

Evolution of Complex Coordinated Behavior

Padmini Rajagopalan
Department of Computer Science
The University of Texas at Austin
Austin, TX, USA
padmini@cs.utexas.edu

Kay E. Holekamp
Department of Integrative Biology
Michigan State University
East Lansing, MI, USA
holekamp@msu.edu

Risto Miikkulainen
Department of Computer Science
The University of Texas at Austin
Austin, TX, USA
risto@cs.utexas.edu

Abstract—Cooperative tasks such as herding and hunting are common among higher animals in nature. A particularly complex example is that of mobbing by spotted hyenas. Through careful coordination, a large number of spotted hyenas can attack a group of lions and successfully steal a kill from them, even though lions are much bigger and stronger. This behavior is more complex than others that hyenas exhibit, and it appears to be heritable. How such behavioral advance can emerge in evolution is a fascinating question; it is difficult to study in nature, but computational simulations can provide insight. In simulation, hyenas initially evolved different levels of boldness, corresponding to simple behaviors such as solo attack, delayed attack, and delayed approach. These behaviors can be seen as stepping stones in constructing the more complex mobbing behavior in later generations. These results suggest a general stepping-stone-based mechanism through which complex coordinated behaviors can arise in humans and animals. This insight should prove useful in building cognitive architectures and team strategies for artificial agents in the future.

Index Terms—cooperation, neuroevolution, multi-agent systems, artificial life

I. INTRODUCTION

Animals, including humans, collaborate in nature to perform various tasks. Cooperative behaviors benefit the group by helping it achieve rewards that would not be possible for an individual alone. Predators cooperate to hunt prey that are stronger and faster than them, and prey may cooperate to defend themselves against predators [1], [2].

An example of a particularly complex cooperative behavior is the precisely coordinated attack on lions by the spotted hyenas of Eastern Africa [3]–[6]. Lions are larger and stronger than hyenas, and therefore the hyenas are naturally fearful of them and stay away from them. However, if the hyenas stumble upon a group of lions that have just made a major kill, such as a zebra or another large mammal, hyenas can gather in large numbers, overcome their fear, and coordinate a group attack that drives the lions away. This behavior is particularly interesting because it is more complex than anything else the hyenas do. It also appears to be largely genetically determined instead of learned. Therefore, it may be an example of evolution making a breakthrough in complexity, eventually leading to more flexible intelligence.

This paper aims to replicate the emergence of mobbing behavior in simulation. Neural networks are evolved computationally to control the behavior of a population of simulated hyenas, and the conditions under which the mobbing behavior

emerges analyzed. The insights from these simulations suggest how such complexity can be evolved in general.

The main finding is that individualistic traits play an important part in this process. Various emotions, such as fear of lions and affiliation towards teammates, affect the willingness of real hyenas to attack lions [6], [7]. Similarly, simulated hyenas evolve different levels of boldness and exhibit different behaviors near the lions. Some of them may attack the lions alone, which is risky but sometimes successful if enough teammates join. Others may approach to the minimal safest distance but not take the initiative in attacking. Yet others may hang back and approach only if the lions are successfully mobbed, making them likely to survive but not achieve the highest rewards. Even though these behaviors do not themselves establish a mobbing behavior, they make it likely for mobbing to happen occasionally by accident. The mobbing may then eventually be built from these simpler behaviors through mutation and crossover. They therefore serve as stepping stones in the evolution of cooperative mobbing.

These insights suggest a general approach for evolving complex behavior by creating a collection of possible stepping stones. It may be used in the future to evolve strategies in teams of robots or video game characters to overcome powerful adversaries or solve problems that involve high risk.

II. RELATED WORK

The biological data on hyenas mobbing lions is first reviewed, followed by computational simulations of this behavior, and the neuroevolution techniques used in this study.

A. Biological Background

Spotted hyenas (*Crocuta crocuta*) share their habitats and resources with other powerful predators such as lions (*Panthera leo*). Because lions are larger and stronger, in any competition between hyenas and lions, lions are expected to win. However, large numbers of hyenas will occasionally gather to “mob” lions and drive them away in order to gain or retain control over a prey carcass [4], [5], [7]. Mobbing refers to two or more hyenas coordinating a charge against one or more lions, and is very dangerous for the hyenas. In fact, lions are the leading cause of death in many hyena populations [4], [8], [9]. Consequently, hyenas can rarely displace lions from food unless the odds ratio (i.e. the ratio of hyenas to lions) is at least four to one [5].

Hyenas have predominantly instinctive behaviors that are not as plastic as those of primates. Most of these behaviors are specific to the challenges they face in everyday life and, therefore, are common to many predatory species in similar habitats. But coordinated mobbing of lions is a much more complex behavior that involves extreme risk for the hyenas and thus, it may be considered as evolutionarily complex and novel.

