Wolves in Sheep's Clothing: Using Shill Agents to Misdirect Multi-Robot Teams

Michael J. Pettinati, Ronald C. Arkin, and Akshay Krishnan

Georgia Institute of Technology, Atlanta, GA 30332 USA
mpettinati3,arkin,a.krishnan@gatech.edu

Abstract. When robots undertake tasks in adversarial environments in which they must cooperate with one another (e.g., military applications or the RoboCup Competition), they are at risk for being deceived by competitors. Competitors can misdirect the team to gain a positional advantage. Our lab is exploring ways in which teams of robots can be misdirected, in part, so counter-deception strategies can be devised. This paper explores how robot shills can be used to misdirect a multi-robot team. It defines behaviors for the agents to be deceived (the mark agents) using the multi-agent coordination literature as well as behaviors for the deceiving team (the shills and lead agent). These behaviors were implemented and simulations were run for a variety of conditions. The results show how shills can facilitate misdirection in certain circumstances. They provide insights into enhancing multi-robot team deception.

Keywords: Robot deception, Multirobot systems, Team misdirection

1 Introduction

A man is milling around a stadium parking lot before a big event. There seem to be countless people wandering around him waiting for the show later that day. A woman moving very quickly enters the man's view; she moves with purpose away from the stadium. The man finds the fast-moving woman interesting but does not consider her much further until he sees a nearby man begin moving quickly (with urgency) in the same direction as the woman. He wonders if there is some kind of emergency for which assistance might be needed. As he begins to move quickly toward the two, several people behind him, who also had noticed the quickly moving pair, wonder what is going on and begin to move with the group.

Research has shown how in teams of humans and animals a small proportion of the group members are able to sway the behavior or movement of the larger group with simple local interactions [1, 3]. In the story above, the woman did not need to call out to all the followers. Instead, with purposeful movement, she was able to attract attention to herself and begin pulling people with her. The people who she pulled with her, in turn, pulled people with them.

This flocking behavior can be useful for groups of robots and has been incorporated into multi-agent robot teams [e.g. 2]. Our lab showed how a robot behavior inspired by lekking in birds could help to support the formation of meaningful task-oriented robot groups. This behavior, though useful, leaves robot teams susceptible to misdirection.

This paper explores misdirection in these robot teams. Specifically, it tries to understand how shill agents (confederate members of the deception team) can help to misdirect mark agents (targets of the deception). The goal of the deceptive agents is to move the marks from a start position to a goal position that is advantageous to the deceiving team. Feints for example (moving in a direction intended to mislead) are common in sports and the military.

In the story above, the man followed along with the woman when he saw the nearby man begin to move with her. This second individual could have followed the woman out of curiosity, but he could have been a confederate of (shill for) the woman. He could have moved to encourage others to follow along and/or to keep others following along with her. In groups, people take an action when they have seen a sufficient number of others take the same action [5]. People assume if many people are taking this action, then it must be correct or appropriate. The man needed to see two people move with urgency in a certain direction before deciding there was something worth seeing in their direction and moving with them. Robots can similarly follow this threshold model to inform their actions.

Our lab has done extensive work in robot deception [e.g. 9, 10] and even provided the first taxonomy on human-robot deception [10]. This work is building upon that previous work by exploring the misdirection of a multi-agent robot team by a multi-agent robot team. This research is being done in part to develop counter-deceptive practices in future works.

The next section of this paper discusses previous research looking at deception between teams of robots as well as how robots have been used to move groups from one location to another. The third section introduces the models of the mark agents and the agents involved in the deception. The fourth section present simulation results involving implementations of these agents. The paper ends with a conclusion and discussion of future work.

2 Related Work

Previous research into multiagent deception has looked at how a deceptive team of robots can keep adversaries away from a certain area that may harbor valuable resources [8]. Our research, instead, focuses on misdirecting adversaries to a certain area.

Robots have been employed in herding situations [4, 6, 11, 12]. These robots "pushed" animals from one location to another. This included herding ducks [11, 12], which have similar flocking behaviors to sheep, into penned areas, and herding birds away from airports to designated safe zones [4, 6]. These "pushing" approaches are fear-based with the robot acting as a predator-like agent [4, 6, 11, 12]. They are fundamentally different than our deceptive approach. The agents that are moving the marks to the goal location in this paper are indistinguishable from the mark agents themselves.

