

Push and Pull: Shepherding Multi-Agent Robot Teams in Adversarial Situations *

Michael J. Pettinati, *Graduate Student Member, IEEE* and Ronald C. Arkin, *Fellow, IEEE*

Abstract— Teams of robots tasked with making critical decisions in competitive environments are at risk for being shepherded or misdirected to a location that is advantageous for a competing team. Our lab is working to understand how adversarial teams of robots can successfully move their competition to desired locations in part so that we can then devise practices to counter these strategies and help make team functioning more successful and secure. In this paper, preliminary research is presented that studies how a team of robots can be shepherded or misdirected to a disadvantageous location. We draw inspiration from herding practices as well as deceptive practices seen in higher-order primates and humans. We define behaviors for the target (mark) agents to be moved as well as members of the shepherding team (a pushing agent and pulling shills) and present simulation results showing how these behaviors move robots to a desired location. These behaviors were implemented and trialed on hardware platform. A discussion of ongoing research into understanding misdirection in multi-robot teams concludes this paper.

I. INTRODUCTION

As teams of robots enter into environments where they are depended upon to make critical decisions, it is imperative to understand what can hinder effective team functioning. This is particularly true when a team of robots is being tasked with operating in a competitive environment. Whether the robot team is for entertainment as in the RoboCup competitions or has a serious assignment like military operations, adversaries to the team are going to do what they can to disrupt team functions.

In such environments, it is reasonable to expect adversaries to try and shepherd their opponents to locations where they have a tactical advantage. This misdirection will require deceptive acts through which adversaries guide unsuspecting mark agents. Our lab is studying how teams can successfully misdirect groups of other agents, in part, to understand how competitors might counter the practices used against them and continue to function effectively.

Our lab has studied numerous aspects of how multi-robot teams function, for instance: how teams can form meaningful task-based groups [3] and how teams can overcome failures in hostile environments to continue communicating and carrying out their task [12, 15]. In addition to the work on multi-robot teams, previous work has produced extensive research related to robot deception (e.g. [10, 16]), including the first taxonomy of human-robot deceptive activities [10].

Up until this point, we have focused on individual agents deceiving.

This paper is a preliminary work that focuses on how a team of robots can be shepherded to a location that may put them at a disadvantage. Fear-based shepherding approaches have been used to move animals from one location to another [4, 7, 13, 14]. In these works, a robot acts as a predator to move a group of animals toward or away from a certain location. These “push” approaches draw on knowledge garnered from herding practices and are often species specific.

This new research explores how a pushing agent can be combined with pulling confederate agents (shills) that are embedded within the group of mark agents (mark agents) to effectively and efficiently move the group. The pushing agent can emit a signal that pushes mark agents away from its location. The shill agents, which initially blend with the mark agents, can emit a signal that attracts mark agents to their locations (a follow me call). This signal may deceptively signal safety, something of interest, etc. As discussed below, animals have used similar deceptive signals to move conspecific groups and human teams have targeted marks using similar push and pull strategies.

The next section of the paper gives an overview of previous work using robotic agents to push (herd) groups (of animals) from one location to another. The section also briefly discusses deceptive signals that have been observed in animal populations as well as those that have been used by humans to misdirect animals and other humans.

The third section introduces the different types of robot agents, both the marks (the agents to be moved) as well as the agents of the shepherding team (pushing agents and pulling shills). The behaviors of these robots are mathematically defined, and the robotic missions for these agents are given. The fourth section presents simulation results given the behavioral definitions and discusses mission runs with a hardware implementation that mirror what was seen in simulation.

The paper concludes with a section discussing what is next in the project, looking at misdirection in multi-robot teams. This includes additional behaviors that can be incorporated into the agents that are present as well as additional deception strategies that can be used to misdirect multi-robot teams.

II. RELATED WORK

A. Herding Groups of Animals Using Robots

There has been research into moving animals against their will using robotic agents and wearable technology [e.g. 2, 4,

*Research supported by the National Science Foundation under CNS EAGER grant #1848653.

M. J. Pettinati (email: mpettinati3@gatech.edu) and R. C. Arkin (email: arkin@cc.gatech.edu) are with the School of Interactive Computing at the Georgia Institute of Technology, Atlanta, GA 30308.

