Affective Herding Behavior in UAVs

Jillian Greczek jgreczek@cse.tamu.edu
Brittany Duncan bduncan@cse.tamu.edu
Dr. Robin Murphy murphy@cse.tamu.edu

This paper presents biologically-based herding behavior for a small unmanned aerial system (sUAS). The motivation for a herding behavior is to use a small helicopter-like UAS to guide a group of people along a path, such as during emergency evacuation or to handle a crowd during a riot. Currently no robotic herding behaviors appear to exist. This paper synthesizes 15 studies from existing biological behavior literature, psychological literature and human-robot interaction studies into a single behavior exploiting affect. The behavior was implemented with the Unity3D engine and Javascript. One UAV was used, an AirRobot AR-100B, and its actual velocity and movement characteristics were simulation. Testing consisted of 1-to-3 "humans" moving down the hallway of variable width (6-20 m) from random starting locations at least 12 feet away from the end of the 14 m hallway. The tests measured how tightly the robot kept the targets close to the path and the end positions of the targets versus the number of targets and the width of the hallway. After 522 runs, the average deviation from the path was -0.08m for one target, -0.41m for two targets, and -0.75m for three targets in a hallway of average size 12.92m. The narrower the corridor, the more likely the robot was to be unable to position itself to herd one target. The success rate for one target suggests the utility of the biologically-based behavior, but declining effectiveness suggests that multiple UAVs might be more effective for crowds.

0 Introduction

This paper presents a novel approach to aiding evacuation and crowd control using a biologically-inspired behavior for an unmanned aerial vehicle (UAV). Small UAVs have been designed to operate indoors and it is easy to imagine one or more small UAVs flying over the heads of a crowd and descending into key crowd control positions. The use of low cost UAVs to herd crowds or evacuees could reduce risk to human life and allow emergency professionals to have a presence in areas humans could not otherwise get into. During an emergency such as a fire, victims could be guided toward exits. In a riot event, herding behavior could be used to contain crowds of rioters. However, research on herding behavior in UAVs is very limited; see for recent efforts [10, 12].

This paper presents the derivation and implementation of affective herding behavior based on biological behavior literature. Herding is defined as directing a target(s) to a desired location or along a desired path using physical cues, such as blocking, and affective cues, such as proxemics and gaze. Because of the lack of existing research on herding behavior, the behavior presented in this paper was drawn from literature about biological herding behaviors exhibited by animals in the real world. The biological basis of the behavior is advantageous because of the established affective component of animal behaviors.

The paper is organized as follows. Section 1 surveys the literature on herding in animals and robots, and a proposed behavior is described in Section 2. Section 3 discusses the implementation of the derived herding behavior. The behavior was then tested in simulation in the experiment described in Section 4. Section 5 summarizes the results of the experiment and discusses future work.

1 Related Work

Herding has been studied in animals [7, 8, 9, 10, 11, 12, 13, 14, 15], most significantly in 2D in ungulates [14] and 3D in dolphins [9] as well as for human crowd control [11]. Two guide robots [10, 12] have implemented herding behaviors, though not based on animal studies. The studies suggest that there are two mechanisms for herding: the primary mechanism is the relative position of the agents (gathering or driving) and the secondary mechanism is making the target uncomfortable while avoiding panic ("showing eye", silent herding, social gaze, etc.). The literature suggests that herding will terminate if a target gets too far. There was no indication of when a target triggered a herding reaction.

1.1 Ethological Herding Behaviors

Herding dog behaviors take two forms: *gathering* or *driving*. Almost all herding behavior in dogs involves getting between the target and the wrong way to go, as described in [7, 13, 15]. Wakeman [13] describes herding behavior in the ancient *puli* breed of herding dog, which unlike newer breeds described in [7], continually runs around herds of sheep to form a physical barrier around them. The puli will run very close to the sheep it is herding, and uses energetic bouncing and high-pitched barks to signal obedience from the herd. Wiese [15] describes two other forms of herding behavior in dogs, "gathering" and "driving." What the puli does can be described as gathering, whereas driving a herd of livestock in a certain direction is called driving. Wiese [15] further describes different types of leadership within the herding environment, where a dog can be either the "header" or the "heeler." A header dog will lead the herd at the front, whereas the heeler dog will follow from behind.

