Constructing Competitive and Cooperative Agent Behavior Using Coevolution

Aditya Rawal, Padmini Rajagopalan, and Risto Miikkulainen, *Senior Member, IEEE*

Abstract— In nature, multiple agents in teams collaborate and compete with one another at the same time. Replicating such agent interactions in games can make for realistic opponent teams. Yet cooperation and competition have mostly been studied separately so far. This paper focuses on simultaneous cooperative and competitive coevolution in a complex predator-prey domain. Multi-Agent ESP [23] architecture is first used to evolve neural networks to control predator and prey agents, but such a naive combination of otherwise successful architectures turns out not to sustain an arms race. An extended architecture consisting of multiple cooperating neural networks within each agent is therefore introduced. This architecture successfully results in hierarchical cooperation and competition in teams of prey and predators: In sustained coevolution, high-level pursuit-evasion behaviors emerge. In this manner, coevolution of neural networks is shown to scale up to an arms race of multiple competing and cooperating agents, more closely modeling coevolution of complex behavior in nature.

I. INTRODUCTION

A major goal in Game AI is to develop intelligent behaviors that appear natural and believable. Extending complex social interactions to autonomous game playing agents is particularly challenging. In nature, cooperative behavior leads to team strategies that compensate for individual limitations and aid in problem solving. For instance, herds and packs allow animals to perform even challenging tasks of attacking larger prey [3][14]. Similarly, competition between opponents may give rise to an arms race that promotes learning of successive strategies targeting each other's weaknesses. For instance, the *Heliconius* butterfly and the passionflower plant compete in nature to evolve new traits in this manner[1]. At the highest level of natural ecosystems, social structures consist of hierarchical layers of competition and cooperation among individuals.

In this paper, computational coevolution is used to study how such competitive and cooperative behaviors can be constructed for teams of game playing agents. Coevolution refers to simultaneous evolution of two or more distinct species with coupled fitness landscapes. In case of competitive coevolution, the fitness of an individual is based on direct competition with an individual from another population. In cooperative coevolution, individuals from different populations cooperate to solve the problem together. Although cooperative and competitive coevolution co-exist in nature, and each affects the other, there have been few previous attempts to study both of them simultaneously.

Aditya Rawal, Padmini Rajagopalan, and Risto Miikkulainen are with the Department of Computer Sciences, University of Texas at Austin, Austin, TX 78712 USA (email: {aditya, padmini, risto}@cs.utexas.edu).

Interactions among multiple autonomous agents are critical in many games. For example, both cooperative and competitive behaviors are required in robotic soccer as well as in Unreal Tournament's™ *capture the flag* [7]. Such games are typically dynamically changing and open-ended, and such behaviors are difficult to achieve using the traditional scripted agent approach.

In contrast, neuroevolution has been used successfully for agent control in several dynamic, open-ended games such as simulated robot soccer [21], robotic battle [18] and Ms. Pac-Man [2]. Neuroevolution techniques are therefore used as the starting point for evolving team behaviors in this paper. The Multi-Agent ESP architecture [23] is extended to competitively and cooperatively coevolve teams of predators and prey. The predator-prey domain was selected as the experimental platform because it is a good surrogate for agent control problems in team games, and because it is easy to simulate with quantifiable results.

Two experiments of increasing complexity are performed: First, a team of predators is coevolved with a single prey. Second, using the insights from the first experiment a team of predators is coevolved with a team of prey. The main contribution of this paper is to develop methods that sustain both competitive and cooperative coevolution in a complex environment. Placing prey and predators against each other gives rise to an arms race that leads to evolution of more and more complex pursuit and evasion strategies [12].

The paper is organized as follows. Section 2 describes prior work on competitive and cooperative coevolution. Section 3 explains the predator-prey problem domain and Section 4 puts forth hypotheses about sustenance of arms race and emergence of behaviors. Section 5 includes detailed descriptions of the experiments and Section 6 gives the results of these experiments. Section 7 analyses the results and Section 8 presents concluding remarks.

II. BACKGROUND AND RELATED WORK

Coevolution is defined as the simultaneous evolution of two different populations whose fitnesses are measured based on their interactions with each other [11]. In competitive coevolution, the two populations have opposing interests and the success of one population depends on the failure of the other. An arms race emerges as the coevolution proceeds; each population evolves a little more at every step so as to defeat the other. Competitive coevolution is usually used to simulate the behavior of competing forces in nature, such as predators and prey. But it can also be used as a method of improving the fitness of a single population by duplicating

it and coevolving the two populations to outdo each other. After it was first described by Hillis [6] as a host-parasite problem, competitive coevolution has been studied extensively by many researchers [5][10][11][13][17]. Mitchell [10], in particular, compares competitive coevolution and plain evolutionary techniques. Competitive coevolution has many advantages: it does not get stuck at local optima as often, it discovers higher-level strategies, it requires sparser training, and it preserves diversity over longer periods.

