

Surviving the Tragedy of Commons : Emergence of Altruism in a Population of Evolving Autonomous Agents

Jean-Marc Montanier¹ and Nicolas Bredeche¹

¹TAO - Univ. Paris-Sud, INRIA, CNRS, F-91405 Orsay, France

Abstract

This paper explores the following question: how a fixed-size population of autonomous agents (such as a swarm of robotic agents) may evolve altruistic behaviors during open-ended evolution. In particular, we focus on a situation where the tragedy of commons can possibly occur: a situation where individuals must display altruistic behaviors in order for the whole population to avoid extinction. Our approach considers a sub-individual framework, defined at the level of genomes rather than agents, in order to provide an efficient algorithmic solution for the emergent of coordination among the population. Experiments show that the proposed evolutionary adaptation algorithm favors the emergence of altruistic behavior under some assumptions regarding genome relatedness. In-depth experimental studies explore the relation between genotypic diversity and degree of altruism as well as the exact nature of the evolutionary adaptation process.

Introduction

Altruism is a remarkable behavior observed in Nature, where actions of an individual benefit other individuals even though these actions may negatively impact the individual's chances of survival. A well-known example is given by individuals that watch out for a predator and signal danger to the group whenever it is required, thus potentially drawing the predator's attention to them. The reason why some individuals may sacrifice themselves for the benefit of the group has long been studied and there are now some widely accepted theoretical basis regarding the relation between genotypic relatedness among individuals and degree of altruism, as first described by Hamilton (1964). Altruism has long been actively studied from Biology to Economics, from Sociology to Game Theory, to cite a few domains. It differs from cooperation as altruism requires no direct benefit nor reciprocity. Moreover, its benefit can only be measured at the level of the population, as summarized by Lehmann and Keller (2006).

This paper is concerned with the emergence of altruism in a fixed-size population of evolving autonomous agents where the environment is such that selfish behaviors lead to extinction. This situation is known as the tragedy of (unmanaged) commons, as introduced by Hardin (1968, 1994):

individuals must share a common limited resource, and possibly sacrifice their own benefit, so that the population survives through generations.

The main motivation behind this research is to propose a practical implementation of evolutionary adaptation in *a priori* unknown environments in the scope of a fixed-size population of autonomous agents. This assumption is central to our motivation as the long term goal is to provide practical algorithmic solutions that can be deployed in a swarm of virtual agents in complex environments as well as real world autonomous robots. The contribution in this paper is then both fundamental and practical as the emergence of altruism during the course of evolution is experimentally studied, with a particular focus on its causes and consequences, and is considered within an experimental setup that is closely related to the target application: a 2D virtual environment with realistic assumptions inspired from autonomous robotics.

The paper is organized as follow: the definitions of altruism and tragedy of commons are provided in the next section, along with a short description of relevant contributions from the fields of Artificial Life and Evolutionary Robotics. Then, the environment-driven evolutionary adaptation algorithm is described as well as the experimental settings used for the experiment. Results from the experiment are given and discussed, with a particular focus on the nature of altruism observed. Finally, the last section provides a discussion and conclusion and sketches future directions for this work.

Context and Motivation

This section starts with a definition of the Tragedy of Commons, a well-known social dilemma where the population welfare strongly depends on individual behaviors. Then, a definition of altruism is given as well as a brief overview of its theoretical foundations in Biology. The section ends with a short review of related works in the field of Artificial Life and Evolutionary Robotics.

The Tragedy of Commons

The tragedy of (unmanaged) commons ((Hardin, 1968, 1994)) is a particular kind of social dilemma where a pop-

ulation of individuals have access to a finite common resource pool: each individual may temporarily increase its fitness through selfish behavior, but this inevitably leads to exhaust the common resource pool, ultimately ending with population extinction. The classic example describes farmers optimizing their personal benefit by owning as many cows as possible without any regards for the common grazing the cows feed from, which will quickly suffer from over-exploitation, ending with cows dying from starvation.

The tragedy of commons has been widely studied in both Evolutionary Biology and Economics (Mankiw, 2009). Using a terminology from Economics, the conditions for the occurrence of the tragedy of commons requires that the resource must be accessible to anyone ("non-excludable") but in limited quantity, thus implying competition ("rivalry") among individuals. It shares some similarities with the well-known public goods dilemma¹ regarding the condition of unrestricted accessibility to the resource, but also differs as the subtractability of the resource may penalize the survival rate of the population (e.g. because of free-riders). From the Biology viewpoint, the tragedy of commons is known to be responsible of in-group competition among individuals.

A possible explanation for the tragedy of commons is the negative impact of reciprocity, where free-riders are favored as they focus on their own personal fitness gain with no regards to the cost at the level of the population (Sober, 1992). However, several strategies have been identified and discussed in the literature for "solving" the tragedy of commons: kin selection, policing (self-regulated punishment) or diminishing returns (population behavior depends on ecological feedback) are all good candidates observed in Nature (Rankin et al. (2007)).

