Proceedings of the 2003 IEEE/RSJ
Intl. Conference on Intelligent Robots and Systems
Las Vegas, Nevada · October 20033

Co-evolution of cooperation in a Pursuit Evasion 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 different artificial evolution approaches to the design of cooperative behavior in a team of simulated mobile robots co-evolved against a second team. 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 is a pursuit-evasion game in which a team of three robots termed: *pursuers*, collectively work to immobilize one of the three robots of the other team, termed: *evaders*. Results indicate that the multiple pools approach applied within a competitive co-evolution process yields superior performance comparative to the other approaches. Specifically, the co-evolutionary process allows the multiple pools approach to 'bootstrap' complementary behavioral roles, facilitating the evolution of a stable cooperative pursuit strategy.

1 Introduction

The use of competitive co-evolution to facilitate emergent behavior [Hillis, 1990], [Angeline and Pollack, 1993], [Rosin and Belew, 1997] particularly cooperation, remains a relatively unexplored area of research in the pursuit and evasion domain [Koza, 1991], [Koza, 1992], [Reynolds, 1994] and related predator-prey systems [Cliff and Miller, 1994; 1995; 1996], [Nishimura and Ikegami, 1997] studied in an evolutionary robotics context [Bullock, 1995].

Various approaches to the evolution of behavior within the pursuit-evasion domain have been studied within a competitive co-evolutionary context. For example, Koza [1991, 1992] applied genetic programming techniques to the co-evolution of pursuer-evader behaviors in a two-player pursuit-evasion game. In similar experiments, Reynolds [1994] observed the evolution of increasingly sophisticated strategies in the competitive co-evolution of a pursuer and evader. Miller and Cliff [1995; 1996] realised the co-evolution of pursuit-evasion strategies in evolutionary robotics, where simulated robots co-evolved vision morphologies as a means of improving pursuit and evasion strategies. Floreano and Nolfi [1997a, 1997b] evaluated a competitive co-evolutionary predator-prey scenario within evolutionary robotics experiments using two mobile robots. This co-evolution scenario was compared with single agent evolution. A comparatively faster evolution of more diverse behavioral strategies was observed with competitive co-evolution. Floreano *et al.* [1998] extended this research with a physical implementation of this predator-prey co-evolution using *Khepera* robots. With notable exceptions such as [Iba, 1996], [Haynes and Sen, 1997], [Yong and Miikkulainen, 2001], [Stone and Veloso, 1998] few researchers have investigated emergent cooperative behavior in a competitive co-evolutionary context, where two teams co-evolve cooperative behavior that ensures the survival of each team.

This paper describes a comparison of three artificial evolution approaches for the synthesis of cooperative behaviour evaluated within a team of simulated *Khepera* robots [Mondada *et al.* 1993], co-evolved against a second team. For each artificial evolution approach, a team of three pursuers implementing initially random behavior is competitively co-evolved with a team of three evaders implementing an obstacle avoidance behavior. The pursuers and evaders are functionally different in terms of sensor and movement capabilities. The team of evaders is able to move 40 percent faster than the team of pursuers. Pursuers are rewarded fitness proportional to how much they are able to slow down the evader team on average, and the evader team is rewarded fitness proportional to how fast it is able to move and avoid obstacles on average. The collective task was for the pursuers to immobilize one or more of the evaders. Previous experiments demonstrated that at least two pursuers are required in order for an evader with superior speed to be immobilized [Nitschke and Nolfi, 2002]. Only cooperative behavior in the team of pursuers is evolved, since that in the evader team there is no incentive for evaders to cooperate, given that evasion is possible as an individual.

The first approach used for the evolution of cooperative behaviour is termed: *single pool*, in which pursuer and evader teams are specified with an identical genotype, meaning that the corresponding phenotype is also the same. The second approach, termed: *plasticity*, uses pursuer and evader teams also specified with an identical genotype, though the corresponding phenotype implements a recurrent neural network controller allowing adaptive behavior during a teams lifetime. The third approach is termed: *multiple pools*, where pursuer and evader teams are specified using a different genotype, and thus the phenotype at the beginning of each teams lifetime is different. For the team of pursuers, the three approaches are evaluated in terms of fitness scored (averaged for the team), and the time period for which an evader is immobilized. For the team of evaders, the three approaches are evaluated in terms of average evader team fitness scored, and the average time for which the evader team is not immobilized. It is hypothesized that the multiple pools approach would yield a superior performance as a result of 'bootstrapping' of behavioural specialisation by the co-evolutionary process. Experimental results support this hypothesis, where the emergence of specialised behavioural roles proved necessary for the formation of a stable cooperative pursuit strategy.