Lehmann et al. [6] used observational data of lion-hyena encounters of seven hyena clans from over 30 years in order to characterize each such encounter. The analysis included dimensions such as the number of hyenas present, the number of lions, whether mobbing occurred, and whether it was successful. The conclusions they reached were:

- 1) Lions and hyenas interacted more frequently at fresh kill sites than at sites with carcasses older than 24 hours. Mobbing rates were also highest at a fresh kill.
- 2) The probability of lion-hyena interaction increased with increasing prey size.
- 3) The presence of adult male lions at the kill site increased the probability of interactions but decreased the probability of successful mobbing. Male lions are more likely to initiate interaction by approaching the hyenas [10]. However, they can also better protect the kill from hyenas because they are larger and stronger than female lions [8], [11].
- 4) The probability of interaction increased as the number of hyenas present increased.

Videos of the mobbing behavior can be seen at <https://bit.ly/36AXPcA>. While the behavior itself is well characterized, it is less clear what factors determine whether it is successful, and most importantly, how it could have emerged in evolution. Computational simulations are a crucial tool in gaining insight into such questions.

B. Simulations of Mobbing Behavior

A significant body of work exists on computational modeling of cooperation in nature. For example, cooperative behavior of micro-organisms like bacteria and viruses has been modeled with genetic algorithms [12], [13]. Ant and bee colonies have been the subject of many studies involving evolutionary computation as well [14]–[16].

Lion mobbing has also been simulated in a number of studies, focusing on different issues and using various methods. In particular, the conclusions from the observational data above were replicated in computational simulations [17], [18]. The behaviors were represented by neural networks and discovered through genetic algorithms. Varying values of parameters such as the attractiveness of the kill, strength of the lions, distance to the lions, and number of hyenas were varied and found to have the expected effect on the probability that the mobbing is successful.

This prior research also demonstrated that mobbing behavior can be modulated by the interaction of two emotions: fear and affiliation. Fear dominates initially, but affiliation gradually builds up through the interactions between the hyenas, and

eventually allows them to attack the lion in a coordinated manner.

Previous computational work also characterized the effect of different communication strategies in mobbing, evolving the behaviors as a set of rules [19], [20]. The results showed that having a single leader to make all mobbing decisions for the hyena team resulted in the most effective coordination. However, emotions were not included in this study, and the single leader result has not yet been verified in real-life hyenas.

In contrast, this paper focuses on a different but fundamental question: how can the mobbing behavior arise from evolution? The challenge is to understand the innovative leap from simpler behaviors, and also how it is possible to evolve when imperfect attempts lead to elimination. The approach is to simulate the evolution of hyena behaviors and identify the stepping stones that lead to mobbing being discovered.

C. Neuroevolution of Behavior

Neural networks and evolutionary computation may be combined into a learning algorithm, neuroevolution, that can be used to solve difficult sequential decision tasks with continuous state and action spaces, and partially observable states. Each agent on the field is controlled by a neural network whose weights and topology are learned using evolutionary algorithms. Neuroevolution has previously been used to discover dynamic and intelligent behavior in autonomous agents. For example, it has been used in simulated robot soccer [21], robotic battle [22] and Ms. Pac-Man [23], [24].

NeuroEvolution of Augmenting Topologies, or NEAT [25], is a particularly appropriate technique for the present study. NEAT optimizes not only the connection weights, but also the topology of the neural networks. It includes mutations to add and delete both nodes and links. Starting from an initially minimal network, connecting inputs directly to outputs, structure is added through these mutations, thus gradually complexifying the networks and elaborating their performance. Topological innovations are protected through speciation, allowing new structure to be refined through evolution before it competes with other structures in the population. Historical markings allow the same structure to be identified in multiple genomes, making crossover of different network topologies possible and efficient.

NEAT technique was shown to be more effective than traditional neuroevolution methods that modify only the connection weights of neural networks [25]. In particular, it makes it possible to discover the appropriate recurrency in the network, resulting in sequential behaviors that can be used in robotics, game agents, and artificial life.

III. EXPERIMENTAL SETUP

The population consists of 100 neural networks that control the hyenas' behavior. During each evaluation, ten hyenas are picked at random from the population to form a team that tries to mob a lion. The world is a discrete 100×100 grid environment without any obstacles (Figure 1). The lion is added to the environment at a random location that stays fixed

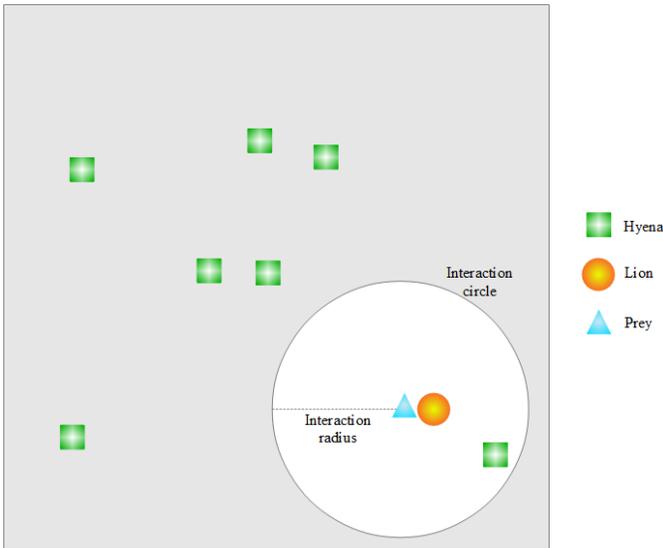


Fig. 1: The simulation environment with multiple hyenas, a lion and an interaction circle. This figure shows a 100×100 grid environment where a lion is in possession of a kill with a number of hyenas around it. The lion is placed at a random location and does not move. The hyenas are initially scattered uniformly randomly around the space but they can move east, west, north or south. If a hyena enters the interaction circle, it may be killed. If a sufficiently large number of hyenas (i.e. at least four) is within the interaction circle, they can mob the lion and drive it away from the kill. As in nature, the hyenas need to coordinate their attack precisely. This mobbing behavior is more complex than other behaviors of the species, yet it appears to be largely genetically determined. It therefore is a good testbed to study how such complex behaviors emerge in evolution.