Definition of Altruism

The emergence of cooperation and altruism has been the focus of a particular attention from many research fields, including of course Biology.

The distinction between cooperation with mutual benefit² (West et al., 2007) and "strong" altruism (termed altruism from now on) depends on the nature of the fitness benefit at the level of *either* the individual *or* the population (Lehmann and Keller, 2006). Cooperation implies that a given individual benefits from its behavior during its lifetime, either through direct or delayed (i.e. through repeated interactions) reciprocity. Altruism, on the other hand, characterizes the sacrifice of (part of) one own's fitness for the benefit of others. Therefore, an altruistic behavior bene-

¹In the public goods dilemma, individuals may choose to invest a part of their benefit for the group welfare.

²Cooperation is also sometimes used as a synonym for altruism (e.g. cooperation in the prisoner's dilemma corresponds to altruism (Sober, 1992)). In this paper, we assume the restricted and well-accepted definition of cooperation as a behavior leading to mutual benefit.

fits other individuals and possibly has a positive impact on longer time-scale (e.g. more than a single lifetime).

Several theories have been identified, covering different kinds of behavior observed in Nature, from mutualism to conditional cooperation. On the one hand, mutualism is the case where cooperation leads to direct benefit even though a single individual displays a cooperative behavior (Maynard Smith J., 1983; Lima, 1989; Packer C., 1988; Dugatkin and Wilson, 1992). On the other hand, the more classic conditional cooperation scheme implies that all individuals share the same cooperative strategy so that the whole population welfare is increased: kin selection (Hamilton, 1964; Maynard Smith, 1964), reciprocity (Trivers, 1971; Axelrod and Hamilton, 1981) or the more controversial group selection (Wynne-Edwards, 1986; Dugatkin, 1994; West et al., 2007) can be accounted for such conditional cooperation.

While the emergence of cooperation can be explained by the fact that every individuals benefit from such a behavior (i.e. no cost to cooperate), the justification for altruism is not as straight-forward. The idea of inclusive fitness proposed by Hamilton (1964) is now widely accepted to account for the emergence of altruism: inclusive fitness considers the fitness of a particular individual to depend both on its own behavior and the behavior of its close relatives. The basic idea is to consider individuals as vehicles for genes, therefore kinship must be taken into account rather than the sole interest of one individual/vehicle. Of course, sacrificing oneself depends on several parameters such as the expected fitness loss (from sacrifice) and benefit (for others) as well as the genotypic relatedness of the individuals concerned (closer relatives may imply increased altruistic behaviors).

Hamilton formalized the relationship between cost, benefit and relatedness in the following equation: $C/B < r$. The Cost C is the amount of fitness lost by an altruistic individual. The benefit B is the amount of fitness gained by the recipient that benefits from the altruistic behavior. And r is the genotypic relatedness between the two individuals. The term *kin selection* has been introduced by Maynard Smith (1964) to illustrate the mechanism and consequences with inclusive fitness: if one's individual is willing to sacrifice itself for closely related individuals, the gene responsible for such an altruistic behavior may spread through natural selection as it is likely to be present also in the genotypic material of its parents.

Models of Altruism in Artificial Life

Altruistic behavior, as well as the emergence of altruism, has also been investigated in the field of Artificial Life. All the major theories have been studied: kin selection (Sober, 1992; Leticia et al., 2004), group selection (Fletcher and Zwick, 2004, 2007) and other mechanisms such as effect of increased environment's viscosity (Mitteldorf and Wilson, 2000), communication (Ackley and Littman, 1994) and tag mechanism (Spector et al., 2004; Spector and Klein,

2006). Previous works have provided studies with various approaches, from game theoretic models to discrete and continuous virtual world simulations. Moreover, kin selection, reciprocity and group selection have been described as variations of a similar mechanism favoring the correlation of interaction between agents (Woodcock and Heath, 2002).

The emergence of altruism under specific condition have also been studied in virtual or real environments, in particular with respect to the public good dilemma (Connelly et al., 2010; Waibel et al., 2009) and to the tragedy of commons (Spector et al., 2004; Scogings and Hawick, 2008), with similar concerns for different selection schemes.

Waibel et al. (2009) discusses the ability to evolve altruism in team of homogeneous robots with group selection in a setup similar to the public good dilemma. Facing the same environmental conditions, Connelly et al. (2010) experimentally show that altruism naturally emerges as long as resources is widely available.

The tragedy of commons has been addressed by Spector et al. (2004), where tag recognition favors the interaction between altruistic agents facing a tragedy of commons, and by Scogings and Hawick (2008) in a prey-predator setup. Even though their work considered population with fixed strategy (rather than evolutionary adaptation), they illustrated the ability of altruistic population to survive in aggressive environment even when confronted to selfish individuals.