This also separates the present work from our lab's recent paper [**7**]. A team of shepherding robots moved a team of mark robots from one location to another. The shepherding team was more effective at moving the marks to the goal location when it combined agents that pulled along with agents that pushed the marks than when the pulling and pushing agents were separate. The pulling and pushing agents, however, were identifiable as different from the marks themselves, contrary to the shills used here.

3 Robot Models

The simulations discussed in the following section replicate the scenario given in the introduction. This section defines the behavioral assemblages that dictate the actions of the agents. The primitive behaviors for each robotic agent are defined in Appendix I. The behavioral assemblages can be seen in Appendix II. Notationally, the behavior assemblages appear bolded throughout this paper and the primitive composing behaviors appear italicized.

There are three types of agents. The first type, the leader agent, plays the role of the quickly moving woman. It leads the other agents toward the goal location. The second agent type is the mark, the agents to be deceived. The mark agents are the crowd outside of the stadium. They wander around and are unresponsive to the surrounding robots until seeing a number of agents moving with intent (moving quickly) at which point they flock to those agents. The third agent type is the shill that act as confederates with the leader agent. They mill among and flock with the mark agents, while also helping to pull the mark agents toward the leader. This is the person who moves as soon as the quickly moving woman appears as illustrated in the introduction.

3.1 Behavior Overviews

The behavioral assemblages for each of the three agent types are summarized in Table 1. The leading agent enacts a **Lead-To-Goal Behavior Assemblage** throughout the simulation that includes three behaviors. The agent is attracted to the goal location (*Go-To-Goal Behavior*); it avoids obstacles (objects) in the environment (*Avoid-Obstacle Behavior*), and it has noise incorporated into its movements (*Wander Behavior*) so that these movements are natural. The leader is the only agent with knowledge of the goal location's position.

The marks are the agents to be relocated from their initial position to the goal location. They begin the simulation wandering slowly around their start location with the **Anchored Wander Behavior Assemblage** active. They avoid crashing into other robots (*Off Robots Behavior*) as well as obstacles (objects) within the environment (*Avoid-Obstacle Behavior*). Otherwise, they simply wander around the area where they begin the simulation (*Wander Behavior and Stay Near Start Behavior*).

Each mark has a set threshold that will cause it to change its behavior to flock. This threshold is the number of agents that the mark needs to recognize as moving with intent. Moving with intent means moving at a speed above a set threshold. As described above, humans will make a decision when they have seen a certain number of others make the same decision [5]. The man in the story from the introduction decided it was prudent to follow the quickly moving woman when a nearby man chose to move toward her. Marks will flock with the robots that show intent once they have seen a sufficient number of agents moving with intent.

Table 1: The behavior assemblages for each agent type along with the composing behaviors.

Robotic Agent	Behavior Assemblage	Composing Behaviors
Leader	Lead To Goal	• *Go-To-Goal* • *Avoid-Obstacles* • *Wander*
Mark	Wander Near Start (Simulation Outset)	• *Wander* • *Stay Near Start* • *Avoid-Obstacle* • *Off Robots*
	Mark Mill Around (Below Flock Threshold)	• *Wander* • *Avoid-Obstacle* • *Off Robots*
	Mark Flock (Above Flock Threshold)	• *Lek Behavior* • *Wander* • *Avoid-Obstacle* • *Off Robots*
Shill	Wander Near Start (Simulation Outset)	• *Wander* • *Stay Near Start* • *Avoid-Obstacle* • *Off Robots*
	Shill Flock (Leader Signaled)	• *Follow Leader* • *Lek Behavior* • *Wander* • *Avoid-Obstacle* • *Off Robots*

All of the shill agents in the simulations are indistinguishable to the marks. Each mark computes the speed of every agent within a certain radius that is not concealed by an object in the environment. The marks measure their speed by considering a robot's current position and its position ten simulation steps before. The agent's speed is how far the agent has moved in that time window.

When a mark is at or above its flocking threshold, the **Mark Flock Behavior Assemblage** is active. This consists of four different behaviors. The mark is attracted to each agent moving with intent that it is able to see within a certain region (*Lek Behavior* [2]). The agent avoids crashing into robots (*Off Robots Behavior*) and obstacles (*Avoid-Obstacle Behavior*) in the environment. There is also noise incorporated into the robot's movement to make it more natural (*Wander Behavior)*.