throughout each simulation. The ten hyenas are also placed at random initial locations. In each timestep, the hyenas decide to either move towards the lion or idle in place. If they decide to move, their movement is translated into a step in one of four directions: east, west, north or south. They move one step at a time, and all the hyena agents in the world take a step simultaneously. Each simulation lasts 200 timesteps, which is sufficient for all the hyenas to reach the lion.

Around the lion there is an area bounded by an interaction circle 20 steps away. A hyena is safe at the circle and beyond it, but inside it may get killed by the lion with a certain probability. The hyenas are aware of whether there are a sufficient number of other hyenas (i.e. at least three) at or within the interaction circle in order to mob the lion. The hyena’s neural network also receives information about whether the lion has already been mobbed. In addition, there are three inputs to each hyena neural network indicating if it is currently inside the interaction circle, at the circle, or outside (Figure 2).

A successful mobbing requires four hyenas to be inside the interaction circle at the same time. Mobbing can happen either

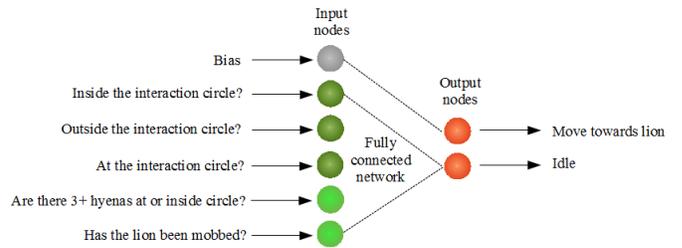


Fig. 2: The initial hyena neural network structure for the lion-mobbing task. The image shows the inputs and outputs of the neural network controlling a simulated hyena. The initial network has no hidden neurons and is fully connected, i.e. all input nodes are connected to all output nodes; hidden nodes are later added and evolved by NEAT. Each hyena receives five binary inputs in addition to the bias. Three inputs indicate whether the hyena is at the interaction circle, inside the circle, or outside. Another binary input is set to one if there are at least three other hyenas either at the interaction circle or within it. This input helps the hyena decide whether mobbing is a safe option. The last input indicates whether the lion has already been successfully mobbed. If so, the hyena can safely enter the interaction circle and partake in the kill. Based on these five inputs, each hyena network has to decide whether to approach the lion or avoid it by idling in place.

by four hyenas stepping into the circle simultaneously, or some of the hyenas stepping in earlier and surviving until a fourth hyena joins them. In nature, lions can kill hyenas easily. In order to simulate this effect, hyenas that enter the interaction circle before there were four of them there are likely to get killed with a high probability. Similarly, a mob consisting of at least four hyenas still has a small probability of dying when they enter the interaction circle together.

Any hyena participating in a successful mobbing event is given a fitness of 10000 points, while hyenas that step into the circle after the lion is successfully mobbed are rewarded 8000 points, simulating the fact that such latecomers or risk-evaders may also get part of the kill. If the lion is never mobbed during the simulation, none of the hyenas can get a fitness reward. Thus, the ideal hyena would go directly to the interaction circle and wait there until three other hyenas also gathered there, and then attack the lion. This behavior requires the other hyenas in the team to also behave the same way, thus ensuring that the attack is perfectly coordinated. This emergence of cooperation among the hyena team is particularly difficult to evolve.

Ten simulations or trials were conducted for every randomly assembled 10-hyena team, and 500 such teams were chosen (randomly, with replacement) for evaluation every generation. The population of hyena neural networks was evolved for 1000 generations in search of teams of hyenas that could successfully mob lions.

IV. CHARACTERIZATION OF FINAL BEHAVIORS

The final generation included several teams that mobbed the lion successfully. The average number of mobbing events

across 500 teams was around 8.5 for every 10 trials. This result indicates that evolution was able to discover complex cooperative behavior with precise coordination. Video of a team that successfully mobbed the lion in all 10 trials is at <https://youtu.be/GKmba9hCxic>.

However, a surprising effect was also found in these simulations; not all the hyenas seen during evolution are perfect mobbers. In fact, four different kinds of behaviors were observed. The first is the perfect mobber. The second is a hyena that is a risk-taker, and commits suicide by running straight to the interaction circle and most of the time dying immediately. The third is a risk-evader, idling outside the interaction circle until the lion has been mobbed, and then joining the feast. The fourth kind of hyena approaches the interaction circle and stays idle there, not participating in any mobbing until the lion has been successfully mobbed by other hyenas in the team, after which it joins in the fray.