Method

In this paper, we are interested in identifying the emergence of altruism in the scope of environment-driven self-adaptation in a population of autonomous agents. The motivation behind this work is two-fold. Firstly, our long-term motivation targets the design of an evolutionary adaptation algorithm for a limited group of autonomous agents that is capable of facing *a priori* unknown situations. An important requirement is that the algorithm should be implementable in a virtual or real-world environment (e.g. multi-agent simulation, agents in virtual worlds, robot swarms).

Secondly, we ask the following question: what can be expected when a population of evolving agents faces the tragedy of commons. This implies to identify if a strategy emerges, but also the nature of this strategy, if any.

In this section, we describe the algorithm and the experimental setting used for this work. In particular, the experimental setting has been designed so that the population faces a setup where the tragedy of commons is expected to occur. Lastly, methodological tools for monitoring altruistic behaviors are introduced at the end of the section.

Algorithm

The mEDEA³ algorithm takes inspiration from the selfish gene metaphor popularized by Dawkins (1976) and performs as an evolutionary adaptation algorithm that can be

³*minimal Environment-driven Distributed Evol. Adaptation.*

distributed over a population of robotic agents (i.e. each agent in the population runs the same algorithm, but carries different genomes). It was first introduced by Bredeche and Montanier (2010) to address robustness issue with dynamic unknown environments and has been successfully validated on real e-puck autonomous robots (Bredeche et al. (2011)).

In this framework, each agent contains an *active* genome, which (indirectly) controls the agent's behavior, and a *reservoir of stored genomes*, which is empty at first. At each time step, each agent *broadcasts* in a limited range (approx. $1/32^{th}$ of the arena's width) a *slightly* mutated copy of its active genome (e.g. with gaussian mutation) and stores genomes received from neighbors, if close enough. At the end of a "lifetime" (i.e. a pre-defined number of time steps), each agent "forgets" its active genome and *randomly* picks one genome from its reservoir of stored genomes (if not empty). Then the reservoir is emptied, and a new lifetime starts. This algorithm is duplicated within each agent in the population, even though agents' behaviors differ depending on each agent's current active genome.

There are three major claims why this algorithm works. Firstly, selection pressure occurs at the population level (the more a genome spreads itself, the higher its fitness) rather than at the individual level (random sampling). Secondly, genomes survive only through spreading (as an active genome is automatically deleted locally at the end of a generation). Thirdly, individual fitness improves over time as conservative variations generate new candidates that explore alternative (but closely related) behavioral strategies.

In practical, this algorithm provides an evolutionary adaptation mechanism, but does not provide a control function. The actual control of the agent behavior shall be performed by a dedicated controller whose parameters are determined from the genome. In other words, the mEDEA algorithm provides evolutionary adaptation by tuning the control architecture. In the rest of this paper, the controller used is a Multilayer Perceptron whose weights are decoded from the genome (more details in the next Section).

The mEDEA algorithm shares some similarity with the basic concepts demonstrated in Tierra (Ray, 1991), AVIDA (Adami et al., 1994) and followers, but also differs as it was originally designed for real world environments with a limited number of moving autonomous agents such as mobile robots. It can also be related to Embodied Evolutionary Robotics (Watson et al., 2002) regarding the possible implementation on physical agents, but with the major difference that it is not meant to optimize a pre-defined objective function.

Experimental Setup

In order to account for the existence of altruism, we have defined a foraging task where a population of autonomous agents must eat food items to maintain a positive energy level. The experimental setup used in the next section is il-

illustrated in figure 1, with food items (circles), agents (small dots) and obstacles. The environment and task depends on the following elements: (1) *Self-sustainability*: foraging is necessary to survive, as each food items give a small amount of battery energy. However, an agent's battery is limited to a maximum amount of energy, and foraging may end up in wasting resource. (2) *Foraging behavior*: an agent may choose to harvest *all* or *part* of a food item. (3) *Re-grow rate*: whenever a food item is harvested, it is removed from the environment until it grows back after some delay. The time to grow back depends on the quantity of energy harvested from the food item.

As a consequence, the environment features a common resource pool for which agents compete: a perfect setup for the Tragedy of Commons to occur. Indeed, it is then enough to set the appropriate delay before a given food item would grow back. This is achieved by setting the maximum re-grow delay for a food item (EP_{LagMax} , with EP as in "Energy Point"), which in turn will be used to compute on-the-fly the re-grow delay of a food item that was just harvested (EP_{Lag}). This is described in equation 1, which also takes into account the amount of energy harvested by an agent from the food item ($E_{harvested}$) and the amount of energy available in each food item (EP_{eMax}).

$$EP_{Lag} = E_{harvested} / EP_{eMax} * EP_{LagMax} \quad (1)$$

Within this setup, it is expected that altruistic agents in aggressive environments shall harvest the minimum amount of energy from each food items, therefore increasing the availability of the resource (short re-grow delay, no wasted energy). On the other hand, selfish behaviors are likely to be fitted for small values of EP_{LagMax} , but are expected to become more and more critical as the value of EP_{LagMax} increases.

