

The Role of Reward Structure, Coordination Mechanism and Net Return in the Evolution of Cooperation

Padmini Rajagopalan, Aditya Rawal, Risto Miikkulainen, *Senior Member, IEEE*,
Marc A. Wiseman, and Kay E. Holekamp

Abstract—Biology can offer insight into how realistic artificial agents and complex interactions between them can be created. For instance, spotted hyenas of Western Africa typically hunt alone, but cooperate once in a while to catch zebras. Using hyenas as motivation, this paper evaluates three potential factors that affect the evolution of cooperation in a team of game agents: reward structure (i.e. whether fitness reward is given to an individual agent or shared by the team), coordination mechanism (i.e. stigmergic vs. direct communication), and net return (i.e. the size of reward relative to the difficulty of obtaining it). Through neuroevolution using Multi-Component ESP [1], three predators were evolved in a field containing multiple fixed-behavior prey. Six such experiments show that (1) although shared reward strongly promotes cooperation, it is not strictly necessary for evolving cooperation; (2) although stigmergic coordination works well in simple, unambiguous tasks (i.e. with a single prey), direct communication is more effective in more complex tasks (i.e. with multiple prey); and (3) the predators evolve to hunt alone or cooperatively based on which approach results in higher net return. Insights from these computational simulations can be used to develop more intelligent behaviors for game agents in the future.

I. INTRODUCTION

Scripting cooperative behavior in multi-agent games is challenging, especially when the agent behaviors are required to be adaptive. Nature provides abundant examples of such behaviors whenever animals coexist as a team, such as an ant colony, a flock of birds, or a pack of wolves. Individual and team interests often conflict in such teams; sometimes the individuals act by themselves, in other times they cooperate explicitly to achieve a common goal. The origins of such cooperation are not fully understood. While experimental psychological research has uncovered many factors affecting cooperation [2], the ones that can most easily be controlled in video games are: how the rewards are distributed, how well the individuals can communicate, how difficult the task is and how it affects the individual's fitness.

In particular, spotted hyenas demonstrate sophisticated hunting behaviors, often catching small prey such as gazelles individually and sometimes coordinating to hunt zebras (which they cannot catch alone) [3], [4]. The rewards from the hunt are not distributed equally but depend upon social rank of an individual. Hyenas also use refined vocal, visual and tactile signals to communicate during the hunt as well

as other social interactions [4], [5], [6], [7]. The social and hunting behaviors of hyenas are complex, however, they are easier to analyze than those of primates. Whereas much of primate behavior is learned and flexible, the hyena behaviors, though not innate, are less plastic [8], [9]. Thus, evolutionary simulations of hyena behaviors make an ideal test bed for studying the emergence of cooperative behaviors in artificial teams.

As a research setting, a Predator-Prey simulation environment was constructed to model hunting behaviors of hyenas. This environment is easy to simulate with quantifiable results and provides an abstract approximation to the hunting environment of hyenas in the wild. The predator-prey domain has previously been used to study the evolution of team behaviors in game agents [10], [11]. Complex agent behaviors have emerged through neuroevolution in games such as simulated robot soccer [12], robotic battle [13] and Ms. Pac-Man [14]. Multi-Component ESP, a neuroevolutionary architecture, was used to coevolve the predators in this domain. This architecture was previously shown effective in evolving cooperation and competition among predator-prey teams [1]. Here, prey capture was considered a cooperative move on the part of the predators because the prey cannot be caught by a single predator on its own.

The study consisted of six experiments, where the reward structure (i.e. whether the reward was given to the individual who caught the prey or shared among the whole team), the coordination mechanism (i.e. through stigmergic or direct communication), and net return (i.e. the difficulty and payoff of prey capture) were varied (Table I). The experiments were motivated by hyenas, which may or may not share kills equally, communicate in many different ways, and hunt different kinds of prey either individually or as a team.

There were two types of prey in the experiments: zebras and gazelles. The zebras had the same speed as the predators, and thus were more difficult to catch. They could not be caught by a single predator, but had to be surrounded before capture. Hence, the capture of a zebra constituted an act of cooperation among the predators, and the number of zebras caught was a useful measure of the extent of predator cooperation. Zebras gave more reward upon capture than gazelles. But in these simulations, the gazelles were made slower than the predators, and thus, could be caught by a single predator chasing it.

First, three predators were coevolved to catch four zebras without communication and with individual fitness rewards. Cooperation did eventually emerge, but it was very difficult

Padmini Rajagopalan, Aditya Rawal, and Risto Miikkulainen are with the Department of Computer Science, University of Texas at Austin, Austin, TX 78712 USA (email: {aditya, padmini, risto}@cs.utexas.edu). Marc A. Wiseman, and Kay E. Holekamp are with the Department of Zoology, Michigan State University, East Lansing, MI 48824 USA (email: {holekamp, wisema39}@msu.edu).