7, 13, 14]. Researchers have focused on using fear, as opposed to deception, to drive certain flocks or herds of animals from one location to another or away from a certain location. This often involves developing algorithms that depend on the flocking behaviors of the animals that the robots are attempting to corral or divert.

For example, studies have been conducted on both physical and simulation systems that divert flocks of birds approaching an airport to a safe zone [4, 7]. Some of this research attempts to describe how the actual species of birds that threaten airports move [7] while others use more general flocking models [4]. In both cases, the modeled birds move together; the pursuer uses the geometric information about how they move together to keep them grouped and to direct them away from the airport to a safe zone.

Vaughan et al. [13, 14] used a sheepdog robot to herd ducks to a desired area. The behaviors of the herding robot were very simple. One robot moved toward the flock with a velocity based on how far the flock was from the goal and was repelled by the goal location [14]. Another robot pushed from behind the flock toward the goal location [13]. To successfully move the ducks (acting in place of sheep) to a goal location, these algorithms depend on the selfish flocking behaviors of these animals [1, 11].

An animal, when threatened by a predator, will try to limit the distance between itself and a conspecific or seek the center of a group of conspecifics to limit the probability the predator attacks it [1]. This “selfish-herding” behavior causes the animals to organize into groups or subgroups that can be pushed. Even with this tendency to congregate, however, the simple algorithms often split the groups of ducks and more complex behaviors were suggested to maintain the group formation [14]. Shepherds use a two-step process to herd animals; they group the animals and push the animals in the direction of a goal location. When the animals begin to split, they repeat the process [11].

Groups of robots are not going to necessarily demonstrate these types of selfish-herding behaviors or flocking behaviors that are present in animal populations. We are interested in how a multi-agent team of robots can be moved from one location to another. We draw inspiration from deceptive tactics introduced in the next subsection that are found in humans as well as other social species.

B. Deceptive Signals to Push and Pull Marks

Researchers have found evidence of primates, such as the wild tufted capuchin monkeys, using deception to take advantage of higher-power conspecifics [17]. Subordinates use predator calls to procure food when it is most highly contested. The call sends conspecifics away from the location of the subordinate and allows the subordinate to procure the food. An adversarial robot can push other agents away from a location like the subordinate monkey pushes conspecifics away from the location (using some type of alert call). This would scatter robots from the location.

If this deceptive agent were working with a team, however, other agents might be able to draw the mark robots to them. Many animals that are to be slaughtered follow a “Judas” animal into the slaughterhouse [8]. The “Judas”

sheep, which lives at the slaughterhouse, uses the fact that sheep have strong flocking instincts (an instinctual trust) to lead the others to slaughter. A robot could use a call to attract or pull others to a location as the other robot pushes them from behind. This is the approach introduced in this paper.

This push and pull process is used in human con games such as the three shell game [6]. Someone who appears to be an outside observer to the game pushes a mark to play. The “inside man” engages the mark with the game and pulls him in. The outside observer stays with the mark continuing to push him to play and assuring him of a big take while the inside man entertains the mark, pulling him into playing with his big personality and spiel. This continues until the mark places a bet for all he has with him (which he loses).

Potential marks (individuals to be deceived) for the three-shell con game have different levels of gullibility. While one player might be swayed to play by the push of a single person and the pull of the game’s host, others may be more skeptical. The con men use “shill” players to populate the game while targeting their next mark to pull people into the game [6]. As individuals see others participating and winning, they are drawn to try it. This is related to Granovetter’s [5] model of collective behavior. People have different thresholds at which point they follow the crowd.

Similarly, robotic agents might need to see a certain number of other agents moving in a certain direction (being pushed and pulled in a certain direction) before they begin to go. The robots within individual teams can have different threshold to determine their response to repulsive agents or pulling agents that try to drive them to a certain end location. It is only when the agent sees enough agents moving that it begins to move along with the group.

III. AGENT MODELS

In this paper, three distinct types of robotic agents are considered when exploring team misdirection. There are the mark agents that the shepherding agents are trying to move from a start location where they are initially grouped (falling within a certain radius from a central point) to a goal location where they are similarly grouped. There is a repulsive or pushing agent that is trying to push the mark agents toward the goal location. There are shill agents that pull the mark agents toward the goal. A general overview of the behaviors of the specific agents is given in the following subsection. The mathematical models are then presented in the second subsection. The mission definitions are given in the third subsection.

