

Evolving a Follower in the Presence of a Potential Leader

Brent E. Eskridge

Department of Computer Science and Network Engineering, Southern Nazarene University, Bethany, OK 73008
beskridge@snu.edu

Abstract

In many real-world tasks, the ability to use a group of autonomous agents provides significant benefits over a single agent. However, these benefits come at the cost of greater complexity, particularly in the areas of cooperation and coordination. While many approaches address this problem, of particular interest is the use of leaders that emerge through action, and not group deliberation. Other agents follow these “emergent leaders” through the use of environmental cues, rather than explicit communication. While there have been many observations of emergent leadership both in natural and artificial systems, there is a lack of understanding into how this behavior can be reproduced and fostered in artificial systems. In the work presented here, experiments inspired by studies of natural systems were performed to evaluate the ease with which following behaviors could be evolved. Agent controllers were evolved both in isolation and in the presence of a potential leader. Results show that the controllers evolved in the presence of a potential leader exhibited following behaviors when there was an evolutionary advantage and did not incur a fitness penalty when doing so. In fact, agents that followed a leader agent were able to achieve higher fitness than agents acting alone in comparable situations.

Introduction

The ability to use groups of autonomous agents, or **multi-agent systems (MASs)**, in interesting, real-world tasks such as exploration, reconnaissance, and search and rescue, depends on effective coordination in complex, dynamic environments. Striking a balance between the effort required to coordinate the group and the effort required to accomplish the task is a significant problem when groups of agents are scaled beyond a few individual agents (Rosenfeld et al., 2008). Despite this problem, the use of groups provides significant performance and adaptive advantages, whether it is the improved protection from predation in natural systems or the increased fault tolerance in artificial systems.

Since coordination of large groups is frequently observed in nature, researchers are increasingly taking inspiration from the mechanisms observed in and models describing collective movement in natural systems. One mechanism that is particularly interesting is that of leadership. While leadership is not a novel idea in MASs, most leaders in

MASs are chosen *a priori* and frequently act as managers. However, leaders in natural systems often emerge from within the group based on the current situation. This is especially the case in fission-fusion societies where group membership changes frequently and long-term relationships are rare. As a result, group leadership in these systems, and the group as a whole, are able to adapt to dynamic environments and complex tasks much easier than current MASs. While most research investigating leader-follower relationships has historically focused on the aspects of leadership, a deeper understanding of followership would result in more significant progress in understanding the leader-follower relationship and in its application in MASs.

In the work described here, the ability to evolve a following behavior in the presence of both effective and ineffective potential leaders was evaluated. It was hypothesized that an agent controller which exhibited a following behavior could be evolved if it provided a clear fitness benefit. Two experiments, inspired by observations of natural systems, were performed to evaluate this hypothesis. First, a single, evolved agent was tasked with reaching maturity while avoiding a predator. The maturation task was complicated by the addition of varying levels of noise in sensing a predator. The first task acted as a baseline of comparison for the second experiment in which an evolved agent was again tasked with reaching maturity while avoiding a predator, but, in this case, had the opportunity to use the actions of a potential leader as an additional indicator of the presence of a predator. The results of these experiments confirm the hypothesis as agent controllers exhibiting following behavior were successfully evolved in relatively few generations. Furthermore, when compared to the results from the single-agent experiment, individuals that followed did not incur a fitness penalty by following a leader and, in some cases, achieved higher fitness by following.

Background

While a number of mechanisms can be used to facilitate coordination of a MAS, of particular interest is the concept of leadership. In both natural and artificial systems, the use of

a leader is frequently found to improve the cooperation and coordination of a group (Couzin, 2009; Yu et al., 2010).

Leadership

While a variety of approaches to leadership exist, this work defines a leader as an agent that initiates actions and communicates motivation implicitly through the actions it takes. In this context, the term “initiator” might be more appropriate than “leader” (Petit and Bon, 2010; Conradt and Roper, 2005), but the term “leader” is consistently used throughout the literature. Traditionally, leaders in MASs are frequently more “managers” that direct other agents, rather than the “initiator” model used here (Farinelli et al., 2004). Not only does this managerial model of leadership usually require explicit communication, but it also often makes *a priori* assumptions about the distribution of knowledge and capabilities within the MAS. These assumptions present problems in the dynamic environments and complex tasks in which MASs can provide the most benefit. In contrast, in the initiator model of leadership, leaders are not chosen, rather, they emerge from within the group by virtue of the fact that others choose to follow them. There are a variety of reasons for a particular agent emerging as a leader, including a behavioral trait, a morphological trait, or unique access to information that increases the emergent leader’s motivation to act first (King et al., 2009). As a result, other agents observe the actions of the emergent leader and determine that it is in their best interests to follow. Emergent leadership is observed frequently in natural systems, including fish (Harcourt et al., 2009), sheep (Pillot et al., 2009), ravens (Marzluff et al., 1996), and crows (Sonerud et al., 2001). It has even been shown to emerge in MASs (Nouyan et al., 2009; Ghijsen et al., 2010), but those studies did not focus their investigation on the emergent aspect of leadership.

