Emergence of Cooperation in a Multiple Predator, Single Prey Game

Geoff Nitschke

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

Abstract

This research concerns the comparison of three different artificial evolution approaches for the design of cooperative behavior in a group of simulated mobile robots. The first and second approaches, termed: *single pools* 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 is a pursuit-evasion game in which a team of robots, termed: *predators*, collectively work to capture or slow a single robot, termed: *the prey*. Results indicate that the multiple pools approach is superior comparative to the other two approaches in deriving robust and consistently effective prey-capture strategies, given that this approach facilitates the evolution of specialized behavioral roles in the predator team.

1. Introduction

The study of the synthesis of collective behavior, particularly the emergence of cooperation, remains a relatively unexplored area of research in the pursuit and evasion domain (Benda *et al.* 1985) and related predator-prey systems (Nishimura and Takashi, 1997) that use multiple predators and multiple 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 via artificial evolution approaches 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 behavior evaluated within a team of three simulated *Khepera* robots (Mondada *et al.* 1993). Cooperative behavior is only evolved for the predators, and the prey is able to move 20 percent faster than the predators. The behavior of the 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 predator team is rewarded fitness proportional to how much it is able to slow down the prey, where the collective task was for the three predators to immobilize the prey. A control experiment using a single predator, demonstrated that at least two predators are needed to accomplish this task. The first approach used for the design of cooperative behavior is termed: *single pool*, in which each predator is specified with an identical genotype. Thus the corresponding behavior (phenotype) for each of the three predators is the same. The second approach is termed: *plasticity*, in which each predator is also specified with an identical genotype, though the corresponding phenotype implements a recurrent neural network controller allowing adaptive behavior during a predator's lifetime. The third approach is termed: *multiple pools*, where each predator is specified using a different genotype, and thus the phenotype at the beginning of each predator's lifetime is different. The three approaches are evaluated in terms of predator fitness scored, the geometrical stability of evolved prey-capture strategies, and the time period for which a prey is immobilized. Experimental results support a hypothesis that the multiple pools approach, which utilized a genetic based behavioral specialization, yielded a superior performance in terms of the three measures defined to quantify evolved prey-capture strategy performance. That is, complementary behavioral roles were evolved, allowing the predators to more effectively execute the task of collective prey capture.

2. Artificial Evolution Approaches

Three artificial evolution approaches were comparatively tested and evaluated for the task of having a team of three predators cooperatively capture a single prey. The evolutionary approaches tested were termed: *single pool*, *plasticity*, and *multiple-pools*. As illustrated in *figures 1*, *2* and *3* respectively, both the single pool and plasticity approaches employed a single population of genotypes, while the multiple pools approach employed three separate populations of genotypes. Each genotype was encoded with the parameters necessary to instantiate a single predator controller (phenotype). Three sets of experiments were run testing each of the three approaches, where each experiment ran for 500 generations, and 10 replications were made. The team of predators 'lived' for either: *5*, *10*, *25* or *50 epochs*. Each epoch constituted a test scenario running for 1000 cycles of simulation time, where all predators and prey were tested for different randomly generated orientations and starting positions.

Comparison of Approaches

Single Pool Approach: As illustrated in *figure 1* this approach generates and tests 3 copies of a single genotype, meaning that

Copyright © 2003, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

the predator team is homogenous. In this approach there is no plasticity so the predators cannot adapt during their lifetime. The fitness assigned to each predator is the 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.

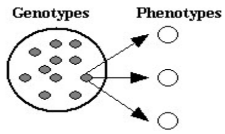

Figure 1. Single Pool – Each predator phenotype corresponds to a genotype selected from a population and copied *3* times.

Plasticity Approach: As illustrated in *figure 2* this approach generates and tests three 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 behavior 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 two randomly selected genotypes from the same population. This process is repeated for all epochs of a predator's lifetime.

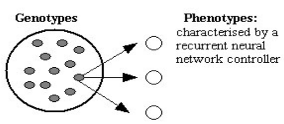

Figure 2. Plasticity. Same as Single Pool, though phenotypes implement a recurrent neural network controller.

Multiple Pools Approach: As illustrated in *figure 3* this approach takes a single genotype from each of the three populations of genotypes. Each genotype is then decoded into a separate phenotype, where this set of three phenotypes then comes to represent the team of predators. In each generation, every individual genotype in a population is tested against two other genotypes, randomly selected from one of the other populations of genotypes. This process is then repeated for all epochs of a predator's lifetime. The advantage of the multiple pools approach is that it encourages behavioral specialization in the group of predators, in that the artificial evolution setup provides for more genetic diversity in the three predator genotypes.

