Evolving an Agent Collective for Cooperative Mine Sweeping

A.E. Eiben, G.S. Nitschke, M.C. Schut, gusz@cs.vu.nl, nitschke@cs.vu.nl, schut@cs.vu.nl,
Computational Intelligence Group, Department of Computer Science, Vrije Universiteit Amsterdam,
De Boelelaan 1081a, 1081 HV Amsterdam, Netherlands

Abstract—**The research goal was to engineer agent collectives that most effectively accomplish a cooperative gathering task. In view of this, we compared reproduction schemes for the artificial evolution of agent controller parameters for a cooperative minesweeping task. Agents utilized cooperative behavior to improve task performance in a simulated environment where different types of mines with different fitness rewards were randomly distributed. We compared the evolution of agent controller parameters with respect to temporal and spatial dimensions of agent reproduction schemes. The first dimension concerned agents reproducing only once at the end of their lifetime or multiple times during their lifetime. The second dimension concerned agents reproducing only with agents in adjacent positions (locally restricted) or with agents located anywhere else in the environment (panmictic). Results indicated that the single reproduction at the end of an agent's lifetime and the locally restricted reproduction schemes afforded the agent collective a higher level of performance in its cooperative gathering task.**

Index Terms — **Emergence, cooperative behavior, artificial evolution.**

I. INTRODUCTION

The research theme of this paper is described by the term: *Emergent Collective Intelligence* (ECI). The end goal of ECI research is to combine and exceed achievements in *multi-agent systems*, *swarm intelligence*, and *evolutionary computation* research via developing synthetic methodologies such that groups of computationally complex agents produce desired emergent collective behaviors resulting from the bottom-up development of certain individual properties and social interactions. This paper investigates certain technical aspects of artificial evolution as means of achieving adaptability at the local level and desired emergent behavior at the global level.

The applications we envision include engineering tasks. For example, in social robotics, the individuals are robots that have to perform certain tasks collectively. In such a group, individuals can be relatively simple but adaptive, where it is the group as a whole that develops the ability to carry out complex collective tasks in unknown environments. We proclaim that adaptability at the local level is attainable by means of an evolutionary process. The technical research goal of this paper was to establish what types of agent reproduction mechanisms operating within an artificial evolution process lead to good solutions for multi-agent task accomplishment.

Our application domain is the gathering of renewable resources from a virtual environment. This gathering task is divided into *locating*, *retrieving*, and *transporting* the resources in question. It is an essential assumption that this task is interfaced to the population of agents via fitness rewards that are given after delivering the resources to a given 'home area'. Additionally, we distinguish resources with different values and postulate that gathering of higher value (more complex) resources necessitates a higher degree of cooperative behavior (more agents). The performance evaluation criterion for the agent collective as a whole is then the *total value gathered cooperatively*, measured at the final generation of the simulation.

In this case the task was to *locate*, *extract* and *transport* different types of mines within an artificial mine field, where cooperative behavior was needed for 'good' solutions. A good solution refers to maximizing the value of mines extracted and transported within a given time period. A mines value was determined by its complexity and hence how difficult it was to disable, extract and transport. We investigated two dimensions of agent reproduction. First, *temporal* reproduction settings, termed: *Single Reproduction at the End of the Agent's Lifetime* (SREL) and *Multiple Reproductions During an Agent's Lifetime* (MRDL). Second, *spatial* reproduction settings, termed: locally restricted reproduction and panmictic reproduction. In the SREL reproduction setting agents reproduce only at the end of their lifetimes, and in the MRDL reproduction setting agents were able to reproduce multiple times during their lifetime. Using the *locally restricted* spatial reproduction setting an agent was only able to reproduce with other agents situated in adjacent positions in the environment. Where as, using the *panmictic* reproduction setting agents were able to reproduce with other agents situated anywhere in the environment.