Communication

Traditional models of leadership in MASs usually rely on significant communication, which is frequently categorized as either being explicit or implicit. Explicit communication can be defined as the intentional signaling of information through a defined protocol. Implicit communication, on the other hand, can be defined as an indirect method of communication that uses the individual’s actions, and the resulting changes in the environment, to communicate information. Environmental cues, frequently observed in natural systems, are an example of implicit communication, in which an individual learns “to associate a behavior, a trace, or an object with the occurrence of a given event” (Drapier et al., 2002). Using implicit communication, an informed individual communicates information through its actions, or how its actions modify the environment (Sumpter, 2010).

While explicit communication is commonly used and can be an effective method of improving cooperation and coordination in a MAS (Balch and Arkin, 1994; Ampatzis et al.,

2008), there can be significant problems with its use. Not only is explicit communication less flexible and more sensitive to environmental conditions than implicit communication, it has problems scaling to large numbers of agents (Anderson and Papanikolopoulos, 2008). Dependence on explicit communication is a significant point of failure and leaves the system vulnerable to a new set of problems such as interference and message authentication. Furthermore, there are many situations where explicit communication in a MAS is physically not possible, not practical, too expensive, too complex to be effective in large teams, or may compromise the system’s ability to accomplish a task. Since explicit communication presents significant problems and is not necessary when implicit communication is available (Balch and Arkin, 1994), implicit communication is increasingly being used instead (Pereira et al., 2002; Ampatzis et al., 2008; de Greeff and Nolfi, 2010). While implicit communication has the potential for lower performance than explicit communication, it is a far more practical choice for larger MASs. Implicit communication is simpler, is more robust to change, has lower power consumption, and is stealthier than explicit communication (Pereira et al., 2002; Anderson and Papanikolopoulos, 2008).

Motivation

There are two components to the development of a follower behavior that warrant investigation. The first is the development of the follower behavior itself, and was the focus of the work presented here. The second is the development of the decision-making process that results in choosing to use the follower behavior, and will be the subject of future work. In an effort to evaluate the development of a follower behavior, the simulations described here use a highly abstract evaluation environment to model a maturation problem inspired by a number of experiments in and observations of natural systems. The intent of using an abstract environment is to minimize the number of confounding variables, while retaining the essential aspects of the systems observed in nature.

Simple Maturation Experiment

For this experiment, a simulated agent was tasked with maturing into adulthood, while avoiding predators. In this first experiment, a single agent acted alone, without the benefit of a potential leader. As such, this experiment provided a baseline against which the results of the subsequent MAS experiment with a potential leader could be compared. It was modeled after studies of activity levels of larval anurans (i.e., tadpoles), salamanders, and fish in the presence of predators (Richardson, 2001; Sih et al., 2003; Harcourt et al., 2009) and game-theoretic models of leadership-follower decisions (Rands et al., 2003). As in the natural systems, the agent in this experiment had to risk capture by predators to forage for food, which both ensured its continued survival and enabled the maturation process. Once the agent reached

Parameter	Value
$e_{consumed}$	0.04
e_{exist}	0.01
e_{gain}	0.01
$e_{initial}$	0.50
e_{max}	1.00
m_{energy}	0.02
$m_{threshold}$	0.60
p_{period}	80
Max timesteps	500
Population	100
Generations	100
Mutation rate	1%

Table 1: Experimental parameters used for the maturation experiment are shown.

full maturity, it was no longer considered to be at risk of capture by the predators, which is consistent with natural systems where mature individuals are generally not vulnerable to the same predators as they were during maturation.

Experimental Setup

For this simulation, the act of foraging for food was abstracted into a single value representing an agent's activity level. At each timestep, the agent gained energy based on its activity level, calculated by the following

$$E_{gain} = a \cdot e_{gain} \quad (1)$$

where a was the agent's activity level lying in the range $[0, 1]$ and e_{gain} was the maximum amount of energy that could be gained by foraging (see Table 1). When the agent had a high activity level, it was considered to be foraging and gained energy. When the agent had zero activity, it was considered to be at rest. The energy gained was then added to the agent's energy reserves, referred to as the energy level and denoted E_{total} with a range of $[0, 1]$. Also, at each timestep, the agent consumed energy as a result of its activity level and the energy costs associated with living. The amount of energy consumed by an agent at each timestep was calculated by the following

$$E_{consumed} = a \cdot e_{consumed} + e_{exist} \quad (2)$$

where a was, again, the agent's activity level, $e_{consumed}$ was the maximum amount of energy consumed by foraging, and e_{exist} was the energy costs for the agent's existence. The energy consumed was then subtracted from the agent's energy level. If the agent's energy level ever dropped below zero, it was considered to have died and the trial was terminated.