Dolphins use a form of gathering to herd prey in three dimensions. Beniot-Bird [9] describes the formations that dolphins form in order to concentrate prey into a small area where they can feed. To concentrate prey density, pelagic dolphins will move in a V-shaped pattern around their prey of up to 16-24 dolphins.

Gaze is also a very important component of herding behavior in dogs [7, 13, 15] and ungulates (hoofed animals) [14]. Hafez [7] describes the herding dog behavior known as "showing eye," where a sheep dog will stare down the sheep it is trying to herd, which is enough to intimidate an already skittish sheep without any barking or movement. Too much activity could scare the sheep into panic. Walther [14] describes a similar behavior in ungulates, where a male deer will keep females within his territory by getting in front of them and staring them down, which is called *silent herding*. A male deer will also herd females back into his territory by focusing only on one female at a time.

1.2 Human Herding Behaviors

Herding in humans is more challenging than with animals because humans may panic more readily. Raafat [11] describes how panic can subconsciously spread through a crowd of people being herded. Helbing [8] shows that herding a crowd of people into a more confined area can actually cause people to become more panicked, which is an example of how herding people differs from herding animals.

1.3 Robot Herding Implementations

Two museum guide robots exhibited ad hoc herding behaviors and discovered gaze, proximity, and orientation were essential. Pandey [10] describes a museum guide robot that led guests around and attempted to keep them engaged by following them and looking at them. In Pandey [10] a museum guide robot maintains human attention by using eye contact. The guide robot also had different states based on the human's engagement, indicated by proximity. The guide robot would also give up after the human moved far enough away from it, suggesting distance as a metric for premature termination of the behavior. Shiomi [12] describes a mall tour guide robot that had more success in keeping guests' attention when it "walked backwards" like a human tour guide would, enabling it to keep eye contact with the people it was leading. People standing around out of range of the tour guide robot were also more likely to join in and pay attention to the robot when the robot was oriented toward its guests.

1.4 Insights from Other Robots

Other robot studies indirectly suggest that familiarity [6], proxemics [2, 4, 6], gaze [2, 3, 5] and audio [5] may be effective affective herding mechanisms. Takayama and Pantofaru [6] show that the more unfamiliar a subject was with a robot, the more nervous they were around it, which could be exploited for dominance. Bethel and Murphy [2], Saerbeck and Bartneck [4], and Takayama and Pantofaru [6] describe how different proximities and motions can create negative reactions, again an effect that could be exploited to drive or gather targets. Shell and Matarić [5] used directional audio alarms to give evacuees guidance with orientation.

2 Approach

An ecological approach to behavior design following [2] considered the herding task and the unique abilities of a UAV for indoor spaces, leading to a heeler type of driving behavior. The herding behavior uses affect to adapt the robot's behavior based on how far the target has strayed from the path. The behavior is based on one target; if multiple targets deviate from the path, the nearest target is selected.

The ecology of the behavior consists of the robot's capabilities, the environment and what it affords, and the task itself. For the purposes of this behavior, helicopter-like UAVs were considered the representative robot, providing four degrees of freedom: *change hovering height, rotate (yaw), move side to side (roll),* and *move backwards and forward*. With these degrees of freedom, a robot could fly over the crowd, change height relative to human height, and maintain "eye contact" while flying backward. The UAV is assumed to be able to detect a target and determine relative range to the target. It is also assumed that the UAV has a recognizable "front," such as a camera payload or LED. The environment consists of the space and the targets. The space is any open area with sufficient overhead height for a UAV to safely fly over a crowd, such as a stadium field, stadium corridors, or streets. The targets are assumed to be moving at normal human walking speed and will turn 180 degrees to avoid a robot that has entered its personal space (within 1.22m of relative distance). The task was to keep the human "targets" moving forward along an imaginary line. The line represented the centerline of a stadium corridor or street, the general direction of a group, or an intended path and also provided a reference for measuring the efficacy of the herding.

