

A Comparison of a Communication Strategies in Cooperative Learning

Michael Solomon
Department of Computer
Science
University of Idaho
83844-1010
solo4857@vandals.
uidaho.edu

Robert B. Heckendorn
Department of Computer
Science
University of Idaho
83844-1010
heckendo@cs.
uidaho.edu

Terence Soule
Department of Computer
Science
University of Idaho
83844-1010
tsoule@cs.uidaho.edu

ABSTRACT

In the wild, groups of hyenas are often observed to cooperate in driving lions away from a kill in order to claim it for themselves. Because lions are much larger and can easily injure or kill a single hyena, this is a difficult, high risk/high reward behavior requiring complex cooperation by the hyenas. The hyenas depend on communication to coordinate their attack. In this paper, we attempt to evolve hyena behaviors that successfully drive-off simulated lions. We are particularly interested in how the type of communication influences the evolution of successful strategies. Several forms of communication are tested including two inspired by hyena behavior. The first is a generalized vocalization or “call” that can be either local or global and used by any hyena. The second mimics the recognition of a special hyena, which we refer to as the “flag-bearer”. Our results show that the presence of a flag-bearer leads to the evolution of significantly more effective coordination than either purely local or name based communication. These results suggest that there may be a “sweet spot” between too little information, which makes coordination difficult, and too much information which makes both evolutionary learning and coordination difficult.

Categories and Subject Descriptors

I.2.9 [Robotics]: Autonomous vehicles

Keywords

Cooperation, communication, multi-agent

1. INTRODUCTION

To successfully solve difficult problems requiring cooperation between multiple agents, agents must be able to effectively coordinate their actions. Communication is a critical component of successful coordination. Fundamentally, communication is the exchange of information. An effective

method of communication must make sure that the right information is exchanged between the right individuals at the right time. This introduces a problem of information filtering, either the sender or the receiver must decide which information is most useful to send and when to send it. Too little information impedes cooperation and, especially in the context of evolutionary learning, too much information might make learning more difficult by introducing a signal-to-noise problem.

Evolutionary methods have proven effective at evolving behaviors exhibiting coordination and cooperation. However, research on comparative communication strategies among agents has been surprisingly limited, focusing on only a few, standard communication forms. To date, most research has focused on either fully distributed swarms in which members can communicate locally, or fully centralized systems in which a central controller coercively directs all of the agents (see [4, 13, 12] for summaries of recent results with multi-agent swarms). In both models, communication plays a critical role and the range and density of the communication network has a significant effect on the success of the evolutionary process. In contrast, many different types of communication are observed in nature, where evolution has evolved specific types of communication to meet specific problem requirements.

In this paper, we examine how the amount of information and range of communication effects the evolution of cooperation. We focus on two very different forms of communication inspired by African spotted hyenas (*Crocuta crocuta*). The first is a generalized vocalization or “call” which can be either local or global and can be used by any hyena. The second mimics the recognition of a special hyena, which we refer to as the “flag-bearer”. The flag-bearer is heard by the entire group, but cannot coerce agent actions. Our results show that the presence of a flag-bearer leads to the evolution of significantly more effective cooperation and coordination than either the local or global call based communication. These results suggest that there may be a “sweet spot” between too little information, which makes coordination difficult, and too much information which makes both evolutionary learning and coordination difficult thereby causing the evolution of successful solutions to take significantly longer or to get stuck in local optima. Interestingly, a similar relationship has been suggested between the amount and quality of information available early in the learning process and the effectiveness of the learning process [6].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '12, July 7-11, 2012, Philadelphia, Pennsylvania, USA.

Copyright 2012 ACM 978-1-4503-1177-9/12/07 ...\$10.00.

2. BACKGROUND

Often communication is interpreted as implying *intent* by the sender. For example, if agent A learns the location of agent B because agent B intentionally broadcast its location, that is typically considered communication. In contrast if agent A learns the location of agent B through a passive medium, e.g. agent A sees agent B, that is often considered *sensing* as distinct from communication. However, this distinction is often unclear. If every agent has an RFID chip or a transponder that continually broadcasts its location is it communicating its location, are other agents sensing its location, or both? Further, from an information flow perspective this distinction may be irrelevant, one agent is learning information about another agent regardless of whether the exchange was intentional or not.