If the agent's energy level exceeded a threshold value, specified by $m_{threshold}$, a portion of the energy was used

to mature the agent. The amount of energy used for maturation was calculated by the following

$$E_{maturation} = \min\{m_{energy}, E_{total} - m_{threshold}\} \quad (3)$$

where m_{energy} was the default amount of energy used, E_{total} was the current energy level of the agent, and $m_{threshold}$ was, again, the threshold value for maturation. This ensured that the agent only used the energy exceeding the threshold for maturation and did not use energy that was reserved for maintenance (i.e., foraging and existence). This is consistent with observations of energy allocation in natural systems (Heino and Kaitala, 1999). The maturation energy was transferred from the agent's energy level to its maturation level, denoted M with a range of $[0, 1]$. When the agent's maturation level met or exceeded 1.0, it was considered to have fully matured and the trial was terminated.

Predation was modeled as a single value, denoted p_{level} , that indicated the current level of predation and cycled between periods of high and low predation with a value in the range $[0, 1]$. This was considered to be a general indication of the activity level of predators in the vicinity of the agent, and did not represent a specific predator. The predation level at timestep t was calculated by the following

$$p_{level} = \left(\frac{\sin\left(\frac{2\pi t}{p_{period}}\right) + 1}{2} \right)^2 \quad (4)$$

where p_{period} was the period of the predation cycle. The predation value was squared to ensure that there were enough opportunities to forage with minimal predation, while still retaining times of high predation. The agent's probability of being captured by a predator at a given timestep was calculated by the following

$$P_{capture} = a \cdot p_{level} \quad (5)$$

where a , again, was the agent's activity level and p_{level} was the predation level. Thus, an agent was free to forage and have a high activity level if the predation level was low, while it was risky to forage when the predation level was high. At each timestep, a random number drawn from the uniform distribution in the range $[0, 1]$ was generated to determine if the agent had been captured. If the random number exceeded $P_{capture}$, the agent was classified as having been captured by a predator and the trial was terminated.

To introduce uncertainty into the simulation, the agent's ability to sense the current predation level was restricted by the addition of sensor "noise." Given the abstract nature of this experiment, this simplified sensing model provided the ability to tune the agent's uncertainty in its knowledge of the environment without introducing too many confounding variables. Sensor noise was modeled as an offset applied to the actual predation level, p_{level} , and was a random number, drawn from a Gaussian distribution with mean 0 and

standard deviation of 1, denoted $N(0,1)$. In each treatment, the randomly generated offset was multiplied by a value less than or equal to 1, denoted p_{noise} , which served to reduce the standard deviation of the Gaussian distribution from which the offset was drawn. The sensed predation level at each timestep was calculated by the following

$$p_{sensed} = p_{level} + p_{noise} \cdot N(0,1) \quad (6)$$

where p_{sensed} was the sensed predation level, p_{level} was, again, the actual predation level, p_{noise} was the sensor noise level, and $N(0,1)$ is a random value drawn from the previously described Gaussian distribution. Hereafter, this noise value is referred to as a sensor noise percentage (i.e., “X%”), with $p_{noise} = 1$ being referred to as “100%.” In some treatments, the agent’s predation sensor was modeled as having completely failed. In these situations, which are denoted by a sensor noise level of “Random,” the sensor produced random values drawn from a uniform distribution in the range $[0,1]$, denoted $U(0,1)$. This represented a worst-case scenario in which the sensed predation level was completely unpredictable, unlike the other treatments in which random noise was added to a known good predation level.

An agent’s decision-making was performed by an **artificial neural network** (ANN) that was evolved using FS-NEAT (Whiteson et al., 2005), a variation on the standard NEAT algorithm (Stanley and Miikkulainen, 2002) in which ANNs in the initial population have no hidden nodes and no connections between nodes except those added by an initial mutation. The inputs to the ANN were a bias signal, the agent’s energy level, the agent’s maturation level, and the current predation level. The output was the agent’s activity level and was also normalized to the range $[0,1]$. The weights of the evolved ANNs were fixed once created.