A basic 2D *driving* pattern of action [15] was selected, with the robot as a heeler, following four phases of escalation: calm, cautious, aggressive, and ignore. Although the UAV can operate in 3D, targets work in 2D, thus a dolphin-like gathering of targets was not appropriate. As with animals, the primary mechanism is the relative position of the robot to the target. The presence of the robot creates a virtual repulsive vector on the target; if the target is behind the robot, the target goes forward, if the robot is on the left side of the target, the target moves to the right. The calm, cautious, and aggressive phases use affective cues as a secondary mechanism to amplify the primary mechanism. The phases were chosen ad hoc but capture the general effect of the silent herding and showing eye tactics.

State	Distance of Target from Line	UAV Max Speed	Min Distance Between UAV and Target	UAV Min Hover Height
Calm	$d <= D_c$	1 m/s	1.0 m	above
Caution	$D_c < d <= D_a$	1 m/s	0.7 m	above
Aggressive	$D_a < d <= D_i$	1 m/s	0.4 m	eye level
Ignore	$> D_i$	0.4 m/s	n/a	above

Table 1 Four phases of the proposed herding behavior.

Table 2 shows the four phases of escalating affective response. The robot begins by being assigned a *path line,* or the direction in which it should herd its targets. The path line is assumed to have a defined end point, such as a doorway, or in the case of this simulation, the end of a hallway. The robot begins behind the target. In the proposed pattern, the UAV stays behind the targets until they deviate a distance *d* from the intended path. Following Wiese [15], the robot acts as a "heeler" and begins calmly herding from the back of the group of targets staying 1.0 m behind. Here the target is aware of the intentional presence of the robot but is not threatened by it, but the robot is also close enough to react to a deviation.

When a target deviates by a distance D_c from the intended path, the robot cautiously attempts to get closer by speeding up and moving in front of the target. Here the robot seeks to get in front of the target in order to drive it back toward the path. The robot does not move in a straight line toward the target, but rather attempts to curve around the target a little, by giving priority to lateral movement before forward movement when coming up behind the target. Only when the robot's lateral movement is within range (shown by min follow distance above) will it then move forward toward the target. This somewhat curbed the robot's tendency to push targets away, but not entirely, as discussed in Section 4.

If the target exceeds a distance D_a, where $D_a > D_c$, then the robot switches to an aggressive combination of silent herding [14] and showing eye [7] tactics, lowering its hovering height to eye level, turning the robot to fly backwards and maintain apparent gaze, and entering the target's intimate space as it moves in front. If the target goes a distance D_i, where $D_i>D_a>D_c$, then the robot gives up on that target and ignores it, moving on to closer targets if there are any.

3 Implementation

The affective herding behavior was implemented using a finite state machine and programmed for an AirRobot AR-100B in the Unity3D engine using Javascript.

A diagram of the herding behavior is shown in Figure 3.1. The parameters D_c, D_a, and D_i, which modulate the escalation of the behavior as a function of the target's distance off course, were arbitrarily chosen to match proxemic zones (D_c = 0.46m, D_a = 1.22m, and D_i =3 .22m). For example, D_c = 0.46m; if a target was within 0.46m of the path line (an intimate distance to the path line), the robot responded calmly and stayed at the edge of the social zone. It should be emphasized that use of proxemic distances for path distances is separate from the use of proxemic distances to generate escalating discomfort. Other options for selecting the parameters certainly exist and could be chosen without changing the general framework proposed in Section 2.