A. Behaviors Overview

The pushing agent is meant to continually push the group of marks toward the goal location. The role of this agent is much like a sheepdog trying to move a herd of sheep or the outside man in the three-shell game urging a mark to play. The agent pushes the mark agent toward the location the mark agent would not otherwise go. The mission of this agent consists of three simple behaviors:

- It needs to pursue the group from behind.
- It needs to avoid other robots.
- It needs to avoid environmental obstacles.

These three behaviors are explicitly defined in Appendix I as the *Pursue Behavior*, the *Repel Behavior*, and the *Avoid-Obstacle Behavior*. The pursue behavior is similar to Vaughan et al.'s herding robot [14]. The pursuing agent computes the centroid the mark agents it sees and is attracted to a point behind this centroid (in the direction from the goal to the centroid). Environmental objects as well as other robots repel the pushing agent when it is within a certain radius of them.

The mark agents are the agents to be moved from one location to another. These could be thought of as the sheep to be herded or the player to be moved to play the three-shell game. The mark agents are not responsive to the other agents in the simulation (aside from not "crashing into them") until a threshold has been reached.

As described above, each agent has a threshold, which determines when it starts to respond to other agents in this simulation. This threshold is the number of agents the mark agent has to "see" that are urging it to move (including the pushing agent(s) and skill(s)) or other mark agents that are moving in response to these pushing/pulling agents to begin responding to the other agents. This threshold model is based on collective behavior seen in humans [5].

A crowd grows as more and more individuals begin to move collectively in a certain direction. They respond to the people around them moving in the direction (the individual is moving or acting in a certain way because the people around him/her are moving or acting in a certain way).

These marks agents, when they are not above their threshold, are:

- Wandering.
- Avoiding other Robots.
- Avoiding obstacles in the environment.
- Staying near the start/goal if close to these points.

These four behaviors are defined below in the *Repel Behavior*, the *Avoid-Obstacle Behavior*, the *Stay Near Start/Goal Behavior*, and the *Go To Goal Behavior*. The agents begin grouped and wander without crashing into other robots or objects when they are below their threshold. They are not responsive to the "call" of skill agents or the push of the repelling agent until their threshold is reached.

Once the threshold is reached, the mark agents are:

- Following the closest skill.
- Avoiding the repulsive agent.
- Avoiding other robots.
- Avoiding obstacles in the environment.
- Wandering.

The mathematics for these five behaviors is defined in the *Repel Behavior*, the *Avoid-Obstacle Behavior*, the *Wander Behavior*, and the *Follow Skill Behavior* below. Mark agents are attracted to the closest skill agent; objects in the environment and other agents in the environment repel them. The repulsive sphere of the pushing agent is larger than other

marks as shown below in the simulation setup section. There is noise introduced into how the mark agents move as well (the wander behavior).

Finally, skill agents play the role of the skills in the three-shell game. The skills in this game act as any other player; they populate the game to draw the mark into playing. Skill agents in the simulation are agents that "know" where they are going and draw the mark agents to the goal location with a call. They mirror the mark behavior by staying near the start location at the beginning. At the time the pushing agent begins to try and push the mark agents, the skill agent will:

- Head directly to the goal.
- Avoid Objects.
- Wander.

The mathematics of these three behaviors is shown in *Go To Goal*, *Avoid-Obstacle* and *Wander*. When the pushing agent begins to try to push the group, the skills move toward the goal with noise incorporated into the path to that goal.

B. Mathematical Models

Each behavior described in the previous section outputs a motion vector for the particular agent. A final motion vector for the agent is a weighted sum of each behavior vector. See Appendix I for the mathematical behavior definitions.

When the pushing agent is trying to move the mark agents from the start to the goal, the pushing agent's final motion vector is a weighted sum of a vector that points at a point behind the mark agents the pushing agent is able to see (*Pursue Behavior*), a vector that is the sum of repel vectors from surrounding robots (*Repel Behavior*), and a vector that is the sum of the repel vectors from surrounding objects (*Avoid-Obstacle Behavior*).