TABLE I
SETUP AND RESULTS OF THE EXPERIMENTS

Experiment Number	Reward Structure	Coordination Mechanism	Number and Type of Prey	Reward for Catching Each Prey	Results
1	Individual	Stigmergic	4 Zebras	150	Little cooperation, evolves slowly
2	Shared	Stigmergic	4 Zebras	150	Systematic cooperation, evolves early
3	Individual	Direct communication	4 Zebras	150	Improved cooperation, evolves gradually
4	Shared	Direct communication	4 Zebras	150	Effective, fast cooperation
5	For gazelles - Individual, For zebras - Shared	Direct communication	4 Gazelles, 1 Zebra	Gazelle - 100, Zebra - 150	Cooperation does not evolve, gazelle-hunting preferred
6	For gazelles - Individual, For zebras - Shared	Direct communication	4 Gazelles, 1 Zebra	Gazelle - 100, Zebra - 450	Cooperation evolves, zebra-hunting preferred

to evolve, suggesting that there was neither sufficient pressure nor sufficient means to establish it. Second, the predators were coevolved to share fitness in the same task; such pressure was sufficient incentive to establish cooperation quickly. The predator agents did not need to communicate directly (i.e. by observing one another) to implement this cooperative behavior; communication through stigmergy [15]), i.e. by observing the prey's behavior, was enough. Next, an experiment was performed with communication turned on, but with individual fitness. Cooperation emerged very easily, suggesting that direct communication is a better coordination mechanism than stigmergy. In addition, sharing the fitness among the predators in the next experiment made the effect even stronger. Interestingly, this result contradicts that of Yong and Miikkulainen [15], where cooperation without communication was found more efficient. However, their simulations contained only one prey as opposed to the four prey in the experiments in this paper, which makes stigmergic signals harder to interpret.

Two further experiments were performed with both zebras and gazelles on the field while varying the reward for catching the zebras. The reward from zebra capture was shared equally among the predators, but the reward from gazelle capture went to the predator that caught it. Cooperation (i.e. zebra hunting) turned out to emerge reliably in the predator agents when the payoff for the zebras was sufficient to make it worth their while; however, these agents were also able to hunt the smaller prey (gazelles) individually when there were no zebras available. Thus the simulations revealed a principled contribution of the reward structure, coordination mechanism, and net return in determining whether cooperation emerged or not.

Section 2 reviews prior work on evolution of cooperation, and section 3 defines the predator-prey problem domain. Section 4 puts forth hypotheses about the factors affecting evolution of cooperative behaviors. Section 5 includes detailed descriptions of the experiments and the results. A discussion of implications and future work in sections 6 and 7 concludes the paper.

II. RELATED WORK

A significant body of work exists on computational modeling of cooperation in nature. For instance, flocking behaviors of birds and shoaling of fish have been modeled extensively using rule-based approaches [16], [17], [18]. Cooperative behavior of micro-organisms like bacteria and viruses has been modeled with genetic algorithms [19], [20]. Ant and bee colonies have been the subject of many studies involving evolutionary computation as well [21], [22], [23]. Perez-Uribe et al. [22] studied the effects of group composition and level of selection in artificial ants. In general, ant-colony optimization [21], [24] is a probabilistic technique that can be used to model ants seeking shortest paths between their colony and a food source. At a higher level, the conflicting goals of individual versus team rewards were studied by Keller and Floreano [25], focusing on the importance of relatedness (kin selection) in robotic agents. Although many mammals like hyenas demonstrate intelligent team behaviors, their behaviors have not yet been analyzed computationally to our knowledge.

Coevolution is defined as the simultaneous evolution of two or more individuals whose fitnesses are measured based on their interactions with each other [26]. In cooperative coevolution, the individuals have to evolve to cooperate to perform a task. They share the rewards and punishments of their individual actions equally. It turns out that it is often easier to coevolve components that cooperate to form a solution, rather than evolve the complete solution directly [27], [28]. The components will thus evolve different roles in the cooperative task.

For example, in the Enforced SubPopulations (ESP) architecture [27], neurons selected from different subpopulations are required to form a neural network whose fitness is then shared equally among them. Such an approach breaks a complex task into easier sub-tasks, avoids competing conventions among the component neurons and makes the search space smaller. These effects make neuroevolution faster and more efficient.

Similarly, Multi-Component ESP extends this approach to evolve a team of agents (Figure 1). Each agent comprises multiple ESP-type neural networks to sense different objects

