Emergence of Cooperation in a Pursuit-Evasion Game

Geoff Nitschke

Artificial Intelligence Laboratory
Department of Information Technology
University of Zurich, Wintcrthurstr. 190
Zurich, Switzerland
nitschke@ifi.unizh.ch

Abstract

This research concerns the comparison of three different artificial evolution approaches to the design of cooperative behavior in a group of simulated mobile robots. The first and second approaches, termed: *single pool* and *plasticity,* are characterized by robots that share a single genotype, though the plasticity approach includes a learning mechanism. The third approach, termed: *multiple pools,* is characterized by robots that use different genotypes. The application domain implements a pursuit-evasion game in which teams of robots of various sizes, termed: *predators,* collectively work to capture either one or two others, termed: *prey.* These artificial evolution approaches are also compared with a static rule based cooperative pursuit strategy specified *a priori.* Results indicate that the multiple pools approach is superior comparative to the other approaches in terms of measures defined for prey-capture strategy performance. That is, this approach facilitated specialization of behavioral roles allowing it to be effective for all predator team sizes tested.

1 Introduction

The use of evolutionary computation as a design process for control of robot reams has received much attention in both real world [Quinn, 2000], [Schultz and Grefenstette, 2000], and simulated problem domains [Kitano, 1997]. Such attention is due to the difficulty in specifying *a priori* an effective means for multiple interacting robot controllers, and the inherent adaptability of artificially evolved controllers. The use of evolutionary computation to exploit emergent cooperation remains a relatively unexplored area of research in the pursuit and evasion domain [Benda *et ai* 1985] and related predator-prey systems [Nishimura and Takashi, 1997] using multiple predators and prey. Various approaches have been used to study the pursuit evasion domain, where the task is for multiple predators to capture a prey by surrounding it [Korf, 1992], [Levy and Rosenschein, 1992], though few researchers have investigated emergent cooperation in these systems, with notable exceptions such as Denzinger and Fuchs [1996], Haynes and Sen [1996] and Yong and Miikkulainen [2001].

This paper describes a comparison of three artificial evolution approaches for the synthesis of cooperative behaviour evaluated within teams of simulated *Khepera* robots [Mondada *et al.* 1993]. The three approaches are compared with a rule based pursuit strategy and each is evaluated in terms of predator fitness scored, and the time period for which a prey is immobilized. For each of the four approaches teams of various sizes are compared. Cooperative behavior is only evolved for the predators, and each prey is able to move 20 percent faster than the predators. The behavior of each prey is not evolved, but instead uses static obstacle avoidance behavior. Functionally, each predator is the same in terms of movement and sensor capabilities. The collective task was for the predators to immobilize a prey. The predator team was rewarded fitness proportional to how much it was able to slow down a prey, where maximum fitness was rewarded if a prey was immobilized. A control experiment using a single predator, demonstrated that at least two predators are needed to accomplish this task. Experimental results support a hypothesis that the multiple pools approach, which encourages behavioural specialization, would yield a superior performance in terms of the two measures defined to quantify evolved prey-capture strategy performance and this superior performance would prove consistent for all group sizes tested.

2 Artificial Evolution Approaches

Four sets of experiments were run in order to test each of the three artificial evolution approaches, and the rule-based approach. Ten replications of each experiment were made, where each ran for 500 generations. The predator team 'lived' for *10 epochs,* where each epoch consisted of *1000 cycles* of simulation time. Each epoch constituted a test scenario where all predators and prey were tested for different, randomly generated orientations and starting positions in the environment. Six different group configurations of predators and prey were tested for each of the four approaches. The group types were as follows. *Group type* 3 predators and 1 prey; *Group type 2:* 3 predators and 3 prey; *Group type 3: 5* predators and 1 prey; *Group type 4:* 5 predators and 3 prey; *Group type 5:* 8 predators and 1 prey; *Group type 6:* 8 predators and 3 prey.

Single Pool Approach: As illustrated in *figure 1 (left)* this approach generates and tests *n* copies of a single genotype, meaning that the predator team is homogenous.

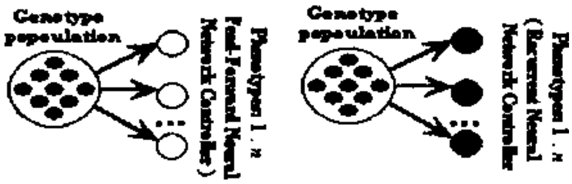

Figure 1. Left: Single Pool Each predator phenotypc corresponds to a genotype selected from a population and copied *n* times. Right: Plasticity - As with Single Pool, though phenotypes implement a recurrent neural network controller.