The key research question addressed within this paper is to determine which combination of the temporal and spatial reproduction setting, that is, which of four possible schemes, leads to superior performance with respect to the agent collective accomplishing its task. The task of each agent was to *locate*, *extract*, and *transport* as high a value of mines as possible to a 'home' area in the environment. The successful delivery of a mine to the home area was termed *gathering*, where gathered mines were equated with a fitness reward. Fitness rewards were proportional to the type and amount of mines gathered. Gathering of higher value (more complex) mines necessitated a higher degree of cooperative behavior (more agents) to extract. Two quality measures were used as the evaluation criteria for the evolved behavior of the agent collective. The first was the *total value gathered cooperatively*. This measure was taken at the final generation of the simulation, where the highest possible values could only be attained if the agents worked cooperatively to accomplish the mine gathering task. The second quality measure was the *stability* of the agent collective, which was number of simulation runs in which all agents survived until the final generation of the simulation.

II. LITERATURE REVIEW

The study of the synthesis of collective behavior, including the emergence of cooperation, has been investigated in both simulated [1], [2], [3], [4], [12], [15] and real world [5], [6], [7], [17], [21], [18] problem domains. Traditionally collective behaviors in agent collectives (multi-agent systems) have been studied using a top down classical approach [10], [16]. Such approaches have achieved limited success given that it is extremely difficult to specify the mechanisms for cooperation or collective intelligence in all but the simplest problem domains. The utilization of evolutionary computation as a mechanism for agent controller design such that the local interactions of many controllers produce a desired collective behavior, has been highlighted as a promising area of research [13], [14]. This is especially true in large agent collectives, which potentially contain thousands of individuals.

Within simulated agent collectives there has been a significant concentration of research on the study of emergent behavior in artificial ant colonies [6], [7], [8], [9], [22].

Drogoul *et al.* [8], [9] presented a simulation model of social organization in an ant colony termed: *MANTA* (*M*odel of an *ANT*-hill *A*ctivity), which was designed to explore the contribution of emergent functionality such as division of labor on emergent cooperation. Results elucidated that emergent division of labor improved the efficiency of emergent functionality in the population. Such emergent functionality included cooperative foraging and sorting behavior. The authors concluded that many of the behaviors viewed as cooperative emerged as a result of the competitive interaction that occurs between individuals in a constrained environment with limited resources.

Perez *et al.* [22] conducted experiments in the context of an artificial evolution process, in order to study the impact of genetic relatedness and different types of genetic selection in the evolution of cooperation for a foraging task. The transportation of certain large food items required that two ants cooperate in order to achieve the task. Artificial ants were rewarded differing fitness scores for either individual or cooperative transportation of food items, such that the total performance of the colony was maximized if ants cooperatively transported food items as opposed to acting individually. In the experimental setup, groups of ants tested were either homogenous or heterogeneous, where the method of genetic selection, which either reproduced the next generation via selecting individuals from different colonies or via selecting different colonies as a whole, delineated homogenous and heterogeneous colonies. Results indicated that the colony-based form of genetic selection and reproduction favored emergent cooperative behaviors, and that cooperative behavior had a low probability of emerging in heterogeneous colonies, where an individual-based form of genetic selection and reproduction was used.

Nolfi *et al.* [19] conducted several experiments to address the problem of how a group of simulated robots (s-bots) could coordinate their movements and actions so as to cooperatively move objects in the environment as far as possible within a given period of time. Nolfi *et al.* [20] conducted a set of experiments designed to facilitate emergent cooperative behavior, where a group of eight s-bots were connected to an object, or connected so as to form a closed structure around an object, and were given the task of moving the object as far as possible in the least amount of time. In a set of experiments the eight s-bots used what the authors termed the *ant formation*, which connected all s-bots to the object, but there were no links between the s-bots themselves. The result was dependent upon the weight of the object, such that the s-bots cooperatively negotiated to either push or pull the object to their destination.

From this overview of different research efforts, associable by gathering and transportation tasks and the general research topic of emergent cooperation, it is clear that some formalization of mechanisms for the design and analysis of emergent cooperation is needed. Specifically, if emergent cooperative behavior in agent collectives was sufficiently understood, purposeful design of cooperative behavior could be applied to benefit a variety of applications in social robotics including resource mining, transportation, surveillance, construction, and mine sweeping [5].