Before the pushing agent begins its pursuit of the marks, the skill agents mirror the mark agents. The final motion vector of the skills is a weighted sum of a vector pointing to the start point (*Stay Near Start/Goal Behavior*), a noise vector (*Wander Behavior*), a vector that is the sum of the repel vectors from other robots (*Spread Behavior*) and objects (*Avoid-Obstacle Behavior*).

As soon as the pushing agent begins to pursue the marks, the skill agents have a final motion vector that is a weighted sum between a vector point toward the goal location (*Go To Goal Behavior*), a vector avoiding objects (*Avoid-Object Behavior*), and a noise vector (*Wander Behavior*).

The marks agent's final motion vector when it is below its threshold is the same as the skills before the pushing agent begins its pursuit of the marks. After agents have their threshold attained, the motion vector is a weighted average of a vector away from the pushing agent (*Repel Behavior*), toward the closest skill (*Follow Closest Skill Behavior*), a vector that is the sum of the repel vectors from other robots (*Spread Behavior*) and objects (*Avoid-Obstacle Behavior*), and a noise vector (*Wander Behavior*).

C. Robot Missions

The FSA diagrams defining the missions of each of the three types of robots are shown in Figures 1, 2, and 3. The

pushing robot (Figure 3) is set behind the group of marks and shills, which enact the Anchored Wander Behavioral Assemblage. This has them wandering around a starting point. After a brief pause, the pushing agent enacts the Push to Goal Behavioral Assemblage until all of the marks have been driven in to the goal location.

The mark agents (Fig. 1), when below their thresholds, will be enacting their Anchored Wander Behavioral Assemblage. They will wander randomly while avoiding objects/other robots when in the area between the start and goal. They will wander in the area of the start/goal if they are in that area. When the threshold has been satisfied, they will enter the Follow Shill Behavioral Assemblage, which includes being repelled by the pushing agent, following the closest shill and avoiding other robots/objects.

Finally, the shill agents (Fig. 2) initially enact the Anchored Wander Behavioral Assemblage to blend with the




Figure 1: The FSA defining the mission for the mark robots.




Figure 2: The FSA defining the mission for a pulling/shill agent.




Figure 3: The FSA defining the mission for the pushing robot.

marks. As soon as the pushing agent begins its pursuit it moves toward the goal with noise and obstacle avoidance incorporated into its movement.

IV. RESULTS

For the simulation, there were four independent variables that were manipulated to generate thirty-six different independent conditions. The independent variables were:

- The number of the mark agents to move.
- The number of pulling (shill) agents.
- The number of pushing agents.
- The complexity of the environment.

There were either four or twelve mark agents. There were zero, one, two, or three pulling agents (shills). There was zero or one pushing agent. There was no obstacle in the environment, a small obstacle in the environment or a large obstacle in the environment.

There were twenty trials run for each of the 36 conditions. The trials were run until all of the mark agents reached the goal location or until 2000 simulation steps had finished. The measures collected were the number of steps to successfully move all of the agents from the start to the goal (in trials when all agents were successfully moved from start to goal) and the proportion of agents that were successfully moved from start to goal when 2000 time steps had passed.

The two-dimensional simulation environment was 60 meters by 240 meters. The mark and shill agents began by wandering grouped in a starting area that had a radius of 10 meters that was centered at (20, 40) in the mission area. The goal location was centered at (220, 40) and had the same radius of 10 meters. The large obstacle had a radius of 10 meters and was centered at (140, 40). The small obstacle had a radius of 3 meters and was centered at (140, 40).

Robots only contributed to each other's motion if they could "see" each other. That is, if a line segment between the centers of robots intersected the circular obstacle, the robots did not influence one another's motion or the active behavior.

Snapshots from a trial with one shill, one pushing agent, a small object, and four marks are shown in Figure 4 (next page). A video of the full trial is available¹. The pushing agent begins blue (all other agents are yellow in the Anchored Wander Behavioral Assemblage). The pushing agent turns orange (begins pursuing the marks) turning the shill agents brown and sending them toward the goal location. The mark agents are green when above their threshold and enacting the Follow Shill Behavioral Assemblage.