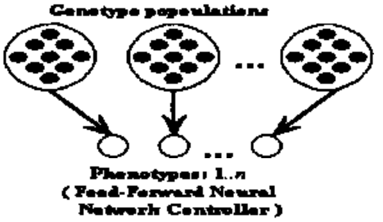

Figure 2. Multiple Pools Approach - Predator phenotypes correspond to *n* different genotypes selected from *n* separate pools of genotypes.

In this approach there is no plasticity so the predators cannot adapt during their lifetime. The fitness assigned to each predator is simply the fitness calculated for the single genotype that specifies the predator team. The main advantage of this approach is its simplicity in terms of behavioral encoding and calculation of team fitness.

Plasticity Approach: As illustrated in *figure J (right)* this approach generates and tests *n* copies of a single genotype, so that as with the *single pool* approach, the predator team is homogenous. The difference is that individual phenotypes are able to adapt during their lifetime as a result of a recurrent neural network learning process. The advantage of the plasticity approach is that it allows for specialization of behaviour by individual predators without being affected by the problem of needing to estimate fitness contribution of different predators to the team as a whole. For both the single pool and plasticity approaches, every individual genotype in the population is tested against *n*-\ randomly selected genotypes from the same population. This process is repeated for all 10 epochs of a predator's lifetime.

Multiple Pools Approach: As illustrated in *figure 2* this approach takes a single genotype from every population of genotypes where the number of populations corresponds to the number of predators. Each genotype is then decoded into a separate phenotype, where this set of phenotypes then represents the team of predators. In each generation, every individual genotype in a population is tested against *n*-\ other genotypes, randomly selected from one of the other populations of genotypes. This process is then repeated for every epoch. The advantage of the multiple pools approach is that it encourages behavioural specialization in the group of predators, in that the artificial evolution setup provides for more genetic diversity.

Rule Based Approach: This approach takes a predator previously evolved for obstacle avoidance and then extends this behavior with a simple set of rules designed to encourage a cooperative encirclement pursuit strategy when at least two predators collectively approach a single prey. The rules cause a predator to move towards the right side of detected prey, and to push to the left against that side, while moving about the prey's periphery. If an obstacle is detected on the preys right side, the predator will attempt to keep moving about the preys periphery until a Tree-space' is found.

Evaluation of Approaches

In order to quantify the effectiveness of emergent prey capture strategies, two different measures were used to evaluate performance. The first was predator team fitness, where fitness awarded to the team was proportional to how much a prey was slowed during the team's lifetime. The second was prey capture time, which was the time period for which a prey was immobilized by the collective efforts of at least two predators. For both the *single pool* and *plasticity* approaches a single genotype specifies the entire predator team. That is, predators are clones of each other, so evaluation of team performance in this case is not problematic. The performance of a predator team executed under either of these approaches is simply measured as the fitness value assigned to the genotype that specifies the team. For the rule-based approach, a pseudo fitness function that emulates the single pool and plasticity approaches was implemented. In contrast to these approaches a predator team using the multiple pools approach is specified by *n* genotypes selected from *n* different populations. Hence, each genotype must be assigned an individual fitness score, and team performance evaluation needs to be computed by estimating the fitness contribution of each genotype to the team as a whole. A method of evaluation widely known as: *fitness sharing* [Bull and Holland, 1997] was implemented for the multiple pools approach, where an equal fitness score is assigned to each individual genotype, thereby assuming that each individual contributed to team performance equally. The advantage of this method is that fitness for individual genotypes is easily calculated and there is no disparity between team fitness and the fitness of individual team members.

Agents, Environment and Artificial Evolution

For all experiments a generational evolutionary algorithm using linear rank-based selection was used [Goldberg, 1989]. Each population contained 100 genotypes, where initial populations consisted of randomly generated genotypes. Genotype length was set to 24 genes, where each gene consisted of several bits encoding each neuron type and connection weights. At the turn of each generation, the 20 genotypes that have accumulated the highest fitness were allowed to reproduce. The total fitness of an individual genotype was the sum of all its fitness for all epochs of its life. Reproduction was done via generating five copies of each genotype in order to create the next generation.

