

Evolution of Altruism and Spatial Dispersion: an Artificial Evolutionary Ecology Approach

Jean-Marc Montanier¹, Nicolas Bredeche^{2,3}

¹NTNU, Trondheim, Norway
jean-marc.montanier@idi.ntnu.no

²UPMC Univ Paris 06, UMR 7222, ISIR, F-75005, Paris, France – ³CNRS, UMR 7222, ISIR, F-75005, Paris, France
nicolas.bredeche@isir.upmc.fr

Abstract

This paper concerns the evolution of altruism in a population of autonomous agents. It explores the relation between altruistic behaviours and spatial dispersion in open-ended evolution whenever energetic constraints must be addressed. A method derived from Embodied Evolution is used to model the spatial interactions between agents from an individual perspective. Firstly, results show that spatial dispersion and levels of altruism are strongly correlated, which confirms theoretical results from biology, but also that this relation may be overshadowed by the complex interactions at work in the ecosystem. Secondly, this paper investigates how robust altruistic behaviours able to cope with various environmental pressures may be evolved. In particular, it is shown that there is a trade-off between efficiency and versatility: the ability to perform well across a wide range of environmental conditions often comes at the cost of sub-optimal performance in terms of survival, especially when compared to more constrained (and less versatile) evolved strategies.

Introduction

Cooperative behaviours are defined by the realisation of actions by an agent which brings benefits to other agents. This paper focuses on altruism, which is a special type of cooperative behaviour characterised by the sacrifice of one agent for the benefit of others (Lehmann and Keller, 2006). This is different from mutual cooperation as the fitness of an agent is permanently impaired by such acts. This type of behaviour may look counter-intuitive from the viewpoint of the theory of evolution as an altruistic agent reduces its own chance of survival in order to increase the chance of survival of other agents. However, such altruistic behaviour is found in multiple biological species, and multiple works have investigated the conditions and mechanisms at play in its evolution.

A common way to explain the evolution of altruism in biology is to consider the survival of genes rather than individuals. In this context, individuals are vehicles for genes that try to survive (Dawkins, 1976). This aspect is captured by the idea of the inclusive fitness which considers that the fitness of a particular individual depends both from its own action and from the actions of its related kin (Maynard Smith,

1964).

Classic approaches in game theory (Maynard Smith, 1974) and adaptive dynamics (Diekmann, 2004) have been used to explore multiple causes that favour the evolution of altruism. The most studied mechanisms are kin selection (Maynard Smith, 1964), group selection (Wynne-Edwards, 1986), tag recognition (Holland, 1996) and environment viscosity (Hamilton, 1964). In particular, kin selection stresses that genes responsible for altruistic behaviours can increase in frequency when there is a chance that beneficiaries of such altruistic acts also carry such genes. In other words, kin selection hypothesizes that the inclusive fitness of an individual is increased if it is genetically close to its neighbours. Kin selection has long been a central idea in the evolution of altruism, and recent works have shown that several other mechanisms (such as group selection) are actually much more related to kin selection than originally expected (West et al., 2007; Grafen, 1984; Queller, 1994; van Baalen and Rand, 1998). Moreover, explicit behavioral strategies have been shown to increase kin selection, such as kin recognition and spatial dispersion (West et al., 2007). In particular, a low spatial dispersion naturally favors reproduction among kins.

From the perspective of artificial evolution, several works have previously addressed the evolution of altruism (Waibel et al., 2009), and of communication (Floreano et al., 2007) (a particular kind of cooperative behaviour) with regards to the level of selection (at the level of the team or the individual), and the composition of teams (homogeneous or heterogeneous). These works successfully show cooperative behaviours could evolve from team level selection or by enforcing homogeneous teams, which is coherent with results previously established in theoretical biology (cf. (Hamilton, 1964)). However, these works rely on a fixed selection scheme (rather than letting it evolve) which prevents studies of particular dispersion strategies that could influence the level of homogeneity and relatedness in the population.

This paper addresses the evolution of spatial dispersion behaviour, in the context of a harvesting task that requires altruistic cooperation among individuals. The question under

scrutiny is to understand how spatial dispersion may evolve when altruistic behaviour comes as a requirement for the population to survive. In particular, it is expected that there is a correlation between the consumption strategy and the spatial dispersion evolved: the more altruistic the individuals, the less spatial dispersion should be observed (Taylor, 1992)). However, questions remain open as to whether such behaviours may be observed easily in nature, and what kind of behaviours may be evolved in term of consumption and spatial dispersion strategies.

The approach followed in this work builds on an existing framework for *in silico* experimental evolution for individual-based modeling and simulation undergoing an open-ended evolutionary process (i.e. long term adaptation in an open environment). In this context, the ability for an individual to survive and pass its genotypic material depends solely on its interaction with other individuals and with the environment, comparably to Dawkins' selfish gene metaphor (Dawkins, 1976) or TIERRA's open-ended evolutionary process (Ray, 1992). Therefore, it is possible to investigate the particular dispersion strategies which comes from a trade-off between harvesting and genotypic material diffusion.