As described above, each agent has a threshold, which determines when it starts to respond to other agents in this simulation. This threshold is the number of agents the mark agent has to "see" that are urging it to move or other mark agents that are moving in response to these pushing/pulling agents to begin responding to the other agents. Half of the mark agents had a threshold of 1 and half had a threshold of 2

¹ https://www.cc.gatech.edu/ai/robot-lab/Deception_2019/Videos/Object.mov

in all simulation trials. The parameters for each of the three agents are shown in the tables in Appendix 2.




Figure 4: Snapshots showing trial where four mark agents were successfully moved from the start to the goal location with one skill and the pushing agent.

A. Simulation Results

To begin, we note that the push only trials (the trials with no pulling skill) were not successful at moving the mark agents to the goal location. In all six conditions (three with 4 marks and three with 12 marks), the pushing agent was not able to move any of the mark agents from start to goal in 2000 simulation steps.

The pushing agent divided the mark agents when it began to pursue them. This is similar to what was seen in with the trials with Vaughn et al. [14] except the mark agents in our simulation did not incorporate a selfish-herd behavior so they split even more easily. The push agent fell into a local minimum at the center of the mark agents it was pursuing. It should be noted that adding additional noise and persistence would allow this pushing agent to break out of this local minimum faster and push at least a single agent to the goal location. Given sufficient steps (beyond the 2000), the push agent did break out of the local minimum and push a single agent to the goal location.

Because the trials without skills had no variation, we compared the conditions that had skills with one another. We began by running Kruskal-Wallis Tests for each environment and each number of marks. This test was used because there were non-normal distributions within the groups as well as unequal variances between groups. There was a significant difference ($p < .05$) between groups in the proportion of marks that were successfully moved from the start location to the goal location. Post hoc testing revealed that the push and pull method, regardless of the number of skills, was significantly better than the pure pull method at moving marks from start to goal. There was not a significant difference in the efficacy with which the marks were moved when looking at different numbers of skills being used. These results can be seen in Figures 5 and 6 as well as tables 1 and 2.

We paired the conditions with push and pull agents with respect to the environment's complexity and for the number of mark agents. We ran two-sample, two-tailed t-tests in cases where the assumptions of normality and equal variance

held. In cases when the normality did not hold, we ran the Mann-Whitney U test, a nonparametric test.

In conditions with four mark agents, there was evidence that the single skill outperformed two skills with respect to how efficiently the marks could be moved from start to goal.

Table 1: Summarizes the proportion of marks successfully moved from the start to goal location when 4 marks were used. Pushing and pulling was better able to shepherd agents than either alone.

Environment	Proportion of 4 Marks Moved to Goal Location				
	1 Pushing / 2 Pulling Agents Mean/Median (Standard Deviation) (n = 20)	1 Pushing / 1 Pulling Agents Mean/Median (Standard Deviation) (n = 20)	1 Pulling Agent Mean/Median (Standard Deviation) (n = 20)	2 Pulling Agents Mean/Median (Standard Deviation) (n = 20)	3 Pulling Agents Mean/Median (Standard Deviation) (n = 20)
Big Object	.575/.75 (.381)	.7875/1.0 (.327)	.1/.0 (.150)	.15/.25 (.150)	.1/.0 (.170)
Small Object	.825/1.0 (.245)	.9/1.0 (.150)	.175/.0 (.216)	.138/.0 (.172)	.213/.25 (.186)
No Object	.925/1.0 (.200)	1.0/1.0 (0.0)	.213/.25 (.186)	.175/.25 (.183)	.25/.25 (.181)




Figure 5: This shows the mean proportion of marks that were successfully moved from the start location in each condition with 4. Conditions with pushing and pulling agents were better able to guide agents than pushing or pulling alone.

Table 2: Summarizes the proportion of marks successfully moved from the start to goal location when 12 marks were used. Pushing and pulling was better able to shepherd agents than either alone.