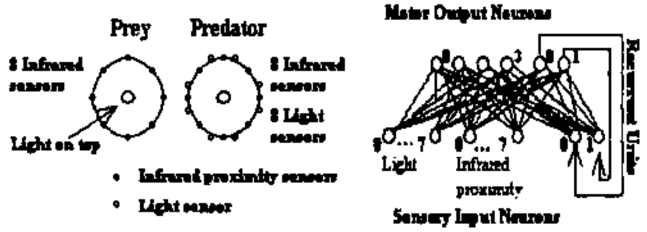

Figure 3. Left: Each predator and prey is a simulated *Khepera* robot. Predators have 8 infrared proximity sensors, and 8 light sensors. Prey have 8 infrared proximity sensors, and a light. Right: A feed forward neural network controller is used for the single pool and multiple pools approaches, and a recurrent neural network controller is used for the plasticity approach.

During this copying process 10 percent of the connection weights were mutated. Mutation added a random value between -1.0 and +1.0 to the weights current value. This process was repeated for the 500 generations that each simulation ran for. As illustrated in *figure 3* (left), the body of each predator and prey is assumed to be a Khepera mobile robot [Mondada *et al.* 1993].

The robots used as predators were equipped with 8 infrared proximity sensors, and 8 light sensors positioned on the periphery of the Khepera. The robots used as prey were equipped with 8 infrared proximity sensors, as well as a light on its top. This light could be detected by the predator light sensors and was used so as each predator robot could distinguish fellow predators from a prey. *Figure 3* (right) illustrates the feed-forward neural network consisting of an input and output layer with no hidden units [Nolfi and Parisi, 1997] that controlled all robots executed under the single pool and plasticity approaches. In the case of the predators, the input layer consisted of 16 units that encoded the activation level of the robots 16 sensors. These 16 input units were connected to 4 output units. *Figure 3* (right) also illustrates the recurrent neural network controller used in the plasticity experiments, where the activation level of two additional output units was copied back into two additional input units. The first two output units represented the two motors of the robot and encoded the speed of the two wheels. These motor units controlled the robots behavior in the environment. The next two output units represented two teaching units that encoded a teaching input for the first two output units. The two motor units used this teaching input in order to learn using the back propagation procedure [Rumelhart *et al.* 1986]. In the plasticity experiments there were an additional two output units that were the recurrent units and contained activation values for the motors from the previous cycle. For the robots that were the prey, a network connecting 8 sensory input units to 4 motor output units was trained for an obstacle avoidance behavior before being placed in the environment.

The environment corresponded to a 1000cm x 1000cm arena with no obstacles. When a predator robot was placed in the environment, sensory input was received via the input units, and activation values were passed to the two motor units, and the teaching units. The activation value of the two motor units was used to move the robot, thus changing the sensor input for the next simulation cycle. The activation value of the two teaching units was used to change the weights that connected the input units to the motor units using back propagation. This cycle was then repeated.

3 Results

Four sets of experiments were executed to test each of the three artificial evolution approaches and the rule-based approach. Each set of experiments tested the six group types and ten replications of each experiment were made.

Figure 4 illustrates the average fitness attained for all groups of predators using the rule-based, single pool, plasticity and multiple pools approaches. *Figure 5* presents the average *prey capture time* attained for the four approaches. *Prey capture* refers to the instance when the prey is immobilized for a time interval $t_0..t_h$

Evolved Behavior

Evolved behavior observed for each of the three artificial evolution approaches is described in this section. For the rule-based approach, even though behavior was not evolved as for the other approaches, a single cooperative pursuit strategy emerged as the result of interaction of multiple predators following a simple set of rules.

Single Pool: For all six group types tested only two cooperative prey-capture strategies consistently emerged. These strategies, termed: *encirclement* and *entrapment* are briefly described in the following. In the encirclement strategy at least three predators move to circle the prey, each moving in the same direction in close proximity to the prey, for some period of time. The strategy was rarely able to immobilize the prey and was only effective in slowing the prey for a short period of time, given that the predators were not able to coordinate their movements for an extended period. The entrapment strategy, illustrated in *figure* 6, also used at least three predators, where either one or two predators moved to each side of the prey, while another, termed: a *blocker,* moved around the flanking predators, to approach the prey from the front, in order to trap the prey in a triangular formation. In the example depicted in *figure 6,* Pi and P_2 are the flanking predators and P_3 is the blocker. Note that at simulation time step w+/ when the prey moved to escape, the flanking predators moved also, turning one way to force the prey in a specific direction. The blocker then moved around in order to affront the prey again. This system of entrapment, movement, and then entrapment continued for several times before the prey was able to evade the predators.