In the following, the experimental setup is described, along with methodological tools and implementation details. A statement of the working hypotheses and outline of the experiments follows. Then, the experiments are described and discussed. Firstly, the possible correlation between spatial dispersion and level of altruism is investigated. Secondly, the trade-off between evolving either efficient or versatile strategies is studied. Finally, the last Section concludes this work and takes a broader perspective from this work, considering implications both from the theoretical viewpoint wrt. biology and from the practical viewpoint wrt. collective adaptive systems.

Method

Open-ended Evolution with mEDEA

The **mEDEA** algorithm, as in *minimal* Environment-driven Distributed Evolutionary Adaptation, was initially introduced in (Bredeche and Montanier, 2010). It performs as an evolutionary adaptation algorithm that can be distributed over a population of agents (i.e. each agent in the population runs the same algorithm, but carries different genomes). While it has been originally designed for collective robotic systems, it can be (and has been) used as a modeling and simulation tool for studying spatial interactions between agents. In previous works (cf. Montanier and Bredeche (2011)), mEDEA has been used to study the impact of genotypic relatedness on altruistic cooperation, in particular whenever genotypic relatedness between individuals is enforced through kin recognition (i.e. explicitly favoring the reproduction of closely related individuals).

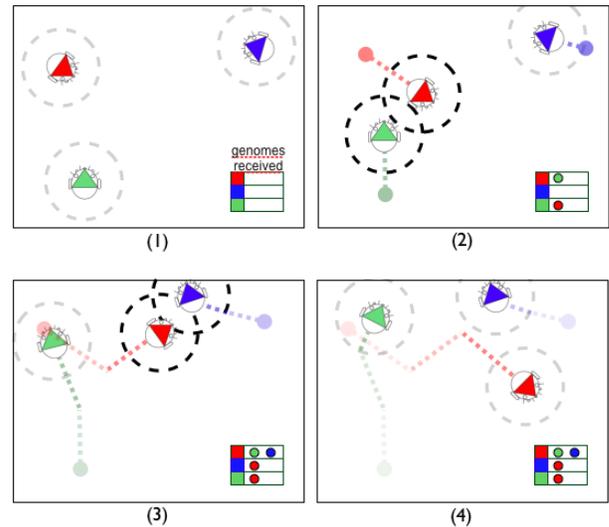


Figure 1: The mEDEA algorithm: a simplified illustration. (1): generation starts, genome reservoirs are empty; (2) and (3): agents move around (each agent is controlled by its own active genome) and exchange mutated genomes when close enough; (4): generation ends - the red genome has spread more and thus have higher probability of being selected (in this case, probability is indeed $p = 1$ in two agents while the two other genomes only get $p = 0.5$ in one single agent). After selection, all the reservoirs are emptied. Note that the next generation will contain slightly mutated copies of the original genomes.

Figure 1 provides an illustrative example of how mEDEA works (see (Bredeche and Montanier, 2010) for a complete description of the algorithm). Each robotic agent contains an *active* genome, which (indirectly) controls the agent's behaviour, and a *reservoir of stored genomes*, which is empty at first. At each time step (or *iterations*), each agent *broadcasts* in a limited range a *slightly* mutated copy of its active genome (gaussian mutation) and stores genomes received from neighbours, if not already stored. At the end of a generation (i.e. a pre-defined number of iterations), 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 generation starts. This algorithm is running independently within each agent in the population. By this mean, agents' behaviours differ depending on each agent's current active genome.

Therefore, selection pressure occurs at the population level (the more a genome spreads itself, the higher the probability it will generate offsprings) rather than at the individual level (random sampling). Genomes survive only through spreading (as an active genome is automatically deleted locally at the end of a generation) and individual may get better over time as conservative mutations generate new candidates that explore alternative behavioural strategies.

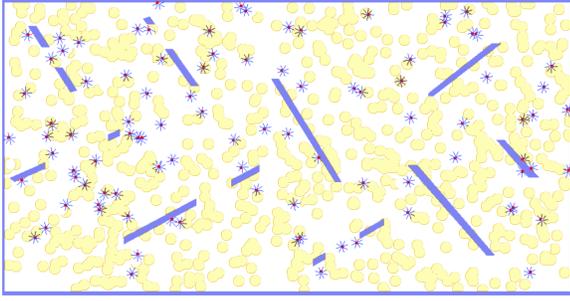


Figure 2: Snapshot from the simulator: food items (circles), agents (dots) and obstacles

On the one hand, mEDEA resembles other open-ended evolutionary setups such as TIERRA (Ray (1991)) as selection pressure occurs through interactions between agents in their environment rather than by explicitly computing a fitness value (such as a metabolic function). On the other hand, relying on a maximum number of active mobile agents makes it possible to set an upper bound in term of computational time required for the simulation, in a similar fashion as it is by setting a maximum number of cells in AVIDA (Ofria and Wilke, 2004). In mEDEA, each genome competes for accessing a limited set of resources: the population of robotic agents.