We explore how the type and quantity of *information* that one agent receives regarding another agent affects the evolution of successful cooperation. Because we are primarily interested in information exchange, we choose not to distinguish between sensing and communication and we will refer to both as communication.

2.1 Background: Multi-Agent Communication

Early research suggests that for evolved cooperation name based approaches to communication are more successful than deictic approaches, particularly for teams of heterogeneous agents [10]. In name based communication agents receive information from (or about) other agents by name: “Joe is at location x,y ”, “Agent 7 is attacking”, etc. In contrast, in deictic communication agents receive information based on relative roles or positions that can change over time: “the closest agent is at location x,y ”, “the highest agent is attacking”, etc.

Despite the early promise of name based communication, more recent research has primarily focused on local communication or centralized communication. In systems using local communication, such as fully distributed systems, agents only exchange information with their nearest neighbors. This is a form of deictic communication, as the communication links change with location. These systems have produced a number of impressive results (see, for example [7, 9, 12, 13]). However, there also is growing empirical and theoretical evidence that for many problems centralized communication and control can significantly improve performance (see for example [11, 8, 2]). Unfortunately, in most centralized communication systems the central agent forms both a bottleneck and weak point largely nullifying the advantages in robustness and scalability of fully distributed systems.

The shift away from name based communication towards either local and distributed communication or centralized communication is largely due to the increasing size of the teams being studied. In a named based system any agent has the potential to communicate with any other agent, thus the number of possible interactions grows exponentially with the number of participating agents. Learning which interactions are beneficial appears to become increasingly infeasible as team size increases, although this has not been rigorously tested.

Here we compare local communication, name based communication, and two novel communication strategies based on the behavior of spotted hyenas. The first novel strategy allows one agent to learn how many other agents are in a

particular state, in this case the number of “calling” or “vocalizing” agents. The second strategy denotes one agent as a “flag-bearer”, all of the other agents can immediately hear when the flag-bearer calls and then become aware of its location. However, the flag-bearer does not have any coercive capabilities typical associated with a centralized control and communication strategy. Further, because the flag-bearer can only broadcast its own location, it does not create a communication bottleneck.

2.2 Background: Hyena-Lion Interactions

Hyenas and lions exhibit a number of complex interactions involving multiple members of both species [15, 3, 14]. We are modeling hyenas cooperative attempts to drive-off one or more lions from a kill. In the wild, a single lion can easily defeat several hyenas and often several lions will defend a kill. So, safely taking a kill requires a coordinated effort by a large number of hyenas.

In the process of taking a kill the hyenas exhibit considerable vocalization, which is assumed to aid the cooperation and coordination required to drive off the lions. Our “number calling” strategy (described below) reflects the hyenas ability to estimate the number of participating hyenas based on the overall volume of vocalization. Rank is also known to play an important role in the cooperative process for the hyenas, with lower rank members often being more active participants and taking higher risks in approaching the lions [15]. Our flag-bearer is a simplified approach that can be interpreted as assigning a rank to one of the hyenas, although given the more aggressive behavior of lower ranked hyenas in the wild it is unclear whether our flag-bearer more accurately represents a high or low ranked individual.

Overall the process of hyenas stealing a lion kill is interesting because it represents a difficult, high risk/high reward, cooperative behavior. Individual lions can easily injure or kill individual hyenas and are fairly aggressive in defending their kills. Thus, the hyenas must be very careful and well coordinated in attempting to drive the lions away.

3. MODEL

Our model is designed to reflect the basic hyena-lion interaction. In a 2 dimensional world with real valued coordinates, lions are guarding a kill (e.g. a zebra) and hyenas are attempting to drive the lions away without undue risk to themselves. The lions’ behavior is fixed, whereas the hyenas’ behaviors evolve. Each hyena evolves its own, potentially unique, behavior, so the hyena clan is heterogeneous.