2 Artificial Evolution Approaches

For the team of pursuers the three artificial evolution approaches were tested and evaluated against co-evolved evader strategies for the task of the pursuer team cooperatively immobilizing an evader. Where as, for the evader team the three approaches were tested and evaluated against co-evolved pursuer strategies in terms of the time for which an evader (averaged for the team) was not immobilized.

Comparison of Approaches

To execute the co-evolutionary scheme under which the three artificial evolution approaches were tested, an extension of realistic Khepera robot simulation software was implemented [Floreano and Nolfi, 1997a], [Floreano and Nolfi, 1997b]. The simulator was based on real sensory values sampled from two functionally different Khepera robots used in previous predator-prey experiments [Miglino *et al.* 1996].

Single Pool Approach: As illustrated in *figure 1* this approach generates and tests *3* copies of a single genotype, meaning that the pursuer and evader teams are homogenous. In this approach there is no plasticity so individual pursuers and evaders cannot adapt during their lifetimes. The fitness assigned to the pursuer and evader teams is the simply the fitness calculated for the genotypes that specify the pursuer and evader teams. 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 2* this approach generates and tests *3* copies of a single genotype, so that as with the *single pool* approach, the pursuer and evader teams are homogenous. The difference is that individual phenotypes, representing the behaviors of individual pursuers and evaders are able to adapt during their lifetime as a result of a recurrent neural network learning process. This learning process is affected by both genetic and environmental factors. Thus, pursuer and evader phenotypes adapt to environmental influences throughout their respective lifetimes, which affects fitness calculated for team genotypes, which in turn influences the selection process in successive generations. The advantage of the plasticity approach is that it allows for specialization of behavior by individual team members without the need to estimate fitness contribution of different team members to the team as a whole. For these co-evolution experiments, both pursuers and evaders implemented a two-layer neural network controller of sigmoid neurons with re-current connections at the output layer. This type of controller is detailed in previous co-evolution research [Floreano and Mondada, 1994], though its structure as relates to the experiments of this paper is briefly described in the following section.

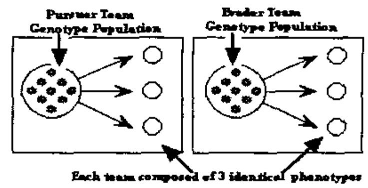

Figure 1. Single Pool: Pursuer and evader teams correspond to a genotype selected from the respective genotype populations and copied *3* times (to represent the 3 pursuers and evaders comprising each team).

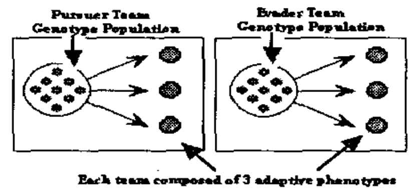

Figure 2. Plasticity: Same as Single Pool, though phenotypes implement a recurrent neural network controller for adaptability during a pursuers or evaders lifetime.

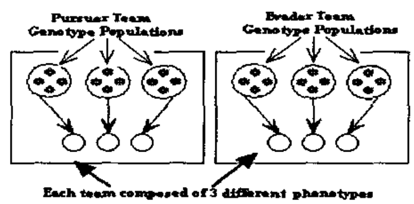

Figure 3. Multiple Pools: Pursuer and evader teams correspond to 3 genotypes selected from *3* separate pools of genotypes.

Multiple Pools Approach: As illustrated in *figure 3* this approach selects a single genotype from each of the three populations of genotypes. Each genotype was then decoded into 3 separate phenotypes representing the pursuer and evader teams. In each generation, every individual genotype in a population is tested against *n* other genotypes, randomly selected from the two other populations of genotypes, where each individual genotype was also tested against the best individuals from the most recent 10 generations. The advantage of the multiple pools approach is that selection operates within each genotype population and each pursuer and evader corresponds to genotypes from different populations, so behavioral specialization in the team is encouraged. The disadvantage is that the assignment of individual fitness is an approximation. Specifically, an equal fitness score is assigned to each of the genotypes, as a means of deriving the contribution of each pursuer or evader to the performance of the team as a whole. Thus the estimation of the fitness for pursuer and evader teams constitutes a problem that may prevent the selection of the best individuals across successive generations.