When the mark is below its threshold (i.e., the number of agents it sees moving with intent is below the specified number), the agent will enact the **Mark Mill Around Behavior Assemblage**. The agent will wander around (*Wander Behavior*), avoid other robots (*Off Robots Behavior*) and obstacles *(Avoid Obstacle Behavior*).

Finally, the shill agents in the simulation behave very similarly to the mark agents. At the outset of the simulation, they enact the mark's **Anchored Wander Behavior Assemblage**. When the leader agent begins to move to the goal location, they enact the **Shill Flock Behavioral Assemblage**. This includes the same behaviors as the Mark Flock behavior (*Off Robots Behavior, Avoid-Obstacle Behavior, Wander Behavior and Lek Behavior*) with the addition of being attracted to the leader when the leader is visible *(Follow Leader Behavior*). This helps pull the flock toward the leader's position, which is approaching the goal throughout the simulation.

3.2 Mathematical Models

Each robot's position is updated every simulation step (simulation second). The distance and direction moved by a robot is based on that robot's baseline speed (in meters per second) and the behavioral assemblage that it is currently executing. Each behavior in the behavioral assemblage outputs a vector. The simulation computes a weighted sum of the vectors that are returned by the behaviors and multiplies the resulting vector by the agent's baseline speed to determine how far and in what direction the agent moves. The behaviors appear in Appendix I. This appendix describes the vector returned by each behavior. The weights and parameters associated with each behavior and behavior assemblage appear in Appendix II.

3.3 Robot Missions

The finite state automata that define the missions for the robot appear in Figure 1. The circles show the behavioral assemblages for the agent that are active throughout the course of a simulation. The rectangles show the triggers by which the agent moves between behavioral assemblages.

The leader (Fig. 1A) waits until all of the other agents are contained within the starting area and wandering around. It approaches the area containing the agents and signals the shills that it is heading toward the goal. It then moves toward the goal location using the **Lead-To-Goal Behavior Assemblage**.

A mark (Fig. 1B) wanders around in the start location until its flocking threshold has been satisfied (**Wander Near Start Behavior Assemblage**). The threshold is some number of agents moving with intent (at or above a certain speed). It flocks with the agents that are moving with intent (**Mark Flock Behavior Assemblage**) when at or above this threshold. It wanders when below this threshold (**Mark Mill Around Behavioral Assemblage**).

The shill agent (Fig. 1C) wanders with the mark agents at the simulation's outset (**Wander Near Start Behavior Assemblage**). When the leader signals to the shill that it is heading to the goal location, the shill begins to flock with the agent's that move with intent; it tries to drag the marks along to the leader's location (**Shill Flock Behavior Assemblage**).