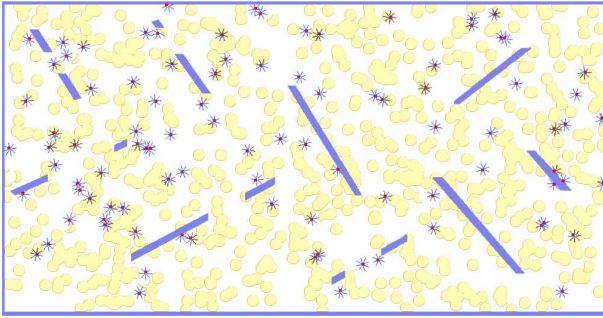


Figure 1: Snapshot from the simulator: food items (circles), agents (dots) and obstacles

Methodology

In order to account for altruism, we define a measure for monitoring the *cost of altruism* for one foraging agent. In the setup described earlier, this corresponds to measuring the

amount of energy that *could be* consumed when harvesting a food item, but which is actually *not consumed* by the agent. This is formally defined in equation 2.

$$Cost = max(0, min(EP_{eMax}, r_{E_{max}} - r_{E_{now}}) - E_{harvested}) \quad (2)$$

Where EP_{eMax} is defined as before (i.e. maximal energy in a food item), $r_{E_{max}}$ is the maximal energy level of an agent, $r_{E_{now}}$ is the current energy level of the agent and $E_{harvested}$ is the energy harvested by the agent from the food item.

While a selfish agent shall have a cost of zero, an altruistic agent should be able to perform a trade-off between its altruistic nature and its survival needs. Therefore, the cost of altruism can be seen as the agent's level of sacrifice which is continuous (a quantity of energy) rather than discrete (eat or dont eat).

Results and Analysis

This section presents results obtained running the mEDEA algorithm in the environment described in the previous section. The organization of the section is as follow: the algorithm is evaluated for its ability to evolve agents with altruistic behavior. Then, the nature of altruistic behavior is investigated, considering the balance between environmental pressure and the algorithm's mechanisms. Finally, the relation between genotypic relatedness and the degree of altruism is explored along with its impact on the survival rate of the population.

All experiments were conducted with 100 robotic agents in the environment described and illustrated in the previous section. The environment contains 800 food items and an agent may harvest a maximum of 50 units from a food item. Each agent consumes 1 unit of energy per step, and can store up to 800 energy units (harvesting surplus is lost). If the agent's battery level drops to zero, the agent stops and its genome is lost. It is then refilled with a small portion of energy, but remains still until it receives a new genome.

The control architecture is a Multilayer Perceptron (MLP) with 5 hidden neurons, 11 inputs (8 proximity sensors, battery level and orientation/distance to the closest food item) and 3 outputs (left/right motor and proportion of energy to be harvested from a food item, if any). The weights of the MLP are decoded from the active genome of the agent. Each agent broadcasts a mutated copy of its own genome and receives genomes from neighbors within a limited range (roughly 1/10th of the length of the larger side of the environment). The mutation operator used in the Medea algorithm is defined as a gaussian mutation with a σ parameter. σ is included into the genome (i.e. similar to a self-adaptive Evolution Strategy) and ranges from 0.01 (low mutation rate) to 0.5 (large mutation rate).

All results shown here have been achieved in ROBOROBO,

a fast 2D simulation for robotic agents, originally introduced by Bredeche and Montanier (2010). The source code for reproducing the experiments is freely available for download (<http://www.lri.fr/~montanier/roborobo-ecal>). For each experimental settings, a set of 600 independent runs have been performed during 320000 iterations (= 800 generations) to provide statistically significant data.

Emergence of Altruism in Medea

A large set of experiments was performed under various environmental pressures by setting a specific value of EP_{LagMax} for each run, ranging from 25 steps (easy environment) to 400 steps (aggressive environment), for a total of 16 setups. For each setup (i.e. a fixed value of EP_{LagMax}), 600 independent runs were performed and results were aggregated to extract various indicators: number of active agents, average cost measure and energy balance (i.e. a positive value means agents harvest more than the minimal requirement). In all experiments, the course of evolution is similar: the number of active agents quickly increases to a stable value while costs start from random values and stabilize to (possibly) positive values. While the increasing number of active agents is expected from evolutionary adaptation, the second observation is of primary importance regarding the possibility of altruistic behavior: a positive cost value would imply that agents do not systematically harvest all possible energy from the food items.