The Unity3D game engine was chosen to simulate the AirRobot platform because of the ease of modeling the AirRobot. Other simulation software packages, such as Microsoft Robotics Studio, did not have the AirRobot included in their existing repertoire of simulated robots, so a platform had to be chosen in which the AirRobot could be quickly mocked up and added in to. Unity3D is also platform-independent, and, since it is used for creating 3D games, also has an easy to use, visual-based GUI. The behavior in this simulation was coded in Javascript, although more complex systems could be created using C#, which is also supported by Unity. The AirRobot AR-100B simulated in this study is a commercially available sUAS designed for aerial surveillance.

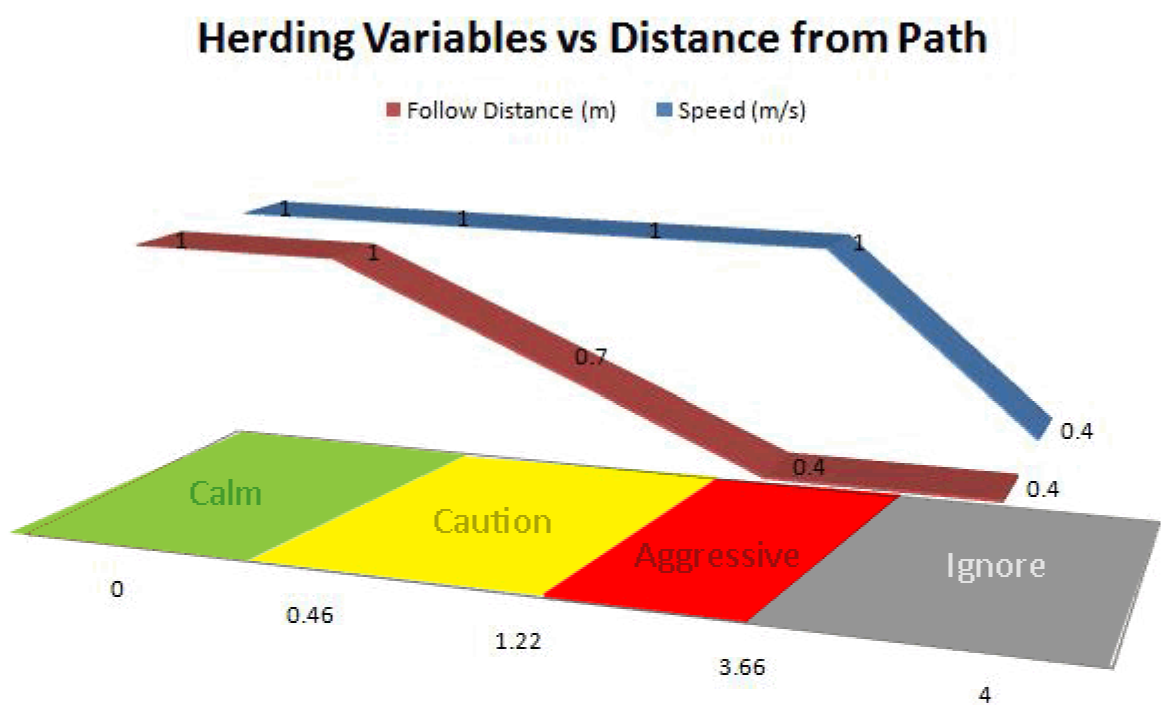

Figure 3.1: Herding behavior variables in each proximity zone.

