

Characterizing Human Perception of Emergent Swarm Behaviors

Phillip Walker & Michael Lewis
School of Information Sciences
University of Pittsburgh
Pittsburgh, Pennsylvania, 15213, USA
Emails: pmwalk@gmail.com, ml@sis.pitt.edu

Katia Sycara
Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania, 15260, USA
Email: katia@cs.cmu.edu

Abstract—Human swarm interaction (HSI) involves operators gathering information about a swarm’s state as it evolves, and using it to make informed decisions on how to influence the collective behavior of the swarm. In order to determine the proper input, an operator must have an accurate representation and understanding of the current swarm state, including what emergent behavior is currently happening. In this paper, we investigate how human operators perceive three types of common, emergent swarm behaviors: rendezvous, flocking, and dispersion. Particularly, we investigate how recognition of these behaviors differ from each other in the presence of background noise. Our results show that, while participants were good at recognizing all behaviors, there are indeed differences between the three, with rendezvous being easier to recognize than flocking or dispersion. Furthermore, differences in recognition are also affected by viewing time for flocking. Feedback from participants was also especially insightful for understanding how participants went about recognizing behaviors—allowing for potential avenues of research in future studies.

I. INTRODUCTION

Robot swarms consist of multiple robots that coordinate autonomously via local control laws based on the robot’s current state and nearby environment, including neighboring robots. Key advantages of robotic swarms are robustness to failure of individual robots and scalability with growing numbers, both of which are due to the distributed nature of their coordination. Multi-robot systems that are not swarms often have explicitly represented goals, form and execute group plans, have heterogeneous capabilities, and assume different roles [1], [2]. Robots in multi-robot systems could act independently without coordinating, or they could cooperate as a team in which all members work towards shared goals. Swarms, on the other hand, necessarily involve coordination between robots and rely on distributed algorithms and information processing. Because of this, global behaviors are not stated explicitly, and instead emerge from local interactions.

There are two different capabilities needed before an operator can successfully supervise a semi-autonomous swarm: *comprehension of swarm state* and *prediction of the effects of human inputs on swarm behavior*. Comprehension of swarm state, which is the main focus of this paper, requires the operator to correctly understand the data being returned from the swarm, and identify the patterns present so as to recognize what global behavior is emerging from the local robot

interactions. This data typically includes position and velocity data, but can also include swarm density, connectivity of the sensing graph, or local environmental variables. Identifying the emergent behavior is often challenging, as deficiencies in the robot hardware or communication capabilities limit the amount of data that can be returned from the swarm. Furthermore, when there is significant noise or error in the robot data, the behaviors may not be readily apparent.

The second capability, predicting effects of human inputs, requires the former capability. An operator cannot accurately predict the effect of their input if they do not have a current and accurate comprehension of the swarm’s state. Having an internal model can help the operator predict what the swarm will do next, given different potential inputs. It is likely, although not necessary, that this internal model is based on the model used to dictate the local interaction rules of the robot. For instance, a common model for designing swarms is the biologically inspired model [3], [4], which uses rules common to those found in collective animal groups. Because humans are often familiar with such collective motion, this is often an easy way for them to understand the swarm as operators. Even for intuitive relationships, such as issuing new heading commands to the leader robot(s) of a flocking swarm, issues in communication, hardware capabilities, or environmental features require the operator to fall back to their model of the swarm dynamics in order to effectively control the swarm. Therefore, one can easily see how comprehension of the swarm state is the first necessary component for proper human swarm interaction (HSI).

We believe that the current behavior of the swarm is likely the most important piece of information in helping an operator make predictions about future states. We hypothesize that some of even the most basic swarm behaviors will be easier or harder to recognize based on the features of that behavior. To that end, the study presented herein tests this hypothesis, and solicits feedback from participants that can also be used to help designers create effective interfaces for HSI systems. In Section II, we introduce previous research that relates to and motivates our work. In Section III, we present the structure and methodology of the study, and follow up with the results in Section IV. Finally, we discuss these results in their context, possible future applications, and conclusions in section V.

II. RELATED WORK

While swarms can operate with full autonomy, schemes requiring a human operator in the loop with robotic swarms are required for many more complicated tasks, such as surveillance and search and rescue. This gives rise to the field of human swarm interaction (HSI), which has been studied primarily for foraging tasks, as in [5], [6], [7]. Study of human involvement in formation control includes work investigating transitions between a flock and torus [8], as well as network configuration [9].