3.1 Lion Behavior

The lions’ behaviors are fixed. Initially the lions are placed randomly within 1 unit of the kill, which we’ll refer to as a zebra. Lions do not move unless forced out of position by being outnumbered by more than 3 to 1, a ratio suggested by field studies. Specifically, if the hyenas within 5 units of a lion outnumber the lions within that range by more than 3 to 1 then the lion moves directly away from the nearest hyena. E.g., for a given lion L if there is one other lion (for a total of 2 lions) and at least 7 hyenas within 5 units of lion L then lion L will move directly away from the nearest hyena which typically will force it away from the zebra. If a hyena ends its move within 3 units of a lion then that hyena risks injury by the lion and is penalized (see Section 4 below).

This represents the potential of being injured by getting too close to a lion.

3.2 Hyena Behavior

Hyenas begin at random locations. All but one hyena start at least 8 units from the zebra, i.e. outside of the range in which they can sense the zebra (see Figure 1). One hyena, but not the flag-bearer if one is used, starts within sensing range, 8 units, of the zebra. Each hyena’s behavior is determined by that hyena’s evolved control structure. Each control structure is a vector expression tree representing a vector function. Inputs to the function (e.g. leaves of the vector expression tree) represent environmental factors the hyena can “sense” (see Table 1). For example, hyenas can “sense” the distance and direction to the nearest other hyena and to the nearest other hyena that is “calling”. The output of the vector function (also a vector) is the hyena’s move for that timestep.

Note that the hyenas can only move at most 1.0 units in a single timestep. If the returned movement vector is longer than 1.0, the hyena moves a distance of 1.0 in the given direction. Each hyena’s vector function is re-evaluated every timestep to determine the hyena’s movement in that timestep. The functions used in the hyenas’ control structure are given in Table 2. Each function takes 2 to 4 vector inputs and returns a vector. Hyenas within 8 units of the zebra are within the “calling area” and call automatically.

As noted in Table 1 some of the input vectors have a maximum range. For example, hyenas can only sense the zebra if they are within 8 units of it and lions if they are within 5. If objects are out of range (e.g. a hyena is more than 8 units from the zebra or more than 10 units from the next nearest hyena) then that input vector is set to the zero vector: zero angle and zero magnitude.

Two features of our model make this problem particularly difficult. First, because the hyenas start each trial at random locations and distances from the lions they cannot simply charge the lions. If they did then they would arrive at different times and the early arriving hyenas would suffer large penalties from being near the lions, at least until enough hyenas converged to drive the lions away. Second, if the hyenas surround the lions the lions will simply “bounce” back and forth between the hyenas inflicting considerable damage. Thus, to be maximally successful the hyenas must coordinate both the timing of their rush, to make sure sufficient numbers of hyenas rush at the same time, and the direction of the rush, to make sure they are not simply pushing the lions into other hyenas.

3.2.1 Hyena Communication - Calling Behavior

Hyenas can “hear” calling hyenas at unlimited distance, i.e. we assume that all hyenas begin and remain within hearing range of each other. A hyena can sense the zebra when it is within 8 units of the zebra (see Table 1). If a hyena is within sensing range of the zebra then that hyena automatically begins “calling”. This is a fixed behavior. How the other hyenas can interpret the calling hyena is varied across the trials and represents the major communication variable being tested. Several possible inputs based on calling are tested, each allowing a different form of communication:

- Nearest hyena - hyenas can “hear” the nearest calling hyena. This is represented as an available input vector




Figure 1: Example layout, with hyenas, lions, and key areas labeled. There are 14 hyenas (smallest circles), two lions, and one zebra kill. Hyenas can only “see” the zebra from within the 8 unit calling area. Hyenas within this area call automatically. In each simulation one hyena (randomly selected, but not the flag-bearer) starts within the calling circle, all other hyenas start outside of the calling circle and within the starting area, i.e. close enough to hear a calling hyena, but too far away to see the zebra. The circles around the lions in the figure are the area in which they attack hyenas.

that points from the active hyena to the nearest calling hyena. This is local, deictic communication.