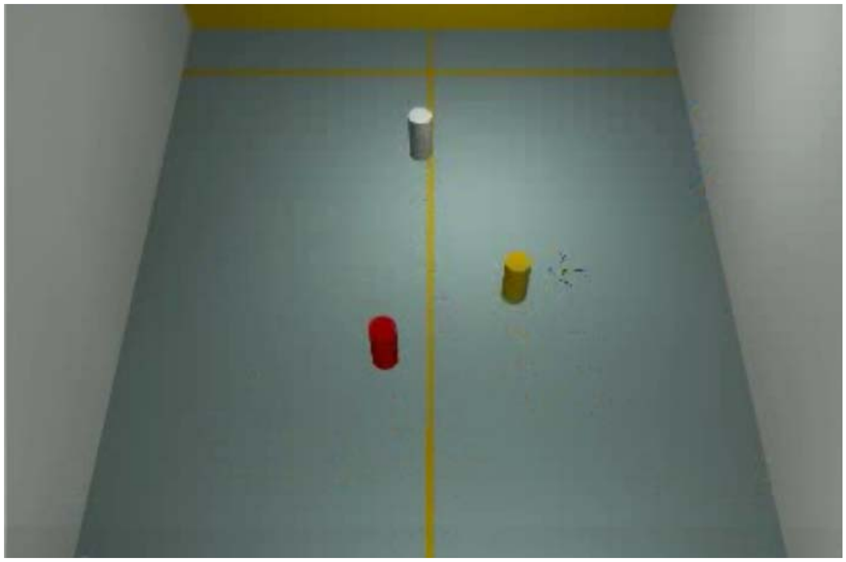

Figure 3.2: Simulation running with three targets.

4 Simulation Trials and Results

522 trials were conducted using computer-based simulation to assess the utility of the affective herding behavior for one robot and up to three targets. The simulation results were analyzed to determine the *average deviation of the target(s) from the path line* at any given point along the path, which provides an indication of how well the robot was able to keep the target(s) on track, individually and collectively. The average deviation data hides an important phenomenon: *losing a target*, where a target deviates so far from the path line that it exceeds D_i and essentially becomes "lost" to the robot. In addition, the use of a hallway in the simulation produced a second noteworthy phenomenon: *hang-ups* where the target would drift away and follow the wall too closely for the robot to impart any effective herding stimulus.

4.1 Simulation Methodology

The simulation methodology consisted of having a single robot employing its real-world speed and movement restrictions operating in a corridor herding 1 to 3 targets under randomized conditions. Although the affective herding behavior is intended for outdoor and indoor spaces, a corridor was selected because it provided natural limits on how far a target could drift off course and also might illuminate any unique issues with indoor spaces. The corridor had a fixed length of 14 m. In order to simplify duplicating human behavior, all targets followed a randomly generated intended path (which can be different than the path line) unless influenced by the robot (as detailed below) or by a wall (when it collided with a wall, the target would hug the wall as it moved forward). The targets' movement speed began at human walking speed (1 m/s) moving forward in the z-direction down the hallway. The robot started behind the targets in the calm zone of behavior. The simulation terminated when all the targets reached the end of the hallway (i.e. the back wall).

Four aspects of the simulation were randomly selected for each trial:

- *the width of the corridor.* The width of the corridor was an integer between 6-20 m.
- *the number of targets*. The number of targets varied between 1 and 3. At the end of 522 trials, 194 trials used 1 target, 175 used 2 targets, and 153 used 3 targets.
- *the starting position of each target*. Each target was placed in a randomized starting x-position and z-position within a 2 m by 2 m bounding box at the start of the hallway.
- *the path of each target*. The movement path of each target consisted of four segments. The 12m active area of the corridor was divided into four regions, each 3m long. At the starting position and whenever the target crossed into in a region, a random direction within a +/- 45 degree cone would be generated. This provided somewhat random movement in the targets while ensuring that the targets would reach the end of the hallway, albeit by hugging the walls, even if the robot completely ignored them. Figure 4.1 provides a visual model of the targets' movement.

The influence of the robot on the targets was a function of relative distance. As described in Section 3, the robot exerts a repulsive vector on a nearby target. In these simulations, the repulsive magnitude was computed as a function of relative distance; recall the direction is the straight line between the center of the robot and the target. The intensity of the repulsion are based again on proximity zones, so when the robot is within 3.66 m, the social zone, the targets begin to slowly move away from the robot (at best 0.25 m/s). When the robot moves within 1.22 m, personal space, the targets move more quickly away from the robot (up to 0.5 m/s). If the robot enters their intimate space, less than 0.46 m away, the targets move back even faster away from the robot at up to1.0 m/s.