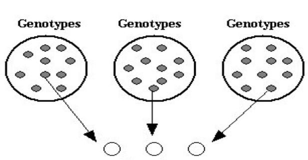

Figure 3. Multiple Pools. Predator phenotypes correspond to *3* different genotypes selected from *3* separate pools of genotypes.

2.2 Evaluation of Approaches

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. In contrast to these approaches a predator team using the multiple pools approach is specified by three genotypes selected from three different populations.

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.

2.3 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. 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 experiment was executed for.

The body of each predator and prey is a simulation of 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 circular periphery of the robot. 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 the prey. A recurrent neural network consisting of input and output layers with no hidden units (Nolfi and Parisi, 1997) controlled all robots. 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. In the plasticity experiments, 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.

2.4 Evaluation of Cooperative Behavior

In order to quantify the effectiveness of emergent prey capture strategies, three 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, and the third measure was a statistical index termed: *group stability index*. Adapted from Baldassarre *et al.* (2002), this index measured how stable a particular geometric formation of a group of predators was with respect to the prey for a given time period. For example, if the predators form a circle about the position of the prey, the index will indicate for how long the predators maintain this circle formation. If the predators are able to hold a certain formation for an extended period of time, the index will be high indicating high group stability.

3. Results

Figures 4, 5, and *6* illustrate the *average fitness*, *group stability*, and *prey-capture time*, for *single pool*, *plasticity* and *multiple pools* approaches.

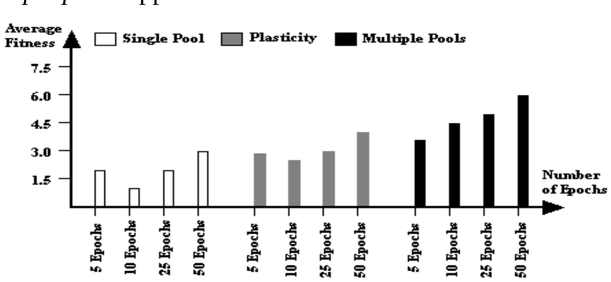

Figure 4. Average fitness for epoch settings tested under the *single pool*, *plasticity*, and *multiple pools* approaches.

3.1 Evolved Behavior

Single Pool: Three emergent cooperative strategies were consistently observed. These were termed: *encirclement*, *entrapment* and *knocker* and are illustrated in figures *7*, *8* and *6* respectively. The three-predator entrapment strategy emerged in *60* percent of replications, while the three-predator encirclement strategy emerged in *40* percent of replications. The most frequently observed was the two-predator *knocker* strategy, which emerged in *80* percent of replications.

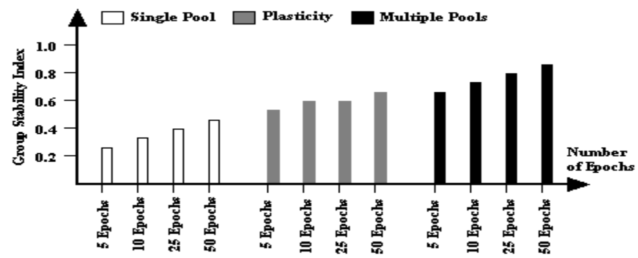

Figure 5. Average *Group Stability Index* for epoch settings tested under the *Single Pool*, *Plasticity,* and *Multiple Pools* approaches.

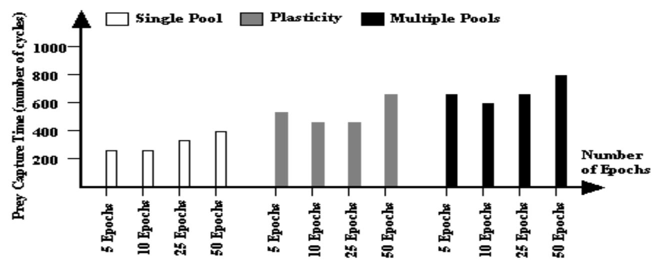

Figure 6. Average *Prey-Capture Time* for epoch settings tested under the *Single Pools*, *Plasticity,* and *Multiple Pools* approaches.

Figure 7 illustrates the encirclement strategy, where 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 slowed the prey significantly and scored a high fitness, though was only effective for a short period of time, as the predators were not able to coordinate their movements for an extended period or able to immobilize the prey. This strategy was also observed using two predators, though prey capture time was longer.

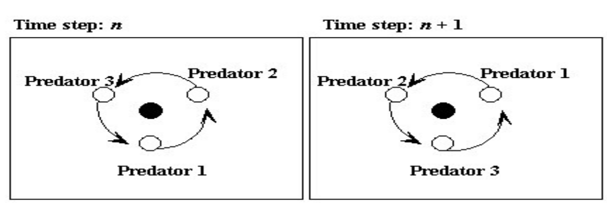

Figure 7. A cooperative strategy, termed *encirclement*, observed in the *single pool* experiments.