III. AGENTS, ENVIRONMENT AND EVOLUTION

A. *The Agent Collective and their Task Environment*

The experiments utilized a simulated minefield and an initial population of 1000 agents, placed at random positions on a grid-cell environment with a 50 x 50 resolution. A maximum of four agents could occupy any given grid-cell within the environment. A home area spanning 4 x 4 grid-cells was randomly placed within the environment. In the simulated environment the resources to be cooperatively gathered were mines, and the home area was where gathered mines were taken. Gathering was the term applied to the process of locating, extracting, transporting, and delivering a mine to the home area. Within the simulated minefield there were three types of mines: *type A*, *type B* and *type C*. The different types of mines had differing *values* to reflect the difficulty (degree of cooperation) associated with gathering it.

The cost of gathering mines comprised two sub-costs: the cost of extracting a mine from its location in the environment, and the cost of transporting a mine to the home area. The costs of extracting and transporting one unit of each of the three mine types are presented in *table 1*. The transport cost was applied per unit being transported, and per grid-cell traversed. Initially, a quantity of between 0 and 3 mines of each type were randomly initialized and placed within each grid-cell. It is assumed that a long-term process of gathering and replenishment in a minefield is being simulated, where mines are considered a renewable resource, and each mine type is renewed at a rate of 3 per simulation iteration. That is, the simulation is of a long-term process of collective gathering behavior being evolved, whilst an unseen competitor renews gathered mines. Additionally, it is assumed that an agent never triggered a mine to detonate.

In order to gather the different mine types a degree of cooperative behavior was necessitated. Cooperation was necessary when at least one agent was attempting to extract a given mine type, and the value of the prevalent agent controller parameter was too low for the agent to individually gather the mine. These prevalent agent controller parameters were termed: *Mine type A capacity*, *Mine type B capacity*, *Mine type C capacity* and *transport capacity*, and provided an indication of the capability of an agent for gathering a particular mine type. Specifically, to gather *one unit* of a particular mine type, the sum of the values of the *capacity* parameter for that mine type (for all agents simultaneously attempting to extract the mine) must exceed a given *capacity threshold*. These capacity thresholds are presented for each mine type in *table 1*. The task of each agent was to gather the highest possible value of mines during the course of its lifetime. This task was interfaced to the agent collective by providing fitness rewards for gathered mines.

The fitness rewards for gathering one unit of the different mine types are presented in *table 1*. The total value of mines that all agents gathered in cooperation with at least one other agent during the course of its lifetime was termed the *value gathered cooperatively*. Further to playing its conventional role in survivor selection, fitness was also used as a metaphor of energy (actions cost fitness). An agent was able to move one grid-cell in any direction per simulation iteration at a cost of one unit of fitness.

B. Evolution Approach

For the evolution of agent controller parameter values, a standard evolutionary algorithm was used [11]. When an agent initiated reproduction, the fittest partner (with the highest energy) of *m* potential partner agents was selected for reproduction. The population initially contained 1000 individuals (agents), and the genotype of each agent was its set of gathering and transport capacities (evolvable parameters illustrated in *figure 1*). These parameter values directly influenced the heuristic agent lifetime behavior, though the behavioral heuristics (*figure 2*) remained static over the course of the evolutionary process. That is, once an agent had gathered as many mines as it could transport, it would begin transporting the mines back to the home area. During reproduction, agent controller heuristics (*figure 2*) were copied from parent to child, and the fitness inherited by a child was the average fitness of the two parent agents. Ninety percent of the inherited fitness was then subtracted from each parent's fitness.

C. Artificial Evolution: Agent Reproduction Scheme

An agent reproduction scheme was devised, in temporal and spatial terms, by distinguishing when and where agents reproduced. For the temporal dimension, we tested a setting termed: *Single Reproduction at the End of the Agent's Lifetime* (SREL). For the spatial dimension, we tested a setting termed: locally restricted. That is, agents reproduce only with agents in adjacent positions. According to previous results this agent reproduction scheme was found to yield superior performance in collective gathering tasks, comparative to other reproduction schemes using different temporal and spatial settings. For example: multiple reproductions during an agent's lifetime, and panmictic reproduction.