Results are summarized in figures 3(a), 3(b) and 2 (resp. number of active agents, cost measure and energy balance), by taking into consideration the last 10 generations of all runs for each setup (i.e. after convergence to stable behaviors). Altruistic behavior in the context of increasing environmental pressure can be observed by looking at the cost, which converges to a stable value, while the energy balance converges to zero (i.e. the limit for survival). Indeed, altruistic behaviors are observed starting with environments with $EP_{LagMax} = 100$, and remains afterwards. With stronger environmental pressures (larger values of EP_{LagMax}), the number of active agents decreases, which confirms that the environment is becoming more and more challenging.

Several observations can be drawn from these results. Firstly, altruistic behaviors are difficult to observe when environmental pressure is low and tragedy of commons not bound to occur (median values are close to zero for values of EP_{LagMax} under 100 steps). This tends to reveal the greedy nature of the algorithm: without environmental pressure, altruism does not emerge spontaneously. In fact, it is possible to classify the individuals' behavioral patterns with respect to (a) their fellow agents (*selfish vs. altruistic* behavior) and (b) the environment (*frugal vs. greedy* behavior). the mEDEA algorithm tends to generate *greedy* but *altruistic* agents depending on the environment at hand. Secondly, altruistic behaviors remain stable in the population even though the environmental pressure increases and

the number of active agents starts to drop, implying limited correlation between the level of altruism and environmental pressure. This is explored in the following.

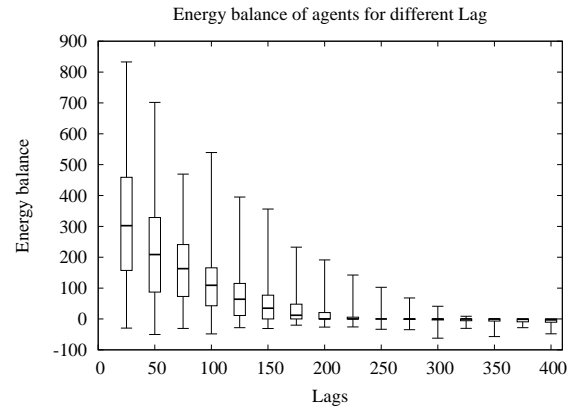


Figure 2: Results with EP_{LagMax} between 25 and 400 : Energy Balance (data: boxplots are drawn from the median values from each run, i.e. for each run, some agents (not shown) are likely to have larger positive energy balances)

Investigating the Nature of Altruism

In order to explore the dynamics of the algorithm, a first experiment is designed to evaluate its ability to converge towards the same results from different initial conditions. Starting with a population of agents already evolved in a challenging setup ($EP_{LagMax} = 400$, strong pressure, used during 1000 generations), the population is abruptly changed to a smoother environment ($EP_{LagMax} = 200$, moderate pressure) and re-adaptation (if any) is studied. The expected outcome is that the number of active agents and the cost measure should converge back to the expected values (shown before). This is indeed what is observed, as shown in figure 4, advocating for the robustness with regards to initial conditions, at least in this case (i.e. starting from already evolved genomes rather than pure random genomes). This is also confirmed by a Mann-Whitney's statistical test.

However, a careful analysis of the results reveals a surprising feature occurring when the environmental pressure is changed: the number of active agents rises significantly before going back down to its final stable value. The same holds for the cost measure, as a sudden drop is observed, preceding a slow convergence to the expected, higher, value. This is indeed a surprise as, for a brief moment, individuals actually have a better survival rate even though more egoistic behaviors are monitored. A closer look at the results in the close vicinity of the change in the environment (not visible at this resolution) actually confirms this: after the environmental change, the number of active agents (resp. cost measure) quickly rises (resp. drops), before slowly converging back

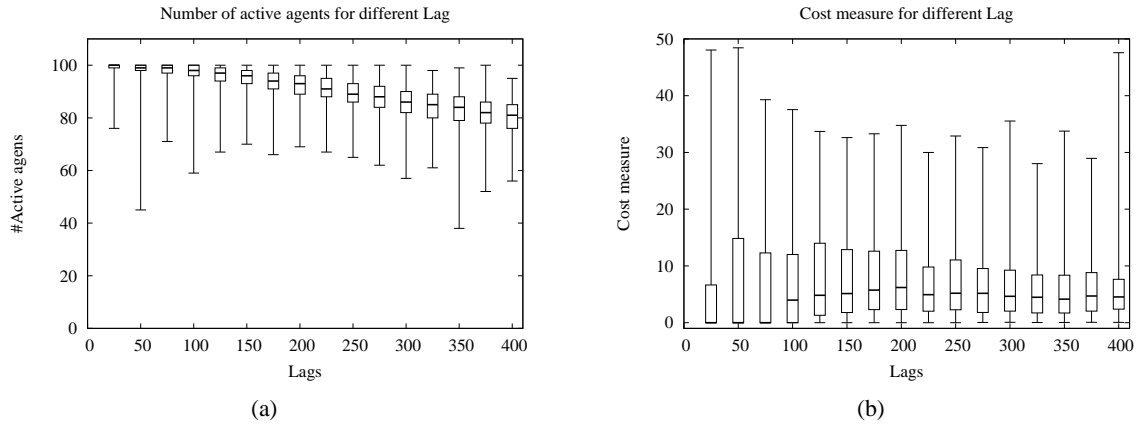


Figure 3: Results with EP_{LagMax} between 25 and 400: *a*) Number of active robots (data: value from each run) ; *b*) Cost measure (data: median values from each run)

to its final expected value.