The arms race can be hard to sustain for extended periods because of over-specialization, red queen dynamics and loss of gradients [11][12]. Over-specialization happens when the two competing populations learn to defeat each other easily but cannot generalize to new environments. Red queen dynamics refers to stagnation caused by oscillation of the two populations among a set of states none of which is an improvement over another. Loss of gradients happens when evolution cannot proceed because all the population members are equally good at losing to or winning over members of the opposing population. Many attempts have been made to overcome these obstacles. For example, in spatial coevolution the hosts and parasites are distributed on a grid, and each of them interacts only with the hosts/parasites that are located close to it on the grid [10][11][22]. Another way to avoid stagnation is to preserve the good behaviors of previous generations in a Hall of Fame so that diversity is not lost [12][16][17]. Resource sharing, also known as competitive fitness sharing [6][13][17], is another tactic where a population member is considered fitter if it defeats an opposing population member that few others have defeated.

In cooperative coevolution, different agents have to evolve to cooperate to perform a task. They share all the rewards and punishments of their individual actions equally. Cooperative coevolution is easier to achieve if the agents are components of the same system in which case they can learn different roles [23]. For example, in the Enforced SubPopulations (ESP) architecture [4], neurons selected from different subpopulations are required to form a neural network whose fitness is then shared equally among them. Such an approach avoids competing conventions among the component neurons and limits their individual search space. This makes neuroevolution faster and more efficient, and helps establish cooperation between the components. Similarly, Potter and De Jong [15] describe an architecture for evolving subcomponents as a collection of cooperating species.

Simultaneous cooperative and competitive coevolution was implemented in an experiment with soccer playing robots by Uchibe and Asada [20]. Their work is significantly different from that in this paper in many ways. While Uchibe and Asada coevolve two cooperating players against one competing player, the experiments in this paper simulate an environment comprising teams of prey and predators. The multiple levels of cooperation and competition in these experiments are closer to the complexity of such processes in nature. Furthermore, Uchibe and Asada use Genetic Programming to evolve decision trees for their agents but are not able to

sustain an arms race. In contrast, Neuroevolution is used in this paper, particularly because it has supported an arms race successfully in games [18].

III. THE PREDATOR-PREY DOMAIN

In the predator-prey domain, predators chase and try to capture prey in a simulated environment. The domain is a special case of the well-known pursuit-evasion problem in mathematics and computer science. Pursuit-evasion problems are common in game agents. They pose a formidable challenge for learning algorithms because of their dynamically changing environments. The predator-prey domain is open-ended and requires continuous discovery of good behaviors on the part of both the predators and the prey. The agents in the simulation should be able to adapt to this changing environment using a supervised training algorithm because the outcome of any single action of the agent is typically not known. The predator-prey domain can easily be extended to include multiple agents, teams of agents having similar goals and other more complex scenarios. It is thus a good problem to study both competitive and cooperative coevolution in games [9].
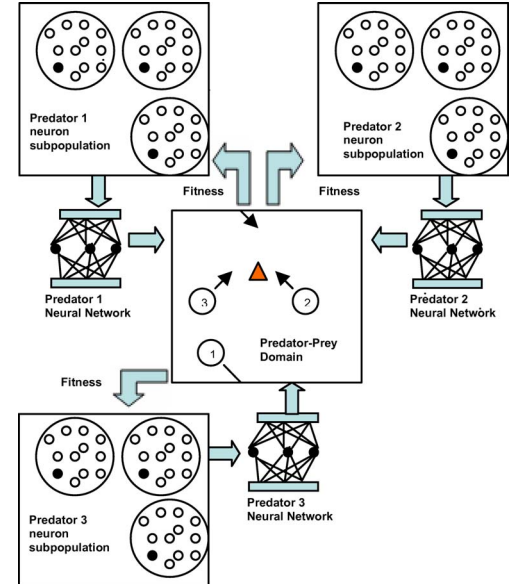

Fig. 1. Multi-Agent ESP architecture; numbered circles are predators and the triangle is the prey. Each Predator-Prey agent is controlled by a single neural network consisting of several hidden neurons. Each neuron in the hidden layer is drawn from its own subpopulation. Agent neural networks are evaluated in the domain at the same time, and fitness for the team is passed back to the participating neurons [23].

In one such study, Yong and Miikkulainen [23] extended the ESP neuroevolution method to the level of networks (see Figure 1). In their Multi-Agent ESP architecture, three neural networks were evolved in parallel to control three predators for the prey-capture task. These predators had to learn to cooperate to capture a single non-evolving fixed-behavior prey that none of them could catch on their own. Yong and Miikkulainen showed that this approach is more efficient than evolving a single central controller for all

predators. Also, they found that cooperation is most efficient through rolebased responses to the environment (i.e through stigmergy), rather than direct communication between the agents.