Treatments with sensor noise percentages ranging from 0% to 50% were evaluated. Two additional treatments were used that represented worst case scenarios for the agent. In the first, the sensed predation level was completely random. In the second, the predation level sensor was missing and the ANN received a constant input of 0 for each timestep. Forty experimental runs were performed for each treatment with ANNs evaluated in five trials. Fitness was calculated as the mean final maturation level of the agent in each trial when the trial ended. While the time required for an agent to mature was not a part of the fitness calculation, there was an implicit benefit for faster maturation since it resulted in lower maintenance costs and afforded fewer opportunities for a predator to capture the agent. Furthermore, a capture did not preclude an ANN from being selected as a parent for the next generation. While this would be the case in a natural system, it is important to remember that it was the evolved ANN that was evaluated and received fitness, not the agent itself. Table 1 shows the experiment-specific parameter settings that were used. NEAT-specific parameter settings were based on standard NEAT defaults and are re-

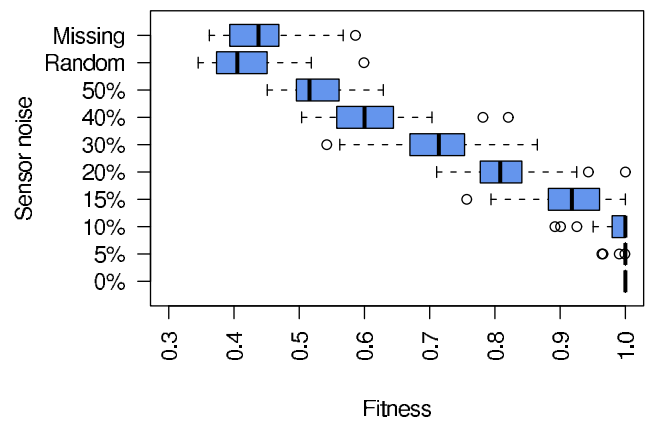


Figure 1: The mean best-of-run fitness values over all five trials for varying levels of sensor noise in the single agent maturation experiment treatments are shown. For each treatment, the box depicts the interquartile range (IQR) from the first quartile to the third quartile over all 40 trials, the vertical line represents the median fitness value, the whiskers represent the ± 1.5 IQR, and the circles represent outlier values.

ported elsewhere (Stanley and Miikkulainen, 2002).

Results

Figure 1 shows the mean best-of-run fitness values over five trials for the single-agent experiment using boxplots. For each treatment, the box depicts the interquartile range (IQR) from the first quartile to the third quartile over all 40 trials, the vertical line represents the median fitness value, the whiskers represent the ± 1.5 IQR, and the circles represent outlier values (Robbins, 2005). All experimental runs were able to evolve an ANN that resulted in the successful maturation of an agent in all five trials when no sensor noise was added to the agent’s predation level precept (i.e., 0% sensor noise). The performance of the ANNs dropped dramatically as sensor noise was added. However, consistently successful ANNs were evolved even with the addition of small amounts of sensor noise. Results of the bootstrapped Kolmogorov-Smirnov¹ test show that there was no statistically significant difference in the fitness values between between the 0% and 5% sensor noise treatments ($p = 0.081$), while there was a statistically significant difference between the 0% and the remaining sensor noise treatments ($p \approx 0$).

Leader Experiment

One reason for the emergence of leaders in natural systems is an individual possessing knowledge that other individuals in the group do not (King et al., 2009). Since incorrect information, in the form of a “noisy” predation sensor, was

¹Unlike other standard significance tests, the KS test does not assume that the data has a normal distribution. Since these results suffer from ceiling effects, this assumption cannot be made.

already present in the simulation, potential leaders were introduced as potentially having more accurate information on the current level of predation.

As in the first experiment, an evolved agent was tasked with surviving until maturity. However, in this experiment, a potential leader, whose actions the evolved agent could observe, was added to the environment. The evolved agent was, therefore, able to observe the activity of the potential leader in response to the current predation level. As a result, the evolved agent could exhibit a following strategy if the leader's actions provided a better indicator of the current predation level than its own sensor. This provided an opportunity to determine the ease with which effective followers could be produced in response to potential leaders.

Experimental Setup

To enable the evolved agent to sense the potential leader's activity level, the potential leader's activity level was added as an input to the ANN configuration used in the single-agent experiment. Note that the use of an input for the activity level of the potential leader represents an observation by the agent of the potential leader, and not explicit communication between the potential leader and the agent. The leader's actions were controlled by a randomly chosen, best-of-run ANN from the single-agent experiment with no sensor noise. As a result, the potential leader was unable to sense the activity level of the agent under evaluation.

In each treatment, Gaussian noise was added to the agent's predation level sensor, the potential leader's predation level sensor, or both. The addition of sensor noise to the potential leader's sensor was used to evaluate the evolved agent's performance in the absence of a perfectly accurate leader. To ensure that the evolved agent always received some environmental cues from the potential leader, the potential leader was ineligible for capture by the predator, regardless of the quality of its actions. While this may not be a realistic or viable long-term assumption, it proved to be an effective simplification given the highly abstract nature of the current experiment and will be revisited in future experiments that use more realistic, high-fidelity environments. The experiment-specific parameter settings used in this experiment were the same as the single-agent experiment and are shown in Table 1.