Evaluation of Approaches

For both the *single pool* and *plasticity* approaches two genotypes, one specifying each the pursuer team and evader teams are used. In both cases pursuers and evaders are clones, so evaluation of team performance is not problematic, as a single value is assigned as the fitness of each team. In contrast, the *multiple pools* approach uses *3* different populations of genotypes, so each genotype must be assigned a 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.

Co-evolutionary Teams

For all experiments two populations, each initially consisting of 100 randomly generated genotypes, were co-evolved where each individual genotype was tested against the best individuals from the most recent 10 generations. This type of co-evolutionary scheme was adapted from that used by Sims [1994], Reynolds [1994], Cliff and Miller [1995], and Floreano *et al.* [1998] and was selected in order to improve co-evolutionary stability.

Each pursuer and evader is a simulated Khepera mobile robot [Mondada *et al.* 1993]. As illustrated in *figure 4*, the robots used as pursuers were equipped with 8 infrared proximity sensors, and 8 light sensors positioned on the periphery of the Khepera. The robots used as evaders were equipped with 8 infrared proximity sensors, as well as a light on its top. This light could be detected by pursuer light sensors and was used so each pursuer could distinguish fellow pursuers from evaders.

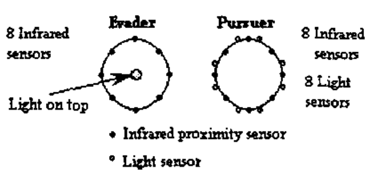

Figure 4. Each pursuer and evader is a simulated *Khepera* robot. Pursuers have 8 infrared proximity sensors, and 8 light sensors. Evaders have 8 infrared proximity sensors, and a light so that pursuers can distinguish them from fellow pursuers and obstacles.

Figure 5 (right-side) illustrates the neural network controller implemented for pursuer and evader teams. This controller is adapted from that described by Floreano *et al.* [1998] and Nolfi and Floreano [1999]. In the case of the pursuers, the input layer consisted of 16 sensory units that encoded the activation level of 8 infrared proximity sensors and 8 light sensors. These 16 input units were connected to 4 output units. As *figure 5* illustrates 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, illustrated in *figure 5* (left-side). 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], where only the connection weights were evolved.

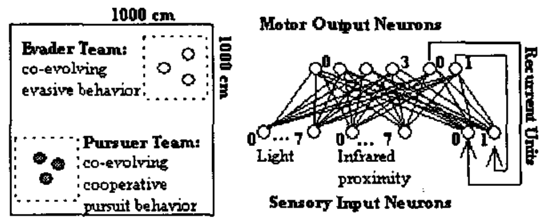

Figure 5. Details of the co-evolutionary experimental setup – A team of 3 pursuers and a team of 3 evaders are placed into a 1000cm x 1000cm arena (left-side). Cooperative pursuit strategies are evolved for the pursuers. A two-layered neural network controller comprising sigmoid neurons is implemented for each pursuer and evader (right-side).

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. The activation values of these two additional output units were copied back into an additional two input units.

In the case of the evaders, a neural network controller connecting 8 sensory input units (representing 8 infrared proximity sensors) to 4 motor output units was trained for an obstacle avoidance behavior before being placed in a co-evolutionary run. Given the simple nature of the neural controllers, direct genetic coding of connection weights was used. In the case of the pursuers genotype length was set to 24 genes, where each gene consisted of 5 bits. That is, 16 genes represented the 8 infrared proximity sensors, and the 8 light sensors, another 4 genes represented the motor output units, and an additional 4 recurrent units used in the plasticity experiments. The 5 bits of each gene encoded connection weights, where the first bit determined the sign of the connection weight and the remaining four bits its strength. In the case of the evaders, genotype length was set to 16 genes, where 8 genes represented the 8 infrared proximity sensors, 4 genes for the motor output units, and an additional 4 genes for the recurrent units used in the plasticity experiments.