- All hyenas - hyenas can “hear” any calling hyena. Effectively, each hyena’s call is a broadcast signal that can be received by any interested hyena. This is represented as 14 available input vectors, one for each hyena (e.g. $calling_1$, $calling_2$, etc.), the input vector points towards the numbered hyena if it is calling and is zero otherwise. This is high density, name based communication.
- Flag-bearer - there is one unique hyena, the flag-bearer, that all other hyenas can hear as soon as it starts calling. This is represented by an available input vector ($calling_{flag}$) that points towards the flag-bearer if it is calling and is zero otherwise. This is somewhat similar to a centralized system, but the flag-bearer has no coercive ability and because it does not transmit messages between other hyenas it does not create a communication bottleneck. Arbitrarily the first hyena in the clan is the flag-bearer.
- Number calling - hyenas know how many other hyenas are calling. This is represented by an available input vector whose magnitude is the number of calling hyenas and whose direction is fixed.
- Landmark - all hyenas know the location of (have a vector to) a fixed landmark (e.g. a stump) located

somewhere within the calling radius - it is placed randomly in each trial. This is used, as discussed below, as a comparison to the flag-bearer to determine whether the advantage of the flag-bearer is simply that it represents a piece of shared information.

It is important to note that while the act of calling is fixed (hyenas always call if within 8 units of the zebra) the hyenas’ response to hearing another hyena calling is purely evolved: hyenas may evolve to ignore all calling, approach callers, avoid callers, etc. depending on how they evolve to use, or not use, the available inputs.

4. FITNESS

For the evolutionary step each hyena and each clan of hyenas is assigned a fitness. At each timestep a hyena is given a fitness based on its distance (d) to the zebra:

$$f_{zebra}(d) = \frac{1}{1+d} \quad (1)$$

and its distance d to *each* lion:

$$f_{lion}(d) = \sum_{\text{for each lion}} \begin{cases} 3(d-3) & \text{if } d < 3, \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

I.e. if the hyena is within 3 units of a lion it **loses** fitness proportional to how close it is to the lion. The total fitness of a hyena is the sum of its fitness at each timestep. So, the best fitness is achieved by quickly getting close to the zebra without getting too close to the lion. The fitness of a hyena clan is simply the sum of the fitnesses of all members of the clan.

5. EXPERIMENTAL PARAMETERS

As noted previously, each hyena has its own evolved control structure represented by a vector expression tree. These trees evolve following a generational, island model algorithm [5] with a population size of 100. In the island model there is a single population of individuals, each of the individuals actually represents a clan of 14 hyenas.

During the selection step, tournament selection is applied to each hyena role. I.e. a tournament is carried out between the hyenas in the role defined by position number 1, then a tournament is carried out between the hyenas in the number 2 position, etc. until a new “parent” clan is created of the winning hyenas. The process is repeated to create a second parent clan. These two clans then undergo crossover to create two offspring clans. Crossover uses the 90/10 rule [1] and is applied between hyenas in the same position: the two number 1 hyenas cross, the two number 2 hyenas cross, etc.

After crossover each of the hyenas in the two new offspring clans undergoes mutation. Each node in the expression tree has a 1% chance of being mutated. Internal nodes are mutated into other nodes with the same arity. If a leaf node representing a constant is selected for mutation the constant value is mutated by a random real value selected uniformly in the range -2 to 2.

Each hyena clan is evaluated by testing it 5 times. In each test the hyenas start in random positions with one hyena (not the flag-bearer) within sensing range of the zebra (see Figure 1). Each evaluation lasts for 100 timesteps. At the

Table 3: Summary of the GP parameters.

Population Size	100
Clan/Team Size	14
Crossover Rate	100%
Mutation Rate	1%
Selection	Tournament (Size 3)
Trials	40

end of each trial the best clan is retested 2000 times and the final fitness is based on the average performance over those 2000 tests. All results are the average of 40 independent trials. The GP parameters are summarized in Table 3.