During the reproduction action, 90% of the fitness of two parent agents was divided amongst and passed onto *p* offspring agents. During reproduction only one partner agent of *m* potential partner agents was selected for reproduction. An agent's fitness could only be replenished when it delivered a mine to the home area. The precondition for locally restricted reproduction setting was that there was at least one potential partner agent in the same grid-cell or an adjacent grid-cell. Reproduction was only possible when both parents current fitness was greater than the value of the *min fit reproduction* parameter.

When *p* offspring agents were produced using the panmictic reproduction setting, each offspring would be placed in a random free grid-cell adjacent to one of the parents. The chance that an offspring an agent was placed in a grid-cell adjacent to parent 1 was 0.5, and the chance that an offspring was placed in a grid-cell adjacent to parent 2 was 0.5. If no adjacent grid cells were free, then the offspring agent died. Using the locally restricted setting offspring agents were always placed in a random free grid-cell adjacent to the parent agent that initiated reproduction.

The number of offspring to be produced was determined as m = the total amount of fitness to be inherited (x) divided by 10. According to the reproduction scheme setting being used, pairs of agents produced p offspring using the genetic operations of crossover and mutation [11]. The core of agent reproduction was the application of uniform crossover to 'recombine' the controller parameters: *mine type A, B, C* and *transport capacities* of two parent agents in order to derive the agent controller parameter values of a child agent. The uniform crossover operator selected a parameter value to be inherited from either parent agent with a 0.5 probability. Child controller parameter values were mutated by a value of either plus or minus 10 with a probability of 0.05.

Value Ranges	Initial	Minimum to Maximum
Parameters: Not Evolvable		
Sight Death_Age Min_Fit_Reproduction	1 [20..100] 50	1 [20..100] 50
Parameters: Evolvable		
Mine Type A Capacity (CA)	[0..100]	[0..Infinity]
Mine Type B Capacity (CB)	[0..100]	[0..Infinity]
Mine Type C Capacity (CC)	[0..100]	[0..Infinity]
Transport Capacity (CT)	[0..300]	[0..Infinity]

Fig. 1. The evolvable and non-evolvable agent controller parameters.

IF AmA < CA **THEN**
  IF (Holding + AmA) < CT **THEN** extract AmA
 ELSE IF AmB < CB **THEN**
   IF (Holding + AmB) < CT **THEN** extract AmB
  ELSE IF AmC < CC **THEN**
    IF (Holding + AmC) < CT **THEN** extract AmC
   ELSE Look-Ahead

Look-Ahead:

IF end of life and SREL active **THEN** reproduce
IF at home **THEN** unload mines transported
 IF MRDL active **THEN** reproduce
IF transporting a quantity of mines **THEN** move to home
ELSE IF mine type A detected **THEN** move to mine type A
 ELSE IF mine type B detected **THEN** move to mine type B
  ELSE IF mine type C detected **THEN** move to mine type C
   ELSE move to a random cell

Fig. 2. Heuristics utilized by agents operating under the pure-evolution approach. *AmA*, *AmB*, and *AmC* denote the amount of mine type *A*, *B* and *C*, respectively, on a given grid-cell. Holding denotes the current amount of all mine types a given agent is currently transporting. *CA*, *CB*, *CC* and *CT*, denote the gathering capacities for mine types *A*, *B*, and *C*, and the transport capacity, respectively.

	Capacity Threshold	Extraction Cost	Transport Cost	Fitness Reward
Mine type A	**300**	**8**	**0.04**	**20**
Mine type B	**150**	**4**	**0.02**	**10**
Mine type C	**75**	**2**	**0.01**	**5**

Table 1. The capacity thresholds, and the costs for extracting and transporting mines, as well as the fitness reward for gathering one unit of the different mine types.

If mutation occurred, the probability of adding versus subtracting 10 from the inherited parameter value was 0.5.

IV. EXPERIMENTATION AND RESULTS

We designed our experiments along two parameter dimensions. Specifically: *Temporal reproduction scheme:* SREL versus MRDL. *Spatial reproduction scheme:* panmictic versus locally restricted.