Each generation, genotypes were ranked by fitness and the 20 genotypes that accumulated the highest fitness were reproduced, via being copied 5 times in order to keep the population size constant. One-point crossover was applied on randomly paired genotypes with a 0.6 probability and mutation, done via flipping bits, was applied to each bit with a 0.05 probability.

The fitness function for the pursuers rewarded the team based upon how much the evader team was slowed during its lifetime. Hence the pursuers attempted to maximize 'capture time', which was the time for which one or more evaders were immobilized. The fitness function of the evaders rewarded the evaders based on their average speed and obstacles avoided during their lifetime. Hence the evader team attempted to maximize their speed of movement before being immobilized.

3 Results

The co-evolutionary process tested each of the three artificial evolution approaches, where the pursuers initially did not implement any cooperative pursuit strategy. Pursuit strategies were co-evolved with evaders implementing obstacle avoidance behaviors. Ten experimental replications of each artificial evolution approach were made. *Figure 6* illustrates the average team fitness attained for pursuer and evader teams using the single pool, plasticity and multiple pools approaches. *Figure 7* presents average capture time and average *free time* attained for pursuer and evader teams respectively using each of the three evolutionary approaches. *Capture time* refers to a time interval $t_0 .. t_i$ when an evader is immobilized, and *free time* refers to the complementary time interval.

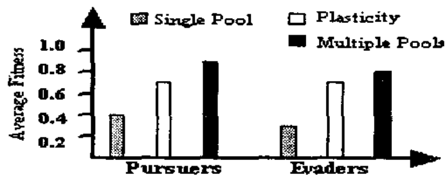

Figure 6. Average fitness for populations representing pursuer and evader teams at the end of the co-evolutionary process (500 generations).

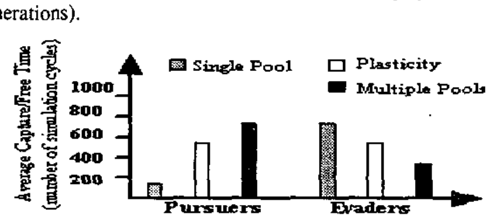

Figure 7. Average capture and free time for populations representing pursuer and evader teams respectively, at the end of the co-evolutionary process (500 generations).

4 Evolved Behavior, Analysis and Discussion

In this section emergent cooperative pursuit strategies observed within the co-evolutionary scheme, using each of the three artificial evolution approaches are discussed. The discussion and analysis is from a behavioral perspective, as fitness comparisons between pursuer and evader teams only illustrate progress and counter progress of pursuit and evasion strategies but do not highlight if evolutionary time corresponds to 'true' progress [Cliff and Miller, 1995] given that the fitness landscape of both teams are continuously changing due to the Red Queen affect [van Valen, 1973].

Single Pool: Given that the pursuers begin with a random behavior, the evaders initially performed very well, though a set of counter-phase oscillations soon emerged in the fitness scores of the pursuer and evader teams. This counter-phase oscillation is supported by other co-evolutionary research [Sims, 1994], [Floreano *et al.* 1998], though neither the pursuers nor evaders maintained dominance throughout the co-evolutionary process. The pursuer team evolved two cooperative pursuit strategies each using three pursuers, termed: *entrapment* and *encirclement*. As illustrated in *figure 8* (left-side), in the encirclement strategy, three pursuers in close proximity to an evader, encircle it, moving in the same direction for some period of time. This caused the evader to spin on its current position as it tried to escape the circle. After approximately 200 generations of the co-evolutionary process, the evaders were able to evolve counter-active evasion strategies, rendering the encirclement strategy less successful. These evasion strategies included an evader closely following a wall or moving slower across the environment, so that it had sufficient time to detect and avoid pursuers.

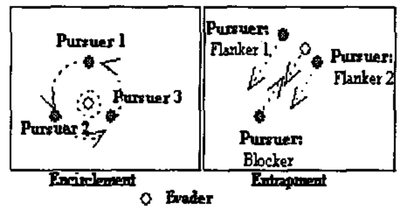

Figure 8. The cooperative encirclement (left-side) and entrapment (right-side) pursuit strategies; each used three pursuers, though neither strategy was successful at immobilizing an evader.