6. RESULTS

Table 4 shows the average fitness of the best clans across all 40 trials. The true or false values in columns 2 through 5 represent whether a particular communication type is available to the hyenas. E.g. in the experiment represented in row 2 hyenas know the number of calling hyenas, but there is no flag-bearer, no landmark, and they can only hear the nearest calling hyena, not all calling hyenas. Column 6 gives the average fitness (and standard deviation) of the best hyena clans from each of the forty trials. The last nine columns show the results of a two-tailed Student’s t-test between each pairs of conditions, *’s represent significant differences at the $P = 0.01$ level, -’s are non-significant. E.g. The difference in average fitness between rows 3 and 1 is significant, but not the difference between rows 2 and 1. $P = 0.01$ is chosen to be conservative because of the multiple tests.

Perhaps most surprisingly, the data in Table 4 suggests that just knowing the number of calling hyenas is not beneficial (e.g. row 1 versus row 2). As noted previously to effectively drive off the lions without injury the hyenas must coordinate their charge so that enough hyenas are present to drive off the lions without injury. This is a numbers game. We had hypothesized that knowing the number of calling hyenas, i.e. the number of hyenas within the calling range of the zebra (Figure 1) would be useful, but the results do not support this conclusion.

In contrast, having a flag-bearer is highly beneficial (rows 5 to 8), unless hear-all is enabled (rows 9 to 10). Observation of the hyenas’ behavior suggests that with a flag-bearer the hyenas usually evolve the same general behavior, which consists of four rules:

1. all hyenas except the flag-bearer remain stationary until the flag-bearer begins calling
2. the flag bearer begins by moving towards the nearest calling hyena and then stops (as noted previously at least one hyena always begins within sensing range of the zebra)
3. once the flag-bearer begins calling all other hyenas converge on it
4. once a *sufficient* number of hyenas have converged on the flag-bearer, all of the hyenas move to drive away the lions

This pattern of behavior meets the timing and positioning requirements described previously for a successful attack.

Table 1: List of allowed inputs to the hyenas’ control structures. Each of these can be a leaf node in the evolved vector expression trees that are used to control the hyenas’ movements. These nodes represent what the hyenas can “sense”.

Name	Description	Maximum Range
Nearest Hyena	Vector to the nearest hyena	10
Nearest Calling	Vector to the nearest calling hyena	No maximum
All Calling	Separate vector for each calling hyena	No maximum
Flag-bearer	Vector to the flag-bear, if its calling	No maximum
Nearest Lion	Vector to the nearest lion	5
Zebra	Vector to the zebra	8
North	North, magnitude 1	No maximum
Num. Calling	Vector whose magnitude is the number of currently calling hyenas and whose direction is always north	No maximum
Random	Vector randomized every time step	No maximum
Last Move	Vector used last time step	No maximum
Constant	Vector randomized exactly once	No maximum
Mirror Nearest	Vector used last time step by the nearest hyena	10

Table 2: List of functions used in the hyenas’ control structure.

Function	Description	Number of Inputs
Sum	Sums 2 input vectors	2
Invert	Inverts a vector	1
LessThanMagnitude	Compares the <i>magnitude</i> of the first 2 input vectors. Returns the 3rd vector if the 1st vector is smaller else returns the 4th vector	4
LessThanClockwise	Compares the <i>direction</i> of the first 2 input vectors. Returns the 3rd vector if the 1st vector is smaller, otherwise returns the 4th vector	4
VectorZero	If the first input has 0 for direction and magnitude, return the 2nd vector, otherwise return the 3rd.	3

The hyenas don’t attack until sufficient numbers have converged and they usually all attack from the same direction. In one instance, we observed hyenas attacking in two groups from roughly opposing sides of the lions, which still successfully forced the lions away from the zebra without inflicting additional damage.