In this paper, their work is extended to the scenario where a team of prey is coevolved together with a team of predators. The world in this simulation is a discrete toroidal environment with three evolving predators which try to catch two evolving prey. The predators are aware of prey positions and the prey are aware of predator positions. However, there is no direct communication within a prey or predator team. The prey and predators move at the same speed and so, in the toroidal world, the predators cannot catch the prey if they use a greedy strategy of just following the prey around. A time limit is placed on the simulation to make sure that the predators do not keep moving at random and capture the prey by accident. Instead, they need to surround the prey from different sides so that they do not have anywhere to escape before they can catch them. This behavior requires high-level cooperative strategies to evolve in the predators. Similarly, the prey can collaborate to evade the predators more effectively. In addition to cooperative behavior, predators and prey compete, and continually evolve to exploit the weaknesses of each other.

IV. Hypotheses

The main goal of these experiments is to sustain simultaneous cooperative and competitive coevolution in a complex environment that includes teams of predators and prey. How different strategies emerge in the prey and predators as they interact with one another in this environment will be characterized in detail. Two hypotheses will be tested:

1) *Can an arms race be sustained in an environment with simultaneous cooperative and competitive coevolution?*
   Although it can be hard to sustain coevolution as discussed in the section 2, the predators and prey should not get stuck at local optima but continue to learn increasingly complex strategies to counter each other. The first hypothesis is that simultaneous cooperative and competitive coevolution can be sustained without stagnation.
2) *Can cooperative and competitive behaviors like baiting and herding emerge in predators and prey in such conditions?*
   As mentioned above, the predators need to cooperate and surround a prey before catching it. If there are multiple prey in the environment, the predators should evolve to catch all of them. An interesting question is whether the predators can group or herd the prey together to make them all occupy the same position before catching them simultaneously. The prey can learn baiting strategies to avoid being captured. Thus the second hypothesis is that the predators and prey will learn such competitive and cooperative behaviors and formations while adapting to each others' strengths and weaknesses.

V. Experiments

This section describes the experiments conducted to observe the interaction between predators and prey in the pursuit-evasion problem. There is cooperative coevolution between the predators as they learn to work as a team to surround the prey and capture them. The prey also cooperate as a team and their goal is to evade the predators. The predator and prey compete against each other. The environment is a 100x100 toroid without any obstacles. The prey and predators can move in four directions: east, west, north and south. They move one step at a time, and all the agents in the world take a step simultaneously. The predators are said to have caught a prey if one of them moves into the same location in the world as the prey.

Multi-Agent ESP is used to evolve both the prey and the predators with the following parameter settings: Each neural network is feedforward with a single layer of 10 hidden neurons and sigmoidal activation functions. Each subpopulation consists of 100 neurons; each neuron (or a chromosome) is a concatenation of real-valued numbers representing full input and output connections of one hidden unit. During each evolutionary generation, 1,000 trials are run wherein the neurons are randomly chosen (with replacement) from their subpopulations to form a neural network. In each trial, the team is evaluated six times. The prey and predators start at random locations each time so that neither of them has an advantage over the other. The fitnesses over the six evaluations are averaged, and assigned to all the neurons that constituted the network. After the trials, the top 25% of neurons within each subpopulation are recombined using one-point crossover. The offspring replace the bottom 50% of the neurons in the corresponding subpopulation, and they are then mutated with a probability of 0.4 on one randomly-chosen weight on each chromosome, by adding a Cauchy-distributed random value to it. Small changes to these parameters lead to similar results.

Throughout the study, there are three predators that form a team that has to cooperate to catch the prey. Each predator has as its inputs the x, y offset distances of all the prey from that predator. Similarly, each prey has as its inputs the x, y offset distances of all the predators from that prey. The output neurons represent different actions that a prey or predator agent can take. Each prey has only four possible output actions in each time step (move east, west, north, or south) and the predators have five (move east, west, north, south, or idle). To evolve blocking strategies in predators, the idle action is often important.

The predator fitness is higher if a predator team catches both prey together rather than one by one. Such a fitness function encourages herding of prey. The average distance of the predators from the prey is also included in the fitness function to help score predator teams that do not catch any prey. This fitness component discourages random predator movements and provides a smooth gradient. More specifically,

$$Z_{\text{predator}} = \begin{cases} 25 & \text{if both prey caught,} \\ \frac{20\gamma_m}{n} + 2\gamma_m + \frac{20d}{n} & \text{otherwise,} \end{cases}$$

where γ_m is the number of prey caught, n is the total number of prey, and d is the normalized sum of distances from the predator to each of the prey at the end of the simulation.

Using a prey fitness function complementary to the predator fitness function (i.e the higher the predator fitness, the lower the prey fitness and vice-versa) [6] did not result in good evasion strategies for prey. Instead, a prey's fitness was defined to be proportional to its lifespan. Even if one of the prey was captured, both prey were punished although less severely than if both were captured. If both prey survived until the end of the simulation, they were assigned the highest possible fitness. The simulation was limited to a maximum of 150 time steps so that the predators would have a finite amount of time in which to capture the prey. More specifically,

$$Z_{\text{prey}} = \begin{cases} 25 & \text{if neither prey caught,} \\ 12.5 & \text{if one prey caught,} \\ \frac{12.5\rho}{R} & \text{if both prey caught,} \end{cases}$$

where ρ is the number of time steps for which at least one prey remained alive, and R is the maximum possible number of time steps.