This led to four different experimental setups, where for each setup we performed 100 independent runs (using different random initialization parameters), where one run was executed for 2000 iterations.

Additionally, we conducted a control experiment in which we fixed the experimental parameters (gathering and transport capacities) as derived by the evolution experiments, where, this control experiment served to elucidate the most appropriate agent controller configuration (parameter values for gathering and transport capacities) for the given collective mine sweeping task. Within each simulation, several experimental monitors were set as objective measures for the performance of the agent collective across successive generations of the evolutionary process.

The monitors for the *average value gathered cooperatively*, and the *number of agents* provided the two objective performance measures for evolved multi-agent system behaviors. The *average fitness* of the population and the *average distance to home*, describe the average energy level and population density of the agents, were additional measures elucidating details of the overall behavior of the agent collective. As previously stated, cooperative behavior was evaluated according to the total value of the mine types (A, B, and C) gathered by the agent collective over the course of a given simulation. Specifically, the measure of cooperative behavior was the total value gathered cooperatively, calculated at the final iteration of the simulation, where cooperative behavior was required to attain the highest values.

V. ANALYSIS AND DISCUSSION

The objective performance measures for evolved agent behaviors were the *total value of mines gathered cooperatively* and the *stability* of the agent collective for a given simulation run. *Table 3* illustrates the average value of mines gathered cooperatively utilizing each of the four reproduction schemes. It is important to note that a high standard deviation (the values in parentheses) was indicative of populations that died out prematurely with some regularity over the 100 simulation runs for a given experimental setup. The second performance measure was the stability of evolved agent behaviors. Here, the term *stability* indicates that, for the gathering and transport parameter values evolved, a particular value gathered cooperatively (plus or minus some variance) was expected.

For each of the four experimental setups we tested how optimal the evolved agent controller parameters were via running four control experiments that used the parameter values attained (for the gathering and transport agent controller parameters) at the end of the evolutionary process as fixed agent controller settings. For each of these four experimental settings, the control experiments yielded superior performance according to the *total value gathered* and *system stability* performance measures. These control experiment results thus indicate that the evolutionary process derived a superior agent controller parameter configuration (when comparing the SREL and locally restricted reproduction scheme combination with other reproduction scheme combinations) for the given agent collective task.

	SREL Panmictic	SREL Local	MRDL Panmictic	MRDL Local
Evolution	23.59 (33.37)	39.10 (17.20)	32.56 (10.00)	22.85 (17.60)
Control	60.25 (1.85)	70.00 (4.88)	43.04 (5.46)	54.28 (0.63)

Table 2. The values attained for the total value gathered cooperatively (standard deviations in parentheses) under both the agent collective utilizing artificial evolution (Evolution) and the agent collective utilizing previously evolved agent controller parameter settings (Control). The upper row at the top of the table refers to the temporal reproduction schemes and the second row refers to the spatial reproduction schemes. Note 'Local' refers to the locally restricted reproduction scheme.

	SREL Panmictic	SREL Local	MRDL Panmictic	MRDL Local
A	2.52 (0.54)	19.67 (2.70)	10.31 (4.66)	17.78 (12.90)
B	21.10 (29.95)	38.26 (5.11)	22.43 (4.41)	49.34 (18.48)
C	42.24 (58.52)	57.98 (8.23)	74.90 (14.72)	50.03 (30.55)
T	19.45 (26.73)	23.27 (2.64)	18.27 (4.21)	39.61 (4.30)

Table 3. Gathering and transport capacities evolved under each of the four agent reproduction schemes. The notation *A*, *B*, *C*, and *T*, denote the mine types *A*, *B*, *C* and *Transport* capacities, respectively.

Table 2 presents the values gathered cooperatively attained for the four agent reproduction schemes. For each scheme, the values attained in the control experiment are presented below the values gathered cooperatively. The value in parentheses presented next to each of the values gathered cooperatively is the standard deviation. A high standard deviation indicates that the agent collective was less stable in its gathering behavior. High standard deviations were the result of many agent populations (of the 100 replications) becoming extinct before the end of a simulation.