These four behaviors evolved consistently across all simulations. Interestingly, each of them represents a different aspect of mobbing, even though on their own they have negative implications for the team. For instance, the suicidal hyena does get to the lion, however too many risk-taking hyenas will leave the team with not enough hyenas to form a mob, thus robbing all the team members of their fitness. In contrast, risk-evaders are the most likely to survive. However, too many risk-evaders in the team prevents mobbing from ever happening. Risk-evaders that idle at the circle have evolved to get close to the lion but not too close. On the other hand, they can fool perfect mobbers into dying because the mobbers expect the idlers to participate in the mob along with them.

However, the roles of risk-taker, risk-evader-outside-circle and risk-evader-at-circle persist in low proportions even in prolonged evolution, i.e. over 1000 generations (Figure 7). Interestingly, the three imperfect behaviors are present in large proportions of the population early on, i.e. during the first five generations (Figures 3 through 6). After that, the population starts to converge towards mobbing behaviors. It is therefore likely that risk-taking and risk-evading tendencies play a role in the discovery of the complex behavior of coordinated mobbing, as will be discussed in the next section.

V. EMERGENCE OF MOBING BEHAVIORS

In order to understand the emergence of mobbing, evolution was run from scratch five times, and each of the hyenas from the first five generations, the 10th generation and the 20th generation in each of the five trials were analyzed in a test bench. In the bench, the hyena was placed in eight different scenarios representing all the various situations in the environment in which it could find itself during evolution. The hyena's reaction to each of these situations was recorded and used to score it along five dimensions corresponding to five behaviors. The fifth behavior, "risk-evader", was a combination of risk-evader-outside-circle and risk-evader-at-circle.

- 1) **Bench Mobbing Score:** The ideal mobber moves towards the lion until it encounters the interaction circle.

It stops there and waits for three other hyenas to join it, at which point it steps into the circle to attack the lion. Each hyena was scored on how similar it was to an ideal mobber in each of the situations in the test bench.

- 2) **Bench Risk-Taker Score:** The ideal risk-taker keeps approaching the lion until it enters the interaction circle. At this point, it has a very high probability of dying and a very small probability of participating in a mobbing event. Each hyena was scored on the extent to which it was an ideal risk-taker in the test bench.
- 3) **Bench Risk-evader-outside-circle Score:** This kind of risk-evader delays its arrival at the kill site by idling somewhere outside the environment. Since the hyena does not know its distance to the lion, but only whether it is inside, outside, or at the interaction circle, the ideal risk-evader-outside-circle idles whenever it is outside the circle until the lion has been mobbed. At that point, it approaches the kill site to partake in the kill. Each hyena was scored on how similar it was to this ideal in the test bench.
- 4) **Bench Risk-evader-at-circle Score:** The ideal risk-evader-at-circle, just like the ideal mobber, moves towards the lion until it encounters the interaction circle. However, it stops there and waits for the lion to be mobbed by other hyenas, at which point it steps into the circle to participate in the feast. Each hyena was scored on how similar it was to an ideal risk-evader-at-circle in the test bench.
- 5) **Bench Risk-evader Score:** An extra dimension was added to the scoring of hyenas representing a generic risk-evader. This dimension combined the behaviors of both kinds of risk-evaders, counting moves that benefit both.

These five scores were normalized and compared for all the hyenas in the population for each of the five runs of evolution. Instead of measuring how similar to a perfect mobber, suicidal risk-taker or risk-evading latecomer a hyena was, each hyena was classified as belonging to one of these five categories based on where it had the highest score, and whether this score was higher than that in the second-highest dimension by a certain threshold (generic risk-evader and the two different kinds of risk-evaders were evaluated separately).

These results are summarized in Figures 3-6. While 256 differently-behaved individuals may exist (2^8 , with two possible decisions in each of eight different scenarios), only 40 different five-dimensional normalized scores are possible, of which around 23 existed in the initial population of neural networks with random weights. Therefore, the initial population was truly diverse.

In generation 0, risk-takers and risk-evaders existed in large numbers, while mobbers were fewer in number (Figure 3). Risk-takers increased for a while in generations 1 to 3 (Figure 4), as this behavior was easy to discover. Mobbers also increased in number, albeit more slowly, at the expense of risk-evaders. Mobbing started to grow further after generation 3 (Figure 5), and risk-evaders persisted because successful

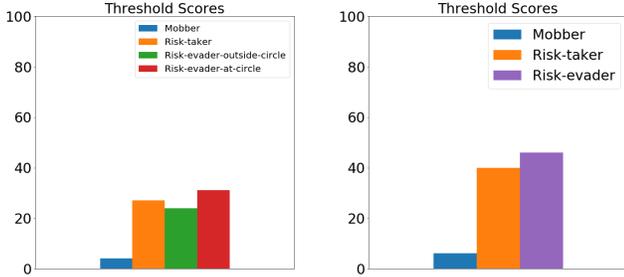


Fig. 3: **Generation 0:** The figure on the left shows the numbers of mobbers, risk-takers, risk-evaders-outside-circle and risk-evaders-at-circle, while the figure on the right shows the numbers of mobbers, risk-takers and generic risk-evaders, classified based on threshold scores. Initially, the simple behaviors of risk-taking and risk-evading of each kind were common, while the more complex behavior of mobbing was rare.