Results

Figure 2 shows the mean best-of-run fitness values for experimental treatments in which a potential leader was present in the environment. In treatments where at least one of the two agents, either the agent under evaluation or the potential leader, had 0% sensor noise, evolved ANNs were consistently able to produce behaviors that resulted in agents successfully maturing in each of the five separate trials. Agents were captured by a predator in only a few trials (see Figure 2a). An analysis using the bootstrapped Kolmogorov-

Smirnov test shows that there was no statistically significant difference between any of the treatments, including treatments with a few outliers in which the agent was captured before reaching full maturation. This indicates that, when appropriate, the evolved ANN was able to use either its own predation level percepts or the activity level of the potential leader with equal effectiveness. One particular treatment of note is the one in which the evolved ANN's predation level sensory input was completely missing. Since the ANNs evolved in this treatment did not differ in fitness from the single-agent treatment, it can be concluded that using the potential leader's activity level as a proxy for the predation level did not incur any inherent fitness penalty.

The results for treatments in which the minimum sensor noise level were 5% are consistent with the previous set in that evolved ANNs were able to use either a direct sensing of the predation level or the potential leader's activity level as a predation level indicator with equal effectiveness (see Figure 2b). However, in this case, evolutionary runs in three treatments achieved higher fitness than similar treatments in the single-agent experiment. In evolutionary runs in which both agents had 5% sensor noise or one agent had 5% sensor noise and the other had 10% sensor noise, evolved ANNs were able to produce behavior that resulted in the successful maturation of an agent in each of the five trials over all forty experimental runs. Although the significance level between fitness values was relatively high ($p < 0.1$) and the difference in fitness was relatively small (0.998 ± 0.008 vs. 1.0 ± 0.0), it bears mentioning since it indicates a trend that will be observed in later treatments.

In the results for treatments in which the minimum sensor noise level was 10%, similar results were found (see Figure 2c). For only the treatment in which the evolved agent had 10% sensor noise and the potential leader's predation level input was random did the addition of the potential leader result in statistically significantly *lower* fitness ($p = 0.001$). Although the relative differences in fitness were slight, the reason for this drop in fitness is unknown and warrants further investigation. Similar to the previous set of treatments, the treatment in which both agents had 10% sensor noise had statistically significantly higher fitness than the single-agent treatment with 10% noise ($p = 0.002$).

Lastly, the results for treatments in which the minimum sensor noise level was 15% were consistent with previous treatments (see Figure 2d). The treatments in which one agent had 15% sensor noise and the other had a random input for the predation level produced results comparable to the single-agent treatment and were not statistically significantly different. The treatment in which both agents had 15% sensor noise had statistically significantly higher fitness, as in the previous treatment sets ($p \ll 0.0001$).

To further investigate the phenomenon in which agents in the leadership experiment were able to achieve higher fitness than a single agent with the same level of noise, the rate at