A low standard deviation indicates a low portion of agent populations dying out prematurely and hence a high stability in the gathering task. The control experiments demonstrated that both the SREL and locally restricted reproduction scheme was operating within a region of the parameter space (defined by the four agent controller parameters) where a high value gathered cooperatively was attainable.

The result of the SREL and locally restricted agent reproduction scheme being most appropriate for both approaches is theorized to be consequent of agents only reproducing at the end of their lifetimes. Using the SREL setting, agents that have performed their task well and have thus survived until the end of allotted lifetime, are allowed reproduce. Given that the reproduction action costs 90% of the parents' energy, agents using the MRDL setting have less of a chance of producing offspring that are well suited to successful task accomplishment.

Using the heuristic controller, child agents inherit only recombined and mutated agent parameter values and an average of parent fitness. However, the nature of the SREL setting holds, in that only agents with appropriate controller parameter settings will have survived until the end of their allotted lifetime (that is, those agents with a high fitness).

Table 3 illustrates that the SREL and locally restricted reproduction scheme was successful in deriving parameter values with a high stability (low standard deviation) comparative to the other reproduction schemes. The efficacy of these parameter values are reflected in the value gathered cooperatively attained by the SREL and locally restricted reproduction scheme (presented in *table 2*).

A high standard deviation in the case of the other reproduction schemes indicates that comparatively, these schemes were less appropriate for guiding the evolutionary process in the derivation of agent controller parameter values that produced a relatively high value gathered cooperatively.

VI. CONCLUSIONS

Results indicated that an agent collective utilizing the single reproduction at end of lifetime (SREL) and the locally restricted reproduction scheme combination yielded superior performance in terms of the evaluation criteria defined. The evaluation criteria were defined as the total value of mines gathered cooperatively, and the stability of the system.

Evolved agent collective behaviors were able to achieve good results, according to the two objective performance measures defined, as well as in comparison to previous experimental results. That is, previous experimental results used an adaptive (though not evolutionary) agent collective in the same task environment, though inferior results were attained, in terms of the same two objective performance measures. The optimality of the evolved agent controller parameter settings was addressed via implementing a control experiment that utilized the evolved values as static agent controller parameters. Agents using these evolved parameter values were able to attain a higher level of performance, according to the evaluation criteria, comparative to the evolutionary process using randomly initialized agent controller parameters and the SREL and locally restricted reproduction schemes.

Two future research objectives have been defined based upon the results presented in this paper. The first is to further investigate the optimality of the evolved agent controller parameters via making a comparison with an advanced control experiment. Such a control experiment would stochastically sample thousands of points within the solution space (agent controller parameter values), and then subsequently test each these agent controller parameters for 100 simulation runs to gather statistics on how the two objective performance measures are addressed. The goal of such an experiment would be to ascertain if evolution discovered an optimal set of agent controller parameters. Such an experiment would illustrate the efficacy of an artificial evolution process for deriving agent controller parameters for a given agent collective task in an unknown environment.

The second future research objective is to increase the complexity of the agent controllers and evolutionary process, giving agents the capacity to learn during their lifetimes, as well as evolution of the capacity to modify genotypes based upon lifetime behaviors (collective or individual). Modifying the evolutionary process such that a greater part of the agent genotype is subject to evolution would also likely yield greater complexity and diversity in emergent behaviors.

REFERENCES

[1] C. Adami, "Evolutionary Learning in the 2D Artificial Life System Avida", Kellogg Radiation Lab, Caltech, Pasadena, Tech. Rep., 1994.

[2] E. Akiyama, K. Kunihiko, "Evolution of Cooperation, Differentiation, Complexity, and Diversity in an Iterated Three-Person Game", *Artificial Life*, 1995, vol. 2(1): 293-304.

[3] R, Axelrod. "The Evolution of Cooperation". New York: Basic Books, 1984.

[4] G. Baldassarre, S. Nolfi, D. Parisi, "Evolving Mobile Robots Able to Display Collective Behaviors", in *Proc. Int. Workshop. Self-Organization and Evolution of Social Behavior*, Zurich, 2002, pp. 11-22.