Rule-Based: A single cooperative strategy emerged using at least two and at most three predators. The strategy was similar to the encirclement strategy that evolved under the single pool approach. As with the single pool version of encirclement, the strategy was rarely successful at immobilizing a prey due to the lack of coordination between predators as they collectively approached a prey.

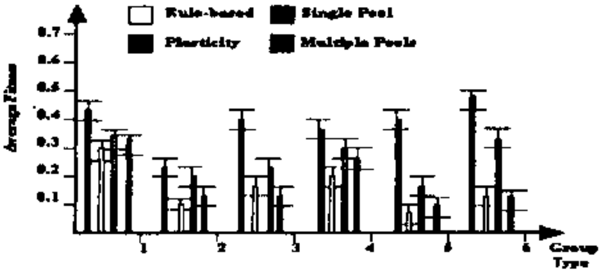

Figure 4. Average fitness presented, with standard error bars, for all group types tested under the three artificial evolution approaches, and the rule-based approach.

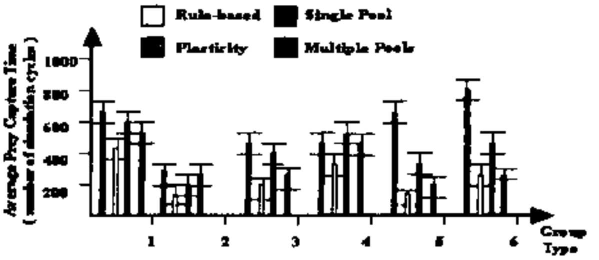

Figure 5. Average *prey-capture time* presented, with standard error bars, for all group types tested under the three artificial approaches, and the rule-based approach.

Plasticity: In these experiments one emergent cooperative prey capture strategy was observed. This strategy, illustrated in *figure* 7, was a derivative of the entrapment strategy and termed: *role switcher.* In this strategy a form of behavioural specialization emerged during the lifetime of the predators. Specifically, two different types of behavioral roles emerged, termed *knocker* and *flanker.* A flanking predator moved along one side and just ahead of the prey forcing the prey to turn in a particular direction, where as a knocking predator 'knocked' into the side of the prey forcing it to slow its movement. The role-switcher strategy was formed by a knocker predator moving along side a flanking predator and role of a flanker and knocker switching whenever a prey attempted evasion. *Figure* 7 depicts predator P_3 formerly flanking assumes the knocking role at simulation time step *n+k* after the prey moves to evade, and predator P_1 formerly knocking assumes a flanking role. This dynamic adoption and switching of roles coordinated the movements of the predators and thus effectively slow a prey's movement. Also, relative to the entrapment and encirclement strategies, role switcher was able to more frequently immobilize a prey.

Multiple Pools: As with the plasticity experiments the *role switcher* strategy was the only cooperative prey-capture strategy that consistently emerged, though a specific difference was noted. Namely, different predators adopted different roles from the beginning of their lifetimes. Specifically, particular predators always assumed the role of a *flanker,* while other predators always assumed the role of a *knocker. Figure 8* depicts predators P_2, P3 and P_5 as maintaining the roles of flankers, and predators P_1 and P_4 maintaining the roles of knockers.

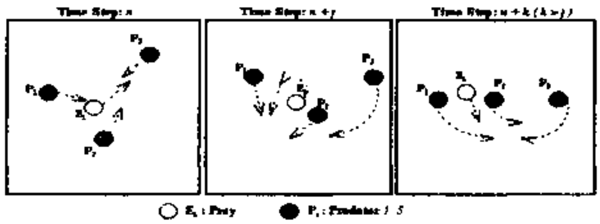

Figure 6. The cooperative entrapment pursuit strategy, illustrated for *group type 1:* 3 predators and 1 prey. The strategy was not successful at immobilizing an evader.