Figure 8 (right-side) also illustrates the entrapment strategy, using three pursuers, where one pursuer moved to flank each side of the evader, while a third, termed: *blocker*, moved so as to approach the evader from the front, in order to trap it in a triangular formation. When the evader moved to escape, the flanking predators moved also, and turning so as to force the evader in a specific direction. The blocker then moved around in order to affront the evader again. This system of entrapment, movement, and entrapment continued several times before evasion was possible. While the entrapment pursuit strategy proved successful in the first 200 generations of the co-evolutionary process, the evaders were able to evolve counter-active evasion strategies similar to those described for encirclement in order to render the entrapment strategy less effective. After 500 generations of the co-evolutionary process, *entrapment* and *encirclement* strategies were only able to immobilize an evader in *20 percent* of single pool experiments. This is reflected in the average capture time and complementary free time presented for pursuer and evader teams, respectively in *figure 7*.

Plasticity: As with experiments run for the single pool approach, the evaders initially scored a high fitness before a similar pattern of counter-phase fitness oscillations emerged as a result of pursuers evolving effective cooperative pursuit behavior and evaders evolving behaviors to counter-act capture. As with the single pool experiments neither the pursuers nor evaders maintained dominance in the co-evolutionary process, though one effective cooperative pursuit strategy emerged. This strategy, termed: *role switcher* was similar to the entrapment strategy observed in the single pool experiments. The *role-switcher* strategy used three pursuers, where one pursuer, termed: a *flanker*, moved to each side of the prey, while a third pursuer, termed: a *blocker*, moved around the flanking predators, to approach the front, in order to immobilize the evader in a triangular formation. The three pursuers then encircled the evader causing it to rotate on its current position. The key difference noted in the role-switcher strategy, was that behavioral specialization evolved in the pursuer team. Each pursuer either assumed the behavioral role of a *flanker* or a *blocker*, and pursuers switched between these roles allowing pursuers to quickly adapt to evaders strategies whilst maintaining the strategy.

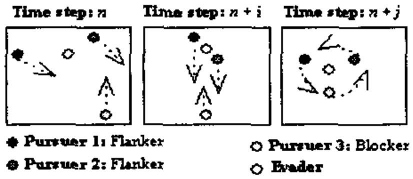

Figure 9. The multiple pools version of the role-switcher pursuit strategy, using three pursuers, emerged in the co-evolutionary process at approximately generation 400.

This dynamic adoption and switching of roles, afforded the pursuer team flexibility in forming and maintaining the entrapment strategy. This is reflected in *figure 6*, which illustrates a higher average fitness for pursuer teams using the plasticity approach, comparative to the single pool approach, when the co-evolutionary process was ended. Though the switching of behavioral roles during the strategy also inhibited the coordination of the three pursuers, meaning that it was difficult for the team to maintain the strategy, and thus immobilize an evader for an extended period. The evaders exploited the lack of coordination between the three pursuers, and were able to evolve a strategy of quick turns when being flanked. This evasion strategy often prevented all three pursuers from being able to maintain close proximity to an evader. As illustrated in *figure 7*, the role-switcher strategy at the end of the co-evolutionary process was able to immobilize evaders (on average) in *50 percent* of plasticity experiments.

Multiple Pools: As with the plasticity experiments the *role switcher* strategy was the only cooperative pursuit strategy that emerged at the end of the co-evolutionary process. *Figure 9* illustrates the multiple pools version of the role-switcher strategy, and its formation in three distinct stages. A specific difference was noted in the multiple pools version of the role-switcher strategy. Namely that different pursuers adopted different behavioral roles from the beginning of their lifetimes. This allowed the pursuers to avoid the interference problem that confounded pursuer teams using the role-switcher strategy under the plasticity approach. Specifically, two pursuers always assumed the role of flankers, while a third always assumed the role of a blocker. In the first 200 generations of the co-evolutionary process, the three pursuers moved about the environment in search of an evader and attempted to capture an evader via remaining in close proximity to each other. Though, the evaders soon developed a counter-evasive strategy where they rapidly and closely followed the walls of the environment often causing pursuers to collide with the walls, given that evader speed of movement was faster. Similar behavior to this has also been observed in the predator-prey experiments of Floreano *et al.* [1998] that used two robots.