Setup

The setup used in this work (displayed in Figure 2) features simulated robotic agents which harvest food items from the environment in order to remain active. Each robotic agent consumes a fixed amount of energy at each iteration, and has a limited energy storage capability.

Each agent can be in three different states: active, dead or listening. At the beginning of a run all agents are in the active state, i.e. they are using an active genome (randomly generated in the range $[-1.0; 1.0]$), and their energy level is greater than 0. If during a generation an agent runs out of energy, it switches to the dead state. In this state the agent has no active genome, remain stationary, and cannot store genomes from robots passing by. The dead state is maintained for one generation, after which the agent switch to a listening state. In this state the agent doesn't move but stores the genomes broadcasted by agents in its neighbourhood. This state is maintained during one generation. If at the end of a generation in listening state the reservoir is still empty, the agent will remain in listening state for another generation, and so on. Any agent with an empty reservoir at the end of a generation switches to the listen state.

Food items are randomly placed in the environment. Once a food item has been harvested it becomes unavailable for some time, termed EP_{Lag} (the regrow delay). This term depends linearly on the energy harvested from it as shown in Equation 1:

$$EP_{Lag} = E_{harvested} / EP_{e_{Max}} * EP_{Lag_{Max}} \quad (1)$$

$EP_{e_{Max}}$ is the maximal amount of energy that can be harvested from a food item by an agent. $E_{harvested}$ is the energy actually harvested from the food item by an agent. $EP_{Lag_{Max}}$ is the maximal regrow delay of a food item.

Environmental pressure can be changed from low to high by setting the value of the $EP_{Lag_{Max}}$ parameter. Large $EP_{Lag_{Max}}$ values result in longer regrow delay (i.e. larger EP_{Lag} values) whenever a food item is completely harvested, which decreases the number of food items available for some time.

Monitoring Consumption Strategy

An agent may display an altruistic behaviour by harvesting only part of a food item. Such a consumption strategy is costly in terms of fitness (as it might run out of energy), and is of benefit to other agents (the food item will regrow faster). On the contrary, selfish agents will completely harvest any food item, which is likely to increase their chance to survive, but also reduces the number of food items available to other agents.

The consumption cost an agent accepts to pay is measured by the difference between how much energy could have been harvested by the agent (in order to completely fill the battery), and how much was really harvested. Equation 2 gives a definition of the consumption cost:

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

$EP_{e_{Max}}$ 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 consumption cost of zero, an altruistic agent should ideally be able to perform a trade-off between its altruistic nature and its survival needs. Therefore, the consumption 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 don't eat).

As a last remark it should be noted that the consumption strategy is but one way to monitor altruistic behaviours. As an example, two different consumption strategies, each combined with a different exploration strategy (travelling speed, area coverage) may well end up with the same number of food items available at any time (slow but greedy vs. fast but frugal agents). The next paragraph investigates how to take into account spatial dispersion strategies, in addition to consumption strategies already considered.

Monitoring Spatial Dispersion

Spatial dispersion may impact harvesting strategies as well as altruistic cooperation in various way as low dispersion

(i.e. remaining in the same region) may both favour exploitation of the same food items as well as increasing kin selection (and therefore impact the level of altruism, as long known since Maynard Smith (1964)). In order to account for spatial dispersion, we devise a measure to approximate the area covered by one agent during its lifetime.

In order to do so, the environment is divided into squared regions. For each agent, the number of regions is counted during its lifetime, and the *AreaCovered* value is computed following Equation 3:

$$AreaCovered = \frac{\#VisitedCells}{\#Cells \times lifetime} \quad (3)$$

In the following, the environment has been divided in 33920 cells of 4 by 4 pixels. The theoretical minimal value of this measure is $1/(\#Cells \times lifetime) = 7.34258 \times 10^{-08}$. The maximal value depends on agents maximal speed ($MaxAgent_{speed}$ given in fraction of cell area), and is given by Equation 4:

$$MaxAreaCovered = \frac{MaxAgent_{speed}}{\#Cells} \quad (4)$$

Working Hypotheses

Firstly, the link between consumption strategies and spatial dispersion will be investigated by monitoring spatial dispersion whenever a pre-defined consumption strategy is used. The expected result is that the level of altruism displayed during consumption of food items should be (negatively) correlated with spatial dispersion (the higher the consumption cost, the lower the dispersion) as lower dispersion theoretically increases genotypic relatedness, which is a key to altruistic behaviour. However, it remains to be investigated if such results can actually be observed as survival becomes more challenging as environmental pressure increases (i.e. spatial dispersion may not be solely driven by altruistic motivation).