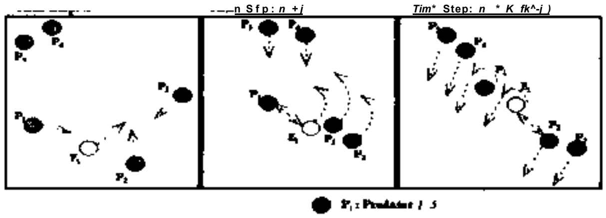

Figure 7. The role-switcher pursuit strategy, illustrated for *group type 3:* 5 predators and 1 prey. Note that a dynamic assumption of behavioral roles in the formation of the strategy is depicted here.

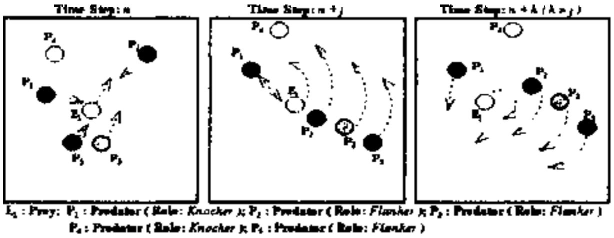

Figure 8. The role-switcher pursuit strategy, illustrated for *group type 3: 5* predators and 1 prey. Here the assumption of behavioral roles is set at the beginning of a predator's lifetime.

Thus, at simulation time step *n+j* predator P_4 does not move to assume the role of a knocker in the strategy as it is already taken by predator P_1 at this stage. This adoption of roles that was maintained throughout the lifetime of the predators served to increase the effectiveness of the role switcher strategy in terms of the measures defined for prey capture strategy performance. Notably the strategy was consistently effective at immobilizing a prey, for all group types tested.

4 Analysis and Discussion

This discussion compares emergent prey-capture strategies with the rule-based pursuit strategy and relates each to the group types tested and the two performance measures.

Single Pool: Two cooperative prey-capture strategics, each using at least three predators, consistently emerged for all six group types tested. Though, relative to the plasticity and multiple pools approaches these emergent strategies performed poorly in terms of average *team fitness* and *prey-capture time.* These results are presented in *figures 4* and 5 respectively and as shown were consistent for all group types tested. The low fitness and prey-capture time of the encirclement and entrapment strategies was found to be a result of interference that occurred as two or more predators collectively approached a prey. This result was found to be due to confused infrared sensor readings of predators in close proximity to each other, and the fact that individual predators did not possess any memory, explicit form of communication, or coordination to facilitate a successful cooperative prey capture strategy. As illustrated in *figures 4* and 5, this result is especially prevalent for *group types 5* and 6, which used eight predators.

Rule-Based: From the interactions of predators following a set of rules, an encirclement strategy emerged, though as with the single pool version of the strategy it was effective for at most three predators. Also, note in *figures 4* and 5, that similar to its single pool counterpart the strategy proved most effective for *group types: 1* and *4.* Given that in these experiments the chance of predators physically interfering with each other as they approached the prey was minimal. The rule-based approach, similar to the single pools approach, also performed best in group types where the predator to prey ratio was medium as in *group types 1* and *4.*

Plasticity: In experiments using this approach the role switcher strategy frequently emerged for all group types tested. In the role-switcher strategy a form of dynamic behavioral specialization emerged in at least three predators. This behavioral specialization was in the form of dynamic role adoption that emerged during the lifetime of a predator and varied from predator to predator depending upon the group type being tested. This dynamic role adoption facilitated cooperation between the predators affording the predator team a high average *fitness* and *prey-capture time* comparative to the single pool and rule-based approaches, for all group types tested. This result is presented in *figures 4* and *5* respectively. In experiments testing *group types* 7, *3* and 5 at least three and at most six predators constituted the role-switcher strategy, where a maximum of three flanked cither side, and the role of a flanking and knocking predator switched whenever the prey turned to evade. In experiments testing group *types* 7 and *6,* at least three and at most eight predators formed the role-switcher strategy, where a maximum of two predators flanked either side of a prey. As with the experiments testing *group types* 7, *3,* and *5* the role of the predator closest to the prey and the knocking predator switched whenever the prey turned so as to escape, though the difference of experimental results from testing *group types J* and 6, was that two sub-groups of predators emerged, meaning that two of the prey were simultaneously engaged by two predator sub-teams.