Although this pattern of behavior appears frequently when a flag-bearer is used, there are still several features of the evolved strategy that are unclear. It is unclear exactly how the hyenas determine that a sufficient number have converged to drive away the lions as the data shows that number calling is not contributing. Our current hypothesis is that the hyenas use the lions themselves as the determinant. That is they ‘push’ up against the lions and only advance as when the lions begin to move away - an indication that sufficient hyenas are present. It is also unclear how the hyenas guide their approach to drive off the lions. In some cases it appears that they approach the zebra and drive off the lions as a side-effect, but in other cases they appear to focus on the lions, briefly moving past the zebra to push the lions away before returning to the zebra.

The data also shows that having a common landmark as a reference point is beneficial (rows 1 and 2 versus rows 3 and 4). However, this benefit is not as great as having a flag-bearer and is not additive with having a flag-bearer. This implies that while there is some benefit to having a

static reference (the landmark) this benefit is superseded by having an “intelligent” reference, i.e. the flag-bearer.

We also observed an interesting intermediate behavior in the earlier generations of many of the runs. In this intermediate behavior the hyenas would approach the zebra, apparently drawn by a calling hyena, but would stop as soon as they were inside the calling radius. This allowed the hyenas to get closer to the zebra, but also kept them a safe distance away from the lions. It was only much later in the evolutionary process that they began to learn to coordinate their attack on the lions and move further toward the zebra. This intermediate behavior emphasizes the importance of the available information about the environment in determining agent behaviors. The calling radius is an identifiable landmark, which is both closer to the zebra than most of the hyenas’ starting points (only one hyena begins within the calling radius), but is also far enough away from the lions to be safe.

To further understand the role of the different inputs we performed “knock-out” experiments on the best clans from each of the 40 trials with the flag-bearer, landmark, and number calling (row 8 in Table 4. In these experiments one or more of the inputs were disabled in the most fit clan from each of the 40 trials and the clans were then retested. Disabled inputs consistently returned a zero vector instead of the correct input. (This is analogous to taking an animal,

Table 4: Results, averaged over 40 trials, with different forms of communication. The true or false values in columns 2 through 5 represent whether a particular communication type is available to the hyenas. E.g. in the experiment represented in row 2 hyenas know the number of calling hyenas, but there is no flag-bearer, no landmark, and they can only hear the nearest calling hyena not all calling hyenas. The best results are observed with a flag-bearer (rows 5 to 8), but not hear-all (rows 9 and 10). The last nine rows show the results of a Student’s two-tailed t-test, * denotes significance at the $P = 0.01$ level, - represents non-significance. For example, rows 5 to 8 are significantly better than rows 1 to 4.

	Flag Bearer	Landmark	Number Calling	Hear All	Avg. Fit (std. dev.)	1	2	3	4	5	6	7	8	9
1	False	False	False	False	860.4 (326.9)	-								
2	False	False	True	False	836.1 (318.5)	-	-							
3	False	True	False	False	991.9 (123.6)	*	*	-						
4	False	True	True	False	992.0 (119.4)	*	*	-	-					
5	True	False	False	False	1056.9 (149.0)	*	*	*	*	-				
6	True	False	True	False	1066.6 (97.9)	*	*	*	*	-	-			
7	True	True	False	False	1066.1 (101.5)	*	*	*	*	-	-	-		
8	True	True	True	False	1055.3 (95.5)	*	*	*	*	-	-	-	-	
9	False	True	False	True	939.7 (151.4)	-	-	-	-	*	*	*	*	-
10	True	True	False	True	954.7 (155.4)	-	*	-	-	*	*	*	*	-

knocking-out a specific gene, e.g. to degrade its hearing, and re-examining its fitness or behavior.) Table 5 shows the results of this experiment.

The knock-out data in Table 5 confirms that the flag-bearer input is the most important input for the evolved clans. When it is disabled, average fitness becomes negative (negative fitness occurs when the lion attacks out-weight the fitness obtained by being close to the zebra). However, decreases in fitness are also observed when either number calling or landmark are disabled. This suggests that at least in some trials the clans are evolving to use those inputs. Interestingly, Table 4 show that the clans can evolve equally successful strategies even when the landmark and number calling inputs are not available during evolution (row 5 versus rows 6, 7, and 8). So, they are often used when available, but are not essential to maximizing fitness.