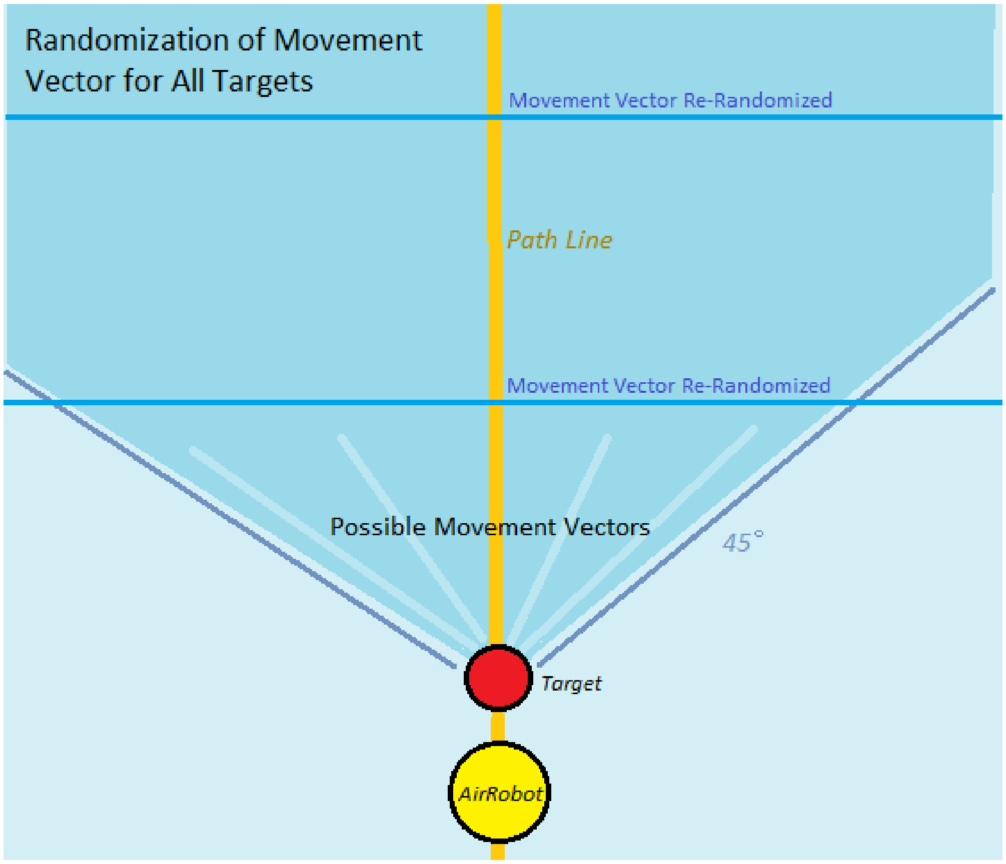

Figure 4.1: Randomization of movement in targets.

4.2 Average Deviation Results

The 522 trials suggest that the affective herding behavior is very effective for one target, but not effective for multiple targets. Figure 4.2 shows visually that the herding behavior keeps a single target well within 0.2m. The average x-deviation was plotted every 1 m along the 12m active region of the hallway, creating a graph of the average path along the path line of each group of targets. The average paths are shown together, where the path line is the x-axis, and x = 1 represents the starting position of the targets. Figure 4.3 shows the statistical properties used to create the graph along with their standard deviation and variance.