The most similar previous work to the study presented herein comes from [10]. In this work, the authors demonstrate through user studies that human perception of biologically-inspired swarms was superior to unstructured motion, but inferior to other types of biological motion with a coherent form or rigid structure, such as a walking human, as demonstrated in [11]. The work herein builds off of this research by determining the differences in operator recognition and discrimination between different types of structured motion in swarms—namely, the different types of behaviors a swarm may be directed to perform.

In [12], the authors also address the problem of recognition of swarm behaviors taking a small sample of the members of the swarm and using a Bayesian classifier to autonomously discern between two behaviors: flocking in a single direction, and flocking in a circle (torus). Therein, the authors found significantly high accuracy even for a sample of one robot. Once the sample is increased to two or higher, classification was near perfect if the simulation was allowed to run for long enough. There are a couple differences between this work and the work in [12] however. First, the behaviors presented herein are distinctly different from another—not just in appearance to a human but also in the underlying control laws. In [12], the behaviors result from a parameter change only, not a change in the control laws. Second, the results presented here are from human participant trials. This is important if we wish to understand how the operator might recognize behaviors or if the operator does not trust an autonomous recognition algorithm. Similar work in [13] represents the swarm as a velocity field. By doing so, the researchers could model the field as a Gaussian distribution, and achieved high rates of autonomous recognition of swarm behavior.

Recognition of behavior is important for another point that we have yet to address: switching between behaviors. In the majority of HSI missions, the operator will need to switch from one behavior to another at some point—especially if mission goals change while the mission is ongoing. A naïve approach would be to switch the behavior immediately upon needing to do so. However, new research of a phenomenon called *neglect benevolence* shows us that the optimal switching time may *not* be at the first possible moment.

A. Neglect Benevolence

A previous study using a foraging scenario [14], found that the performance of a human-swarm system could be strongly affected by the time between two different commands

applied by the human to the robots. In particular, results showed that one group of subjects who performed well waited between initial and corrective command (when changing the goal heading of the swarm). The phenomenon was called *neglect benevolence*, since neglecting the swarm for some amount of time led to improved performance. Further analysis in [14] found that in transient states (i.e. moving from one goal to another), applying another input to change the goal could have differential effects depending on the timing of the input.

To determine whether this phenomenon was unique to the particular situation studied, or a more general concept, [15] reported simulations of swarm systems starting at different configurations and performing rendezvous where the operator inputs changed the rendezvous point. The robots moved with a repulsive virtual force to avoid collisions with neighbors, and an attractive force to maintain cohesion of the swarm. The simulation gave a variety of outcomes depending on *when* the input was given, following the desired change in the goal. In particular, it was observed that giving input immediately after the need to change the rendezvous point arose resulted in several robots becoming disconnected, and never returning to the rendezvous point. If, however, the input was delayed, the swarm often stayed together and completed the rendezvous at the desired point. In [15], an algorithm was reported that computes the optimal time for insertion of the human input. Therefore, neglect benevolence and optimal human input timing is a nuanced, yet quantifiable notion.

In [16], the authors investigated whether operators presented with an HSI reference task that benefited from neglect benevolence could learn to approximate optimal input timing, after gaining experience interacting with the swarm (implicit learning). This study divided subjects into two groups that were each tasked with diverting a swarm headed from an initial state to a first configuration, and then to a different, second configuration (see Figure 1). The goal of the participants was to give the input at a time that would cause the quickest convergence to the second configuration. Because the system exhibited neglect benevolence, the optimal time was neither at the beginning nor end of the swarm state evolution between first and second configuration. One group of the participants were given a simple display, showing each of the robot positions as they moved. The second group received an aided display, which showed the same as the first, but added lines from each robot's current position to its goal in the formation.

The main finding of this study indicates that, although humans had challenges in term of determining optimal input timing, they nevertheless managed to improve their performance over time (i.e. decrease time to convergence to the second configuration). The quality of this approximation began at a higher level for those using the aided display, and developed slowly for participants in the unaided condition. The success of the augmented display (aided condition) in improving the performance of participants supports our hypothesis that challenges in interaction with swarms could be partially due to the perceptual inaccessibility to the human of the variables on which the swarm is coordinating. Therefore,