A candidate hypothesis for explaining the algorithm's behavior is to reconsider the very nature of *what* can be stated as its intrinsic motivation: mEDEA may be performing a trade-off between survival *and* stability of evolutionary dynamics, rather than survival only. In order to investigate this hypothesis, we define a measure of evolutionary stability that takes into account the number of ancestors from a previous generation for individuals of the current generation (i.e. the larger the number, the more the ancestor with one offspring only). Larger numbers imply a more stable population as it means that more genomes actually survived through their offsprings. In other words, a population with many ancestors imply lack of selective pressure. In practical, this is defined as follow: $nbStrains_{gen=N-b}/nbActiveAgents_{gen=N}$, with $nbStrains$ the number of ancestors from b generations ago with at least one descendant in the current generation. The value is normalized in $[0, 1]$. Lower values imply increased selective pressure.

Figure 5 tracks this value for a few generations: for each generation (i.e. each boxplot), the (normalized) number of ancestors from $b = 10$ generations ago with at least one offspring in the current generation is drawn. During the short increase in performance after the environmental change, the number of ancestors decreases for at least 10 generations, which indicates that fewer genomes actually benefited from a stronger selective advantage. However, selective pressure then goes back to a more conservative level, even though behaviors end up being sub-optimal with respect to survival (as shown before). Why the best genomes for survival do not remain in the population is yet to be fully understood. In this context, it is likely that egoistic agents may only temporarily benefit from the change, as they may not be enough in numbers to take over the population before altruistic agents adapt to the new environment. Indeed, very specific initial con-

ditions (forcing egoistic behavior at start-up) or dedicated mechanisms in the algorithm (see next section for a discussion) may be required to obtain the best population wrt. surviving rate.

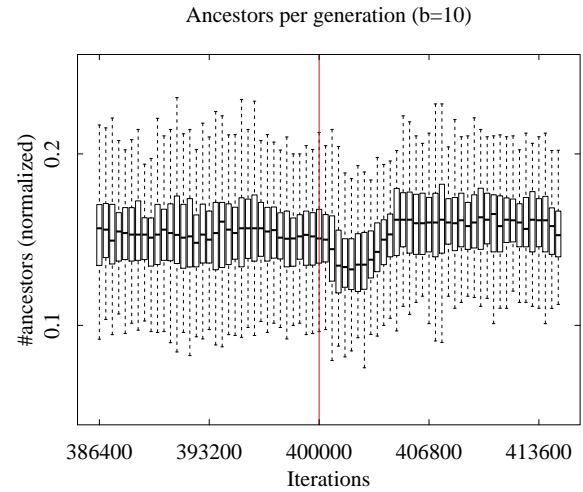


Figure 5: Ancestors from generation $N - 10$ with at least one offspring in the current generation (34 generations before and after the change are shown).

Discussion on Diversity and Altruism

As stated previously, it is likely that selective pressure acts in favor of a trade-off between optimizing survival and algorithmic internal stability. But what happens if one were to deliberately enforce genotypic homogeneity? In the following, we address this question and discuss its possible implications. The motivation is two-fold: firstly, the goal is to explore the relation between genotypic homogeneity, level of altruism and survival rate. Secondly, part of the answer

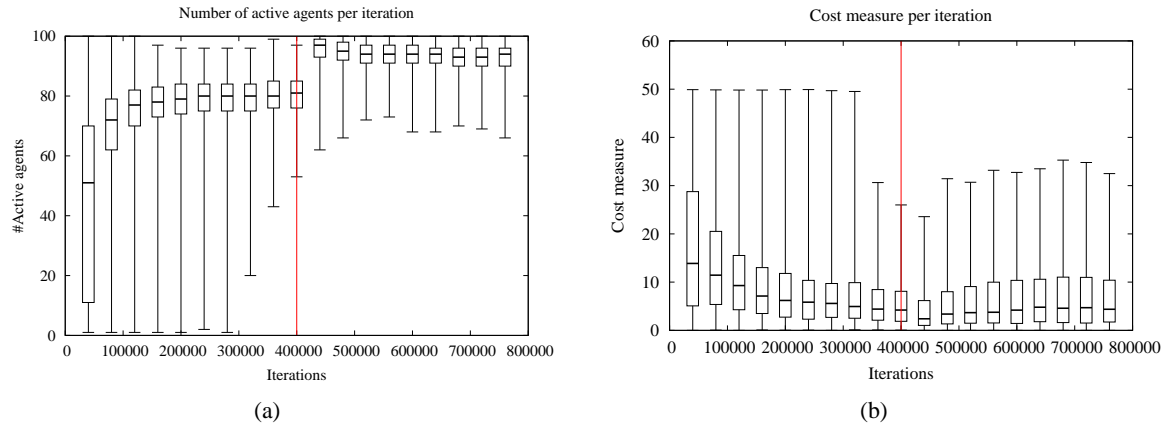


Figure 4: Environment change from strong to moderate pressure (see text). *a)* Number of active robots ; *b)* Cost measure

to this question is a first step towards controlling the evolutionary dynamics at work in the algorithm.