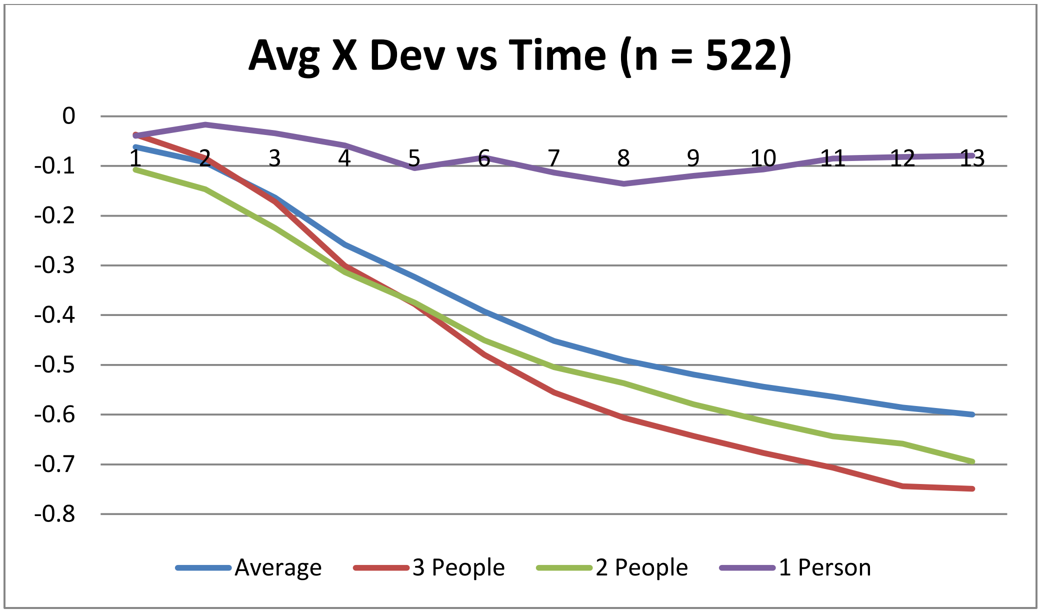

Figure 4.2: Average path of different instances of targets

	Avg. X Dev	Std Dev	Variance
Total	-0.388194143	0.190498109	0.03628953
3 People	-0.47172993	0.25324916	0.06413514
2 People	-0.449835336	0.199671659	0.03986877
1 Person	-0.081557931	0.035810584	0.0012824

Figure 4.3: Average Deviation from Path

4.3 Hang-Ups

Finally, the width of the corridor also affected the accuracy of the herding behavior. Because the target behavior was reactive, a narrow hallway sometimes caused the robot to drive the target against the wall. If the targets' randomization of movement did not cause it to move away from the wall, the target would hug the wall for the rest of the simulation, keeping the robot from getting in front of them and often causing it to get "hung up" on that target, ignoring the rest.

4.4 Losing a Target

These results also show that as the number of targets increases, so does the frequency of the robot "losing" a target, causing the target to go so far off the path that it goes out of range of the behavior or against the wall. If only one target is being herded, the likelihood is low, with a lost target occurring only in 4.7% of the trials. This is what happens to other targets when the robot gets "hung-up" on a target pressed against the wall, but could also happen for other reasons, such as two targets moving quickly in opposite directions so that at least one goes out of range before the robot can get to it. When there is one target, the robot can "lose" it by pushing the target out of range while chasing after it. Although the robot's movement tried to keep it from going straight toward a target and pushing the target away, the modification to the robot's movement was not enough to completely stop it from pushing targets away or out of range. The likelihood of losing a target based on the number of targets is shown in Figure 4.6.

	# of Missed Targets/Total	% Likelihood to lose at least one target
Total	226/522	43.3%
3 People	128/153	83.7%
2 People	89/175	50.9%
1 Person	9/194	4.7%

Figure 4.6: Likelihood of the behavior missing a target based on the number of targets.

5 Conclusions

This paper presented biologically-based affective herding behavior for a small unmanned aerial vehicle, where one robot herds one or more targets (humans). The synthesized behavior has the robot driving the targets from behind as a *heeler*. The behavior considered affect only in the sense of coercing the target to resume the path using proxemics. This appears appropriate for crowd control or for containment to prevent a riot, but ignores the role positive affect in encouraging a smaller, well-mannered group to evacuate following a desired route.