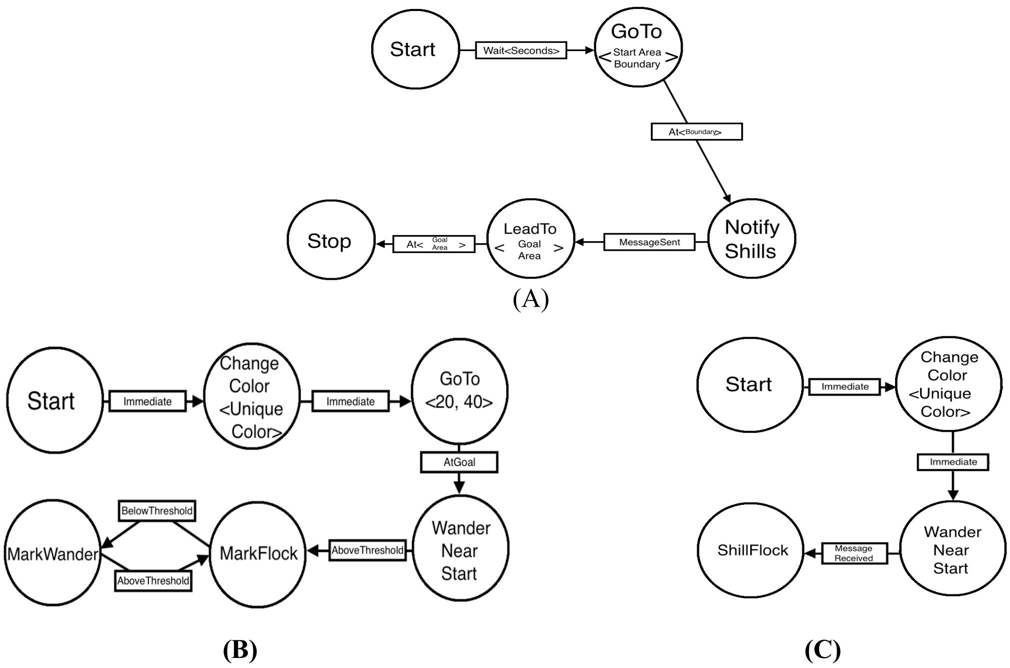

Figure 1: The FSAs for Agents: (A) Leader Agent (B) Mark Agent (C) Shill Agent

4 Misdirection Simulations

Each simulation began when the leader agent signaled the shill agents and it began to move to the goal. The shills and marks began the simulation wandering within a start area of ten-meter radius (20m, 40m). The simulation ended when all of the marks were within ten meters of the goal location (220m, 40m) or when the simulation had run 2000 steps (seconds). The simulation environment was 60m by 240m.

The five independent variables that were manipulated between simulations are summarized in Table 2. Snapshots from a trial with twelve mark agents, two shill agents and a large object appear in Fig. 4. The marks were successfully relocated from the start area to the goal area. The robots have unique colors so marks can identify and track other agents to compute their speeds.

The number of shills was varied from 0 to 2 to understand the conditions under which shills may facilitate misdirection. Shills are attracted toward the leader and flock along with the marks. These agents help to keep the flock of marks moving toward the leader and help to meet the marks' thresholds for flocking. The shill agents move with intent in the general direction of the leader.

Table 2: These are the five independent variables that were manipulated between simulation conditions. There were 30 different conditions tested by running series of simulations with the values indicated.

Independent Variable	Values Tested		
Number of Shills	0, 1, and 2 for all other parameter settings		
Number of Marks	4	12	
Mark Agents Thresholds for Flocking	2 marks -threshold 1 2 marks - threshold 2	6 marks - threshold 1 6 marks - threshold 2	1 mark - threshold 1 2 marks - threshold 2 2 marks -threshold 3 2 marks -threshold 4 2 marks - threshold 5 2 marks - threshold 6 1 mark – threshold 7
Shill *Lek Behavior* Weight	High	High	Low, High
Environment Obstacles	No Object Small Object Large Object	No Object Small Object Large Object	No Object Large Object

The shill *Lek Behavior* could be given a higher or lower weight. A higher *Lek Behavior* weight makes the shill more responsive to all the agents moving with intent in the simulation. A lower weight makes the shill less responsive to all flocking agents and gives greater influence over its movement to the leader agent.

The complexity of the environment was varied as well. Simulation cases included no object present, a small object present, or a large object present (Figure 2). The objects were centered at (140, 40) and had a radius of 3 meters or 10 meters. The objects obscured all agents' lines of sight. During a trial with objects, the marks could lose sight of agents moving with intent, so they would fall below their flocking threshold and simply begin to mill about. The object could also prevent a shill agent from seeing the leader; this removes the shill's ability to move toward the goal.

The number of marks and their thresholds for flocking were varied as well. In certain simulations, all the marks were easily persuaded to flock, they all had low flocking thresholds (of 1 and 2). In other conditions, certain marks had much higher thresholds for flocking.

The degree to which the deceptive team (the leader and shill agents) was successful in misdirecting the flock was assessed by looking at the proportion of the marks moved from the start to the goal region.

4.1 Results

We compared the median number of mark agents successfully moved from the start location to the goal location using the Kruskal-Wallis one-way analysis of variance test.

The number of shills was the independent variable for each test. The environment complexity, the number of marks, and the marks' thresholds varied between the tests but were held constant within tests. This Kruskal-Wallis test was used because the proportion of agents that were successfully misdirected did not follow a normal distribution. There were many trials within conditions where almost all the mark agents were moved to the goal location (the proportion of agents was near 1) or almost none of the mark agents were moved to the goal location (the proportion of agents was near 0).