In practical, this will be done by enforcing the amount of energy that is *left* when harvesting a food item. In Equation 2, this corresponds to setting a value for $E_{harvested}$ so that the *Cost* paid is equal to the "fixed" cost expected. In this paper, two different fixed costs, each close to one particular extreme consumption behaviour, are investigated: whenever a food item is harvested, either 5 (slightly altruist) or 40 (*very* altruist) units of energy are left over on a total of the 50 units of energy a food item can provide, implying different consequences on the food item's delay to regrow. These *fixed cost* consumption strategies will be referred to as $cost = 5$ and $cost = 40$ consumption strategies in the next Section.

Secondly, the possible benefits of leaving to evolution both the consumption strategy *and* the spatial dispersion strategy will be investigated. It is indeed not clear that letting

both strategies evolve should lead to better survival strategies, as evolution may face a more difficult challenge due to an increased number of degrees of freedom. In practical, the consumption cost to be paid when a food item is harvested will be left to the robot to decide and both consumption cost (cf. Equation 2) and spatial dispersion will be monitored. The expected result is that evolving both consumption cost and spatial dispersion may possibly lead to a richer set of behaviours whenever environmental pressure varies, though possible benefits remain to be identified. This consumption strategy will be referred to as *dynamic cost* in the next Section.

Results

Technical Details

A Multi-Layer Perceptron (MLP) is used to encode the controller of each robotic agent. The input layer is composed of 12 inputs (8 for distance sensors, 1 for the direction to the closest energy point, 1 for the distance to the closest energy point, 1 for the battery level of the agent, 1 to detect the presence of an energy point under the agent), the hidden layer is composed of 5 neurons, and the output layer is composed of 3 neurons (rotational speed, translational speed and amount of energy to be harvested (used only if a food item is within reach)). The output neuron for energy harvesting is not taken into consideration when a fixed cost is used. The weights of the MLP are decoded from the active genome of the agent. A gaussian mutation is used, and initial weights are set randomly around zero. The σ parameter for mutation is evolved, and a minimal value (0.01) is fixed to avoid obtaining a population of clones

All experiments are performed with RoboroBo, a fast open-source multi-robot simulator (Bredeche et al., 2013). In order to ensure the reproducibility of the experiments presented in this paper, the full implementation is available online¹ and parameters used are summarized in Table 1. One run takes approximately one hour to be performed using one core of a quad-core 2 CPUs Intel 2.26 GHz processor. All experiments presented in this paper are performed on a computer cluster equipped with such processors. For each setup considered in the next section results, each Figure results from a compilation of 500 independant runs, and statistical significance is tested using the Wilcoxon signed-rank test (Wilcoxon, 1945).

Evolution of Spatial Dispersion Strategies

In order to obtain results on a large range of environmental pressures, the experiment starts with a low environmental pressure ($EP_{LagMax} = 25$ iterations) until the 400000th iteration. After this, the environmental pressure

¹http://pages.isir.upmc.fr/evorob_db/moin.wsgi

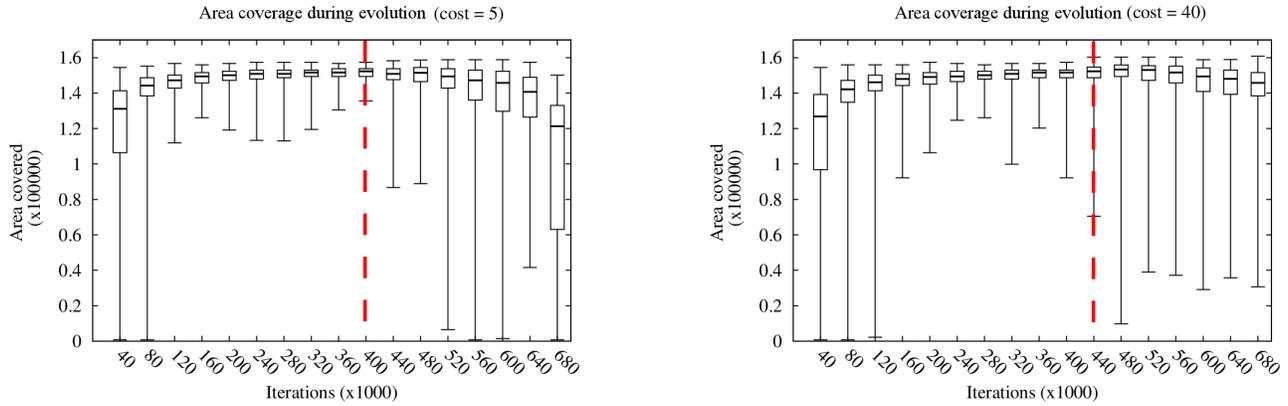


Figure 3: Area coverage measured when the cost of altruism is fixed to 5 (left) and 40 (right) and the pressure of the environment increases by step of 80 iterations every 4000 iterations from iteration 400000 (before the environmental pressure is $EP_{LagMax} = 25$ iterations).