Figure 8 depicts an example of the entrapment strategy, using three predators, where a predator moves to each side of the prey (*predator 1* and *2*), while a third (*predator 3*) moves around one of the flanking predators to approach the prey from the front, in order to halt the prey in a triangular formation. When the prey turns to escape, the two flanking predators move also, turning to force the prey in a specific direction. The third predator (*predator 2* now) then moves about also in order to affront the prey. This system of entrapment, movement, and then entrapment continues several times before the prey is able to evade the predators.

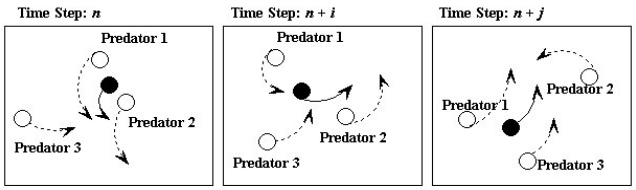

Figure 8. A cooperative strategy, termed *entrapment*, observed in the single pool experiments.

Figure 9 depicts an example of the knocker strategy, where two predators move so as to flank either side of the prey. One predator places itself on the right side of the prey moving in front and forcing the prey to spiral inwards, while the predator to the left moves with the prey also, knocking into the side of the prey forcing it to slow more than usual.

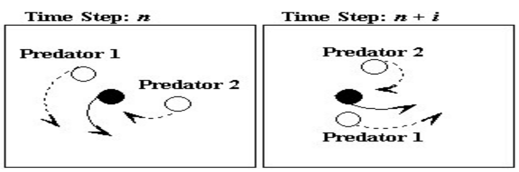

Figure 9. A cooperative strategy, termed *knocker*, observed in the single pool and plasticity experiments.

Plasticity: In these experiments three cooperative prey capture strategies emerged. These included the two-predator strategy: knocker, as well as a derivative of the previously described entrapment strategy. This new derivative strategy, termed: *role switcher*, illustrated in *figure 10* proved superior comparative to other emergent strategies in terms of prey-capture time. *Figure 10* highlights the predators using a form of behavioral specialization so that a third predator moves along side another of the predators and role of the flanker switches between the two predators whenever the prey tries to evade the predators. The knocker strategy emerged most frequently, observed in *80* percent of replications, and the *role switcher* strategy emerged in *50* percent of replications.

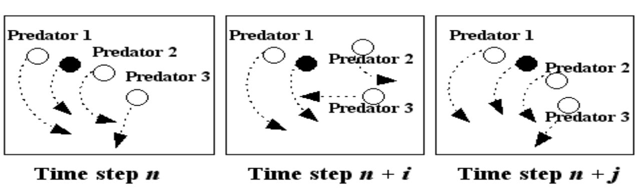

Figure 10. *Role Switcher.* Predator 3 moves aside flanking predator 2, and the flanking role switches when the prey attempts evasion.

Multiple Pools: In these experiments a derivative of the *knocker* strategy was the only cooperative strategy that frequently emerged, observed in *90* percent of replications. In this derivative version, a particular predator always assumed the flanking role, while the other predators assumed the role either the knocker or that of an idle third predator.

4. Analysis and Discussion

In this section the three artificial evolution approaches are discussed in relation in relation to the three performance measures defined for cooperative prey capture.

Single Pool: These experiments yielded prey-capture strategies using two and three predators. The knocker strategy, illustrated in *figure 9*, was an emergent prey-capture strategy using two predators, and the encirclement and entrapment strategies, illustrated in *figures 7* and *8*, were emergent prey-capture strategies using three predators. Comparative to the three-predator strategies, the two-predator knocker strategy proved superior in terms of measures for *prey capture time* and *group stability*. Low group stability and prey-capture time of the encirclement and entrapment strategies was found to be a result of physical interference between the three predators as they collectively approached the prey. The chance of interference was reduced with the simpler knocker strategy, which thus had the advantage of being dominant in the population of individuals, and selected for in the evolutionary process. This is reflected in *figures 4*, *5* and *6*, which illustrate a *higher average fitness*, *group stability* and *prey-capture time* for experiments testing *25* and *50 epochs*. The use of more epochs provides more tests for each individual genotype thus increasing the likelihood that only robust strategies will be propagated. Also, note in *figure 4*, comparing experiments using *5 epochs* and *25 epochs* under the single pool approach, a similarly high fitness was attained. This was due to the encirclement and entrapment strategies emerging with frequency comparable to the knocker strategy, so a high average fitness was scored. Experiments using the single pool approach indicated that evolution had selected for and propagated a simpler and more robust strategy using only two predators that was consistently effective due to minimal interference between these two predators.