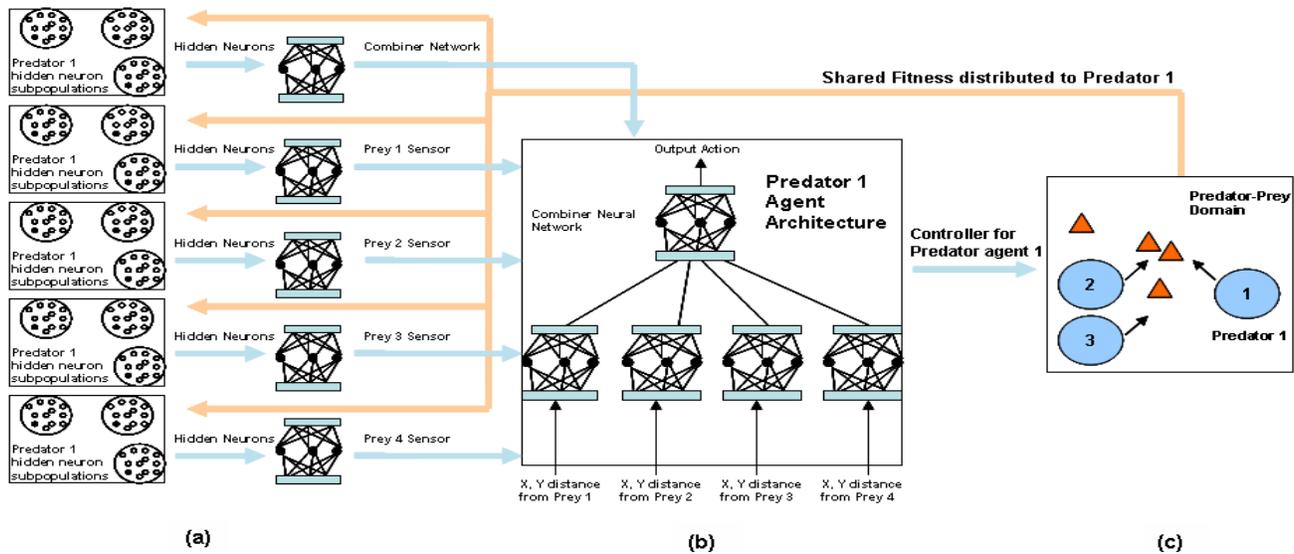


Fig. 1. Multi-Component ESP in the predator-prey domain for predator agent in Experiment 1. A single predator agent (shown in (b)) is composed of five neural networks. Four of these sense one of the prey agents. Their outputs are given to a fifth combiner network that outputs the next move for that predator. Each network is evolved in a separate ESP process, where one subpopulation is evolved for each of the neurons in the network (a). The predator is evaluated in the domain simulation with prey and other predator agents (c). Its fitness is distributed equally among all the networks and among all the neurons that participated in it. In this manner, evolution can discover neurons and networks that cooperate well to form an effective agent.

in the environment. The team's reward from fitness evaluations is shared equally by the component networks of all the agents [1]. The cooperative coevolution approach has been shown to be effective when coevolving teams of agents. First, Yong and Miikkulainen [15] showed that a team of predators that share fitness can evolve to cooperate to catch prey with or without communication. In their experiments, without communication, the roles the predators evolve are more rigid but more effective; with communication, their roles are less efficient but more flexible. Second, Rawal et al. [1] showed that the Multi-Component ESP architecture can coevolve a team of predators with a team of prey. The individuals cooperate within the team, but the predator team competes with the prey team. Therefore, the Multi-Component ESP architecture will be used to evolve the predators in this paper as well. In prior work, the outputs of the neural networks within a predator or prey agent were summed to get the final output action. However, preliminary experiments showed that including a combiner network to combine the outputs of these networks was more powerful and resulted in the emergence of more complex behaviors. Hence, this technique was used in this paper (Figure 1). The combiner network weights were evolved using the same technique as the other networks.

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 [10]. Pursuit-evasion problems are common in the real world as well as in artificial environments like video games. They are challenging because they are highly dynamic and the optimal actions are

generally not known. The predator-prey domain can easily be extended to include multiple agents acting in teams. Factors such as reward structure, coordination mechanism and net return are also modifiable, making this domain an ideal platform to study hunting behaviors of hyenas.

More specifically, a predator-prey domain was constructed in this paper to study constraints under which cooperation emerges. A team of predators (hyenas) is evolved using cooperative coevolution to capture fixed-behavior prey (a gazelle or a zebra). The world in this simulation is a discrete toroidal environment with 100 x 100 grid locations without obstacles, where 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 take a step simultaneously. To move diagonally, an agent has to take two steps (one in the east-west direction and one in the north-south direction). A predator is said to have caught a prey if it moves into the same location in the world as the prey. The predators are aware of prey positions and the prey are aware of predator positions. Direct communication among predators (in terms of knowledge of other predators' positions) is also introduced in some cases. In all other cases, the predator agents can sense only prey movements and have to use that to coordinate their actions (stigmergic communication). There is no direct communication among the prey. Each predator has as its inputs the x and y offsets of all the prey from that predator. In the case of communicating predators, they also get as input the x and y offsets to the other predators. When fitness rewards from prey capture are shared, all the predators gain fitness even when only one of them actually catches the prey. In cases with individual fitness, only the particular predator that captures the prey gets the reward.

There are two types of prey in the environment - a smaller prey (gazelle) that moves with 0.75 times the speed of the predator and a larger prey (zebra) that has the same speed as the predator. The prey behaviors in these experiments are hard-coded and do not evolve. Each prey simply moves directly away from the current nearest predator. The predators can therefore catch the smaller prey individually, but cannot catch the larger prey by just following the prey around, because their grid world is toroidal. The predators have to surround a zebra from different directions before they can catch it. In cases where both types of prey exist in the field simultaneously, the predators need to decide whether to catch the small prey individually or to coordinate and hunt the larger prey together. The larger prey give more reward than the smaller prey, and the relative reward amounts can be varied.

Thus, three parameters are progressively modified in these experiments: (1) whether only the individual actually catching the prey receives the fitness, or whether it is shared by all individuals, (2) whether the predators can observe one another or not (direct vs. stigmergic communication), and (3) the size of the fitness reward from catching a prey. These experiments are used to contrast the role of each of these parameters in the evolution of cooperation.