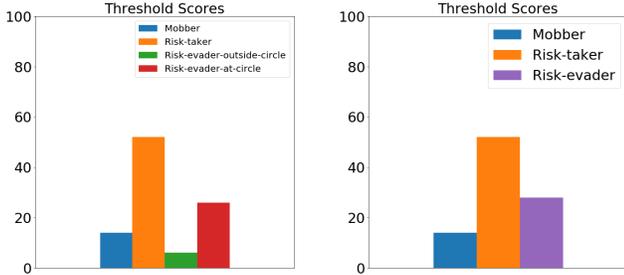


Fig. 4: **Generations 1-3:** Risk-takers became quickly more common because this behavior provides an easily discovered stepping stone that allows the hyenas to attack. Risk-evaders also persist, representing a risk-averse approach. In contrast, the number of mobbers grew more slowly.

mobbing behaviors led to risk-evading behaviors being encouraged as well. Eventually, at about generation 10, almost all the hyenas were mobbers (Figure 6). Interestingly, risk-evaders-at-circle increased slightly. This type of behavior is likely the least common because it is the most detrimental to the team: It deceives some mobbers into committing suicide. However, at this stage it may serve as a stepping stone in discovering the perfect timing for the attack.

Mobbing is the most profitable behavior in this environment, but it is difficult to discover in evolution. Even after generation 10, risk-takers and risk-evaders persisted in a very small proportion of the population and never died out entirely. But the system was strong enough to always favor mobbers after mobbing behavior was discovered. The persistence of risk-takers and risk-evaders is discussed next in Section VI.

VI. BEHAVIORAL INTERACTIONS

In order to characterize how the four behaviors interact in prolonged evolution, and how they interact in a team of hyenas, all the hyenas in the population were given mobbing, risk-taker and risk-evader scores in the actual behavioral simulation. That is, they were not placed in a test bench as in Section V, but their behaviors were observed and assessed in the actual simulation.

- 1) **Simulation Mobbing Score:** The number of times in 10 trials that a hyena participated in a mobbing event, averaged over all the times it was randomly picked for a team. The maximum possible value was 10, but even a perfect mobber would sometimes get a less-than-perfect score because in some trials it could be initially located too far away from the lion to join in the mob.
- 2) **Simulation Risk-Taker Score:** The number of times in 10 trials that a hyena stepped into the interaction circle before the lion was mobbed and died immediately, averaged over all the times it was randomly picked for a team. The maximum possible value was 10, but a hyena may have been located too far away from the lion to reach it before it was mobbed. The risk-taker score was affected also by whether there were any risk-evaders-at-circle that fooled ideal mobbers into stepping into the circle—such deaths were counted as a risk-taker.
- 3) **Simulation Risk-evader-outside-circle Score:** The number of times in 10 trials that a hyena idled for at least one time step at a grid cell outside the interaction circle, averaged over all the times it was randomly picked for a team, provided it eventually joined in the attack after the lion was mobbed. The maximum possible value was 10.
- 4) **Simulation Risk-evader-at-circle Score:** The number of times in 10 trials that a hyena was present at the interaction circle but did not participate in lion-mobbing, averaged over all the times it was randomly picked for a team, provided it joined the attack after the lion had been mobbed. The maximum possible value was 10, but again a lower score was possible if the hyena was initially located too far away to travel to the interaction circle in time for the mobbing.

A. Interaction in Prolonged Evolution

The averages of each of these scores across the population of 100 hyenas over 1000 generations of evolution are depicted in Figure 7. The four behavior scores remain more or less constant across generations and none of them die out even in such prolonged evolution. Mobbing behavior was the most common with an average score around 6, while average risk-evader-outside-circle score was next with a score around 3.75. Risk-takers and risk-evaders-at-circle were less common with scores around 1.6 and 0.09. Since these four scores were calculated and averaged separately for every hyena in every team picked randomly from the population, they may not always add up to 10 exactly for the population as a whole,

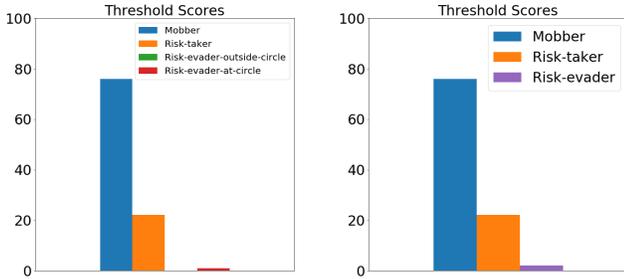


Fig. 5: **Generations 4-5:** With enough risk-takers and risk-evaders in the population (Figure 4), mobbers started to emerge as their combinations. Because of the superior fitness of the mobbers, other behaviors started to subside.