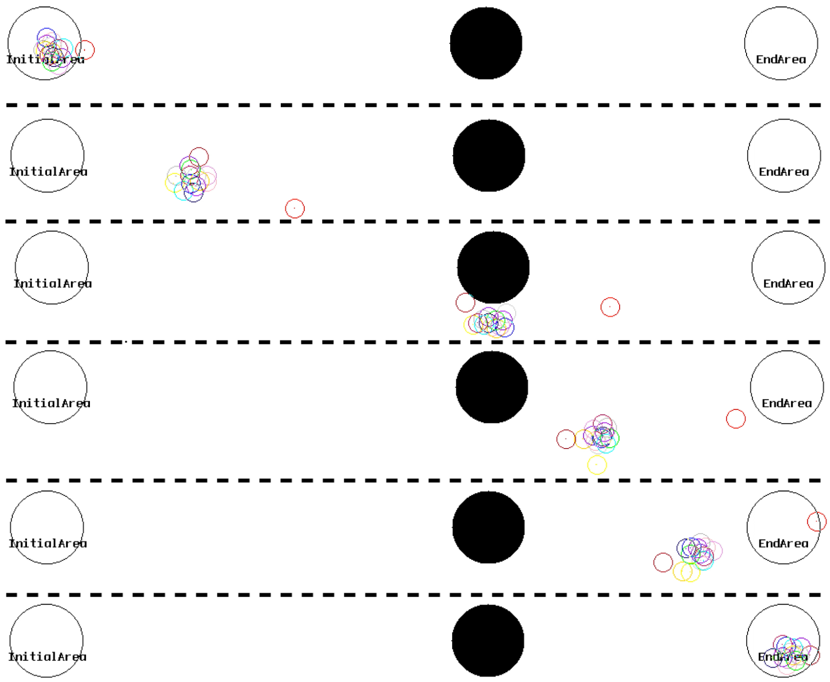

Figure 2: Snapshots from a simulation trial with twelve marks and two shills. The marks are successfully relocated from the start to the goal.

There were no significant differences between the groups with respect to the proportion of mark agents that were successfully moved to the goal location when the *Lek Behavior* weight was high and the flocking thresholds for all marks was low ($p > .05$). Additional shills in these simulations did not change the leading agent's ability to pull mark agents from the start to the goal location in any environment. In these cases, it seems the leader alone was sufficient to pull large proportions of the marks to the goal location (see Table 3 and Figure 3).

The results for when the marks' thresholds were changed to include agents that had high flocking thresholds still had no significant differences between the three conditions with different number of shills ($p > .05$). In these cases, however, it was often true that the marks were not successfully moved to the goal location. In all of the conditions, the median values for the proportion of shills moved to the goal location were below .25, fewer than one in four marks was moved from the start to the goal location.

The weight on the shills' *Lek Behavior* was changed to a lower weight. These results are summarized in Table 4 and Figure 4. There was an extremely significant difference ($p < .001$) between all groups in the no object condition. With two shills, the agent was

Table 3: The results from the simulations run with twelve marks, a high *Lek Behavior* weight, and low flocking thresholds for all mark agents. There were no significant differences between conditions. Marks did not facilitate the misdirection.

Environment	0 Shills Median (Standard Deviation) n = 20	1 Shill Median (Standard Deviation) n = 20	2 Shills Median (Standard Deviation) n = 20
Big Object	1.0 (0.438)	0.708 (0.374)	0.667 (0.392)
Small Object	0.958 (0.390)	0.958 (0.405)	0.583 (0.384)
No Object	1.0 (0.283)	1.0 (0.253)	1.0 (0.293)