Overall these results suggest that in different trials the hyenas evolve different strategies for attacking the lions and that these strategies depend on different inputs. When it is available, the flag-bear based communication always plays a critical role. Without a flag-bearer evolution cannot find as successful solutions (in the given number of generations) and knocking it out has a significant negative effect in teams that evolved with it. In contrast, the landmark and number calling inputs are somewhat useful on their own, but can be superseded by the having a flag-bearer. Finally, when a flag-bearer is used it appears to sometimes take advantage of either the landmark, the number calling, or both - hence the decrease in fitness in Table 5 from row 8 to rows 5, 6, and 7 when landmark or number calling is knocked-out.

Interestingly, Table 4 shows that hear-all (the name based case) is generally detrimental (row 3 versus row 9 and row 7 versus row 10). Further, hear-all negates the benefits of having a flag-bearer (rows 3 and 7 versus row 10). It is important to note that with hear-all the hyenas could evolve to only “listen” to hyena 1, thereby perfectly replicating the flag-bearer case. However, we hypothesize that with 14 hyenas, it takes considerable evolutionary time for the clans to “agree” on which hyena to listen to. Figure 2 shows the evolutionary curves (averaged across all 40 trials) with a flag-bearer and with and without hear-all. The improve-

ment in fitness is much slower with hear-all than without. This supports the hypothesis that learning with name based communication is difficult when there are many named individuals to consider.

7. CONCLUSIONS

In this paper we used a model of hyena-lion interactions to explore several different and nontraditional communication strategies for a cooperative, multi-agent teams. Our results show that having a single, global identifiable individual - a flag-bearer - significantly improves the evolution of cooperative, coordinated behavior. This benefit arises even though the flag-bearer has none of the coercive ability typically associated with a centralized control structure. Importantly the flag-bearer also does not form a communication bottleneck by acting as a central communication node for the other agents. Thus, a limited number of flag-bearers holds promise as an effective coordination method even for very large swarms where a more typical centralized communication structure would be infeasible.

We studied two other sources of shared information: number calling, which gives general information about the state of the swarm; and landmark, which represents a shared reference similar to the flag-bearer, but static. Number calling generally did not improve fitness while the inclusion of a landmark improved fitness, but not as much as flag-bearer. This suggests that while a shared reference point, the landmark, can be useful, if that reference point is also an evolving agent the potential benefits are much greater.

We also studied traditional name based communication, in which each agent could listen *by name* to the calls of any other agent (referred to in the paper as “hear-all”). This increased performance slightly, but not by a statistically significant amount. It also significantly hindered the performance with a flag-bearer. In our model there are 14 agents, we hypothesize that the difficulty with name based calling arises because each agent most learn which of the other agents it is beneficial to listen while the other agents’ behaviors are themselves evolving.

Overall the results support the idea that there is a “sweet spot” in communication bandwidth. Small, tailored increases




Figure 2: Average best teams over 40 trials with a flag-bearer and with and without hear-all. Under the hear-all condition the hyenas can evolve to listen to only one hyena, thereby replicating the flag-bearer cases. However, the data shows that evolution is faster and achieves better results with the more limited communication case

Table 5: Knockout results, averaged over 40 trials. False means that input was “knocked-out” of the best evolved clan and the clan was retested with that input returning a zero vector. Negative fitness is possible with repeated lion attacks. (The results in row 8 are for the same case as in Figure 4 row 8, differences are due to retesting the clans.) The last eight columns show the results of a Student’s 2-tailed t-test between the respective rows, * denotes significance at the conservative 0.01 level, - represents non-significance. For example, the average fitness given in row 8 is significantly better than any of the other rows, but the fitness in row 7 is not significantly better than row 5 or 6. Knocking out the landmark, number calling, or both inputs have a small, but significant affect on fitness. Knocking out the flag-bearer input has a very large affect on fitness (rows 1-4), confirming that it is the most significant form of communication when we tested. Knocking out the other inputs has a smaller, but significant effect when the flag-bearer is available (row 8 versus 5, 6, or 7), but without the flag-bearer changing the other inputs does not significantly influence fitness.