Also, the wall following behavior made it difficult for two pursuers to flank each side of an evader. To counter-act this behavior two of the pursuers also developed a wall following behavior while a third maintained the role of an idle pursuer in one corner of the environment. Thus, evaders following a wall were often trapped by the pursuer team in a corner. After approximately 300 generations the evaders adapted to this pursuit strategy that exploited corners, and evolved the next stable evasion strategy. This was for the evaders to move randomly about the environment, though only at approximately 75 percent of full speed. As evaders moving at full speed often detected pursuers too late (due to the limited range of infrared sensors) to avoid being flanked, and subsequently immobilized. After approximately 400 generations the next stable pursuit strategy emerged, where two pursuers maintained the behavioral roles of flankers, searching the environment as a pair, while the third pursuer maintained the behavioral role of a blocker, waiting idly in one position. The function of the blocker in forming the role-switcher strategy was either to 'chase' an evader towards the two flankers, or to move in order to capture an evader in a triangular formation as the two other flanking pursuers forced the evader towards its own position. As reflected in *figure 7*, multiple pools role-switcher was successful at immobilizing evaders (on average) in *70 percent* of experimental replications. Adoption of behavioral roles was maintained throughout the lifetime of the pursuers, which served to aid in the formation of a stable pursuit strategy. *Figure 6* presents the benefit of the role-switcher strategy at the end of the co-evolutionary process. Specifically, in the comparatively higher fitness of pursuer and evader teams implementing the multiple pools approach.

5 Conclusions

This paper presented a set of experiments testing three different artificial evolution approaches for the synthesis of cooperative pursuit strategies within a team of simulated mobile robots, competitively co-evolved with a second team of robots. Results indicated that the multiple pools approach applied within a competitive co-evolution process yielded superior performance comparative to the single pools and plasticity approaches. In competitive co-evolution, the multiple pools approach implemented within a co-evolutionary context allowed the exploitation of bootstrapping of complementary behavioral roles, facilitating the evolution of a stable cooperative pursuit strategy. Emergent pursuit strategies observed using the other two approaches proved less effective, due physical interference that occurred between pursuers as they collectively approached an evader in attempted formation of a pursuit strategy. Behavioral specialization, as observed under the multiple pools approach, alleviated the problem of physical interference resulting from a lack of coordination, given that the three pursuers maintained three complementary behavioral roles allowing them to form a stable pursuit strategy that effectively immobilized an evader.

A comparison with other research investigating emergent cooperation within a co-evolutionary context in the pursuit-evasion domain is difficult given the limited literature on co-evolving teams within physically realistic environments. That is, with notable exceptions such as the two robot predator-prey co-evolution experiments of Floreano *et al* [1998], the co-evolution of robot behaviors within a pursuit domain has typically employed simulated grid-world environments [Iba, 1996], [Haynes and Sen, 1997] and [Yong and Miikkulainen, 2001]. Also, there is relatively little literature describing the co-evolution of robot teams for the purpose evolving cooperative behavior within the teams themselves.

Though the robot teams in this research were simulated, the robot simulator used a continuous domain and the simulation incorporated noise in sensory data, namely confused infrared sensor readings resulting from two or more robots being in close proximity to each other. This noisy sensor data was a key reason for interference occurring between multiple pursuers as they collectively approached an evader. 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. Finally, experimental results highlighted that artificial evolution applied within a competitive co-evolution context is an effective method for the derivation of cooperative pursuit strategies in a team of robots with no explicit communication, or coordination mechanisms. The advantage of co-evolution in evolving more complex behaviors is supported by other research. For example, the evolution of predators against the fixed behavior of a prey in the case of Nolfi and Floreano [1998] did not attain the same performance levels as in the co-evolutionary case.

References

[Angeline and Pollack, 1993] Angeline, P. J. and Pollack, J. B. Competitive environments evolve better solutions for complex tasks. In Proceedings of the Fifth International Conference on Genetic Algorithms, pages 264--270. Morgan Kaufmann, San Mateo, USA.