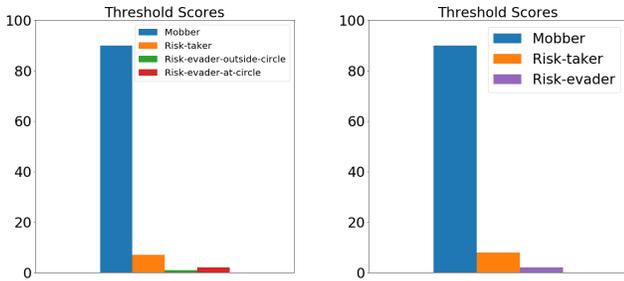


Fig. 6: **Generation 10:** Mobbers made up most of the population at this point, but other behaviors persisted as well in very small numbers. This persistence continues even in prolonged evolution as seen in Figure 7.

but they do give a clear picture of how the hyenas are split four ways into varying degrees of boldness.

B. Interaction in Teams

While an individual hyena can occasionally exhibit different behaviors, each hyena usually had one dominant behavior, i.e. an individual degree of boldness. It is therefore possible to study what makes a successful team. Hyenas from the final generation were put together into a team in four different ways to check their mobbing prowess:

- 1) The team was composed of 10 clones of the hyena that had the top mobbing score. Ten trials were conducted for this team, where each trial had the lion and the hyena clones in random starting positions in the grid world.

Result: The team successfully mobbed the lion 9 times. In some cases, some of the hyenas exhibited risk-evader-at-circle tendencies, resulting in their teammates getting killed, which is a possible outcome in nature.
- 2) The team was composed of six clones of the hyena with the top mobbing score, two clones of the one with the top risk-evader-outside-circle score, and one each of the

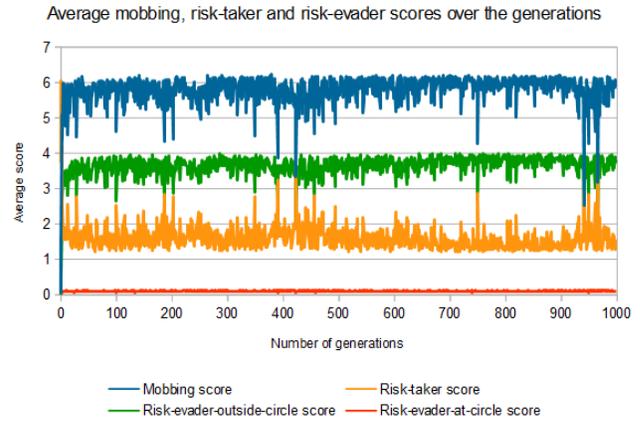


Fig. 7: Average mobbing, risk-taker, risk-evader-outside-circle and risk-evader-at-circle scores (in y) over 1000 generations (in x). Mobbing was discovered within the first 10 generations as shown in Figures 3-6, and stays at the same level throughout. On average, each hyena successfully mobs the lion six out of 10 times. Average risk-taker scores persist at a low level of 1.6. Risk-evader-outside-circle is less common than mobbing, at 3.75 on average. Average risk-evader-at-circle score is very low at around 0.09. These suboptimal behaviors act as stepping stones in discovering mobbing in prolonged evolution; they likely play a role in maintaining it and keeping it robust.

hyenas with top risk-evader-at-circle score and the top risk-taker score. This experiment sought to emulate the actual probability of finding different proportions of each kind of hyena behavior when choosing hyenas at random from the population.

Result: The successful team mobbed the lion 8 times, which is close to the value seen on average in the 500 randomly chosen teams in the final generation. This result shows that a team created in this fashion is a typical team. It tolerates the other behaviors relatively well, but it is not always successful in mobbing, especially when risk-evaders-at-circle are present.

- 3) The team had the six hyenas with the highest mobbing scores (instead of clones of the top mobber), the two hyenas with the highest risk-evader-outside-circle scores, the hyena with the top risk-evader-at-circle score, and the one with the highest risk-taker score. This team is more heterogeneous than the team created with clones, and their average mobbing score is well below that team's. It is therefore expected to be slightly less adept at mobbing the lion.

Result: This heterogeneous team successfully mobbed the lion 6 times. Many times in the 10 trials, several hyenas committed suicide (either by themselves or because of risk-evaders-at-circle) and there weren't enough hyenas left at the circle for mobbing.
- 4) This team was composed of those hyenas that scored about equally in each of the four behaviors. This team

is also a very heterogeneous team, but is not expected to perform any mobbing.

Result: This team mobbed the lion successfully only once in the 10 trials. The hyenas in this team do not cooperate well, but they are still the products of 1000 generations of evolution targeted towards successful mobbing, and therefore they sometimes mob successfully as well.

These results suggests that although the mobbing behavior is relatively robust, more heterogeneous teams are not as effective. The implications of these results will be discussed in the next section.

VII. DISCUSSION

As was described above, the hyenas in the population exhibit different levels of bold behavior (mobbing, suicide) and shy behavior (risk-evaders outside circle, risk-evaders at circle). While some of the behaviors may seem detrimental to the success of a mobbing event, all these behaviors exist in large portions of the population in early generations, and persist in smaller proportions throughout evolution even after the discovery of successful mobbing. This led to the hypothesis that the behaviors, while destructive to mobbing success at the end of evolution, may contribute to its emergence during evolution: they may act as stepping stones that eventually lead to the discovery of mobbing. This is not a new idea; Meyerson et al. [26] found that carefully chosen or learned behavior characterizations can act as local optima that can channel evolution towards successful behaviors in complex domains. Similarly, Woolley and Stanley [27] discovered that the intermediate stepping stones during evolution often do not resemble the final image in the Picbreeder online interactive image evolution service [28].