A set of additional experiments have been performed where genotypic relatedness is favored during the selection process, in order to decrease genotypic distance among individuals in the population. In practical, the algorithm’s random selection that is embedded in each agent is replaced by a tournament selection (Miller and Goldberg, 1995) (also embedded in each agent), where ranking is based on the genotypic (euclidian) distance between the previously active genome and the locally available genomes (the closer, the better). Tournament selection combined with genotypic distance (termed kin-tournament from now on) makes it possible to introduce an explicit pressure towards kin selection, which can easily be tuned by the size of the tournament.

Experiments with a tournament size of 3 (roughly corresponding to medium pressure towards kin selection) have been achieved with two setups, one with moderate environmental pressure ($EP_{Lag_{Max}} = 200$) and the other with a strong pressure ($EP_{Lag_{Max}} = 400$). For each setup, 200 runs were performed, and statistical test are computed with Mann-Whitney’s Test to clearly establish the difference in performance. Performing kin selection increases the level of altruism in both cases (roughly doubling it, $p - value < 10^{-15}$). While the number of runs with extinctions is roughly similar ($p - value = 0.07$ for $EP_{Lag_{Max}} = 200$, and $p - value = 0.71$ for $EP_{Lag_{Max}} = 400$), enforced kin selection suffers from a smaller number of active agents ($p - value < 10^{-15}$).

These results can be put in perspective with Hamilton’s idea of inclusive fitness (Hamilton, 1964). The intrinsic mechanisms in the algorithm, in particular conservative mutation, already imply a strong genotypic relation between one genome and its offsprings. Kin selection is shown to artificially increase the already existing level of altruism, at the cost of a decreased overall performance wrt. to individual survival. This is not a surprise as altruistic behaviors

were already shown previously to lead to sub-optimal survival rate, which is even more critical when environmental pressure is aggressive. Nevertheless, the kin-tournament selection as proposed here actually does provide an interesting tunable mechanism to act on the level of altruism, and could possibly lead to a more competitive, heterogeneous population if kin selection is penalized rather than favored.

Conclusions and Perspectives

In this paper, we investigated evolutionary adaptation in a population of robotic agents whenever altruistic behaviors are mandatory to survive. The algorithm under scrutiny was shown to naturally evolve greedy-altruistic agents within aggressive environments (ie. greedy behavior whenever it does not impact the survival rate of the population). An important message from this paper is that evolutionary adaptation in this context does not automatically lead to the best survival strategy but rather converge towards a trade-off between algorithmic stability and survival. Also, the relation between genotypic relatedness and the level of altruism was confirmed and a possible mechanism to control the level of altruism has been identified.

Perspectives from this work include deeper investigation regarding the exact causes of the sub-optimal survival strategies obtained. Moreover, tuning the level of altruism offers interesting perspectives with regards to modeling environmental-feedback induced altruistic behaviors, such as *diminishing returns*, where altruism may be regulated by the environment (Rankin et al., 2007).

Acknowledgements

This work was made possible by the European Union FET Proactive Initiative: Pervasive Adaptation funding the Symbion project under grant agreement 216342. Some of the experiments presented in this paper were carried out using the Grid’5000 experimental testbed (cf. <https://www.grid5000.fr>). We’d like to thank Simon Carrignon for his last minute (much appreciated) help with generating figure 5.