IV. HYPOTHESES

The main goal is to discover what factors influence whether the predators evolve to cooperate in their hunt or not. Spotted hyenas in the wild usually prefer to hunt small prey alone even though bigger, more difficult-to-catch prey such as zebras would give them more food. Only rarely, do they team up to hunt zebras [3], [4]. What are the reasons behind these seemingly suboptimal hunting habits? The following hypotheses about the evolution of cooperative behaviors in hyenas will be tested in the simulations:

- 1) *Sharing the fitness reward is an incentive for the predators to cooperate.*

If only the predator that catches a prey gets the reward, there is no reason for the other predators to help in the capture. If they all share the rewards, there is a benefit to cooperation, and it should emerge in evolution.

- 2) *If cooperation is easy to establish, e.g. through communication, it is more likely to emerge.*

Evolving to cooperate just through stigmergy (sensing of clues from the environment) is a difficult task. If the predators can communicate directly (i.e. by sensing each others' positions), the cost of cooperation is significantly lower because it becomes easier for the predators to coordinate their actions. Cooperation is thus more likely to emerge.

- 3) *Evolution of cooperation depends on payoff upon prey capture.*

Hyenas can both communicate and share fitness rewards, but they still typically hunt on their own. Perhaps the energy necessary to coordinate their hunting behaviors to catch larger prey is not off-set by

the higher rewards they get upon catching it. If the reward was higher, cooperation would be more likely to emerge.

These hypotheses will be tested in the computational experiments outlined next.

V. EXPERIMENTS

The experiments were designed to observe cooperation in predators hunting prey that are difficult for a single predator to catch on its own. The goal is to test whether predator communication, fitness distribution and size of the fitness reward influence evolution of cooperation in hypothesized ways.

A. Experimental Setup

Multi-Component ESP [1] is used to evolve the predators with the following parameter settings: Each predator consists of one network for the combiner, and one for each prey and predator it tracks: four in Experiments 1 and 2, six in Experiments 3 and 4, and seven in Experiments 5 and 6. The outputs of the predator and prey tracking networks are given as input to the combiner network to decide the next move for that predator. Each network (including the combiner) has a feedforward architecture with a single layer of 10 hidden neurons and sigmoidal activation functions. Each hidden neuron is evolved in a separate subpopulation consisting of 100 neurons; each neuron is represented as a concatenation of real-valued numbers representing full input and output connections. 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 predators are evaluated six times, with the prey and predators starting at random locations each time. The fitnesses over the six evaluations are averaged, and assigned to all the neurons that constituted the network. Fitness is the sum of rewards from all the different prey the predators catch during an evaluation (see last but one column in Table I). 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.

A summary of different experiments was given in Table I. In short, the first, i.e. the control experiment, includes non-communicating predators with individual fitness. The second experiment adds shared fitness rewards, while the third experiment adds communication instead. The predators in the fourth experiment both communicate and share fitness rewards.

The fifth and the sixth experiment focus on the hyena environment more specifically by having both zebras and gazelles on the field. The predators can both communicate and share the reward from zebra capture. They still prefer hunting individually for smaller animals (gazelles) that are