Plasticity: Experiments executed under the plasticity approach yielded the knocker strategy as well as the three-predator *role-switcher* strategy. In role switcher the predators developed a dynamic form of behavioral specialization, meaning that each assumed a role based upon its current position and orientation relative to the prey at a given point in time. *Figure 10* illustrates the predators switching roles during the course of the role-switcher strategy, so that flanking predators switched roles whenever the prey turned to escape. This dynamic role assumption served to make the role-switcher strategy more effective, comparative to the single pools approach, in terms of the defined performance measures. This is reflected in *figure 4*, which illustrates a progressively higher yet similar average fitness, comparative to the single pool approach, for *10, 25 and 50 epochs*. Specifically, physical interference between predators was reduced as they collectively approached the prey, given that they assumed specific, though not always complementary behavioral roles.

In the single pool approach, such interference often caused the entrapment and encirclement strategies to fail prematurely thereby making their selection and propagation by the evolutionary process more unlikely.

Multiple Pools: In these experiments the emergence of only one strategy was observed, which was a variation of the *role-switcher* strategy that emerged in the plasticity experiments. Though, where as the plasticity version of role-switcher used a dynamic form of role assumption, the multiple pools version used predators evolved for a specific role. That is, for all epochs tested in these experiments, each of the three predators evolved a specific behavioral role that allowed the team to collectively overcome the physical interference that confounded the teams tested in the single pool and plasticity approaches. In the plasticity experiments the adoption of specialized roles was dependent upon the positions of the predators at a given time. This dynamic form of role assumption still proved problematic in situations where all three predators were in close proximity to the prey when approaching it, as two or more predators typically attempted to assume the same behavioral role, thus causing physical interference, the role-switcher strategy to fail, and allowing the prey to escape. Where as, in the multiple-pools version of role switcher strategy, different predators evolved so as to assume a specific role during their lifetime so each initially behaved differently. That is, one predator always assumed the role of a flanker, one the role of a knocker and another the role of an idle predator, where these behavioral roles complemented each other in formation of the strategy. The effectiveness this strategy in the multiple pools approach, is evident from the higher fitness, group stability and prey-capture time.

5. Conclusions

This paper presented a set of experiments testing three different artificial evolution approaches for the synthesis of cooperative behaviour in a team of simulated mobile robots, operating within a pursuit-evasion domain. The team of robots, termed predators, was given the task of collectively capturing a single robot termed the prey. The performance of emergent prey capture strategies was quantified in terms of *group stability*, *prey-capture time*, and *predator fitness*.

Results presented indicated the multiple pools approach to be superior in terms of these three measures. In the multiple pools approach only one cooperative prey-capture strategy, which was a variation of the plasticity role-switcher strategy, consistently emerged. The superiority of the multiple pools approach was found to be a result of a genetic based behavioral specialization in the multiple pools role-switcher. That is, in the multiple pools role-switcher, specific roles were consistently assumed by the different predators. This improved the effectiveness of multiple pools knocker strategy compared to the single pool and plasticity approaches, in that it reduced interference between predators as they collectively approached the prey. This is especially evident from a comparison with the single pool experiments, where three-predator strategies with no behavioral specialization emerged and performed less well.

Finally, experimental results highlighted that artificial evolution is an effective method for deriving cooperative prey capture strategies using predator teams with no explicit communication, or coordination mechanisms.

References

Baldassarre, G., Nolfi, S., and Parisi, D. 2002. Evolving Mobile Robots Able to Display Collective Behaviors. In *Proceedings of the International Workshop on Self-Organization and Evolution of Social Behaviour*, pages 11--22. Monte Verità, Switzerland. Addison Wesley.

Benda, M., Jagannanthan, V., and Dodhiawalla, R. 1985. An Optimal Cooperation of Knowledge Sources. Technical BCS-G2010-28, Boeing AI Center.

Bull, L., and Holland, O. 1997. Evolutionary Computing in Multi-Agent Environments: Eusociality. In *Proceedings of the Second Annual Conference on Genetic Programming*, pages 347--352. San Francisco, U.S.A. Morgan Kaufmann.

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

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

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

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

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

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

Nishimura, S. I., and Ikegami, T. 1997. Emergence of Collective Strategies in a Prey-Predator Game Model. *Artificial Life*, 3(1): 243--260. MIT Press, Cambridge, USA.

Nolfi S. and Parisi D. 1997. Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, 5(1): 75--98. MIT Press, Cambridge, USA.

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. 1986. Learning internal representations by error propagation. *Parallel Distributed Processing, Volume 1: Foundations*. MIT Press, Cambridge, USA.

Yong, C. H., and Miikkulainen, R. 2001. Cooperative Co-evolution of Multi-Agent Systems. Technical Report AI01-287. Department of Computer Science, University of Texas, USA.