References

- Ackley, D. H. and Littman, M. L. (1994). Altruism in the evolution of communication. In Brooks, R. and Maes, P., editors, *Artificial Life IV*, pages 40–48, Cambridge, MA. MIT Press.
- Adami, C., Brown, C. T., and Kellogg, W. K. (1994). Evolutionary Learning in the 2D Artificial Life System Avida. In *Artificial Life IV*, pages 377–381.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211(4489):1390–1396.
- Bredeche, N. and Montanier, J.-M. (2010). Environment-driven Embodied Evolution in a Population of Autonomous Agents. In Springer, editor, *The 11th International Conference on Parallel Problem Solving From Nature (PPSN 2010)*, pages 290–299.
- Bredeche, N., Montanier, J.-M., Wenguo, L., and Alan FT, W. (2011). Environment-driven Distributed Evolutionary Adaptation in a Population of Autonomous Robotic Agents. *Mathematical and Computer Modelling of Dynamical Systems*.
- Connelly, B. D., Beckmann, B. E., and McKinley, P. K. (2010). Resource abundance promotes the evolution of public goods cooperation. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 143–150.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.
- Dugatkin, L. A. (1994). Subtle ways of shifting the balance in favor of between-group selection. *Behavioral and Brain Sciences*, 17:618–619.
- Dugatkin, L. A. and Wilson, D. S. (1992). The prerequisites for strategic behaviour in bluegill sunfish, *leptomis macrochirus*. *Animal Behaviour*, 44(Part 2):223 – 230.
- Fletcher, J. A. and Zwick, M. (2004). Strong altruism can evolve in randomly formed groups. *Journal of Theoretical Biology*, 228(3):303 – 313.
- Fletcher, J. A. and Zwick, M. (2007). The evolution of altruism: Game theory in multilevel selection and inclusive fitness. *Journal of Theoretical Biology*, 245(1):26 – 36.
- Hamilton, W. (1964). The genetical evolution of social behaviour. i. *Journal of Theoretical Biology*, 7(1):1–16.
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162:1243–1248.
- Hardin, G. (1994). The tragedy of the unmanaged commons. *Trends in Ecology & Evolution*, 9(5):199–199.
- Lehmann, L. and Keller, L. (2006). The evolution of cooperation and altruism - a general framework and a classification of models. *Journal of Evolutionary Biology*, 19(5):1365–1376.
- Leticia, A., A. F. J., and Cutter, A. D. (2004). The kin composition of social groups: trading group size for degree of altruism. *Am Nat*, 164(2):132–44.
- Lima, S. L. (1989). Iterated Prisoner's Dilemma: An Approach to Evolutionarily Stable Cooperation. *American Naturalist*, 134:828–834.
- Mankiw (2009). *Principles of economics, 3rd edition*. South-Western College Pub, february edition.
- Maynard Smith, J. (1964). Kin selection and group selection. *Nature*, 201:1145.
- Maynard Smith J. (1983). *Game theory and the evolution of cooperation*, page 445 456. Cambridge University Press, Cambridge.
- Miller, B. L. and Goldberg, D. E. (1995). Genetic algorithms, tournament selection, and the effects of noise. *Complex Systems*, 9:193–212.
- Mitteldorf, J. and Wilson, D. S. (2000). Population viscosity and the evolution of altruism. *Journal of Theoretical Biology*, pages 481–496.
- Packer C. (1988). Constraints on the evolution of reciprocity : lessons from cooperative hunting. *Ethology Sociobiology*, 9:137–147.
- Rankin, D., Bargum, K., and Kokko, H. (2007). The tragedy of the commons in evolutionary biology. *Trends in Ecology & Evolution*, 22(12):643–651.
- Ray, T. S. (1991). An approach to the synthesis of life. In Langton, C., Taylor, C., Farmer, J. D., and Rasmussen, S., editors, *Artificial Life II*, volume XI of *Santa Fe Institute. Studies in the Sciences of Complexity*, pages 371–408. Addison-Wesley, Redwood City, CA.
- Scogings, C. and Hawick, K. (2008). Altruism amongst spatial predator-prey animats. In Bullock, S., Noble, J., Watson, R., and Bedau, M. A., editors, *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, pages 537–544. MIT Press, Cambridge, MA.
- Sober, E. (1992). The evolution of altruism: Correlation, cost, and benefit. *Biology and Philosophy*, 7(2).
- Spector, L. and Klein, J. (2006). Genetic stability and territorial structure facilitate the evolution of tag-mediated altruism. *Artificial Life*, 12.
- Spector, L., Klein, J., and Perry, C. (2004). Tags and the evolution of cooperation in complex environments. In *Proceedings of the AAI 2004 Symposium on Artificial Multiagent Learning*. AAI Press.
- Trivers, R. L. (1971). The evolution of reciprocal altruism. *The Quarterly Review of Biology*, 46(1):35.
- Waibel, M., Keller, L., Floreano, D., and Member, S. (2009). Genetic team composition and level of selection in the evolution of multi-agent systems. *IEEE Transactions on Evolutionary Computation*.
- Watson, R., Ficici, S., and Pollack, J. (2002). Embodied evolution: Distributing an evolutionary algorithm in a population of robots.
- West, S. A., Griffin, A. S., and Gardner, A. (2007). Social semantics: altruism, cooperation, mutualism, strong reciprocity and group selection. *Journal of Evolutionary Biology*, 20:415–432.
- Woodcock, S. and Heath, J. (2002). The robustness of altruism as an evolutionary strategy. *Biology and Philosophy*, 17:567–590. 10.1023/A:1020598804674.
- Wynne-Edwards, V. C. (1986). *Evolution through group selection*. Blackwell Scientific, Oxford.