The best predator and prey teams from each generation are saved in a Hall of Fame, and used during crossover in later generations.
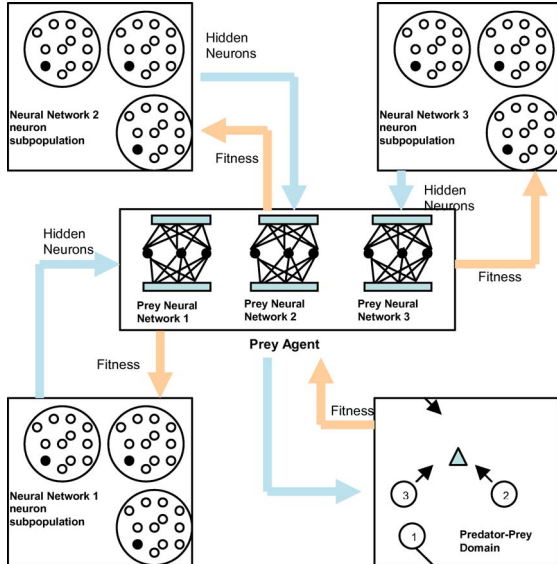

Fig. 2. Split Network Architecture - A single agent consists of multiple networks. Each network is dedicated to one opposing team member. Fitness is distributed equally among the participating hidden neurons of the agent networks.

VI. RESULTS

At first, some preliminary experiments were carried out that helped determine the underlying problems in scaling coevolution to teams of prey and predators. The insights from these experiments were then used to develop new techniques for sustaining coevolution of a predator team against both single and multiple prey. Every experiment consisted of several simulation runs. Similar behaviors evolved during these runs, and the results from a typical run are discussed here.

A. Preliminary Experiments

In the first experiment, a team of two prey was coevolved with a team of three predators. Each prey and predator was controlled by a single neural network and fitness shared within a team through Multi-Agent ESP architecture. The experiment did not result in evolution of smart strategies on the part of predators. This outcome in turn implied that the prey were not required to learn any effective evasive behaviours. A state of stagnation resulted where the prey always win.

In order to study this stagnation in greater detail, the experiment was simplified to involve a single prey. It is important to note that now the problem is much more complex for the prey than the predators, especially because the prey has to keep track of each of the three predator positions. Therefore the predators are easily able to catch the single prey. Overwhelming amounts of input information to the prey seemed to be a possible cause for this situation. However, just increasing the number of hidden neurons in each prey neural network did not yield improved results. This outcome suggests that the components within the existing Multi-Agent ESP architecture (hidden neurons) were not able to effectively cooperate to decompose the prey evasion task. To solve this problem, new prey architecture was introduced.

Another hierarchical layer was added to the Multi-Agent ESP architecture bringing cooperatively coevolving components to the agent level.

In this architecture, each individual prey agent is split into three neural networks evolved with Multi-Agent ESP as shown in Figure 2. Each neural network of the prey has x and y offset distances of one predator as its input, and its outputs represent the confidence values of four possible prey actions. The output values from each neural network corresponding to a given action are added up and the action corresponding to the largest sum is selected (see Figure 3).

With this split network architecture, coevolution was successfully sustained as will be described next. The success rate of prey and predators alternate in cycles and new behaviors emerge in each phase. Videos of these behaviors are available at http://www.dailymotion.com/conference_videos/1. (Note: In these videos, prey are black squares and predators are colored squares)

B. Arms Race in Predator Team vs. Single Prey

Each predator agent assumed the role of either an attacker or blocker. The role of an attacker is to chase the prey, while
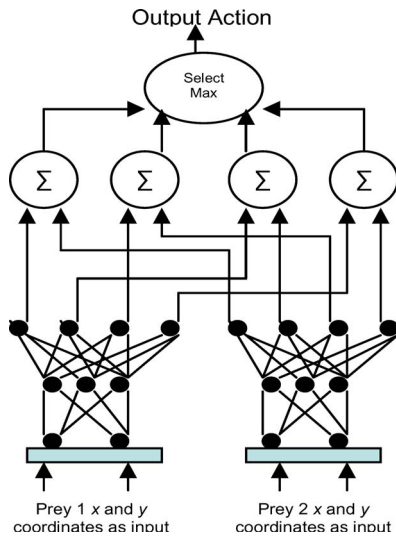
Fig. 3. Neural Network Output Combiner for Predator agent. The output values from each neural network corresponding to a given action are added up and the action corresponding to the largest sum is selected.

the blocker moves in a localized area to obstruct the prey's path.