The dynamic assumption of roles during a predators lifetime allowed for the formation of sub-groups in the predator team, and thus served to yield a higher average team fitness and prey capture time. As presented in *figures 4* and *5* respectively, these two measures were relatively higher in experiments testing *group types* 7, *4,* and *6,* and low in experiments testing *group types 3* and 5. In experiments testing group *types 3* and 5, this result was due to multiple predators attempting to assume the same behavioral role as they collectively approached the single prey. For *group type* 2, the 1 predator to 1 prey ratio increased the chance of the strategy dispersing given an increased probability of distraction from a passing prey. In the other group types tested a higher predator to prey ratio reduced the number of predators that collectively approached the prey at any given time, thus increasing the likelihood of unhindered dynamic adoption of a behavioral role in the role-switcher strategy. It is theorized that behavioral specialization for role switching emerged as an indirect result of interference that occurred when at least three predators collectively approached a prey. Such interference was observed in the single pool experiments and often caused emergent strategies to fail prematurely, thereby making it more difficult for such strategies using at least three predators to be selected for and propagated by the evolutionary process. Where as, the *role-switcher* strategy was able to achieve a high team fitness and prey-capture time, it was at best only able to slow the prey, and never completely immobilizing the prey in experiments testing *group types 2* through to *5.* This is reflected in *figure* 5, which presents a relatively higher average prey capture time for experiments testing *group types* 7 and *6.* Thus, the dynamic adoption of roles was only effective where the predator to prey ratio was sufficiently high, though not so high as in the case of *group type 5* where an excess of physical interference between predators prevented the role-switcher strategy from regularly forming.

Multiple Pools: In experiments using the multiple pools approach only a single cooperative prey-capture strategy emerged. This strategy was classified as a derivative of the role-switcher strategy using at least three predators that was observed in the plasticity experiments. Though the multiple-pools role-switcher used a genetic based specialization. In the plasticity experiments the adoption of specialized behaviour was dependent upon the positions of the predators at a given time, where as in the multiple pools experiments, different predators initially behaved differently and assumed genetically pre-determined roles prior to engaging in the multiple pools version of the *role-switcher* prey-capture strategy. That is, one predator always assumed the role of the knocker while the others always assumed the role of a flanker, or that of an idle predator. Idle predators, for example predator P_4 in *figure 8,* served the purpose of reducing any potential interference between predators as they collectively approach a prey. Idle predators also increased the chance of success of the prey-capture strategy via limiting the number of predators that constituted the strategy. The fact that the predators were genetically different is one explanation for the evolution of specialized behavioural roles that complemented each other in the emergence of the multiple pools role switcher strategy. In the experiments testing *group type* 6_y two specialized sub-groups of predators, comprising two groups of four predators emerged.

The effectiveness of these two specialized sub-groups is illustrated in *figures 4* and 5, which present a relatively high average fitness and prey-capture time for all group types, comparative to the single pool, rule-based and plasticity approaches. Note that as with the plasticity approach, the multiple pools approach is more effective at exploiting an environment containing at least two prey via the formation of two specialized predator sub-groups. Thus, the multiple pools approach facilitated the evolution of specialized predator sub-groups to simultaneously capture two prey, where genetic based specialization proved to be more effective, in terms of measures for *fitness scored* and *prey capture time.* This comparison was made with the dynamic assumption of behavioral roles as done by predators using the plasticity approach.

5 Conclusions

This paper presented a set of experiments testing three different artificial evolution approaches and a rule-based approach for the synthesis of cooperative behaviour within a team of predators, where the task was to cooperatively capture at least two prey. Results compared the performance of these approaches for six different group configurations. The comparison was in terms of two measures designed to quantify the efficacy of emergent cooperative prey capture strategies. These measures were *predator fitness and prey-capture time.* Results presented indicated the multiple pools approach to be superior in terms of these two measures for all group types tested. The superiority of the multiple pools approach was found to be a result of a genetic form of behavioral specialization that assigned behavioral roles at the beginning a predators lifetime. The multiple pools approach also facilitated the evolution of two specialized sub-groups of predators in scenarios using at least two prey. Evolution of these sub-groups improved the effectiveness of these strategies in terms of higher *team fitness* and *prey capture time,* comparative to the rule-based, single pool and plasticity approaches. Specifically, the specialized sub-groups aided in reducing interference between predators as they collectively approached a prey.