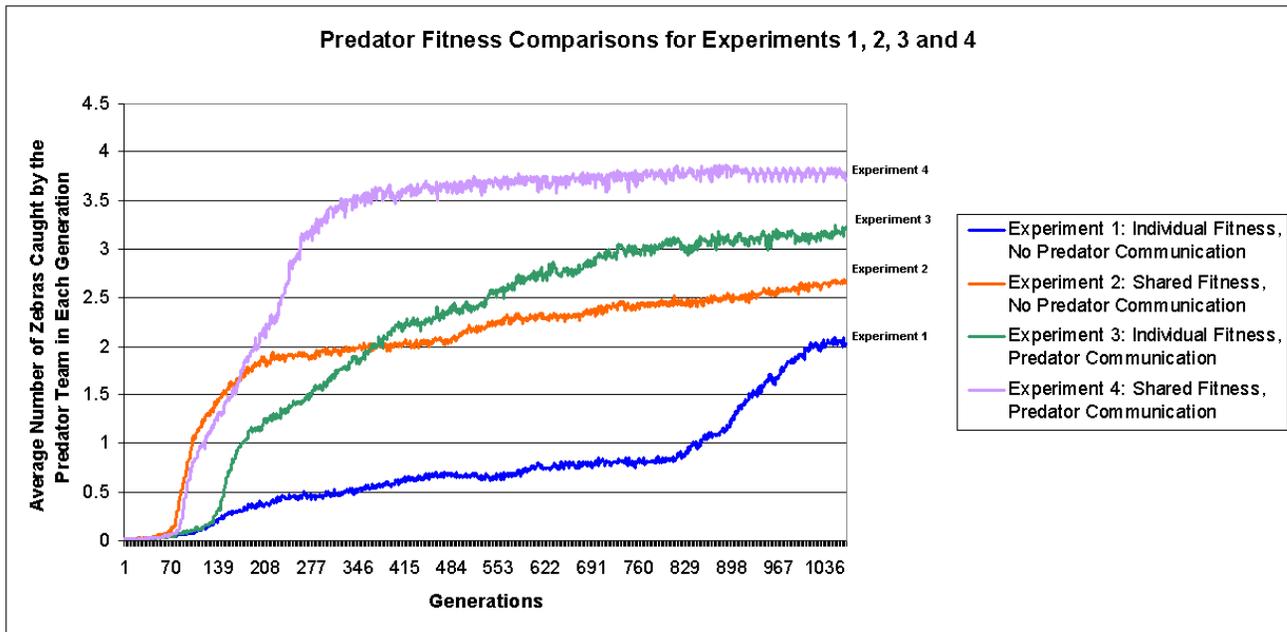


Fig. 2. Average number of prey (zebras) caught (out of four possible) in Experiments 1, 2, 3 and 4. The total number of prey caught by the three predators was averaged over 6000 trials for each generation. Cooperation is slow to evolve with individual rewards and without communication, and is less efficient (Experiment 1). Introduction of reward sharing results in faster and more effective evolution of cooperation (Experiment 2). Knowledge of positions of other predators makes it easier to evolve coordinated hunting strategies (Experiment 3). Evolution of cooperation is strongest when reward sharing and communication are combined (Experiment 4).

less rewarding. Cooperative hunting of big prey such as zebras is rare. Perhaps the cost of cooperation is higher than the reward obtained? Experiments 5 and 6 evaluate this hypothesis.

In each study, there are three predators; the number of prey varies from experiment to experiment. In Experiments 1 through 4, there are four zebras, and each one gives a fitness of 150 points to each of the three predators upon capture. In Experiments 5 and 6, there are four gazelles and one zebra, and each gazelle gives a fitness of 100 to the predator that catches it. In Experiment 5, catching the zebra gives a reward of 150 points to each predator, while in Experiment 6, this reward is increased to 450 points (Table I).

The output neurons represent different actions that a predator agent can take. Each prey has four possible 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.

B. Results

The results are summarized in the rightmost column of Table I. Here, the number of zebras caught by the predators is taken as the metric indicating the extent of cooperation because zebras need to be surrounded by more than one predator before they can be caught. This requires active movement coordination on the part of the predators.

In the control experiment (Experiment 1), the predators neither communicate nor share fitness. Cooperation does not evolve initially and as a result, they rarely catch any zebras.

On the other hand, adding reward sharing (Experiment 2) increases the number of prey caught as the predators efficiently evolve to cooperate over the early generations. The average number of zebras caught in each generation in Experiments 1 and 2 are contrasted in Figure 2.

Similarly, adding communication to predators with individual fitness in Experiment 3 results in the predators easily evolving to cooperate, leading to more prey captures (Figure 2). This effect is even stronger with both communication and fitness sharing enabled (Experiment 4; Figure 2), suggesting that these two factors affect different aspects of the evolution process, i.e. how easy it is to establish cooperation, and how worthwhile it is.

Experiments 5 and 6 were designed to answer the question: If there are both gazelles, which can be caught easily but give a lower fitness, and zebras, which need all the predators to cooperate to catch them but give higher fitness, which one is preferred? In Experiment 5, the predators prefer to hunt gazelles instead of evolving to cooperate to capture the zebra. The reward for catching the zebra is not large enough for cooperative behaviors to be selected during evolution. In contrast, in Experiment 6, it is large enough, and the predators slowly evolve to team up to capture this more difficult prey, thus verifying the hypothesis that net return is important in the evolution of cooperation (Figure 3). Interestingly, they are still able to hunt gazelles as well, but only do it when there are no zebras around even though zebras are still hard to catch. This result is important because it suggests that cooperative strategies include individual strategies as a special case.