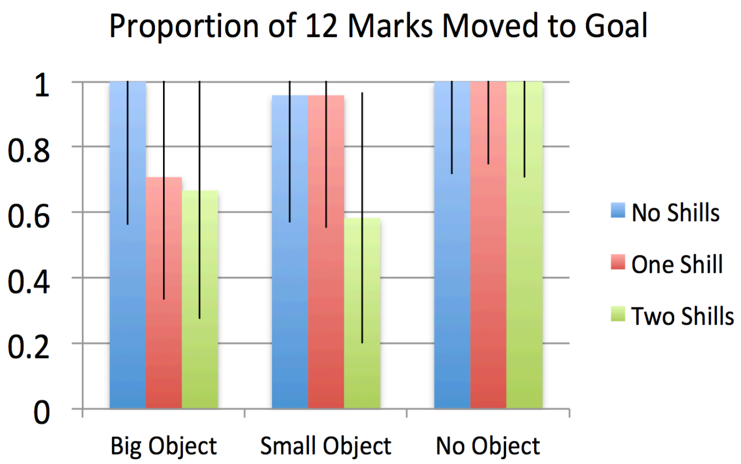

Figure 3: There were no significant differences between groups when all 12 marks had low flocking thresholds. The leader alone was able to misdirect the agents.

Table 4: The results from the simulations run with twelve marks, a low *Lek Behavior* weight, and high flocking thresholds for some mark agents. There was a very significant difference between all groups with no object present. The shills facilitated the misdirection. With the big object present, they did not.

Environment	0 Shills Median (Standard Deviation) (n = 10)	1 Shill Median (Standard Deviation) (n = 10)	2 Shills Median (Standard Deviation) (n = 10)
Big Object	0.0 (0.0)	0.0 (0.478)	0.0 (0.478)
No Object	0.0 (0.167)	0.417 (0.423)	1.0 (0.053)

consistently able to misdirect almost all of the mark agents. With no shills, the leader was consistently not able to misdirect the mark agents. In the big object condition, however, there was no difference between groups. The leader, in all three conditions, was unable to misdirect the marks. The object in a large portion of the trials obscured the leading agent from the shill agents. This meant that the shill agents would not continue toward the goal. The group of flocking marks and shills ended up stalling behind the object while the leader continued on to the goal location.

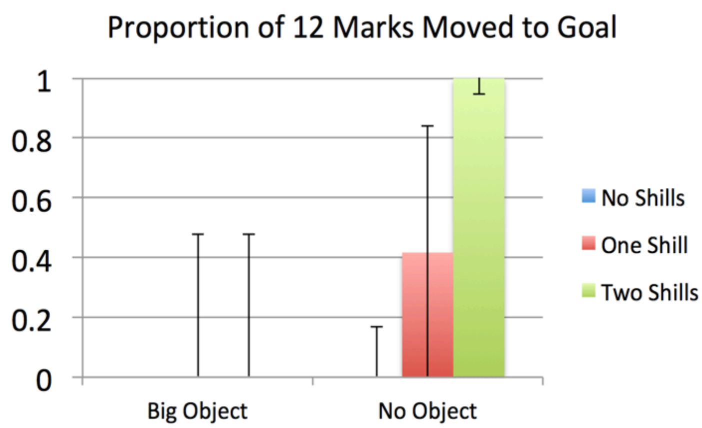

Figure 4: There were significant differences between all groups with no object present. The shill agents facilitated the misdirection. With a large object present, there was not a significant difference between groups. All groups with the large object had a median value of 0. The object obscured the leader agent from the shills.

4.2 Discussion

When teams of marks are "naive" (their thresholds for flocking are universally low), shills are not necessary to successfully misdirect them.

The simulations in which the mark agents all had low thresholds for flocking (thresholds of 1 or 2) shills did not make a difference. The leader alone was able to bring the marks from the start location to the goal location (cf. Pied Piper story). The medians for these conditions were all at or near one, and there were no significant differences between groups.

When teams of marks contain agents with higher flocking thresholds, a leader alone is often not able to successfully misdirect them.

The conditions in which the marks had higher thresholds for flocking had medians of 0 when a leader agent was the only member of the deception team. The use of the shills aided in the misdirection under these conditions.

The weight of a shill's lekking behavior must be low enough to prevent it from dominating the follow the leader behavior.

The conditions in which the shills had a high *Lek Behavior* weight were not significantly different from the conditions in which only a leader agent composed the deceiving team. It appears that in these conditions the *Lek Behavior* vector cancels the vector produced by the *Follow Leader Behavior* when the flock of marks and shills approaches the leader agent. This prevents the shills from pulling the marks all the way to the goal location. Often the group stopped just short of the goal location.

If the deceptive team is going to function effectively, shill agents must be able to see the leader agent throughout the deception or the shill agent must have knowledge of the goal location.