Parameter	Value
arena width and length	1024 * 530 pixels
lifetime (i.e. generation duration)	400 iterations
selection scheme	random
population size	100 agents
agent size	1 pixel
proximity sensor range	64 pixels
radio broadcast signal	32 pixels
agent rotational velocity	0.52 rad/iteration
agent translational velocity	3 pixels/iteration
genome length	84 real values (83 MLP weights + σ)
variation operator	Gaussian mutation with σ parameter
cells size	4 by 4 pixels
theoretical max area	2.21108×10^{-5}
theoretical min area	7.34258×10^{-8}
EP_{LagMax}	50
number of energy points	800
r_{LagMax}	400
energy consumption	1 per iteration

Table 1: Parameters for experiments.

slowly increases every 4000 iterations (10 theoretical generations) until the population goes extinct (i.e. no genome left to exchange). Each increase of the environmental pressure is done by a fix amount of 80 iterations in the regrow delay (EP_{LagMax}). As an example, $EP_{LagMax} = 105$ iterations at the 404000th iteration of the simulation, and $EP_{LagMax} = 185$ iterations at the 408000th iteration.

Results obtained when two harvesting strategies with fixed costs of 5 (less altruistic strategy) and 40 (more altruistic strategy) are used are presented in Figure 3 (500 runs for each setup). With both strategies the area dispersion evolved is increasing until iteration 400000th and decreasing after (p -value < 0.05 for comparison of iteration 400000 and every iteration after 560000). This shows that different spatial dispersion strategies are displayed through evolution depending on the consumption strategy used and the environmental pressure at hand.

The differences between spatial dispersions evolved under different consumption strategies is expected from results obtained in biology (as said before). However, results shown here are contradictory with theory: spatial dispersion is shown to be higher for the more altruistic con-

sumption strategies when challenging environment are considered (while kin selection, favored by lower dispersion, should be paired with an increased altruistic behaviour (Taylor, 1992)). Rather than contradicting well established theoretical results, individual based modeling and simulation actually points out the complex interactions between individuals and the environment. Indeed, dispersion strategies may be influenced by much more than just acting on genotypic relatedness. The number of active agents, the availability of energy points, and the regrow delay are all possible causes to explain a particular dispersion strategies. Therefore, one question remains: in a comparable setup (i.e. removing all other possible causes), how does dispersion strategies compare when evolved with different fixed consumption strategies.

Fair Comparison of Dispersion Strategies

In order to compare results from the two setups considered previously, agents' spatial dispersions are measured in a similar environment. The environment used for comparison features a consumption cost artificially fixed to 0 (whatever the initial consumption cost used during evolution) and a low environmental pressure ($EP_{LagMax} = 25$ iterations). Moreover, genome transmission and selection are shut down, and robots continue to run even if energy is depleted. This makes it possible to compare the different behavioural strategies by *replaying* evolved genomes with all other parameters set to similar values. The following replay procedure is defined: (1) genomes from the 600000th iteration of a given run are randomly sampled to assemble a population of 100 individuals ; (2) this population is embodied in 100 robots (one genome per robot) (3) The spatial dispersion of these robots is measured during 40000 iterations.

For each fixed cost strategy considered earlier, genomes

are extracted from 20 runs selected randomly among the 500 discussed earlier, and replayed following the procedure described above. The replay procedure is repeated 20 times for each setup (total of 800 replays). In addition, we devise a third type of behaviour, which stands as a control experiment: agents using random controllers. Random translational and rotational speeds are assigned to 100 agents at each iteration, and spatial dispersion is measured for 40000 iterations. This is inspired from typical tests in ecology field studies where geographic tracking of animals positions are evaluated wrt. brownian motion (Borger and Fryxell, 2012). This control experiment is termed *random behaviour*, and is also performed 800 times.

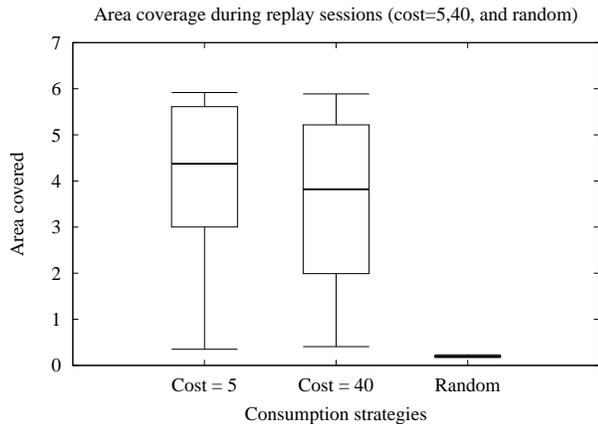


Figure 4: Area coverage for three consumption strategies when replayed, and for runs using random movements. All area coverage are significantly different: $p - value < 0.05$ for comparisons between each boxplot.