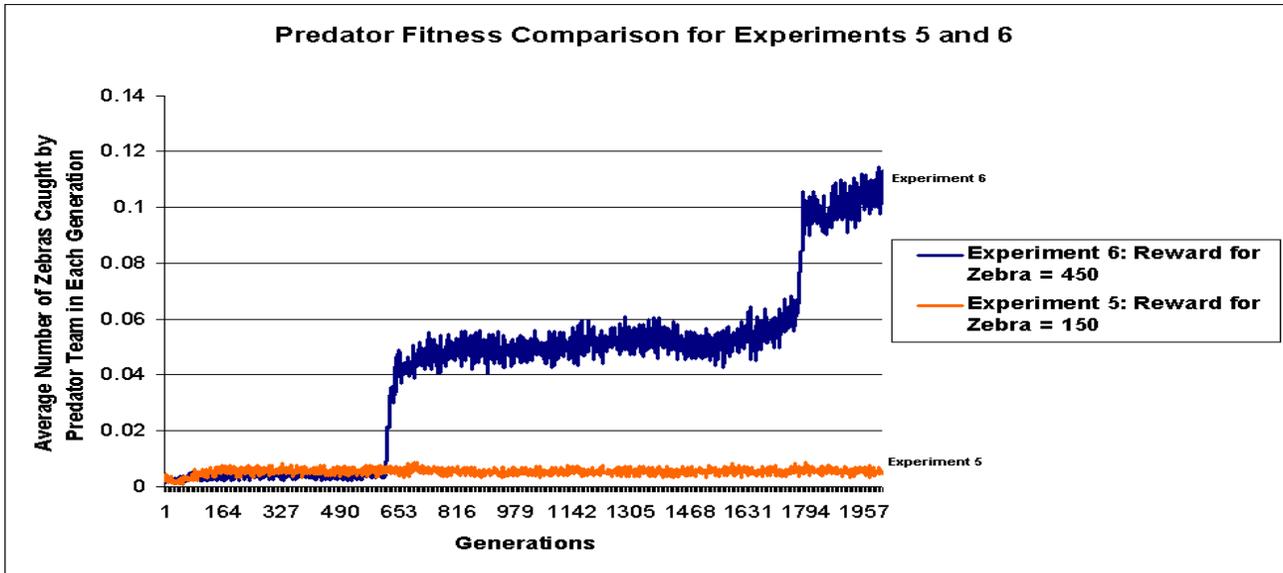


Fig. 3. Average number of zebras caught in Experiments 5 and 6. The total number of prey (out of one possible) caught by the three predators was averaged over 6000 trials for each generation. When the payoff on capturing a zebra is low with respect to the difficulty of catching it (Experiment 5), the predators prefer to hunt the easy-to-catch gazelles individually. When the net return for capturing the zebra is high enough (Experiment 6), the predators evolve to discover cooperative strategies to hunt it. Once it is caught, they continue by hunting gazelles.

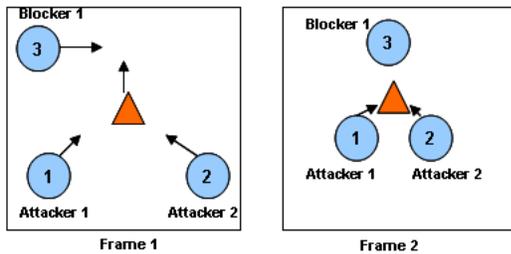


Fig. 4. Collaboration during a zebra hunt (frequently observed in Experiments 2, 3 and 4). Two predators assume the role of attacker and chase the prey towards the third predator, which acts as blocker (Frame 1). They surround the prey from three directions and capture it (Frame 2).

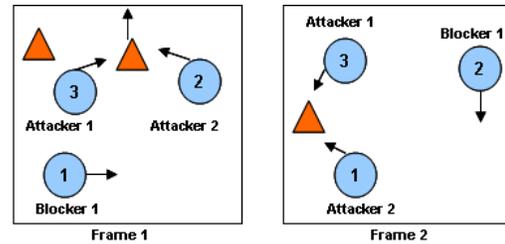


Fig. 5. Role switching between the hunting of two prey (Experiment 4). The predators assume certain roles (attackers, blocker) while hunting a prey (Frame 1). These roles can change when hunting the next prey (Frame 2).

C. Evolved Behaviors

These overall results show that communication among the predators, reward distribution when a prey is caught, and net return from capturing a prey are important factors in the evolution of coordinated hunting, as hypothesized. Examples of evolved behaviors are shown at <http://nn.cs.utexas.edu/?cooperation> and described below. In Figures 4, 5, 6 and 7, the circles represent predators, the triangles represent large prey (zebras) and the squares represent smaller prey (gazelles).

When the predators neither communicated nor shared fitness, they performed very poorly on the prey-capture task. Initially, they did not evolve to cooperate to catch the prey, and the prey easily eluded any individual predators. No predator knows where the other predators are, and there is no fitness incentive to cooperate when it senses (through stigmergy) other predators hunting the same prey (Video 1). Eventually, the predators do evolve to coordinate through

stigmergy, but this is not easy or efficient. Once fitnesses are shared, the predators have a direct incentive to collaborate, and this leads to their quickly evolving specific roles to cooperate to catch the prey. These roles are rigid enough to not require communication. The predators act like a well-trained team that knows exactly how each team member will behave in any given scenario: Usually one team member stays in place while the other two chase the prey towards it (Figure 4, Video 2).

When comparing the control experiment to predators that can communicate, predators are again found to evolve cooperation much faster (Figure 4, Video 2). Communication makes it easier for a predator to sense what the other predators are doing and act accordingly, thus reducing the cost of evolving to cooperate. Further, as expected, evolution of cooperation was strongest in predators that both communicate and share fitness rewards. While the non-communicating predators develop fixed roles, communication allows them to adapt to the current situation (cf. [15]). These predators also

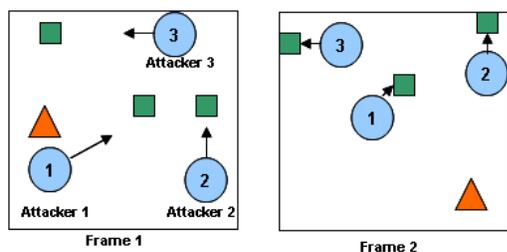


Fig. 6. Predators prefer to hunt gazelles (Experiment 5). When the net return from capturing a zebra is low, the predators prefer chasing the slower gazelles (Frame 1). They successfully catch them on their own (Frame 2).

have more flexible behaviors, i.e. they can change roles in the middle of the hunt (Figure 5, Video 3).