In the set of simulations with no object present in the environment, high flocking thresholds, and a low *Lek Behavior* weight for the shill agents, the marks were successfully moved to the goal location in all trials except one when two shills were present. With a large object inserted into the environment, the median proportion of marks moved to the goal location was 0. The shills lost sight of the leading agent and were unable to help drag the marks to the goal location. The shills may need to incorporate additional behaviors to keep the leader agent within view or may need to have additional knowledge about the goal under these conditions.

5 Conclusions and Ongoing Work

This paper explored how shill agents could be used to facilitate the misdirection of a team of mark agents. The simulations presented here show that in cases where mark agents have low thresholds for flocking together, a leading agent that moves with intent is sufficient to pull agents from a start location to a goal location. In cases where the mark agents have higher thresholds for flocking, shill agents help to carry out the misdirection. The shills facilitate misdirection when they do not lose sight of the leading agent and when the influence of the flocking marks does not dominate the influence of the leading agent on their motion.

Shills could employ more complex behaviors in order to more effectively misdirect mark agents. For example, they could observe and model the movement of the mark agents and coordinate their behavior to optimize the deception. Any additional behaviors employed by the shill agents, however, further differentiates them from the marks (providing opportunities to spoil the deception). In this study, shill agents were designed to be as simple and indistinguishable from the marks as possible.

In ongoing research, we are developing counter-deceptive strategies for these scenarios. We are currently evaluating if counteragents that employ novel strategies to deter misdirection can overcome the deceptive practices of the opposing team.

Acknowledgments

This research is supported by the National Science Foundation under CNS EAGE grant #1848653.

Appendix I: The definitions of the robotic behaviors discussed in the text appear below.

A) Off Robots Behavior: The agent is repelled from a robot with variable sphere of influence and variable gain. The vector returned from the behavior is the sum of the individual vectors for all of the robots. This keeps the robots from crashing into one another in the simulation.

$$V_{magnitude} = \begin{cases} \frac{\mathrm{R} - \mathrm{d}}{\mathrm{R}}, & \mathrm{d} \leq \mathrm{R} \\ 0, & \mathrm{d} > \mathrm{R} \end{cases}$$

$\mathrm{V}_{direction}$ = Direction from the center of the other robot to this robot's center

where:
R = Radius of the repulsion sphere
d = Distance of robot to another robot

B) Lek Behavior: The agent is attracted to a surrounding robot that is moving faster than the speed threshold with a variable gain and a variable region of influence. The vector returned is the sum of the vectors for all agents. This is for group formation in the simulation.

$$V_{magnitude} = \begin{cases} 1 - \frac{(\mathrm{A} - \mathrm{D}) - (\mathrm{d} - \mathrm{D})}{\mathrm{A} - \mathrm{D}}, & \mathrm{D} \leq \mathrm{d} \leq \mathrm{A} \\ 0, & \text{otherwise} \end{cases}$$

$\mathrm{V}_{direction}$ = Direction from the center of this robot to the other robot's center

where:
A = Radius of the attraction sphere
D = Radius of the dead zone sphere
d = Distance of robot to another robot

C) Avoid-Obstacle Behavior: Agent is repelled from obstacles (objects in the environment) with variable gain and sphere of influence. The robot avoids designated obstacles in the environment.

$$V_{magnitude} = \begin{cases} \infty, & \mathrm{d} \leq \mathrm{r} \\ \frac{max - d}{max - r}, & r < d \leq max \\ 0, & \max > \mathrm{d} \end{cases}$$

$\mathrm{V}_{direction}$ = Direction from the center of the object to this robot's center

where:
max = Radius of obstacle detection sphere
r = Radius of the circular obstacle
d = Distance of robot to the center of the obstacle

D) Follow Leader Behavior: A shill agent is attracted to the position of the leader with variable gain and sphere of influence. The leader is pulling it toward the goal.

$$V_{magnitude} = \begin{cases} 1 - \frac{\mathrm{R_L} - \mathrm{d}}{\mathrm{R_L}}, & \mathrm{d} \leq \mathrm{R_L} \\ 0, & \mathrm{d} > \mathrm{R_L} \end{cases}$$

$\mathrm{V}_{direction}$ = Direction from the center of this shill robot to the leader's center

where:
$\mathrm{R_L}$ = Radius of the region in which agents are attracted to the shill
d = Distance of robot to the shill

E) Go To Goal Location Behavior: Agent is attracted to a goal location. This moves the agent in the direction of a designated goal location.