The median results for each fixed consumption strategy and the random behaviour are presented in Figure 4. As expected, both strategies evolved under environmental pressure (cost fixed to 5 and cost fixed to 40), display higher level of dispersion than random movements. Then, comparing the fixed cost consumption strategies, it is shown that the area covered is significantly higher when a lower level of altruism is used during evolution ($p - value < 0.05$ between a cost fixed to 5 and a cost fixed to 40). As a consequence, this confirms expected results from theory: there is indeed a negative correlation between spatial dispersion and consumption behavior: enforcing a consumption strategy which displays altruistic behaviour leads to lower spatial dispersion, which is expected to increase kin selection, and thus altruistic behaviour among closely related individuals.

Typical runs for each behaviour studied (i.e. runs at the median), are shown in Figures 5(a), and 5(b). In these two Figures, which can only be interpreted in light of the previously shown quantitative analysis, illustrate slight but visible differences in spatial dispersion. Trajectories obtained with a highly altruistic consumption strategy ($Cost = 40$) dis-

play more localized behaviours (i.e. robot circling around in the same area) and fewer wandering trajectories. Counting the average number of encounters per agent also advocates for local interactions: there are significantly less encounters for the $Cost = 40$ strategy than for the less altruistic, more exploratory, $Cost = 5$ strategy ($p - value < 0.05$, Wilcoxon test).

Evolving Consumption and Dispersion Strategies

In this last part of the paper, we investigate the impact of evolving both the consumption strategy and the dispersion strategy. By doing so, we intend to address the following questions: (1) What kind of (consumption and dispersion) strategies can be expected when evolved under different environmental pressures ; (2) What are the possible benefits and drawbacks of evolving the consumption strategy rather than enforcing an ad hoc consumption strategy.

As before, 500 runs are performed. The setup is similar to the previous setups for fixed cost, except that the cost of altruism is now chosen by the robot controller. This setup is termed "Dynamic Cost" has the cost paid may change any-time and depends from evolution (i.e. the controller output fixing the amount of energy taken from a food item is actually used). Figure 6 shows the boxplot results for consumption costs paid and area dispersion evolved by all agents throughout evolution. As before, the environment becomes gradually more challenging starting iteration 400000, and stops when all runs have gone extinct.

A notable difference is that during the first part of the runs, the consumption costs paid stick to zero, which is not unexpected as there is no benefit at being altruistic in an environment that represent an easy challenge. The consumption cost paid then abruptly changes as soon as the environmental pressure increases ($p - value < 0.05$ for comparison between iteration 400000 and iteration 440000). It then fluctuates around a value of 5 until iteration 680000, and remains significantly higher than at the beginning of the run ($p - value < 0.05$, comparing results from iteration 400000 and any iteration afterwards). Moreover, the final value (at iteration 680000) is similar to the value obtained at iteration 440000 ($p - value = 0.31$). Hence, there appears to be two stable values (either no altruism ($Cost = 0$) or low altruism ($Cost = \sim 5$)) for consumption cost depending on the challenge posed by the environment.

Regarding spatial dispersion, Figure 6-right shows that the area covered by each agent levels up, and then off, until iteration 400000. Then, as environmental pressure starts to increase, the area covered is continuously decreasing, and ends up as significantly lower beyond iteration 560000 ($p - value < 0.05$, comparing area covered at iteration 400000 and any iteration from 560000).

In order to compare the behaviours obtained with a dynamic cost, replay sessions are performed in the exact same fashion as it was for the fixed cost setups. Quantitative re-

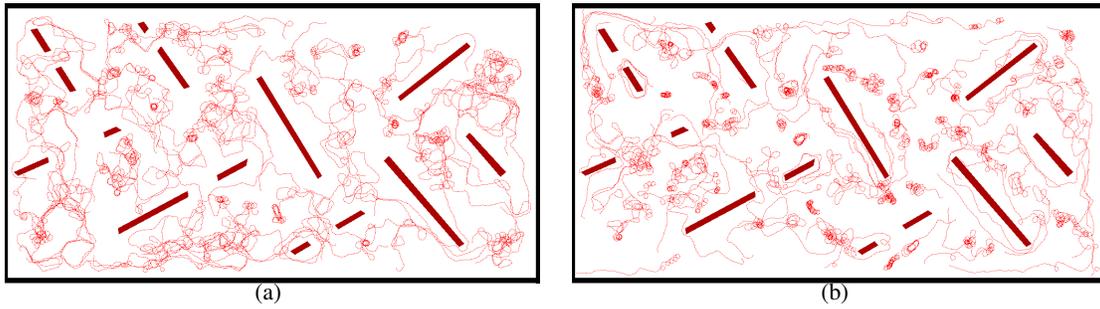


Figure 5: Trajectories from median runs for consumption strategies of cost=5 (left) and cost=40 (right).