Environment	Proportion of 12 Marks Moved to Goal Location				
	1 Pushing / 2 Pulling Agents Mean/Median (Standard Deviation) (n = 20)	1 Pushing / 1 Pulling Agents Mean/Median (Standard Deviation) (n = 20)	1 Pulling Agent Mean/Median (Standard Deviation) (n = 20)	2 Pulling Agents Mean/Median (Standard Deviation) (n = 20)	3 Pulling Agents Mean/Median (Standard Deviation) (n = 20)
Big Object	.442/.458 (.188)	.421/.50 (.192)	.058/.083 (.055)	.046/.042 (.050)	.054/.083 (.056)
Small Object	.404/.417 (.217)	.379/.333 (.220)	.063/.042 (.076)	.075/.083 (.071)	.092/.083 (.10)
No Object	.579/.625 (.196)	.463/.583 (.268)	.075/.083 (.060)	.113/.083 (.062)	.133/.125 (.091)




Figure 6: This shows the mean proportion of marks that were successfully moved from the start location in each condition with 12. Conditions with pushing and pulling agents were better able to guide agents than pushing or pulling alone.

In environments with no object and a small object, there was a significant difference ($p < .05$) between the one skill and two skill cases. In successful trials where all four agents were moved from the start location to the goal location, the one skill did it faster (in fewer steps) than the two skills. The two skills brought the agents to the goal with less cohesion. See Figures 7 and 8 and Tables 3 and 4.

There was no significant difference between the two conditions ($p > .05$) in the large object case. In both conditions, the robots got caught on the large object at the midpoint, which was a larger factor than the cohesion with respect to the number of simulation steps it took to get to the goal location. There was also variation in how long it took for the agents to get around the object. Additional noise in the movements may have helped move around the object more quickly.

There were clear differences in the efficacy with which one and two skills along with the pushing agent could move four marks from the start to the goal location versus the efficacy with which they could move twelve marks from the start to the goal location. In the no object and four mark conditions, there was very little variance. In all twenty trials with one skill, all four marks were successfully moved; in seventeen of twenty two-skill cases the four marks were successfully moved. In the cases with twelve-mark agents, there was not a single trial where all twelve agents made the goal (for the one or the two-skill cases).

In three of the four conditions with variance, statistical tests found extremely significant ($p < .001$) differences between the twelve-mark condition and the four-mark condition. The proportion of marks that could be moved by the skills was higher in the four-mark cases. The only exception was the two-skill condition when the big obstacle was present in the environment.

B. Demonstration on Hardware

We implemented the robotic behaviors described above in the Robotarium at Georgia Tech [9]. The parameters were adapted to the much smaller environment and three exemplary trials were run with no object present and three marks present. The three trials consisted of no skill present, one skill present, and two skills present (with one pushing agent). On the trials run on physical systems, we witnessed the same type of dynamics as in the simulation.

With no skill present, the group of three agents split up and the pushing robot was in a local minimum for a brief period of time before breaking free and pushing a single agent to the goal location. The video of this trial is available for viewing². In each of the trials with the skills, the agents moved smoothly from the start configuration to the goal location. Snapshots from a trial with one skill can be seen in Figure 9 (next page). The video of the full trial is available for viewing³.

Table 3: In an environment without obstacles 1 skill could move four marks to the goal location more efficiently than 2 skills.

Simulation Steps to Move All 4 Marks from Start to Goal	Two-Tailed Mann-Whitney U Test Result			
	1-Skill Median (Standard Deviation) (n = 20)	2-Skill Median (Standard Deviation) (n = 17)	U	p-value
No Object	1053.5 (44.359)	1077.0 (29.794)	U= 94.5	p= .02202




Figure 7: Chart showing how it required fewer simulation steps to shepherd four marks with 1 skill than 2 skills.

Table 4: In an environment with one small obstacle 1 skill moved four marks to the goal location more efficiently than 2 skills.

Simulation Steps to Move All 4 Marks from Start to Goal	Two-Tailed Mann-Whitney U Test Result			
	1-Skill Median (Standard Deviation) (n = 11)	2-Skill Median (Standard Deviation) (n = 13)	U	p-value
Small Object	1052.0 (20.328)	1099.0 (26.868)	U= 16.5	p= .00158




Figure 8: Chart showing how it required fewer simulation steps to shepherd four marks with 1 skill than 2 skills.