[Bullock, 1995] Bullock, S. Co-evolutionary design: Implications for evolutionary robotics. CSRP 384, University of Sussex, United Kingdom.

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

[Cliff and Miller, 1994] Cliff, D., and Miller, G. F. Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In, *From Animals to Animats III: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 411--420. Bradford Books, Cambridge, USA.

[Cliff and Miller, 1995] Cliff, D., and Miller, G. F. Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In, *Proceedings of the Third European Conference on Artificial Life*, pages 200--218. Springer-Verlag, Berlin.

[Cliff and Miller, 1996] Cliff, D., and Miller, G. F. Co-evolution of Pursuit and Evasion II: Simulation Methods and Results. In, *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 506--514. MIT Press, Cambridge, USA.

[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. MIT Press, Cambridge, USA.

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

[Floreano and Nolfi, 1997a] Floreano, D., and Nolfi, S. Adaptive behavior in competing co-evolving species. In, *Proceedings of the Fourth European Conference on Artificial Life*, pages 378--387. MIT Press, Cambridge, USA.

[Floreano and Nolfi, 1997b] Floreano, D., and Nolfi, S. God Save the Red Queen! Competition in Co-evolving Robotics. In, *Proceedings of the Second International Conference on Genetic Programming*, pages 398--406. Morgan Kaufmann, San Mateo, USA.

[Floreano and Mondada, 1998] Floreano, D., and Mondada, F. Hardware solutions for Evolutionary Robotics. In, *Proceedings of the First European Workshop on Evolutionary Robotics*. Springer-Verlag, Berlin.

[Haynes and Sen, 1997] Haynes, T., and Sen, S. Co-adaptation in a Team. International Journal of Computational Intelligence and Organizations. Vol. 1(4): 1-20. Lawrence Erlbaum Associates, New Jersey, USA.

[Hillis, 1990] Hillis, W. Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, vol. 42(1): 228-234. Elsevier Science, Amsterdam, Netherlands.

[Iba, 1996] Iba, H. Emergent cooperation for multiple agents using genetic programming. *Parallel Problem Solving from Nature*. Springer-Verlag, Berlin.

[Koza, 1991] Koza, J. R. Evolution and co-evolution of computer programs to control independently acting agents. In, *Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 366--375. MIT Press, Cambridge, USA.

[Koza, 1992] Koza, J. R. Genetic Programming: On the programming of computers by means of natural selection. MIT Press, Cambridge, USA.

[Miglino *et al.* 1996] Miglino, O., Lund, H., and Nolfi, S. Evolving Mobile Robots in Simulated and Real Environments. Artificial Life, vol. 2(1): 417-434. MIT Press, Cambridge, USA.

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

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

[Nitschke and Nolfi, 2002] Nitschke, G., and Nolfi, S. Emergence of Cooperation in a Pursuit-Evasion Scenario. *Technical Report CNR02-26*. Institute of Cognitive Science and Technologies, C.N.R. Rome, Italy.

[Nolfi and Floreano, 1998] Nolfi, S., and Floreano, D. Co-evolving predator and prey robots: Do 'arms races' arise in artificial evolution? Artificial Life, vol. 1(1): 1-10. MIT Press, Cambridge, USA.

[Reynolds, 1994] Reynolds, C. W. Competition, Co-evolution and the Game of Tag. In, *Proceedings of the Fourth International Workshop on Artificial Life*, pages 59--69. MIT Press, Cambridge, USA.

[Rosin and Belew, 1997] Rosin, C, and Belew, R. New methods for competitive co-evolution. *Evolutionary Computation*, vol. 5(1): 1--29. MIT Press, Cambridge, USA

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

[Sims, 1994] Sims, K. Evolving 3D Morphology and Behavior by Competition. In, *Proceedings of the Fourth International Workshop on Artificial Evolution*, pages 28--39. MIT Press, Cambridge, USA.

[Stone and Veloso, 1998] Stone, P., and Veloso, M. Towards collaborative and adversarial learning: A case study in robotic soccer. *International Journal of Human-Computer Systems*, vol. 48(1): 83-104. Elsevier Science, Amsterdam.

[van Valen, 1973] van Valen, L. A new evolutionary law. *Evolution Theory*, vol. 1(1): 1--30. MIT Press, Cambridge, MA.