A comparison with other evolution of cooperation approaches in the pursuit-evasion domain [Denzinger and Fuchs, 1996], [Haynes and Sen, 1996] and [Yong and Miikkulainen, 2001] is difficult given the real world nature of the experiments described in this paper. Though the robots were simulated, the environment was a continuous domain and the simulation incorporated noise in sensory data, namely confused infrared sensor readings resulting from two or more Khepera's being in close proximity to each other. This noisy sensor data was a key reason for interference between multiple predators as they collectively approached a prey. Also, a continuous environment does not allow for the selection of distinct sets of situation/action values that are possible in grid world implementations [Denzinger and Fuchs, 1996] where a finite set of actions and resultant outcomes can be defined. While, the emergence of cooperation is simpler to analyze in these grid world domains, they are limited by their own implementations, so the study of mechanisms that facilitate emergent cooperation such as behavioral specialization is limited to trivial situations.

References

[Bcnda *et al.* 1985] Benda, M., Jagannanthan, V., and Dodhiawalla, R. An Optimal Cooperation of Knowledge Sources. Technical BCS-G2010-28, Boeing AI Center, August 1985.

[Bull and Holland, 1997] Bull, L., and Holland, 0. Evolutionary Computing in Multi-Agent Environments: Eusociality. In *Proceedings of the Second Annual Conference on Genetic Programming,* pages 347-352. San Francisco, USA, July, 1997.

[Denzinger and Fuchs, 1996] Denzinger, J., and Fuchs, M. Experiments in Learning Prototypical Situations for Variants of the Pursuit Game. In *Proceedings of the Second ICMAS conference,* pages 48-55, Kyoto, Japan, August, 1996.

[Goldberg, 1989] Goldberg, D. Genetic Algorithms. Addison Wesley, New York, USA. 1989.

[Haynes and Sen, 1996] Haynes, T., and Sen, S. Evolving behavioral strategies in predators and prey. *Adaptation and Learning in Multi-Agent Systems,* Springcr-Verlag, Berlin, 1996.

[Kitano *et al.* 1999] Kitano, H., Asada, M., Noda, I., and Veloso, M. (1999). RoboCup: Today and Tomorrow - What we have learned. *Artificial Intelligence,* vol. 110(2): 193-214. Elsevier Science Publishers, Amsterdam, Holland.

[Korf, 1992] Korf, R. E. A simple solution to pursuit games. *Working Papers of the Eleventh International Workshop on DAI,* pages 195-213. Geneva, Switzerland, February, 1992.

[Levy and Rosenschein, 1992] Levy, R., and Rosenschein, J. S. A Game Theoretic Approach to Distributed Artificial Intelligence and the Pursuit Problem. In *Decentralized AI III,* pages 129-146, Kaiscrslautern, Germany, August, 1992. Springer-Verlag.

[Mondada *et al.* 1993] Mondada, F., Franzi, E., and lenne, P. Mobile Robot Miniaturization: A tool for Investigation in Control Algorithms. In *Proceedings of* Third International Symposium on Experimental Robotics, pages 501-513, Kyoto, Japan, October, 1993. Springer-Verlag.

[Nishimura and Ikegami, 1997] Nishimura, S. I., and Ikcgami, T. Emergence of Collective Strategies in a Prey-Predator Game Model. *Artificial Life,* 3(1): 243-260. July, 1997.

[Nolfi and Pansi, 1997] Nolfi S. and Parisi D. Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior,* 5(1): 75-98. July, 1997.

[Quinn, 2000] Quinn, M. Evolving cooperative homogeneous multi-robot teams. In *Proceedings of the International Conference on Intelligent Robots and Systems,* pages 1798-1803, Takamatsu, Japan, October 2000, IEEE Press.

[Rumelhart *et al* 1986] Rumelhart, D. E., Hinton, G. E., and Williams, R. J. Learning internal representations by error propagation. *Parallel Distributed Processing, Volume I: Foundations.* MIT Press, Cambridge, 1986.

[Schultz and Grefenstette, 2000] Schultz, A. C. and Grefenstette, J. J. Continuous and Embedded Learning in Autonomous Vehicles: Adapting to sensor failures. In *Unmanned Ground Vehicle Technology II,* Proceedings of SPIE, vol. 4024(1): 55-62, American Association for Artificial Intelligence, 2000.

[Yong and Miikkulainen, 2001] Yong, C. H., and Miikkulainen, R. Cooperative Co-evolution of Multi-Agent Systems. Technical Report AIO 1-287. Department of Computer Science, University of Texas, July 2001.