Initially in generations 50-75, the prey evolves only a greedy fleeing strategy, where it moves away from the closest approaching predator. Simple predator behavior is enough to catch the prey in this case. Two predators block the prey and the third approaches it from the third direction (Figure 4, Phase 1, and Video 1). The success rate of the predators is high in this phase. At generations 75-100, the prey evolves to selectively use the option of fleeing from the closest predator, and sometimes goes around in a small circle with the closest predator following on its tail (Figure 4, Phase 2 and Video 2). At this stage, the other predators too move between fixed positions without making any new move to catch the prey because they are acting as blockers. The prey survives more often and therefore has high success rate in this phase. For generations 100-150, the predators learn to avoid this deadlock (Figure 4, Phase 3 and Video 3). Two of them now approach the prey from opposing directions (acting as attackers) and the third one assumes the role of blocking. The predators are more successful in this phase of the arms race. In the next phase (generations 150-180), the prey demonstrates intelligent baiting behavior by waiting for the two predators to converge towards it before moving away in a direction opposite to that of the predators (Figure 4, Phase 4 and Video 4). Since the third predator, the blocker, remains mostly stationary, the prey can easily dodge it. To counter this move, the predators learn to dynamically switch roles in generations 180-200 (Figure 4, Phase 5 and Video 5). The blocker also starts to follow the prey when it tries to escape.

In generations 200-250, the prey is captured often by two attackers and one blocker (Figure 4, Phase 6). In generations 250-300, the prey learns to avoid the blocker by sidestepping as it crosses the blocker's path (Figure 4, Phase 7 and Video

6). The blocker counters this (in generations 320-360) by not blocking the path of the prey directly. It stays a couple of steps away from the straight line path of the prey and then moves towards it (in a direction perpendicular to the prey) as soon it comes within a catch able distance (Figure 4, Phase 8 and Video 7). In the next phase (generations 400-450), the prey learns to reverse its direction so that it avoids both the attackers and the blocker (Figure 4, Phase 9 and Video 8).

The observations above clearly demonstrate that the split network architecture is successful in sustaining an arms race. This idea is next applied to teams of predator and prey agents to help sustain simultaneous cooperative and competitive coevolution.

C. Cooperative and Competitive Coevolution in Prey and Predator Teams

Let us now return to the original problem of coevolving three predators versus two prey. To solve the initial challenge of stagnant behaviors in predators, the split network architecture is used for the predators as well. Each predator now consists of two neural networks to keep track of the two prey. As with the single prey experiment, each prey consists of three neural networks (one for each predator). Fitness is shared between multiple neural networks within an agent through the Multi-Agent ESP architecture. At a higher level, each team of agents also shares fitness in a similar fashion.

The hierarchical layers of cooperation and competition are shown in Figure 5. There are three levels of cooperation operating in this system. Hidden neurons selected from separate coevolving subpopulations within a single neural network cooperate to form the Level 1 of this hierarchy. At the next level, multiple neural networks within a single agent cooperate to generate agent behavior. At the third level the individual agents in a team cooperate to defeat the other team. At the highest level, there is competition between the teams of predators and prey. In the previous experiment of predator team vs. single prey, each predator agent has only a single neural network to track the single prey. Similarly, there is no team-level cooperation for the prey.

This scenario of multiple evolving prey is far more complex than that of a single evolving prey. There are multiple predators and prey on the field simultaneously and thus there are far more factors that affect the evolution of both the teams. The hierarchical structure makes it possible to distribute roles effectively, which in turn allows both populations to adapt to the continuously changing environment.