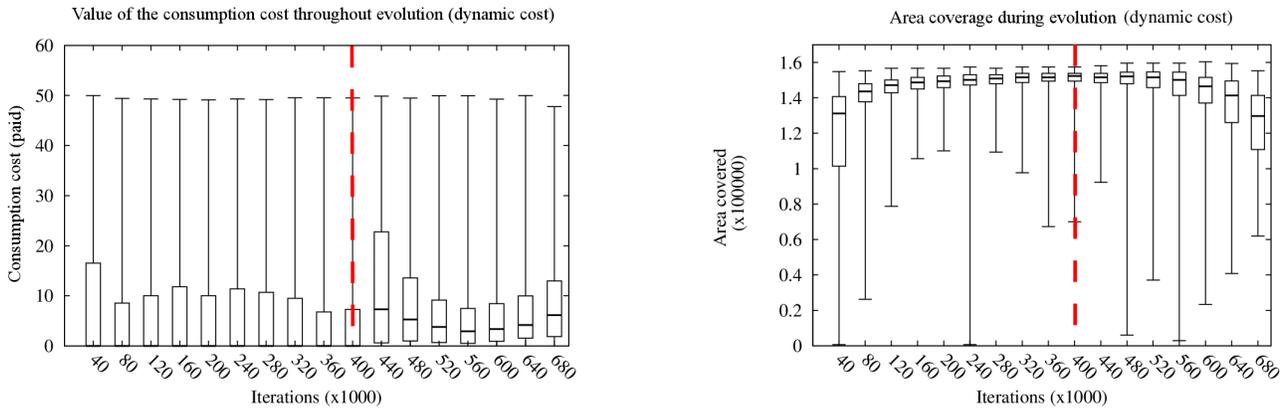


Figure 6: Consumption cost of altruism (left) and area coverage (right) measured when the Consumption cost is dynamic and the pressure of the environment increases by step of 80 iterations every 4000 iterations from iteration 400000 (before the environmental pressure is $EP_{Lag_{Max}} = 25$ iterations).

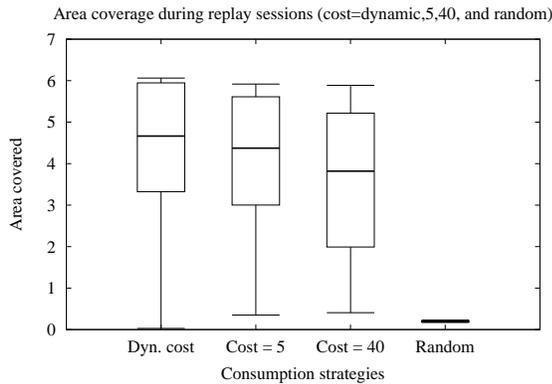


Figure 7: Area coverage for two consumption strategies when replayed, and for runs using random movements. All area coverages are significantly different.

Results obtained by all consumption strategies studied are presented in Figure 7. Results obtained with the dynamic cost strategy is different from all other strategies (p -value < 0.05 for comparisons between each cost strategy). This con-

firms the impact of cost strategies on the evolution of behaviours.

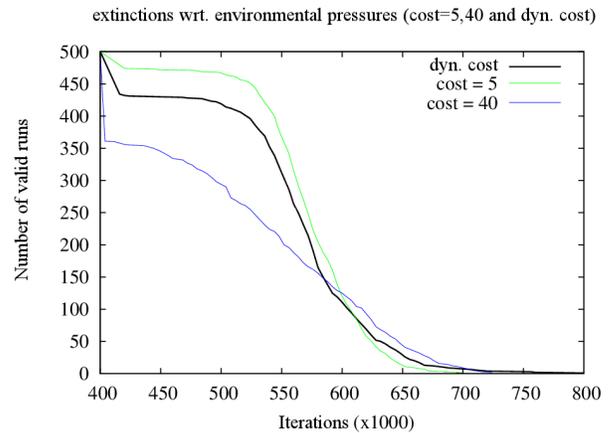


Figure 8: Number of active runs when the environmental pressure is increasing.

Another way to study the differences between different cost strategies is to observe the extinction of runs : that is,

for a given iteration, how many runs is there with at least one active agent. Figure 8 shows the number of active runs for all three setups considered ($Cost = 5$, $Cost = 40$ and dynamic Cost) starting iteration 400000, that is when environmental pressure starts to gradually increase. From this figure, several particular situations arise: the $Cost = 5$ strategy dominates before iteration 580000, and is replaced by the $Cost = 40$ strategy around iteration 600000 and until iteration 700000. Afterwards, only the $Cost = 40$ and dynamic Cost remain for another 100000 iterations until all runs for both go completely extinct. Though the dynamic Cost strategy goes a little further than the $Cost = 40$ strategy, there are too few runs left to make any statistically significant remark between the two setups beyond iteration 700000. Most notably, the dynamic Cost strategy is nearly never completely dominated by the two other strategies. As such, it appears that relying on a dynamic cost strategy provides some kind of trade-off between optimal performance in specific contexts and versatility in various contexts, without having to take the risk of guessing which cost must be paid prior to evolution.