It is interesting to note that this result directly contradicts that of Yong and Miiikkulainen [15]. They found that adding communication to predators that already shared fitness rewards made the evolution of cooperation more difficult and less robust; role-based cooperation based on stigmergy was more efficient. However, their experiment had only a single prey, which means the stigmergic signals based on its behavior are easy to interpret. All the predators attempt to capture the same prey, therefore the movements of this prey can be easily used to infer what the other predators are doing. Four prey, in contrast, all respond to the three predators simultaneously, which makes it hard to infer which ones are being closely chased and which ones are not. Such ambiguity makes cooperative strategies harder to evolve.

With two different types of prey, whether cooperation evolves or not depends on the value of the prey relative to the difficulty of catching it. When the zebra reward was not much higher than the gazelle reward, the predators did not evolve cooperation, preferring to catch gazelles on their own (Figure 6, Video 4). The reward for catching the zebra is much higher in Experiment 6, and the predators evolve to cooperate to catch it first (with some interference from nearby gazelles; Figure 3). Once the zebra is caught, the predators return to hunting the gazelles individually (Figure 7, Video 5). Thus, increasing the net return on the zebra leads to multiple (i.e. both individual and cooperative) hunting strategies in the predators. These results therefore show how reward structure, coordination mechanism, and net return interact in producing cooperative behavior.

VI. DISCUSSION AND FUTURE WORK

The principles of reward structure, coordination mechanism, and net return can be manipulated to evolve more realistic and intelligent agents and teams for video games and robotics applications. Such opportunities constitute a most interesting long-term direction, demonstrating how computational insights into biology can benefit computational engineering as well.

The simulations also lead to insights into cooperation in nature. For instance, the predators in Experiment 3 (individual rewards with communication) evolve cooperative hunting strategies efficiently, but they do not have any fitness

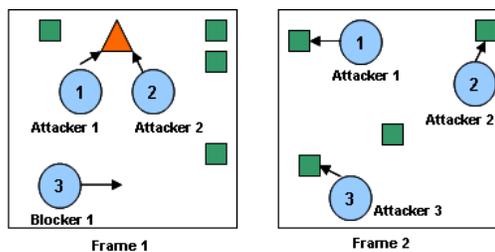


Fig. 7. Predators prefer to hunt the zebra (Experiment 6). When the net return from capturing a zebra is high enough, the predators evolve to cooperatively hunt the zebra first (Frame 1). After the zebra is captured, they return to hunting gazelles individually (Frame 2).

incentive for cooperation. Instead, they use one another to improve individual fitness. Is this real cooperation? In biological literature, a cooperater is defined as an individual who pays a cost for another individual to receive a benefit [29]. This is a useful working definition in artificial settings as well. Thus in Experiment 3, though not all the predators gain by coordinating their behaviors, it is still considered cooperation.

Interestingly, a similar situation arises in hyenas as well. Hyenas have a complex social structure that is based on hereditary social ranking [6]. Higher ranked hyenas have more offspring and are given a bigger share of the prey that is caught through cooperation. Rewards are thus not shared equally, but higher ranked hyenas use lower ranked hyenas to gain individual fitness.

The reasons for such a rank-based society are unknown; if anything, rank-based evolution should be less efficient than evolution based on actual fitness evaluations. To analyze these paradoxical behaviors, future experiments can be designed to simulate rank-based social structure in reproduction and hunting. When fitness evaluations are costly, rank-based evolution may be more efficient, even when it results in paradoxical hunting behaviors. If confirmed, this principle may lead to the construction of better evolutionary computation techniques for domains where fitness evaluation is expensive, such as intelligent agents in games.

Another future direction would be to evaluate these factors affecting cooperation in the presence of evolving prey. Such competitive coevolution between predators and prey could lead to complex behaviors in both. In the wild, hyenas have to coordinate to hunt zebras partly because the zebras themselves cooperate to defend their herd against hyenas. Simulations of evolving predators and prey can help understand such predator-prey dynamics better. They may also lead to techniques that result in more complex behaviors for artificial agents, as well as agents that adapt in response to player actions.

VII. CONCLUSION

The predator-prey domain was used in this paper to determine the reasons for complex behaviors observed in nature. The experiments confirmed that predator coordination mechanism, reward structure, and net return on prey capture

are important factors in the evolution of efficient cooperative hunting behaviors. When hyenas survive on gazelles, they do not need to cooperate. However, if the zebras are available and tasty enough, they will. These results are intuitive, but this is the first time easily replicable experiments were constructed to verify them. The same factors that were established to be important in the evolution of cooperation in this domain can be manipulated in more complex artificial environments to build interesting behaviors for game agents in the future.

ACKNOWLEDGMENTS

This research was supported in part by NSF under grants DBI-0939454, IIS-0915038, and IIS-0757479.