$V_{magnitude}$ = Adjustable gain value

$\mathrm{V}_{direction}$ = Direction to the goal location from the robot's center

Appendix II: The parameters used for the behavior assemblages appear below. The parameters for the set of simulations in which the Lek Behavior was changed appear in parentheses

Leader Agent Parameter	Value	Units
Lead to Goal Behavior Assemblage		
x position of goal	240	m
y position of goal	40	m
Move to Goal Gain	1	
Wander Gain	0.9	
Avoid Obstacle Gain	1	
Avoid Obstacle Sphere	3	m
Avoid Obstacle Safety margin	.5	m

Mark Agent Parameter	Value	Units
Anchored Wander Behavior Assemblage		
Stay Near Start/Goal Gain	.2	
x coordinate of start location	20	m
y coordinate of start location	40	m
Start location attraction radius	10	m
Avoid Obstacle Gain	.2	
Avoid Obstacle Sphere	6	m
Avoid Obstacle Safety margin	.5	m
Avoid Robots Gain	.2	
Repel Sphere	2	m
Wander Gain	.2	
Mark Mill Around Behavior Assemblage		
Avoid Robots Gain	.2	
Repel Sphere	2	m
Avoid Obstacle Gain	.2	
Avoid Obstacle Sphere	6	m
Avoid Obstacle Safety Margin	.5	m
Wander Gain	.2	
Mark Flock Behavior Assemblage		
Avoid Robots Gain	1	
Repel Sphere	2	m
Lek Gain	1	
Attract Sphere	40	m
Dead Zone Sphere	4	m
Wander Gain	.5	
Avoid Obstacle Gain	1	
Avoid Obstacle Sphere	6	m
Avoid Obstacle Safety Margin	.5	m

References

[1] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin, "Effective leadership and decision-making in animal groups on the move." *Nature* 433.7025 (2005): 513.

[2] B. A. Duncan, P. D. Ulam, and R. C. Arkin, "Lek behavior as a model for multi-robot systems." *Proc. 2009 IEEE International Conference on Robotics and Biomimetics (ROBIO)* (2009): 25-32, IEEE.

[3] J. M. Epstein, "Modeling civil violence: An agent-based computational approach." *Proceedings of the National Academy of Sciences* 99.suppl 3 (2002): 7243-7250.

[4] S. Gade, A. A. Paranjape, and S. Chung, "Herding a flock of birds approaching an airport using an unmanned aerial vehicle." *AIAA Guidance, Navigation, and Control Conference.* 2015.

[5] M. Granovetter, "Threshold models of collective behavior." *American journal of sociology* 83.6 (1978): 1420-1443.

[6] A. A. Paranjape, S. Chung, K. Kim, and D. H. Shim, "Robotic herding of a flock of birds using an unmanned aerial vehicle." *IEEE Transactions on Robotics* 34.4 (2018): 901-915.

[7] M. J. Pettinati and R. C. Arkin, "Push and Pull: Shepherding Multi-Agent Robot Teams in Adversarial Situations", *15th IEEE International Conference on Advanced Robotics and Its Social Impacts (ARSO 2019)*, Beijing, China

[8] Sanghvi, N., Nagavalli, S., and Sycara, K. "Exploiting robotic swarm characteristics for adversarial subversion in coverage task". In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 2017.

[9] J. Shim and R. C. Arkin. "Biologically-inspired deceptive behavior for a robot." *International Conference on Simulation of Adaptive Behavior*. Springer, Berlin, Heidelberg, 2012.

[10] J. Shim and R. C. Arkin, "A taxonomy of robot deception and its benefits in HRI." *2013 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, 2013.

[11] R. Vaughan, N. Sumpter, R. Frost, S. Cameron, "Robot sheepdog project achieves automatic flock control." *Proc. Fifth International Conference on the Simulation of Adaptive Behaviour*. Vol. 489. 1998.

[12] R. Vaughan, N. Sumpter, J. Henderson, A. Frost, and S. Cameron, "Experiments in automatic flock control." *Robotics and autonomous systems* 31.1-2 (2000): 109-117.