V. CONCLUSION AND FUTURE WORK

In this paper, we presented simulation results as well as results on physical systems showing how pushing and pulling (skill) agents can shepherd mark robots from a start location to a goal location. *The combined push and pull shepherding approach was superior to both push and pull alone with respect to the number of agents that could be moved from the start location to the goal location successfully.* The pushing agent divided the mark agents; it

²https://www.cc.gatech.edu/ai/robot-lab/Deception_2019/Videos/noshills.mp4

³https://www.cc.gatech.edu/ai/robot-lab/Deception_2019/Videos/1skill.mp4




Figure 9: Snapshots of a pushing agent and single skill successfully moving three marks from a start location to a goal location. Physical hardware implementation

pushed them out of the start area, but it was unable to get any more than one agent into the goal location. It was caught between agents in a local minimum. It bounced out of that local minimum and pursued a single agent to the goal location. The pulling agents (in cases without the pushing agent) were sometimes able to get multiple agents from the start to the goal location, but on average, the push and pull approach was able to move significantly more mark agents from the start to the goal location.

In the future, we will investigate using more complex behaviors for a pushing agent (for example drawing inspiration from shepherds) as well as incorporate behaviors into our mark agents where they are attracted to one another. As noted in the second section of the paper, individual agents have had success pushing groups of agents to a specific location when those agents are herd animals. Our lab has explored situations in which robots are able to use a Lek behavior to form groups [3] as well as situations in hostile environments where robotic agents might be attracted to one another under certain conditions [15].

The number of marks influences the herding efficacy of the push and pull approach. The pushing and pulling agents were able to successfully move all marks from the start to the goal frequently in the four-mark conditions. In the twelve-mark conditions, the group consistently split. As noted above, we need to consider the use of more complex behaviors for the push agent. Skill agents may need to be responsive to mark agents as well.

Additional skills can be problematic for smaller groups of marks. A single skill was more efficient at moving four skills from the start to the goal location. The second skill caused the group to have less cohesion when it was small enough to stay as a single unit. We are going to have to consider in what situations additional skills will be useful.

APPENDIX I

This appendix contains the behavioral formulas for each agent tested in the simulation.

A) Pursue Behavior: Attraction to a point behind the centroid of the mark robots. The push robot computes the centroid of the mark robots within view and moves to the point a set distance from the centroid in the direction from the goal location to that centroid. This puts the pushing agent behind a group of mark agents.

$$V_{magnitude} = \text{Adjustable gain value}$$

$$V_{direction} = \text{Direction from the center of the robot to the point behind the marks}$$

B) Spread or Repel Behaviors: Repel from like or different robot with variable gain and variable sphere of influences. Different types of robotic agents may have different spheres of influence.

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

$$V_{direction} = \text{Direction from the center of the other robot to this robot's center}$$

where:

R = Radius of the repulsion sphere (differs for different robot types)
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.

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

$$V_{direction} = \text{Direction from the center of the object to this robot's center}$$

where:

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

D) Stay Near Start/Goal Behavior: Agent is attracted to the start or the goal location. This allows the mark agents to wander within a certain radius of the start and goal locations.

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

$$V_{direction} = \text{Direction of the start or the goal point}$$

where:

R_R = Radius of the region in which agents are attracted to the start/goal
d = Distance from the robot's center to the start point/goal point

E) Follow Skill Behavior: Agent is attracted to the position of the closest skill with variable gain and sphere of influence. The skill is pulling it toward the goal location.

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

$$V_{direction} = \text{Direction from the center of this robot to the skill's center}$$

where:

R_S = Radius of the region in which agents are attracted to the skill
d = Distance of robot to the skill

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

$$V_{magnitude} = \text{Adjustable gain value}$$

$$V_{direction} = \text{Direction to the goal location from the robot's center}$$

G) Wander (Noise) Behavior: Agent moves in a random direction with a variable gain. This is to allow for agents' movements to incorporate realistic noise into them.

$$V_{magnitude} = \text{Adjustable gain value}$$

$$V_{direction} = \text{Random direction}$$

APPENDIX II

This appendix contains the parameters used for each simulation described in the paper.