The predators must choose between two alternatives - catching the prey one by one or herding them together. Herding of prey is a complex behavior especially because the predators have to sacrifice the immediate gain of catching a single prey to achieve better efficiency by catching them together. In the beginning (generations 0-25) when predators have not yet learned high-level pursuit behaviors, they unsuccessfully attempt to herd the prey before capture. The prey easily evade the predators during this time. At (generations 25-50), predators first attempt to herd the prey, but if their pursuit fails, they switch to catching the prey one
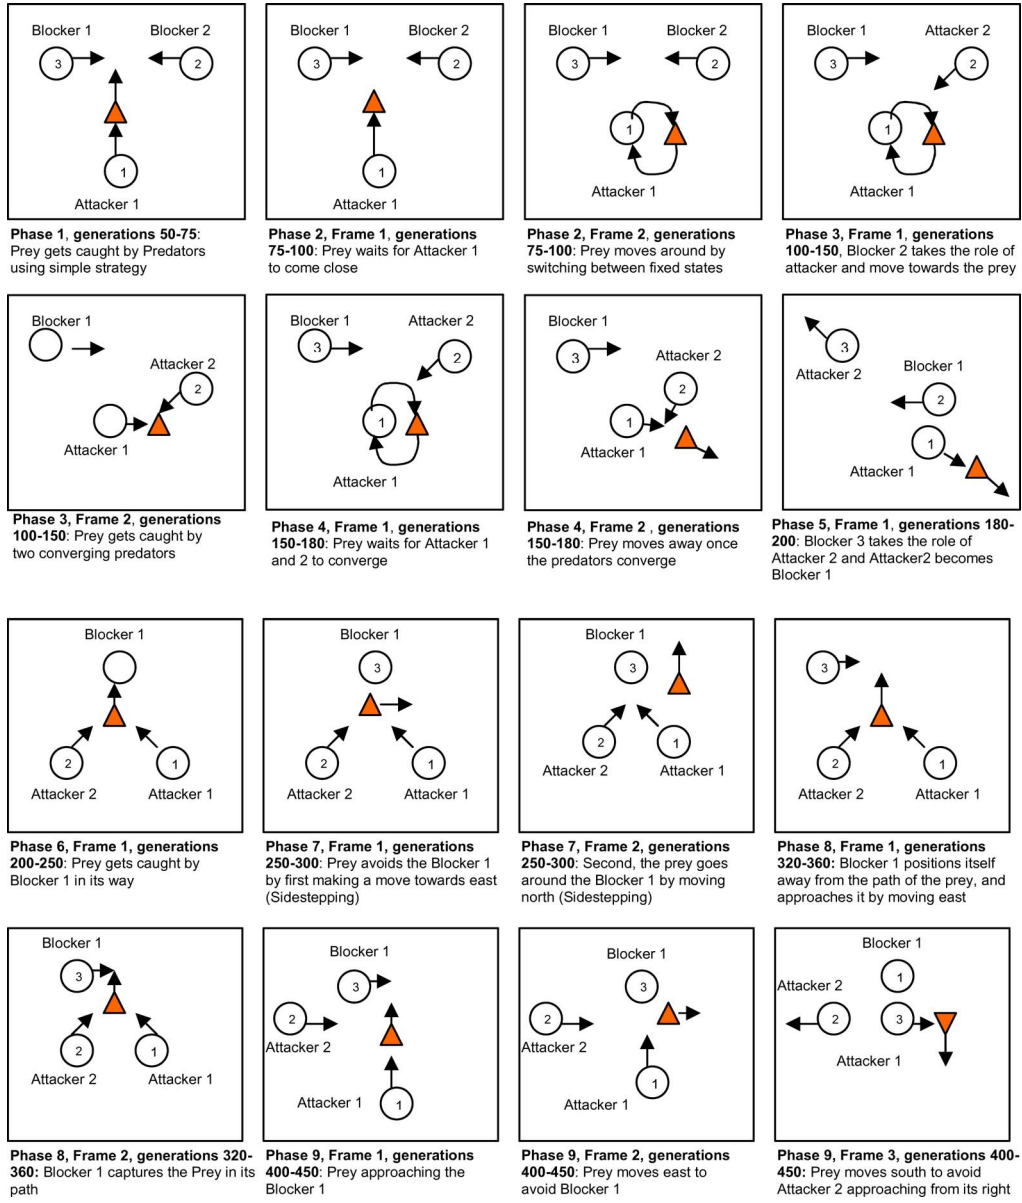
Fig. 4. Arms Race in Team of Predators vs. Single Prey: Emergence of predator-prey behavior in phases.

after the other. At this point most of the behaviors observed in the single prey scenario (like dynamic role switching in predators and baiting by the prey) also evolve in this case. At generations 150- 200, predators are able to succeed in herding the prey and capturing them simultaneously (Figure 6, Phase 1 and Video 9). To counter herding (generations 250-320), the prey evolve to scatter in different directions just before the predators converge on them (see Figure 6, Phase 2 and Video 10). One reason for this last-minute scattering could be that once the predators have almost converged, they are all roughly in the same location, making it easier for the prey to evade them. In this manner, behaviors coevolve in cycles, resulting in complex final behaviors for both predators and prey.

VII. DISCUSSION

The experiments show that it is possible to sustain coevolution of teams of competing and cooperating agents. This result was made possible by a new architecture that consists of cooperating components.

Initially when each agent consisted of a single neural network, the simulation stagnated to fixed behaviors and arms race did not occur. The predators did not learn to catch the prey and hence no smart evasive strategies emerged in the prey. Simplifying the problem to a single prey suggested that the agents apparently did not have enough computing resources to track the actions of the opposing team. However, the problem was not that simple: Simply increasing the number of hidden neurons did not solve it. Only when a new Multi-Component ESP architecture was created, coevolution
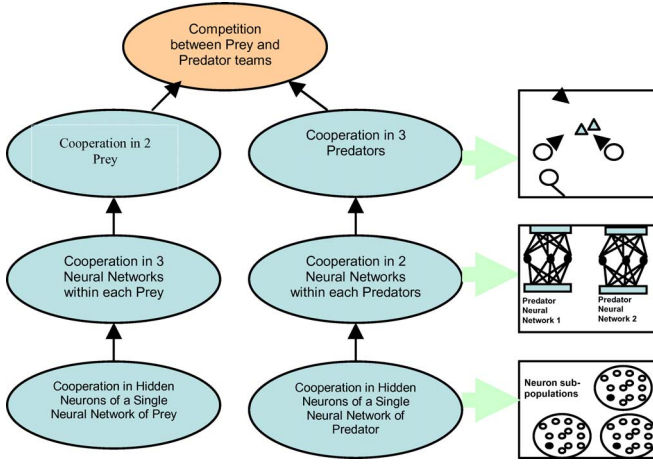
Fig. 5. Hierarchical Levels of Cooperation and Competition in teams of Prey and Predators.

