Emergence of Altruism in Open-ended Evolution in a Population of Autonomous Agents

TRACK : Artificial Life / Robotics / Evolvable Hardware

Jean-Marc Montanier
TAO - Univ. Paris-Sud, INRIA, CNRS
F-91405 Orsay, France
montanier@lri.fr

Nicolas Bredeche
TAO - Univ. Paris-Sud, INRIA, CNRS
F-91405 Orsay, France
bredeche@lri.fr

ABSTRACT

This paper summarizes recent works on the evolution of altruism to solve the tragedy of commons in the context of open-ended evolution with a fixed number of robotic agents.

Categories and Subject Descriptors: I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence

General Terms

Algorithms, Experimentation

Keywords

Artificial Life, Embodied Evolution, Evolutionary Adaptation, Altruism, Tragedy of Commons

1. INTRODUCTION

Altruism can be observed whenever a specific individual in a population deliberately sacrifices part of its own fitness in order to increase the fitness of other individuals. The reason why some individuals may sacrifice themselves for the benefit of the group has been long studied and there are now some widely accepted theoretical basis based on genotypic relatedness, such as the idea of inclusive fitness [4]. It defers from cooperation as altruism requires no direct benefit nor reciprocity, and its benefit can only be measured at the level of the population [6].

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 [5]: individuals must share a common limited resource, and possibly sacrifice their own benefit, so that the population survives through generations. The objective is the following: given a simple environment-driven evolutionary adaptation algorithm distributed over a swarm of agents, can we observe the emergence of altruistic behaviors?

2. METHOD

The mEDEA[1] algorithm, first introduced in [1], takes inspiration from the selfish gene metaphor [3]. Various properties of this algorithm have already been studied with regards to robustness to environmental changes and emergence of behavioral consensus, both in simulation and with real world autonomous robots [2].

This algorithm is defined as follow: 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* a (slightly) mutated copy of its active genome (e.g. with gaussian mutation) and stores genomes received from neighbors, if any. At the end of a "lifetime" (ie. 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.

3. EXPERIMENTS

3.1 Experimental Setup

Experiments were conducted with 100 robotic agents in a 2D simulated environment (see figure 1). 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 (ie. similar to a self-adaptive Evolution Strategy) and ranges from 0.01 (low mutation rate) to 0.5 (large mutation rate).

In order to account for altruism, we introduce the notion of *cost of altruism* for one agent foraging behavior. This corresponds to monitoring the amount of energy that *could*

[1] *minimal Environment-driven Distributed Evol. Adaptation*

GECCO'11, July 12–16, 2011, Dublin, Ireland.
ACM 978-1-4503-0690-4/11/07.

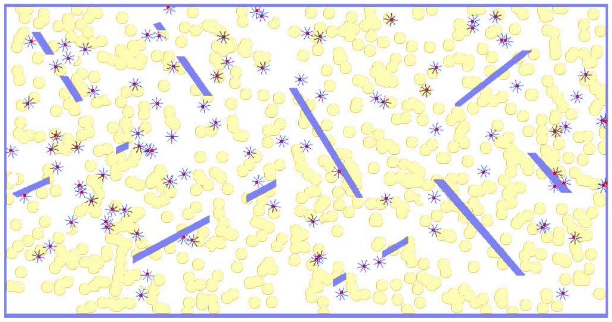

Figure 1: Snapshot from the simulator: food items (circles), agents (dots) and obstacles

be consumed when harvesting a food item, but which is actually *not consumed* by the agent. This is formally defined in equation 1.

$$Cost = max(0, min(EP_{max}, r_{E_{max}} - r_{E_{current}}) - E_{harvested}) \quad (1)$$

Where EP_{max} is defined as the maximal energy available from a food item, $r_{E_{max}}$ is the maximal energy level of an agent, $r_{E_{current}}$ 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 die). Being selfish or egoistic is a critical issue as a food item grow back *faster* if not *all* of its content has been harvested.

3.2 Results

Figures 2 and 3 respectively show the evolution of the number of active agents and the cost through generations in the context of an environment where egoistic individuals negatively impact the population ($EP_{max} = 200$, $r_{E_{max}} = 400$).

These results show that the number of active agents quickly increases to its maximum while the cost starts from a large value and converges to a stable, non-zero, value. 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 non-zero cost value implies that agents do not systematically harvest all possible energy from the food items. In other words, the resulting behavior can be qualified as altruistic as it implies a loss from the individual viewpoint, even after stabilization.

4. CONCLUSIONS AND PERSPECTIVES

mEDEA, an evolutionary adaptation algorithm dedicated to fixed-size population of autonomous agents, was shown to naturally evolve altruistic agents within an aggressive environment. Several other future directions can be identified, from further experiments regarding the exact nature of altruistic behaviors (on-going work) to reproducing the same results with physical robots.

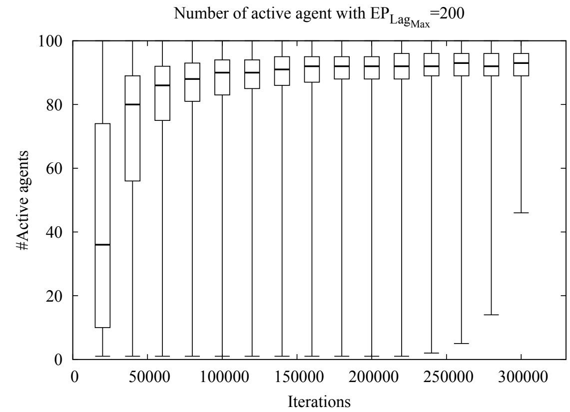

Figure 2: Number of active robots

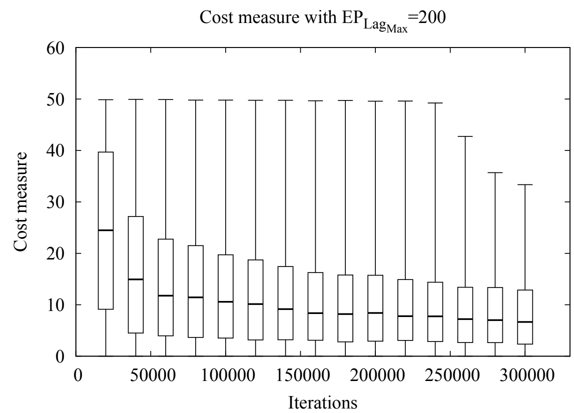

Figure 3: Cost measure

5. REFERENCES

[1] Nicolas Bredeche and Jean-Marc Montanier. 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, 2010.

[2] Nicolas Bredeche, Jean-Marc Montanier, Liu Wenguo, and Winfield Alan FT. Environment-driven Distributed Evolutionary Adaptation in a Population of Autonomous Robotic Agents. *Mathematical and Computer Modelling of Dynamical Systems*, 2011.

[3] Richard Dawkins. *The Selfish Gene*. Oxford University Press, 1976.

[4] W. Hamilton. The genetical evolution of social behaviour. i. *Journal of Theoretical Biology*, 7(1):1–16, July 1964.

[5] Garrett Hardin. The tragedy of the commons. *Science*, 162:1243–1248, December 1968.

[6] L. Lehmann and L. Keller. The evolution of cooperation and altruism - a general framework and a classification of models. *Journal of Evolutionary Biology*, 19(5):1365–1376, 2006.