Conclusion

In this paper, we have studied the importance of spatial dispersion during the evolution of altruistic cooperation in a population of autonomous agents. This work followed an individual-based modelling and simulation approach using the mEDEA algorithmic framework for open-ended evolution, which enables to study the actual evolution of individual spatial dispersion and consumption strategies in the context of a harvesting task.

Firstly, well-established theoretical results on kin selection in the evolution of altruism were confirmed experimentally regarding the negative correlation between the level of altruistic behaviour and spatial dispersion. However, our work revealed that such a confirmation may not come for free when individual modeling and simulation is considered as the complex interactions between the population and its environment may provide contradictory results at first glance. This result is important as it may have an impact on field observation from nature, where the level of spatial dispersion may fail to directly explain the occurrence of altruistic behaviour.

Secondly, we showed that there is a trade-off between setting ad hoc mechanisms (here, the strategy used when eating a food item) and letting such mechanisms to evolution. On the one hand, results show that the more there is left to evolution, the less likely optimal behaviour may be reached (compared to *a priori* carefully crafted strategies). On the other hand, fully evolved strategies turn out to be more versatile, i.e. showing good performance in a larger set of contexts, and require less prior knowledge compared to more constrained evolutionary setups.

Lastly, this paper intends to contribute both to theoretical biology, by providing new results from an individual-based

modeling perspective where spatial dispersion is the product of open-ended evolution, and to collective adaptive systems, as the algorithm used throughout this paper may straightforwardly be implemented onto real robots (and has already been, albeit for a different problem).

Acknowledgments

Experiments presented in this paper were carried out using the Grid'5000 experimental testbed, being developed under the INRIA ALADDIN development action with support from CNRS, RENATER and several Universities as well as other funding bodies (see <https://www.grid5000.fr>). This work was made possible by the Alain Bensoussan Fellowship Programme .

References

- Borger, L. and Fryxell, J. (2012). *Quantifying individual differences in dispersal using the net squared displacement*, chapter 17. Oxford University Press, Oxford (UK).
- Bredeche, N. and Montanier, J.-M. (2010). Environment-driven Embodied Evolution in a Population of Autonomous Agents. In *The 11th International Conference on Parallel Problem Solving From Nature (PPSN 2010)*, pages 290–299.
- Bredeche, N., Montanier, J.-M., Weel, B., and Haasdijk, E. (2013). Roborobo! a fast robot simulator for swarm and collective robotics. *CoRR*, abs/1304.2888.
- Dawkins, R. (1976). *The Selfish Gene*, volume 32. Oxford University Press.
- Diekmann, O. (2004). A beginner's guide to adaptive dynamics. *Mathematical Modelling of Population Dynamics*, 63:47–86.
- Floreano, D., Mitri, S., Magnenat, S., and Keller, L. (2007). Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17(6):514–519.
- Grafen, A. (1984). Natural selection, kin selection and group selection. *Behavioural ecology: an evolutionary approach, 2nd edition*, pages 62–84.
- Hamilton, W. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology*, 7(1):1–16.
- Holland, J. (1996). *Hidden Order: How Adaptation Builds Complexity*. Basic Books.
- 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.
- Maynard Smith, J. (1964). Group selection and kin selection. *Nature*, 201:1145–1147.
- Maynard Smith, J. (1974). The theory of games and the evolution of animal conflicts. *Journal of theoretical biology*, 47(1):209–221.
- Montanier, J.-M. and Bredeche, N. (2011). Surviving the tragedy of commons: Emergence of altruism in a population of evolving autonomous agents. In *Proceedings of the 11th European Conference on Artificial Life (ECAL'11)*, pages 550–557.
- Ofria, C. and Wilke, C. O. (2004). Avida: A software platform for research in computational evolutionary biology. *Artificial Life*, 10(2):191–229.
- Queller, D. (1994). Genetic relatedness in viscous populations. *Evolutionary Ecology*, 8(1):70–73.
- 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*, page 371408. Addison-Wesley, Redwood City, CA.
- Ray, T. S. (1992). Evolution, ecology and optimization of digital organisms. Technical report, Santa Fe Institute.
- Taylor, P. D. (1992). Inclusive fitness in a homogeneous environment. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 249(1326):299–302.
- van Baalen, M. and Rand, D. (1998). The unit of selection in viscous populations and the evolution of altruism. *Journal of theoretical biology*, 193(4):631–648.
- Waibel, M., Keller, L., and Floreano, D. (2009). Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transactions on Evolutionary Computation*, 13(3):648–660.
- 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.
- Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80–83.
- Wynne-Edwards, V. C. (1986). *Evolution through group selection*. Blackwell Scientific, Oxford.