Phase 1, Frame 1, generations 150-200: Predators surround the prey to herd them

Phase 1, Frame 2, generations 150-200: Prey are herded and caught simultaneously

Phase 2, Frame 1, generations 250-300: Prey bait the predators by first moving the close to each other

Phase 2, Frame 2, generations 250-300: Prey scatter in different directions to disturb predator group formation

Fig. 6. Arms Race in Team of Predators vs. Team of Prey: Emergence of predator-prey behavior in phases.

was sustained. As part of this architecture, the number of neural networks was increased within a single agent. This number was matched to the number of members in the opposing team so that each neural network kept track of one of them. As a result, coevolution was sustained even in the case of competing teams of three predators and two prey. The important insight is that it is easier to coevolve components that cooperate to form a solution, rather than evolve the complete solution directly. This idea is, of course, the same that motivated ESP and Multi- Agent ESP. In this paper, it is shown to apply to the level of complex behavior in individual agents, as part of the multi-level hierarchy of collaboration and competition.

As demonstrated by Yong and Miikkulainen [23], the predators initially evolve rigid role-based behaviors in which some of them act as attackers and others as blockers. Each predator reacts to the preys' actions, and direct communication (knowledge of team members' positions) between the predators is unnecessary. The attackers actively pursue the prey from different directions, and the blocker moves into a location that prevents the prey from escaping. However, the prey evolves a smart counter strategy: it waits for two attacking predators to converge before fleeing (Figure 4). Further, it reacts differently to each predator, suggesting that it has learned the possible role definition for the predators. To counter these more complex strategies in prey, the predators learn to switch roles (as shown in Figure 4). This result is interesting because in the original experiments by Yong and Miikkulainen, such dynamic role switching occurred only in communicating agents [23]. However, coevolution is powerful enough to find a way to switch roles even without communication. Coordination is still based on stigmergy, i.e, absorbing the clues in the environment (such as prey behavior), but it is a more complex and dynamic version of it. Communication is without doubt a useful ability, but it is interesting to see that quite complex and flexible team behavior can be achieved efficiently even without it.
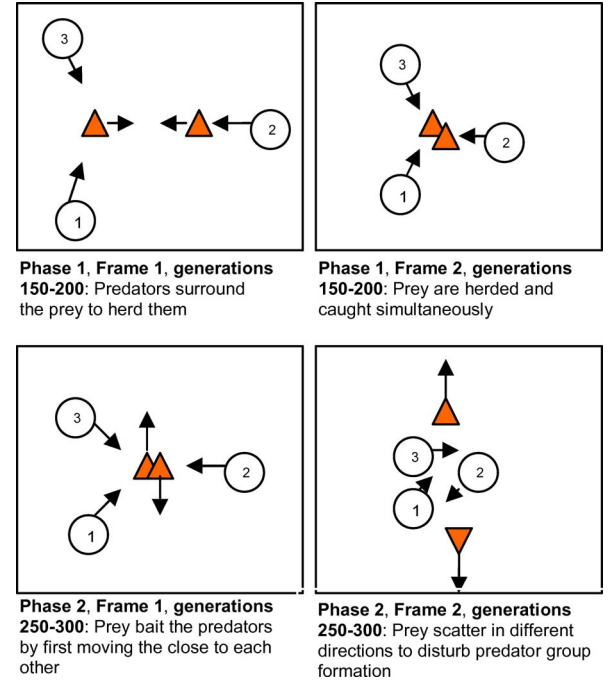
VIII. FUTURE WORK

This research makes many directions of future work possible. First, the arms race between predators and prey should continue beyond the observations made in this paper, and could lead to the emergence of even more complex behaviors. However, as the simulation progresses, these behaviors take more and more time to discover and they are more difficult to analyze. Thus the next step is to develop new methods to look at and analyze such increasingly complex behaviors.

Preliminary experiments show that using a combiner neural network to aggregate the split networks output in the Multi-Component ESP architecture is another interesting research direction. This approach makes it possible to utilize correlations among relative locations of various agents. Different split network topologies can be evolved based upon the complexity of the problem domain.

Application of this work to other domains like Robot Soccer or Unreal Tournament[TM] is an interesting possibility as well. Eventually, simultaneous competitive and cooperative coevolution may thus make it possible to build artificial systems that are comparable in complexity to those seen in nature.

