agents will look like and how they will communicate with other agents in the environment (Huhns, 1999). Although MAS are considered to deploy a collection of agents to solve some particular problem, the perspective remains mainly individual-based, i.e. each agent is engineered in isolation. Once the agent architecture is defined, developing a MAS is difficult in the sense that the designer needs to be aware of all the parts that constitute the system and how these parts will interact. All interactions and their protocols are defined a priori.

There are a lot of applications for which MAS are the right approach, but which are too complex to be engineered. Since other agents are acting in the environment, it becomes inherently non-stationary. This causes an inability to define all different percepts and actions in advance. Therefore the agent has to learn how to react to changes in signals and how to interact with other agents in the environment (Sen, 1999).

Most learning is concerned with the construction of internal models of the environment which provide the agent with a mechanism for anticipation and prediction. Current approaches use existing Machine Learning(ML) techniques which are to some extent adapted to work in a MAS environment. In this context questions like how to extend the ML algorithm to handle communication and cooperation have to be answered. We believe that answers to these questions can be found by looking at biological systems.

In Evolutionary Computation(EC), models of biological systems are used to perform optimisation or learning tasks (Mitchell, 1997). EC is used when it is difficult to obtain an exact solution due to the complexity of the application. As opposed to MAS, this approach is population-based, i.e. EC evolves sets of simple individuals which have to behave optimally in some environmental setting. EC populations contain a large amount of simple, sometimes similar (redundant) individuals instead of a few hand-crafted, highly intelligent agents.

In EC there exists some god-like central control-mechanism which directs the entire evolutionary process. This mechanism is in contradiction with the autonomous nature of an agent in a MAS. Further, this central control mechanism is artificial in the sense that there is no biological counterpart which performs the same operation. In Biology, fitness and selection are an average result of the reproductive success of a species instead of a calculated property. Moreover, in stan-

standard EC algorithms, reciprocal interactions between the individuals in the population are ignored, except for some research on relative fitness calculation in the context of coevolutionary algorithms (Paredis, 1997) and niching (Mahfoud, 1995). These are only three reasons why the existing EC model needs to be adapted when using EC in MAS. Hence, in order to study the merits of evolution in MAS we need evolutionary systems which incorporate interactions and biological fitness. Such systems are similar to Complex Adaptive Systems (CAS) described in (Holland, 1995) and they define a context for MAS research from an EC point of view.

Our research group is interested in cooperation and transitions in CAS. Particularly we are interested in the origin and role of cooperation in an environment which is inherently competitive, i.e. every agent tries to reach its personal goals and since this happens in a common environment there will be competition if the number of resources is limited. For short, how can agents know with whom and when to cooperate? This topic is studied from a RL perspective and from an EC perspective. The answer to the question of the origin and role of cooperation is also important in modeling Evolutionary Transitions since cooperation provides the leverage for the construction of higher-level units. This topic will be covered in the second part of this abstract.

## 2 Niching as a Key to Cooperation

MAS are inherently decentralised, which makes the coordination of those systems usually very challenging. The agents or controllers can only access limited information about each other and the overall system. They have to base their decisions on partial information regarding the state of the system and/or on some private local knowledge.

In the Reinforcement Learning community this problem is theoretically studied on a small scale and the Markov game model (Litmann, 1994) (Hu, 1999) also called the stochastic game model is proposed as the underlying system. A MAS has characteristics of both a Markov Decision Problem (MDP) and Game Theory. The system can be in several states as in a MDP and has multiple agents whose actions collectively influence the reward as in a game. The Markov game model<sup>1</sup> is a direct extension of the above models for

<sup>1</sup>The Markov Game Model is defined by a set of states  $S$ , and a collection of action sets  $A_1, \dots, A_n$  (one set for *every* agent). The state transition function  $S \times A_1 \times \dots \times A_n \rightarrow P(\mathfrak{R})$  maps a state and an action from *every* agent to a probability distribution on  $S$ . Each agent has an associated

MAS. This model augments the MDP with actions that are distributed over the different agents as in the game theory model. Every step in the process, the system is in a certain state and a corresponding game has to be played. Although this model gives a natural mapping of the problem, learning in it is not trivial since the Markovian property is no longer valid.

In MDPs agents learn a value for an action in a certain state, while in stochastic games values are learned for combinations of actions, and learning is thus done in a product space. Since the size of this product space grows exponentially with the number of agents, this approach is not scalable. In our view the dimensionality of the product space can be limited by discarding agents which have little influence on the given reward. In other words, we are interested in finding out which combinations of agents have a relevant impact on the reward function. In examples from Biology, these are the individuals that are in competition for the same resources.

The concept of competing for resources is directly mappable to biological species with overlapping niches. A biological niche can be defined as *the unique position occupied by a particular species, conceived both in terms of actual physical area that it inhabits and the function that it performs in the environment* (Dictionary of Science). Species with non-overlapping niches may coexist side by side without competition, in a stable way. But two species with overlapping niches will be direct competitors. An obvious consequence is that the weaker species eventually will go extinct. However if there is a possibility for genetic variation in resource utilisation in one of the two species, divergence of resource use may emerge, and thus leaving the two species to co-exist (Mahfoud, 1995).