	Flag Bearer	Land-mark	Number Calling	Avg. Fit (std. dev.)	1	2	3	4	5	6	7
1	False	False	False	-162.5 (658.2)	-						
2	False	False	True	-157.4 (816.2)	-	-					
3	False	True	False	-128.6 (788.8)	-	-	-				
4	False	True	True	-123.7 (664.4)	-	-	-	-			
5	True	False	False	668.3 (510.5)	*	*	*	*	-		
6	True	False	True	888.9 (382.8)	*	*	*	*	*	-	
7	True	True	False	842.4 (387.5)	*	*	*	*	-	-	-
8	True	True	True	1064 (95.12)	*	*	*	*	*	*	*

in communication, e.g. the flag-bear, can improve the rapid evolution of cooperation and coordination. In contrast, allowing unlimited communication, e.g. hear-all, has much smaller benefits, if any, and can interfere with the evolution of more limited and successful communication methods by overwhelming the useful signal with noise from other agents and thereby impeding the evolution of successful solutions.

8. REFERENCES

- [1] W. Banzhaf, P. Nordin, R. E. Keller, and F. D. Francone. *Genetic Programming: An Introduction*. 1998.
- [2] D. Braess, A. Nagurney, and T. Wakolbinger. On a paradox of traffic planning. *Transportation Science*, 39(4):446–450, 2005.
- [3] S. Cooper. Optimal hunting group size: the need for lions to defend their kills against loss to spotted hyaenas. *African Journal of Ecology*, 29(2):130–136, 1991.
- [4] G. Dudek, M. Jenkin, E. Milius, and D. Wilkes. A taxonomy for multi-agent robotics. *Autonomous Robots*, 3(4):375–397, 1996.
- [5] A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer, 2003.
- [6] J. Elman. Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1):71–99, 1993.
- [7] M. Ghavamzadeh, S. Mahadevan, and R. Makar. Hierarchical multi-agent reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 13(2):197–229, 2006.
- [8] M. Land and R. Belew. No perfect two-state cellular automata for density classification exists. *Physical Review Letters*, 74(25):5148–5150, 1995.
- [9] V. Lesser. Reflections on the Nature of Multi-Agent Coordination and Its Implications for an Agent Architecture. *Autonomous Agents and Multi-Agent Systems*, 1(1):89–111, 1998.
- [10] S. Luke and L. Spector. Evolving teamwork and coordination with genetic programming. In *Proceedings of the first annual conference on genetic programming*, pages 150–156. MIT Press, 1996.
- [11] M. Mitchell, J. Crutchfield, and R. Das. Evolving cellular automata with genetic algorithms: A review of recent work. In *Proceedings of the First International Conference on Evolutionary Computation and Its Applications (EvCAS96)*. Citeseer, 1996.
- [12] L. Panait and S. Luke. Cooperative Multi-Agent Learning: The State of the Art. *Autonomous Agents and Multi-Agent Systems*, 11(3):387–434, 2005.
- [13] L. E. Parker. Current state of the art in distributed autonomous mobile robotics. In *Distributed Autonomous Robotic Systems*, pages 3–12. Springer, 2000.
- [14] J. E. Smith, R. C. Van Horn, K. S. Powning, A. R. Cole, K. E. Graham, S. K. Memenis, and K. E. Holekamp. Evolutionary forces favoring intragroup coalitions among spotted hyenas and other animals. *Behavioral Ecology*, 21(2):284–303, 2010.
- [15] H. Watts, L. Blankenship, S. Dawes, and K. Holekamp. Responses of spotted hyenas to lions reflect individual differences in behavior. *Ethology*, 2010.