IX. CONCLUSIONS

The experiments in this paper confirmed two hypotheses. First, competitive and cooperative coevolution were successfully sustained in the predator-prey domain. Second, a hierarchy of cooperation and competition similar to that in nature was observed to emerge, including various high-level competitive and cooperative strategies in both predators

and prey. This process was made possible by a new Multi-Component ESP architecture for a single agent, where each agent controller consists of multiple cooperating neural network modules. As a result, the predators learned to switch roles dynamically based on stigmergy, and to herd the prey together before capturing them. To counter these predator behaviors, the prey learned high-level strategies such as baiting, scattering, direction reversal and sidestepping. These behaviors were learned in an arms race with predators and prey each being the successful population in turn. Such a competitive and cooperative coevolution is a possible way to construct complex behaviors for games in the future.

X. Acknowledgments

References

[1] W. W. Benson, K. S. Brown Jr., and L. E. Gilbert, "Coevolution of plants and herbivores: Passionflower butterflies," *Evolution*, pp. 659–680, 1976.

[2] P. Burrow and S. M. Lucas, "Evolution versus Temporal Difference Learning for Learning to Play Ms. Pac-Man," in *IEEE Symposium on Computational Intelligence and Games*, 2009.

[3] R. Gadagkar, *Survival Strategies: cooperation and conflict in animal societies.* Harvard University Press, 1997.

[4] F. Gomez and R. Miikkulainen, "Incremental evolution of complex general behavior," *Adaptive Behavior*, pp. 317–342, 1997.

[5] J. Grefenstette and R. Daley, "Methods for Competitive and Cooperative Evolution," in *ICMAS*, 1995, pp. 276–282.

[6] D. W. Hillis, "Co-evolving parasites improve simulated evolution as an optimization procedure," *Physica D*, pp. 228–234, 1990.

[7] P. Hingston, "A Turing Test for Computer Game Bots," in *IEEE Transactions on Computational Intelligence and AI in Games vol.1, Issue: 3*, 2009, pp. 169–186.

[8] H. Juill and J. B. Pollack, "Coevolutionary learning: A case study," in *Proceedings of the Fifteenth International Conference on Machine Learning*, 1998, pp. 251–259.

[9] G. Miller and D. Cliff, "Co-evolution of pursuit and evasion I: Biological and game-theoretic foundations," *Technical Report CSRP311*, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.

[10] M. Mitchell, "Coevolutionary learning with spatially distributed populations," *Computational Intelligence: Principles and Practice*, New York: IEEE Computational Intelligence Society.

[11] M. Mitchell, M. D. Thomure, and N. L. Williams, "The role of space in the success of coevolutionary learning," in *Artificial Life X: Proceedings of the Tenth International Conference on the Simulation and Synthesis of Living Systems*, 2006, pp. 118–124.

[12] S. Nolfi, D. Floreano, "Co-evolving predator and prey robots: Do Arms Races arise in artificial evolution?," *Artificial Life*, pp. 311–335, 1998.

[13] L. Pagie and M. Mitchell, "A comparison of evolutionary and coevolutionary search," *International Journal of Computational Intelligence and Applications*, pp. 53–69, 2002.

[14] A. Perez-Uribe, D. Floreano, and L. Keller, "Effects of group composition and level of selection in the evolution of cooperation in artificial ants," in *7th European Conference on Artificial Life*, 2003, pp. 128–137.

[15] M. A. Potter and K. A. D. Jong, "Cooperative coevolution: An architecture for evolving coadapted subcomponents," *Evolutionary Computation*, pp. 1–29, 2000.

[16] C. D. Rosin and R. K. Belew, "Methods for competitive coevolution: Finding opponents worth beating," in *Proceedings of the Sixth International Conference on Genetic Algorithms*, 1995, pp. 373–381.

[17] C. D. Rosin and R. K. Belew, "New methods for competitive coevolution," *Evolutionary Computation*, 1996.

[18] K. O. Stanley and R. Miikkulainen 'Competitive coevolution through evolutionary complexification'," *Journal of Artificial Intelligence Research*, pp. 63–100, 2004.

[19] K. O. Stanley, B. D. Bryant, and R. Miikkulainen, "Real-time neuroevolution in the NERO video game," in *IEEE Transactions on Evolutionary Computation Special Issue on Evolutionary Computation and Games*, 2005, pp. 653–668.

[20] E. Uchibe and M. Asada, "Incremental Coevolution with Competive and Cooperative Tasks in a Multirobot Environment," in *Proceedings of the IEEE*, 2006.

[21] S. Whiteson, N. Kohl, R. Miikkulainen, and P. Stone, "Evolving Soccer Keepaway Players Through Task Decomposition," *Machine Learning*, pp. 5–30, 2005.

[22] N. Williams and M. Mitchell, "Investigating the success of spatial coevolutionary learning," in *Proceedings of the 2005 Genetic and Evolutionary Computation Conference*, 2005, pp. 523–530.

[23] C. Yong and R. Miikkulainen, "Coevolution of role-based cooperation in Multi-Agent systems," *IEEE Transactions on Autonomous Mental Development*, 2010.