In this context it is important to introduce a way of measuring the niche overlap between agents, in order to decide whether or not to incorporate the actions of these agents in the product space. Once we know this measure, it can help an agent to find possible and interesting cooperators. More precisely, agents with a high degree of niche overlap are likely to share resources and goals and for this reason cooperation can benefit them.

We want to test these ideas by using the following experimental setup. Suppose a model including a fixed number ( $> 1$ ) of different resources (food) distributed over a two-dimensional grid, and a number of different agent species. Each cell can contain a quantity of each of the different resources, and each agent species

---

reward function  $R_i : S \times A_1 \times \dots \times A_n \rightarrow \mathfrak{R}$ .

needs one specific kind of resource to survive. Agents move over the grid and each time-step they try to acquire as much as possible of their species-specific resource from the cell where they are at that moment. When more than one agent is present in a cell during the same time-step, a *common pool resource game* (CPRG) (Gintis, 2000) with multiple resources is played.

Our idea is for an agent to learn from the CPRGs the agents that belong to the same niche, i.e., the agents needing the same resources. Then only this subset of agents which belong to the same niche will be modeled by the agent. It is our opinion that cooperative behavior may emerge between individuals belonging to the same niche, this as a consequence of the increased knowledge agents have concerning agents sharing the same resources.

The group behaviour of the agents (the turn-based bidding) is cooperative. Although this cooperative behaviour decreases the short-term gain of the agents, it can emerge in such a selfish context. This introduces the topic of evolutionary transitions which is discussed next.

### 3 Interactions, Transitions and Individuality

CAS are defined as systems made up of large collections of active agents that are diverse both in form and capabilities (Holland, 1995). These active agents are able to adapt to their environment and can interact with the other agents in the collection. The idea is that these interactions can result in evolutionary stable groups of agents in the sense that they evolve as if they were a new entity in the collection. The origin of such groups has a biological metaphor which is called an *Evolutionary Transition*. The easiest way to explain Evolutionary Transitions is by giving a well-known biological example used by Michod to describe the dynamics of evolutionary transitions (Michod, 1997): the transitions from molecular replicators to hypercycles, cfr. also (Eigen, 1979).

Imagine a well-mixed system of molecular replicators or genes. These replicators struggle to survive in an environment with a limited amount of resources, i.e. those elements which are well-designed for that environment will survive while the others will go extinct. Although this system is inherently competitive, networks of molecular replicators can emerge through cooperative interaction if the circumstances are right. In this example, interactions occur through catalytic proteins or through the surfaces of the other replicators

in the system. These two types of interaction allow cooperation to emerge.

However, these replicators will have the tendency to defect since cooperation will result in a decrease of their fitness and hence in their reproductive success. This decrease results from the fact that when a replicator is acting as an catalyst, it cannot reproduce or it loses energy required for reproduction. This leads to a lower density of that replicator which results in a decrease of its fitness. A decrease will not occur when the molecular replicators keep their solitary state. Hence, they have to gain more by behaving selfishly. Furthermore, in a system of interacting genes, the reproductive success of each gene also depends on the frequency of the other genes in the network.

Even though defection seems the way to go, these networks are able to prosper in their environment. So, how can these cooperative gene networks increase in number or how can these frequency-dependent effects be neutralised. Generally, Michod claims that natural selection is not able to overcome these effects and suggests spatial structuring and kin selection as a solution. Through the introduction of spatial effects, the spread of selfish replicators can be limited and conditions are provided to maintain genetic relatedness among the replicators. Kin selection can help to reduce the amount of defection in a group of replicators.

Even if a network emerges as a result of cooperation it always can disintegrate again because competition between the individual replicators never ends. To protect the network against parasitic genes, some mechanism has to evolve to protect it against disintegration. A biological example of such a mechanism is the fixation of the location of proteins through the construction of a protective cellular membrane resulting in a cellular structure. This results in a new entity or individual at a higher level which in turn try to prosper in the environment.

All biological examples of transitions use the same abstract processes; A collection of competing elements in which cooperation emerges and produces new selective units. It is our opinion that this process abstraction can be formulated as an algorithm for the creation of complex structures. This algorithm will be guided by the principles of conflict and cooperation inherent to the biological transition examples.

To build this abstract algorithm we first need to capture the mechanisms underlying these transitions in a model of cooperative problem solving. We divided the construction of the model in four steps:

**Step 1:** Create a system of interacting organisms which can evolve in a biological manner. Examine whether cooperation can emerge without explicit partner selection.

**Step 2:** Make the interaction between the organisms explicit in order to identify the networks of connected organisms. Again analyse whether cooperating groups can emerge in this setting.

**Step 3:** Introduce Michod's concept of group functionality, i.e. the reason groups are able to increase in number is a result of the idea that a group of agents has more functionalities than one simple agent

**Step 4:** Examine whether the cooperating groups remain stable in the sense that they can constitute a new entity at a higher level which is capable of differential survival of its own.