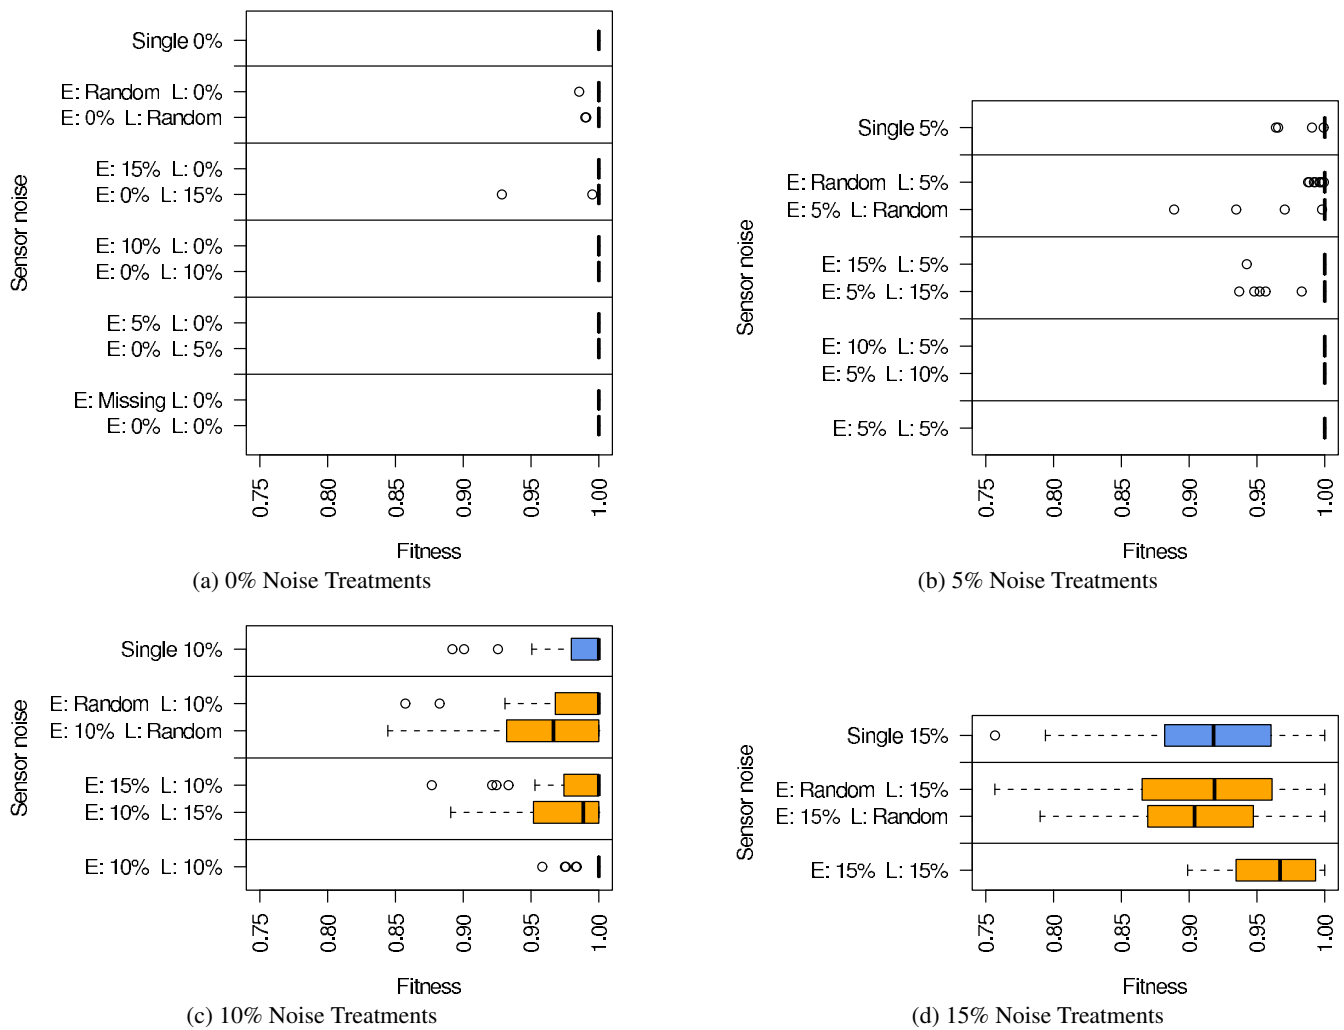


Figure 2: The mean best-of-run fitness over five trials and forty runs for the maturation experiment in the presence of a potential leader are shown. Individual treatments are organized by sensor noise levels and are compared to results from the single-agent experiment with similar sensor noise levels. The “E” term represents the sensor noise of the evolved agent and the “L” term represents the sensor noise level of the potential leader.

which fit ANNs evolved in each treatment were compared using the randomized two-way ANOVA test (Piater et al., 1998). In each case, evolved ANNs from the leadership experiment treatment achieved higher fitness faster than the treatment from the single-agent experiment with $p = 0.05$. Figure 3 illustrates these results in a comparison of the fitness curves between treatments.

Discussion

There are two main conclusions that can be drawn from these results. First, ANNs were evolved in a few generations that were capable of using the observed leader’s actions as an indicator of the current predation level if their own sensor percepts were unreliable. Although increased noise in the predation sensor made the evolution of follow-

ing behaviors more difficult, the fitness curves in Figure 3 show that effective controllers exhibiting following behaviors were evolved in under 100 generations for up to 15% sensor noise. In treatments for which following was a viable strategy, the mean generation at which the best-of-run ANN was found was 34.9 with a standard deviation of 16.5. Furthermore, this following behavior did not incur a fitness penalty as some of the treatments in the 0% sensor noise treatments shown in Figure 3a illustrate.

Second, these results demonstrate that an effective following behavior can result in performance that is superior to performance in the single-agent experiment. This superior performance is shown not only in the statistically significantly higher fitness of the evolved ANNs, but also in the statistically significantly faster rate in which fit ANNs were