REFERENCES

- [1] A. Rawal, P. Rajagopalan, and R. Miikkulainen, "Constructing competitive and cooperative agent behavior using coevolution," in *IEEE Conference on Computational Intelligence and Games (CIG '10)*, August 2010.
- [2] S. Kopelman, J. M. Weber, and D. M. Messick, "Factors influencing cooperation in commons dilemmas: A review of experimental psychological research," *P. C. Stern, S. Stonich, and E. U. Weber (Eds.), The drama of the commons*, p. 113156, 2002.
- [3] K. E. Holekamp, L. Smale, R. Berg, and S. M. Cooper, "Hunting rates and hunting success in the spotted hyaena," *J. Zool.*, vol. 242, pp. 1–15, 1997.
- [4] H. Kruuk, *The spotted hyena: a study of predation and social behaviour*. Chicago, IL: University of Chicago Press, 1972.
- [5] J. R. Henschel and J. D. Skinner, "Territorial behaviour by a clan of spotted hyenas (*Crocuta crocuta*)," *Ethology*, vol. 88, pp. 223–235, 1991.
- [6] K. E. Holekamp, S. T. Sakai, and B. L. Lundrigan, "Social intelligence in the spotted hyena (*Crocuta crocuta*)," *Phil. Trans. R. Soc. B.*, vol. 362, pp. 523–528, 2007.
- [7] M. G. L. Mills, *Kalahari hyaenas: the behavioural ecology of two species*. London, UK: Unwin Hyman, 1990.
- [8] K. E. Holekamp, "Questioning the social intelligence hypothesis," *Trends in Cognitive Science*, vol. 11, pp. 65–69, 2006.
- [9] K. E. Holekamp, E. E. Boydston, and L. Smale, "Group travel in social carnivores," *On the move: how and why animals travel in groups* (eds S. Boinski and P. Garber), University of Chicago Press, Chicago, IL, pp. 587–627, 2000.
- [10] G. Miller and D. Cliff, *Co-evolution of pursuit and evasion I: Biological and game-theoretic foundations*. School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK: Technical Report CSR311, 1994.
- [11] G. N. Yannakakis and J. Hallam, "Evolving opponents for interesting interactive computer games," in *Schaal, S., Ijspeert, A., Billard, A., Vijayakumar, S., Hallam, J., Meyer, J.A. (eds.) From Animals to Animats 8: Proceedings of the 8th International Conference on Simulation of Adaptive Behavior (SAB-04)*, 2004, p. 499508.
- [12] E. Uchibe and M. Asada, "Incremental coevolution with competitive and cooperative tasks in a multirobot environment," in *Proceedings of the IEEE*, 2006, pp. 1412–1424.
- [13] K. O. Stanley and R. Miikkulainen, "Competitive coevolution through evolutionary complexification," *Journal of Artificial Intelligence Research*, vol. 21, pp. 63–100, 2004.
- [14] 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, pp. 53–60.
- [15] C. Yong and R. Miikkulainen, "Coevolution of role-based cooperation in multi-agent systems," *IEEE Transactions on Autonomous Mental Development*, 2010.
- [16] A. Czirok and T. Vicsek, "Collective behavior of interacting self-propelled particles," *Physica A*, vol. 281, pp. 17–29, 2000.
- [17] C. W. Reynolds, "Flocks, herds, and schools: A distributed behavioral model," in *Computer Graphics, 21(4) (SIGGRAPH '87 Conference Proceedings)*, 1987, pp. 25–34.
- [18] H. Seno, "A density-dependent diffusion model of shoaling of nesting fish," *Ecol. Modell.*, vol. 51, pp. 217–226, 1990.
- [19] N. Kubota, K. Shimojima, and T. Fukuda, "Virus-evolutionary genetic algorithm -coevolution of planar grid model," in *Proceedings of the Fifth IEEE International Conference on Fuzzy Systems (FUZZIEEE '96)*, 1996, pp. 8–11.
- [20] O. Roeva, T. Pencheva, S. Tzonkov, M. Arndt, B. Hitzmann, S. Kleist, G. Miksch, K. Friehs, and E. Flaschel, "Multiple model approach to modelling of escherichia coli fed-batch cultivation extracellular production of bacterial phytase," *Electronic Journal of Biotechnology*, vol. 10, no. 4, pp. 592–603, 2007.
- [21] M. Dorigo, V. Maniezzo, and A. Colomi, "Ant system: optimization by a colony of cooperating agents," *IEEE Transactions on Systems, Man, and Cybernetics-Part B*, vol. 26, no. 1, pp. 29–41, 1996.
- [22] A. Perez-Urbe, 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 (ECAL'2003)*, 2003, pp. 128–137.
- [23] M. Waibel, D. Floreano, S. Magnenat, and L. Keller, "Division of labour and colony efficiency in social insects: effects of interactions between genetic architecture, colony kin structure and rate of perturbations," in *Proceedings of the Royal Society B*, vol. 273, 2006, pp. 1815–1823.
- [24] S. Iredi, D. Merkle, and M. Middendorf, "Bi-criterion optimization with multi colony ant algorithms," in *Evolutionary Multi-Criterion Optimization, First International Conference (EMO'01)*, 2001, pp. 359–372.
- [25] L. Keller and D. Floreano, "Methods for artificial evolution of truly cooperative robots," *Biologically Inspired Systems for Intelligent Computing*, pp. 768–772, 2009.
- [26] 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.
- [27] F. Gomez and R. Miikkulainen, "Incremental evolution of complex general behavior," *Adaptive Behavior*, pp. 317–342, 1997.
- [28] D. E. Moriarty and R. Miikkulainen, "Forming neural networks through efficient and adaptive coevolution," *Evolutionary Computation*, vol. 5, no. 4, pp. 373–399, 1997.
- [29] M. A. Nowak, "Five rules for the evolution of cooperation," *Science*, vol. 314, pp. 1560–1563, 2006.