In answer to Step 1 we rebuild a model suggested by N. Packard (Packard, 1988) in which simple agents follow a foodgradient, consume food, reproduce and struggle to survive. For more detailed information we refer to (Packard, 1988). In this model, we introduced indirect interaction in the sense that instead of allowing the agent to consume all resources, their consumption was determined through a *n-player common pool resource game* (Gintis, 2000). The behaviour of the agent is genetically determined, i.e. an agent can be selfish or cooperative. The goal of this experiment was to examine whether it is true that cooperating agents can survive in a competitive environment. The results showed that this is the case. The difficulty in this simple experiment is the observation of the groups, i.e. Who is forming a cooperative group with whom.

Therefore we made the interaction and partner choice explicit in Step 2. The interaction with the chosen partner will remain as long as the agent is in the neighbourhood of its partner. When the social link exists and the agent consumes food, it will give a part (here 1/2) of its food to its partner. Hence it behaves altruistically. The goal of these experiments was again to examine whether in this setup altruistic agents can survive. Again the results showed that these altruists could maintain themselves.

To examine Step 3, we need variation in the behavior of bugs to introduce the idea of agent and group functionality. To do this, we removed the gradient search from the Packard model and allowed the bugs to use individual strategies for finding food. In our model, we added only one rule to the bug's genotype and moreover these rules were more simple: each time a bug

wants to move, the gradient is calculated and the result is compared to the if-clause of the bug's rule. If the clause is true, the bug moves accordingly otherwise it does not move at all. A bug pays both a movement and a metabolic tax. Thus, when a bug is at a location with few resources and does not move then it will die. In order to survive, bugs need to cooperate, i.e. to share their rules. An altruistic bug provides its partner not only with a part of its food but also with the rule it uses to survive. The partner thus gains extra rules which it can use to move to regions with higher amounts of resources. When the partner behaves selfishly it will gain extra resources and information from the others in the short run but it might not survive in the long run. Hence, survival is not a property of a single bug anymore but of the group of cooperating bugs. For the moment, we don't have enough results to draw reliable conclusions but we are convinced that we will obtain them in the near future.

Once a group of cooperating individuals emerged in Step 3 that group has to be protected against defecting members or external parasites which reap off the benefits of the group. This step, Step 4, is not completed yet, but once coordinated group movement is established we can examine the emergence of new evolutionary entities. In this context we are examining aggregation techniques which merge the group into a single unit. This aggregation should occur when the group remains active over a period of time, i.e. all components remain connected, eat enough to allow survival of the network and move in union.

Once the abstract algorithm is completely captured we will try to embed the model into Learning Classifier Systems and (Evolutionary) Reinforcement Learning systems and perform some experiments on typical MAS problems.

## References

- M. Eigen and P. Schuster (1979). *The Hypercycle: A Principle of Natural Self-Organization*, Springer-Verlag, Berlin.
- H. Gintis (2000). *Game Theory Evolving; A problem-centered introduction to modeling strategic interaction*, Princeton University Press, Princeton, New Jersey.
- J. H. Holland (1995) *Hidden Order, How Adaptation Builds Complexity*, Helix Books, Addison-Wesley Publishing Company, Inc
- Hu J. and Wellman M. P. (1999) Multi Agent Reinforcement Learning in Stochastic Games.
- M. N. Huhns and L. M. Stephens (1999). Multiagent

Systems and Societies of Agents. In Gerhard Weiss (ed.) *Multiagent systems, A modern approach to Distributed Artificial Intelligence*, 79-120, Cambridge MA: MIT Press.

M.L. Litmann (1994) Markov games as a framework for multi-agent reinforcement learning. In Proceedings of the Eleventh International Conference on Machine Learning, 157-163.

S. Mahfoud (1995). *Niching Methods for Genetic Algorithms*, Phd thesis, University of Illinois at Urbana-Champaign. Urbana, Illinois.

R. E. Michod (1997). *Darwinian Dynamics; Evolutionary Transitions in Fitness and Individuality*, Princeton Paperbacks, New Jersey.

M. Mitchell (1997). *An Introduction to Genetic Algorithms*. A Bradford Book, MIT Press, 3th edition.

N. H. Packard (1988). Intrinsic Adaptation in a Simple Model for Evolution. In C. Langton (ed.) *Artificial Life, SFI Studies in the Sciences of Complexity*, 141-155.

J. Paredis (1997). Coevolutionary Algorithms. In T. Bäck, D. Fogel, Z. Michalewicz (eds.) *The Handbook of Evolutionary Computation*, Oxford University Press.

S. Sen and G. Weiss (1999). Learning in Multiagent Systems. In Gerhard Weiss (ed.) *Multiagent systems, A modern approach to Distributed Artificial Intelligence*, 259-298, Cambridge MA: MIT Press.

J.M. Smith and Eörs Szathmáry (1997). *The Major Transitions in Evolution*, Oxford University Press

Michael Wooldridge (1999) . Intelligent Agents . In Gerhard Weiss (ed.), *Multiagent systems, A modern approach to Distributed Artificial Intelligence*, 27-78. Cambridge MA: MIT Press