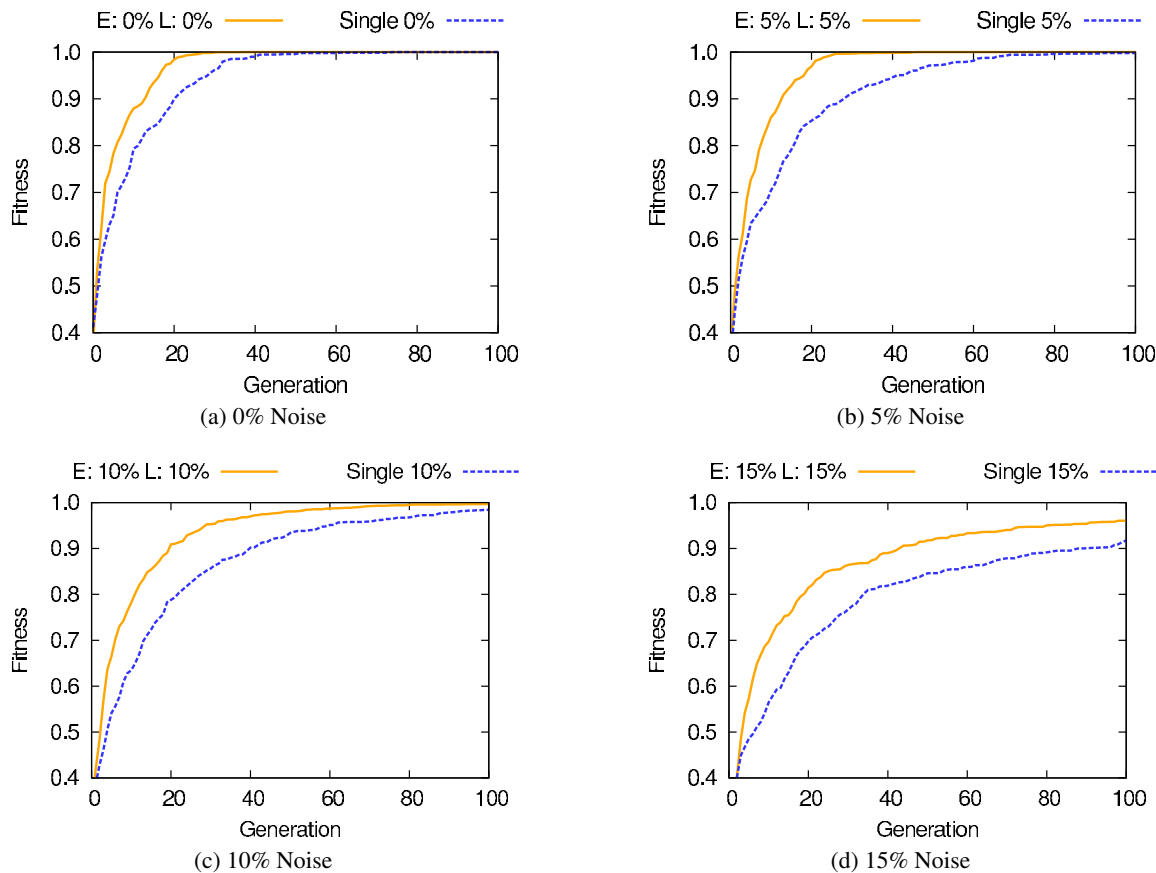


Figure 3: Plots of the mean best-of-run fitness at each generation over all forty experimental runs for selected treatments of the single-agent and leadership experiments are shown. Selected treatments compare the fitness of treatments in the single-agent experiment with treatments in the leadership experiment in which both agents had the same sensor noise level.

evolved. While this phenomenon of higher group performance, referred to as the “many wrongs principle,” is observed in group navigation found in nature (Simons, 2004), its observation was not expected in this highly abstract environment, as there were only two agents present and the potential leader was ignorant of the evolved agent’s presence.

Conclusions

For MASs to be useful in the dynamic, real-world tasks for which they can provide the most benefit, the problem of effectively coordinating even moderate numbers of agents must be solved. One promising approach is through the use of “emergent leadership.” In emergent leadership, leaders arise from within a group by virtue of the actions that they take, and do not require extensive communication. For emergent leadership to work, however, other agents within the group must decide to follow the leader. In the work presented here, the ability to evolve agent controllers capable of exhibiting following behaviors was evaluated. The experimental results demonstrate that the effective controllers were evolved in relatively few generations, even in the pres-

ence of sensor noise, and without the benefit of explicit communication between the leader and follower. Furthermore, following did not incur a loss in performance when compared to the single-agent simulations and even resulted in a performance increase in some simulations.

There are a variety of directions for future work. First, the work presented here used a static environmental configuration in which it was either beneficial or not beneficial to follow the leader. As noted above, the second component of interest in followership is the development of an effective decision-making process. Additional experiments indicate that attempting to evolve both a following behavior and the decision-making involved in deciding to follow at once can be too complex to evolve in a single ANN. Further work is necessary to ensure that agents can adapt to dynamic environments in which the benefits to following vary with time and the decision to follow is less clear-cut. Also, in these experiments, the agent was presented with a simple choice: follow a single known agent or follow no one. When an agent is a member of a much larger group, it is presented with the much more difficult choice of deciding *which* agent

to follow, if one at all. Further work is also required to determine if the relative ease with which following behaviors were evolved persists as the number of agents is scaled up.

Acknowledgments

The author wishes to thank Ingo Schlupp and Dean Hougen for their contributions and insights. This work was supported by NSF grant No. BCS-1124837. Some of the computing for this project was performed at the OU Supercomputing Center for Education & Research (OSKER) at the University of Oklahoma.