[5] Y. Cao, A. Fukunaga, A. Kahng, "Cooperative Mobile Robotics: Antecedents and Directions". *Autonomous Robots*, 1997, vol. 4(1): 7-27.

[6] J. Deneubourg, S. Goss, J. Pasteels, D., Fresneau, J. Lachaud, "Self-Organization Mechanisms in Ant Societies (II): Learning in Foraging and Division of Labor". *Behavior in Social Insects*, 1987, vol. 54(1): 177-196.

[7] G. Di Caro, M. Dorigo, "AntNet: Distributed Stigmergic Control for Communications Networks". *Journal of Artificial Intelligence Research*, 1998, vol. 9(1): 317-365.

[8] A. Drogoul, J. Ferber, B. Corbara, D. Fresneau, "A behavioral simulation model for the study of emergent social structures", in *Toward a practice of autonomous systems*, F. Varela, P. Bourgine, Eds. Brussels: Kluwer Academic Publishers, 1992, pp. 161-170.

[9] A. Drogoul, B. Corbara, S. Lalande, "MANTA: New Experimental Results on the Emergence of (Artificial) Ant Societies", in *Artificial Societies: the computer simulation social life*, N. Gilbert, N., R. Conte. Eds. London: University College of London Press, 1995, pp. 190-211.

[10] L. Dugatkin, "N-Person games and the evolution of cooperation: A model based on predator inspection in fish". *Journal of Theoretical Biology*, 1990, vol. 142(1), 123-135.

[11] A. Eiben, J. Smith, "Introduction to Evolutionary Computing", Berlin: Springer-Verlag, 2003.

[12] J. Epstein, J., R. Axtell, "Growing Artificial Societies: Social Science From The Bottom Up". New York: Brookings Institute Press, 1998.

[13] D. Floreano, S. Nolfi, "Evolutionary Robotics", Cambridge: MIT Press, 2000.

[14] I, Harvey, "Artificial Evolution and Real Robots", in *Proceedings of International Symposium on Artificial Life and Robotics*, S. Masanori. Ed. Berlin: Springer-Verlag, 1997, pp. 18-20.

[15] H, Iba, "Emergent cooperation for multiple agents using genetic programming", in, *Parallel Problem Solving from Nature, PPSN IV*. Berlin: Springer-Verlag, 1996, pp. 32-41.

[16] N, Jennings. Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence*, 1995, vol. 75(2), 195-240.

[17] C, Kube, H, Zhang, "Collective Robotics: From Social Insects to Robots". *Adaptive Behavior*, 1993, vol. 2(1), 189-218.

[18] N, Mitsumoto, T., Fukuda, K., Shimojima, A., Ogawa, "Micro Autonomous Robotic System and Biologically Inspired Swarm Strategy as a Multi Agent Robotic System", in, *Proceedings of the IEEE International Conference on Robotics and Automation*. Tokyo: IEEE Press, 1995, pp. 63-67.

[19] S. Nolfi, G. Baldassarre, D. Parisi, "Evolution of collective behaviour in a team of physically linked robots", in *Applications of Evolutionary Computing*, R. Gunther, A., Guillot, J. Meyer. Eds. Heidelberg: Springer Verlag, 2003, pp. 581-592.

[20] S. Nolfi, J. Deneubourg, D. Floreano, L, Gambardella, F. Mondada, M. Dorigo, "Swarm-Bots: Swarm of Mobile Robots able to Self-Assemble and Self-Organize". *Ecrim News*., vol. 53, no. 1, pp. 25-26, Apr. 2003.

[21] V. Trianni, R. Gross, T. Labella, E. Sahin, M. Dorigo, "Evolving Aggregation Behaviors in a Swarm of Robots", in, *Proc. Seventh European Conf. Artificial Life,* Dortmund, 2003, pp. 865-874.

[22] A. Perez-Uribe, D. Floreano, L., Keller, "Effects of group composition and level of selection in the evolution of cooperation in artificial ants", in, *Proc. Seventh European Conf. Artificial Life,* Dortmund, 2003, pp. 128-137.