Mark Agent Parameter	Value	Units
Anchored Wander Behavior Assemblage		
Stay Near Start/Goal Gain	1	
x coordinate of start location	20	m
y coordinate of start location	40	m
x coordinate of goal location	220	m
y coordinate of goal location	40	m
Start location attraction radius	10	m
Goal location attraction radius	10	m
Avoid Obstacle Gain	1	
Avoid Obstacle Sphere	6	m
Avoid Obstacle Safety margin	.5	m
Avoid Like Robots Gain	1	
Repel Sphere	2	m
Wander Gain	.75	
Follow Shill Behavior Assemblage		
Follow Shill Gain	1	
Attraction Sphere	10	m
Avoid Repulsive Agent Gain	1	
Repulsive Sphere	10	m
Avoid Object Gain	1	
Avoid Obstacle Sphere	6	m
Avoid Obstacle Safety margin	.5	m
Avoid Like Robots Gain	1	
Repel Sphere	3.5	m
Wander Gain	.25	
Shill Agent Parameter		
Anchored Wander Behavior Assemblage		
Stay Near Start/Goal Gain	1	
x coordinate of start location	20	m
y coordinate of start location	40	m
x coordinate of goal location	220	m
y coordinate of goal location	40	m
Start location attraction radius	10	m
Goal location attraction radius	10	m
Avoid Obstacle Gain	1	
Avoid Obstacle Sphere	6	m
Avoid Obstacle Safety margin	.5	m
Avoid Like Robots Gain	1	
Repel Sphere	2	m
Wander Gain	.75	
Go To Goal Location Assemblage		
Move to goal gain	1	
x coordinate of goal location	220	m
y coordinate of goal location	40	m
Wander gain	.7	
Avoid Obstacle Gain	1	
Avoid Obstacle Sphere	3	m
Avoid Obstacle Safety margin	.5	m
Repulsive/Pushing Agent Parameter		
Push to Goal Behavior Assemblage		
Pursue Gain	1	
Distance of pursuit	6	m
Attraction Sphere	15	m
x position of goal	220	m
y position of goal	40	m
Spread Gain	1	
Repel Sphere Other Repulsive Agents	5	m
Repel Sphere Mark Agents	2.5	m
Avoid Obstacle Gain	1	
Avoid Obstacle Sphere	6	m
Avoid Obstacle Safety margin	.5	m

REFERENCES

- [1] G. Beauchamp, "The Selfish Herd." In *Social Predation* (2014): 157-175, ELSEVIER.
- [2] Z. Butler, P. Corke, R. Peterson, and D. Rus, "Virtual fences for controlling cows." *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004.* Vol. 5. IEEE, 2004.
- [3] B. A. Duncan, P. D. Ulam, and R. C. Arkin, "Lek behavior as a model for multi-robot systems." In *2009 IEEE International Conference on Robotics and Biomimetics (ROBIO)* (2009): 25-32, IEEE.
- [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] D. W. Maurer, "The argot of the three-shell game." *American Speech*, 22(3) (1947): 161-170.
- [7] 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.
- [8] H. Pedersen, "Follow the Judas sheep: Materializing post-qualitative methodology in zooethnographic space." *International Journal of Qualitative Studies in Education* 26.6 (2013): 717-731.
- [9] D. Pickem, P. Glotfelter, L. Wang, M. Mote, A. Ames, E. Feron, and M. Egerstedt, "The Robotarium: A remotely accessible swarm robotics research testbed." In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, Singapore (submitted, available as *arXiv preprint*).
- [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] D. Strömbom, R. P. Mann, A. M. Wilson, S. Hailes, A. J. Morton, D. J. T. Sumpter, and A. J. King, "Solving the shepherd problem: heuristics for herding autonomous, interacting agents." *Journal of the royal society interface* 11 (2014): 20140719.
- [12] P. Ulam, and R. C. Arkin. "When good communication go bad: communications recovery for multi-robot teams." *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004.* Vol. 4. IEEE, 2004.
- [13] 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.
- [14] 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.
- [15] A. Wagner and R. Arkin. "Multi-robot communication-sensitive reconnaissance." *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004.* Vol. 5. IEEE, 2004.
- [16] A. R. Wagner and R. C. Arkin. "Acting deceptively: Providing robots with the capacity for deception." *International Journal of Social Robotics* 3.1 (2011): 5-26.
- [17] B. C. Wheeler, "Monkeys crying wolf? Tufted capuchin monkeys use anti-predator calls to usurp resources from conspecifics." *Proceedings of the Royal Society B: Biological Sciences* 276.1669 (2009): 3013-3018.

ACKNOWLEDGMENT

The authors thank Akshay Krishnan for completing the hardware implementation discussed above.