Specifically in the lion-mobbing case, risk-taking hyenas that go straight to the interaction circle and commit suicide may be needed in order to discover that the interaction circle exists and that any hyena that steps into it is killed. This discovery may lead to the emergence of hyenas that either do not approach the interaction circle quickly or idle once they arrive at the circle in order to avoid being killed. But risk-taking hyenas are still necessary for the mob to eventually enter the interaction circle in order to attack the lion and gain fitness. Enough suicidal hyenas may eventually discover mobbing by accident.

Similarly, risk-evaders outside the circle ensure that they arrive at the circle after all the risk-taking hyenas are killed off. In many cases, the lion has already been mobbed, and the hyena can now gain some fitness for joining in the attack late. While the reward for being a latecomer is not as high as that given to a mobber, the risk of dying is also much lower than that faced by a mobber. They may act as a counterbalance to risk-takers, eventually leading to hyenas that wait until it is safe to enter.

Some risk-evaders are more efficient than others, waiting at the interaction circle itself for the lion to be mobbed. In this way, they make sure that they never lose out on the risk-evader

reward just because the simulation time ran out. They are also in a perfect position to coordinate their action: All it takes is a mutation that causes them to step in when there is enough of them at the circle. This behavior can thus be a stepping stone for discovering proper timing for the mobbing.

In this manner, the three unsuccessful behaviors can each be seen as a constituent of the complex mobbing behavior. Once they are discovered, evolution is in a position to cross them over to create successful mobbing. Similarly, they stay in the population because mobbing is still fragile, and a deleterious mutation or crossover results in one of these constituent behaviors. A dynamic equilibrium exists between them, resulting in proportions of behaviors observed in the simulations. Moreover, such diverse behaviors serve to make mobbing more robust. The mobbers cannot count on their teammates always being perfect mobbers, and they evolve to be more flexible. It is also possible that the communicative behaviors that the hyenas in nature engage in before mobbing serve a useful purpose: they make the hyenas more uniformly bold and coordinated in their responses, and therefore increase the chances of success.

These observations therefore suggest how behavioral stepping stones are discovered and utilized to discover something as complex as mobbing, and how such behavior is maintained over prolonged evolution. These conclusions can be seen as a general approach for constructing complex behaviors in science and engineering.

The computational model developed in this work to study lion-hyena interactions implemented behaviors through neural networks. In order to replicate mobbing behaviors from nature, various parameters such as mobbing rewards and probability of injury or death had to be set carefully and systematically. The resulting successful settings suggest principles that make such behaviors possible. Surprisingly, even though the neural networks are deterministic, the hyenas exhibited multiple behaviors resulting in a variety of outcomes. The expectation was that if they evolve a good mobbing strategy, they should always use it. If the net return from mobbing is very low, they should evolve to never mob the lions. Instead, a variety of behaviors evolved and persisted, demonstrating a dynamic equilibrium that can serve as a foundation for adaptation in a changing world.

While emotions were not explicitly modeled in this study, the four behaviors can be seen to express different levels of fear and affiliation. The role of emotions, as well the importance of individualistic traits in lion-mobbing has not been studied before. It is not clear exactly what information emotions provide to the hyenas and how they regulate behavior. One effect identified in this study is that they may serve as a regularizer, making the responses of the hyenas more uniform so that their mobbing attempts can be more successful. It is difficult to simulate emotion inputs to hyena neural networks. Similarly, the different roles of individual hyenas are also hard to replicate in simulation when not much about these roles has been observed in nature. Extending the current work with emotions is therefore a most interesting

direction of future work.

Including communication between hyenas is another interesting direction. Communication may play a role in precise coordination as well as in establishing uniform emotions. If successful, such a study may result in conclusions about the role of emotions and communication in general in constructing complex cooperative behavior.

VIII. CONCLUSION

Experiments on evolving mobbing behavior on simulated hyenas demonstrated that behaviors with different levels of boldness emerge reliably in early generations and persist even in prolonged evolution. While these behaviors are simple and even detrimental in the mobbing situation by themselves, they can be seen as stepping stones in making the innovative leap to mobbing. Encouraging such imperfect individual traits constitutes a general approach for evolving complex innovative behavior, and may be useful in constructing cooperative teams of heterogeneous robots, game agents, and artificial life simulations in the future.

ACKNOWLEDGMENT

This research was supported in part by National Science Foundation Bio/computation Evolution in Action CONSORTIUM (BEACON) Science and Technology Center grant DBI-0939454, and in part by National Institutes of Health grant 1R01-GM105042.