References

- Ampatzis, C., Tuci, E., Trianni, V., and Dorigo, M. (2008). Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behavior*, 16(1):5–26.
- Anderson, M. and Papanikolopoulos, N. (2008). Implicit cooperation strategies for multi-robot search of unknown areas. *Journal of Intelligent and Robotic Systems*, 53(4):381–397.
- Balch, T. and Arkin, R. C. (1994). Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52.
- Conradt, L. and Roper, T. (2005). Consensus decision making in animals. *Trends in Ecology & Evolution*, 20(8):449–456.
- Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, 13(1):36–43.
- de Greeff, J. and Nolfi, S. (2010). Evolution of implicit and explicit communication in mobile robots. In *Evolution of Communication and Language in Embodied Agents*, pages 179–214. Springer.
- Drapier, M., Chauvin, C., and Thierry, B. (2002). Tonkean macaques (*Macaca tonkeana*) find food sources from cues conveyed by group-mates. *Animal Cognition*, 5(3):159–165.
- Farinelli, A., Iocchi, L., and Nardi, D. (2004). Multirobot systems: a classification focused on coordination. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 34(5):2015–2028.
- Ghijsen, M., Jansweijer, W. N. H., and Wielinga, B. J. (2010). *Adaptive Hierarchical Multi-agent Organizations*, volume 281 of *Studies in Computational Intelligence*, pages 375–400. Springer.
- Harcourt, J. L., Ang, T. Z., Sweetman, G., Johnstone, R. A., and Manica, A. (2009). Social feedback and the emergence of leaders and followers. *Current Biology*, 19(3):248–252.
- Heino, M. and Kaitala, V. (1999). Evolution of resource allocation between growth and reproduction in animals with indeterminate growth. *Journal of Evolutionary Biology*, 12(3):423–429.
- King, A. J., Johnson, D. D., and Van Vugt, M. (2009). The origins and evolution of leadership. *Current Biology*, 19(19):R911–R916.
- Marzluff, J. M., Heinrich, B., and Marzluff, C. S. (1996). Raven roosts are mobile information centres. *Animal Behaviour*, 51(1):89–103.
- Nouyan, S., Groß, R., Bonani, M., Mondada, F., and Dorigo, M. (2009). Teamwork in self-organized robot colonies. *IEEE Transactions on Evolutionary Computation*, 13(4):695–711.
- Pereira, G., Pimentel, B., Chaimowicz, L., and Campos, M. (2002). Coordination of multiple mobile robots in an object carrying task using implicit communication. In *International Conference on Robotics and Automation*, pages 281–286. IEEE.
- Petit, O. and Bon, R. (2010). Decision-making processes: The case of collective movements. *Behavioural Processes*, 64(3):635–647.
- Piater, J. H., Cohen, P. R., Zhang, X., and Atighetchi, M. (1998). A randomized ANOVA procedure for comparing performance curves. *International Conference on Machine Learning*, pages 430–438.
- Pillot, M. H., Gautrais, J., Gouello, J., Michelena, P., Sibbald, A., and Bon, R. (2009). Moving together: Incidental leaders and naïve followers. *Behavioural Processes*, pages 235–241.
- Rands, S., Cowlshaw, G., Pettifor, R., Rowcliffe, J., and Johnstone, R. (2003). Spontaneous emergence of leaders and followers in foraging pairs. *Nature*, 423(6938):432–434.
- Richardson, J. M. (2001). A comparative study of activity levels in larval anurans and response to the presence of different predators. *Behavioral Ecology*, 12(1):51–58.
- Robbins, N. (2005). *Creating more effective graphs*. Wiley-Interscience.
- Rosenfeld, A., Kaminka, G., Kraus, S., and Shehory, O. (2008). A study of mechanisms for improving robotic group performance. *Artificial Intelligence*, 172(6-7):633–655.
- Sih, A., Kats, L. B., and Maurer, E. F. (2003). Behavioural correlations across situations and the evolution of antipredator behaviour in a sunfish-salamander system. *Animal Behaviour*, 65(1):29–44.
- Simons, A. (2004). Many wrongs: the advantage of group navigation. *Trends in Ecology & Evolution*, 19(9):453–455.
- Sonerud, G., Smedshaug, C., and Bråthen, Ø. (2001). Ignorant hooded crows follow knowledgeable roost-mates to food: support for the information centre hypothesis. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 268(1469):827–831.
- Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127.
- Sumpter, D. J. (2010). *Collective Animal Behavior*. Princeton University Press.
- Whiteson, S., Stone, P., Stanley, K. O., Miikkulainen, R., and Kohl, N. (2005). Automatic feature selection in neuroevolution. In *Genetic and Evolutionary Computation Conference*, pages 1225–1232.
- Yu, C.-H., Werfel, J., and Nagpal, R. (2010). Collective decision-making in multi-agent systems by implicit leadership. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 1189–1196. IFAAMAS.