The proposed affective herding behavior was found to be effective for one target, but broke down with two or more targets. Assuming that the degeneration was not an artifact of the simulation or test conditions, one solution is to consider multiple robots using the affective herding behavior. However, the behavior is a serial behavior in that it looked only for one deviate target and did not perceive other targets or the group motion (e.g., was the whole herd spreading out? Where was the centroid of the group?). This suggests that this behavior will not scale beyond 1 robot to 1 target and is unlikely to produce a satisfactory emergent behavior. The average deviation results point to re-evaluating the biological literature and synthesizing the insights differently or focusing on gathering or on the robot as a *header*, not a *heeler*.

The *losing a target* and *hang-up* phenomena could be due to unrealistic simulated human behavior, but they raise an important question: is there a fundamental difference in herding in open spaces, such as fields, or in closed spaces, such as corridors? If so, one behavior may not be sufficient for both cases.

6 Acknowledgements

This work was supported in part by CRA-W Distributed Research Experiences for Undergraduates and the Texas A&M Computer Science and Engineering Department.

REFERENCES

1. CHATTARA, U., A. SEYFRIED, and P. CHAKROBORTY (2009) *COMPARISON OF PEDESTRIAN FUNDAMENTAL DIAGRAM ACROSS CULTURES.* Advances in Complex Systems **12**, 393-405.
2. Bethel, C.L. and R.R. Murphy, *Non-facial/non-verbal methods of affective expression as applied to robot-assisted victim assessment*, in *ACM/IEEE International Conference on Human-Robot Interaction* 2007: Arlington, Virginia. p. 287-294.
3. Mutlu, B., et al., *Footing in human-robot conversations: how robots might shape participant roles using gaze cues*, in *ACM/IEEE International Conference on Human-Robot Interaction Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*. 2009: La Jolla, CA. p. 61-68.
4. Saerbeck, M. and C. Bartneck, *Perception of affect elicited by robot motion*, in *ACM/IEEE International Conference on Human-Robot Interaction* 2010: Osaka, Japan. p. 53-60.
5. Shell, D. and M. Matarić (2005) *Insights Toward Robot-Assisted Evacuation* The International Journal of the Robotics Society of Japan **19**, 797-818.
6. Takayama, L. and C. Pantofaru, *Influences on proxemic behaviors in human-robot interaction*, in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on* 2009: St. Louis, MO. p. 5495 - 5502
7. Hafez, E.S.E., *The Behavior of Domestic Animals*. 1962, Baltimore: The Williams & Wilkins Company.
8. Helbing, D., I. Farkas, and T. Vicsek (2000) *Simulating dynamical features of escape panic.* Nature **407(6803)**, 487-490 DOI: 10.1038/35035014.
9. Benoit-Bird, K.J. and W.W.L. Au, *Cooperative prey herding by the pelagic dolphin, Stenella longirostris.* Acoust. Soc. Am., 2009. **125**(1): p. 125-137.
10. Pandey, A.K. and R. Alami, *A step towards a sociable robot guide which monitors and adapts to the person's activities*, in *Advanced Robotics, 2009. ICAR 2009. International Conference on* 2009: Munich. p. 1-8.
11. Raafat, R.M., N. Chater, and C. Frith (2009) *Herding in humans* Trends in Cognitive Sciences **13**, 420-428 DOI: 10.1016/j.tics.2009.08.002.
12. Shiomi, M., et al., *A larger audience, please! — Encouraging people to listen to a guide robot*, in *Human-Robot Interaction (HRI), 2010 5th ACM/IEEE International Conference on* 2010: Osaka. p. 31-38.
13. Wakeman, M.C. *The Herding Puli*. [web page] 2002 2002 [cited 2010 6-21-10]; Herding behavior and training of Pulik]. Available from: http://www.showdogsupersite.com/puliworld/herding.html.
14. Walther, F.R. (1991) *On herding behavior*. Applied Animal Behaviour Science **20**, 5-13 DOI: 10.1016/0168-1591(91)90145-P
15. Wiese, W.H. (2008) *Farm Dogs and Other Dogs of Agriculture.* Journal of Agricultural & Food Information **9**, 77-91 DOI: 10.1080/10496500802113936.