REFERENCES

- [1] C. Packer and L. Rutten, "The evolution of cooperative hunting," *American Naturalist*, pp. 159–198, 1988.
- [2] B. Penzhorn, "A long-term study of social organisation and behaviour of Cape Mountain zebras *Equus zebra zebra*," *Zeitschrift für Tierpsychologie*, vol. 64, no. 2, pp. 97–146, 1984.
- [3] K. E. Holekamp, J. E. Smith, C. C. Strelhoff, R. C. V. Horn, and H. E. Watts, "Society, demography and genetic structure in the spotted hyena," *Molec. Ecol.*, vol. 21, no. 3, pp. 613–632, 2012.
- [4] M. Trinkel and G. Kastberger, "Competitive interactions between spotted hyenas and lions in the Etosha National Park, Namibia," *African Journal of Ecology*, vol. 43, no. 3, pp. 220–224, 2005.
- [5] H. Kruuk, "The spotted hyena: a study of predation and social behavior," *Wildlife Behavior and Ecology*, pp. 315–325, 1972.
- [6] K. D. S. Lehmann, T. M. Montgomery, S. M. MacLachlan, J. M. Parker, O. S. Spagnuolo, K. J. VandeWetering, P. S. Bills, and K. E. Holekamp, "Lions, hyenas and mobs (oh my!)," *Current Zoology*, 2016.
- [7] H. Watts and K. E. Holekamp, "Interspecific competition influences reproduction in spotted hyenas," *Journal of Zoology*, vol. 276, no. 4, pp. 402–410, 2008.
- [8] S. Cooper, "Optimal hunting group size: the need for lions to defend their kills against loss to spotted hyenas," *African Journal of Ecology*, vol. 29, no. 2, pp. 130–136, 1991.
- [9] H. Hofer and M. East, "Population dynamics, population size, and the commuting system of Serengeti spotted hyenas," *Serengeti II: dynamics, management, and conservation of an ecosystem*, vol. 2, p. 332, 1995.
- [10] J. P. Elliott and I. M. Cowan, "Territoriality, density, and prey of the lion in Ngorongoro Crater, Tanzania," *Canadian Journal of Zoology*, vol. 56, no. 8, pp. 1726–1734, 1978.
- [11] B. M. Kissui and C. Packer, "Top-down population regulation of a top predator: lions in the Ngorongoro Crater," *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 271, no. 1550, pp. 1867–1874, 2004.
- [12] N. Kubota, K. Shimojima, and T. Fukuda, "Virus-evolutionary genetic algorithm-coevolution of planar grid model," in *Fuzzy Systems, 1996., Proceedings of the Fifth IEEE International Conference on*, vol. 1. IEEE, 1996, pp. 232–238.
- [13] 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.
- [14] M. Dorigo, V. Maniezzo, and A. Colomi, "Ant system: optimization by a colony of cooperating agents," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 26, no. 1, pp. 29–41, 1996.
- [15] 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 *Advances in Artificial Life*. Springer, 2003, pp. 128–137.
- [16] 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," *Proceedings of the Royal Society B: Biological Sciences*, vol. 273, no. 1595, pp. 1815–1823, 2006.
- [17] P. Rajagopalan, "The evolution of coordinated cooperative behaviors," Ph.D. dissertation, Department of Computer Science, University of Texas at Austin, Austin, TX, August 2016. [Online]. Available: <http://nn.cs.utexas.edu/?rajagopalan:thesis>
- [18] P. Rajagopalan, K. E. Holekamp, and R. Miikkulainen, "Factors that affect the evolution of complex cooperative behavior," in *The 2019 Conference on Artificial Life (ALIFE)*. MIT Press, 2019, pp. 333–340.
- [19] M. Solomon, T. Soule, and R. B. Heckendorn, "A comparison of a communication strategies in cooperative learning," in *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference*. ACM, 2012, pp. 153–160.
- [20] J. Fairey and T. Soule, "Evolution of communication and cooperation," in *Proceedings of the 2014 conference on Genetic and evolutionary computation*. ACM, 2014, pp. 169–176.
- [21] S. Whiteson, N. Kohl, R. Miikkulainen, and P. Stone, "Evolving soccer keepaway players through task decomposition," *Machine Learning*, vol. 59, no. 1-2, pp. 5–30, 2005.
- [22] K. O. Stanley and R. Miikkulainen, "Competitive coevolution through evolutionary complexification," *J. Artif. Intell. Res.*, vol. 21, pp. 63–100, 2004.
- [23] P. Burrow and S. M. Lucas, "Evolution versus temporal difference learning for learning to play Ms. Pac-Man," in *Proceedings of the IEEE Symposium on Computational Intelligence and Games (CIG 2009)*, 2009, pp. 53–60.
- [24] J. Schrum and R. Miikkulainen, "Evolving multimodal behavior with modular neural networks in Ms. Pac-Man," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2014)*. ACM, 2014, pp. 325–332.
- [25] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary Computation*, vol. 10, no. 2, pp. 99–127, 2002.
- [26] E. Meyerson, J. Lehman, and R. Miikkulainen, "Learning behavior characterizations for novelty search," in *Proceedings of the Genetic and Evolutionary Computation Conference 2016*. ACM, 2016, pp. 149–156.
- [27] B. G. Woolley and K. O. Stanley, "On the deleterious effects of a priori objectives on evolution and representation," in *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. ACM, 2011, pp. 957–964.
- [28] J. Secretan, N. Beato, D. B. D'Ambrosio, A. Rodriguez, A. Campbell, J. T. Folsom-Kovarik, and K. O. Stanley, "Picbreeder: A case study in collaborative evolutionary exploration of design space," *Evolutionary Computation*, vol. 19, no. 3, pp. 373–403, 2011.