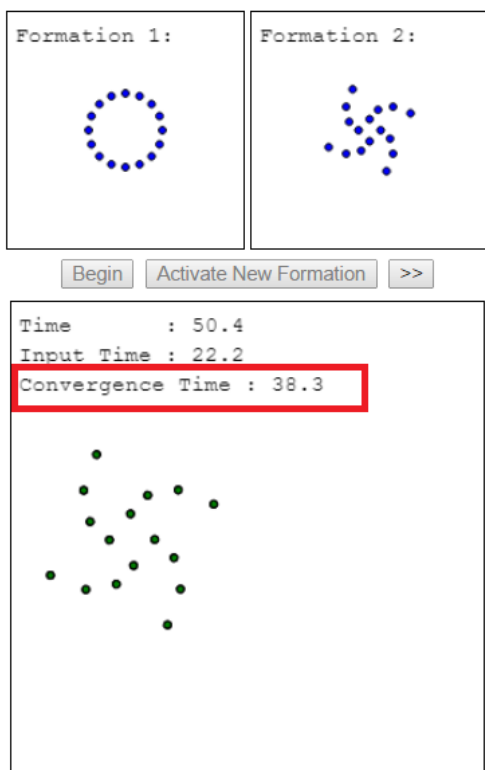


Fig. 1. For each trial in the motivating study, participants were asked to apply the input at whatever time they thought would minimize the total time required for the robots to converge to Formation 2, after initially moving to Formation 1. This figure is taken from [16].

the study reported in this paper is aimed at evaluating human perceptual recognition and discrimination between different swarm behaviors. We believe understanding what allows better recognition for the operator can aid in future studies that deal with input timing in an HSI system.

III. BEHAVIOR RECOGNITION STUDY

A. Overview and Hypotheses

The study involving participants dealing with a system that exhibited neglect benevolence [16] showed that operators could learn to better time their inputs over the course of interacting with the swarm. Because we believe that having a better internal model of the swarm, and recognizing how the swarm is evolving, can help improve timing of inputs, this naturally begs the question of how well an operator can recognize these emergent swarm behaviors in general. If an operator is better able to recognize what behavior is being performed, he or she should be able to better recognize the proper timing of input. This is especially the case if the behaviors take some time to reach consensus. Furthermore, if we are able to understand which behaviors are easier or harder than others to recognize, this could help designers of interfaces for human-swarm systems know what properties of the swarm state to display, and how much information is needed for the operator to best learn to time their inputs.

To that end, we designed a study to investigate how well operators recognized three of the most common algorithmically generated swarm behaviors as an investigation into how we might improve human control of swarms in general. Our study contrasts and complements [10], who studied perception of biological swarm behaviors and subsequently compared their perception with that of simpler displays of rotating dots. Similar to the aided condition [16], the display used in this study shows robot positions and velocities to inform the operator about the current state of the swarm. We ask participants to discriminate between three types of behavior in the presence of background noise: rendezvous, flocking, and dispersion. Because of humans' innate ability to recognize biological motion and common fate [11], [17], we believe participants in general will be good at recognizing all behaviors, even when background noise is high ([17] reports discrimination at S/N as high as 97%). However, we hypothesize that there will be significant differences between the behaviors in terms of recognizability, and in the factors, including individual differences, that give rise to recognizability. Specifically, we hypothesize that the flocking behavior will benefit from longer viewing times before a response is given, due to the time it takes for consensus to emerge and the collective movement to begin before the behavior becomes apparent.

B. Task Description

This study involved the human participant viewing a series of videos with one of three types of swarm behaviors: rendezvous, where each robot moved to the center of the bounding x- and y-axis values of its neighbors (the parallel circumcenter algorithm [18]); flocking, where each robot tried to match the velocity of its neighbors, while also maintaining a minimum distance from close neighbors, and maximum distance from far neighbors; and dispersion, where each robot moved away from the average position of its neighbors. In each video, there was some amount of background noise, i.e., some of the swarm members moved randomly, ignoring all neighbors. The goal of the study was to determine how much background noise could be present before the participants would stop recognizing the behavior being performed. Each participant started at 50% background noise for each behavior. If the participant answered correctly for a behavior, the next time they viewed a video with that same behavior, the noise level would be increased according to the following formula:

$$e_1 = (100 + e_0)/2 \quad (1)$$

Where e_0 and e_1 are the current and new noise percentages, respectively. In other words, the noise would increase to the halfway point between the previous noise level and 100% noise. If the participant answered incorrectly, the next time they viewed that behavior the noise level would be set as follows:

$$e_1 = (50 + e_0)/2 \quad (2)$$

The participants viewed each behavior six times, along with six videos with 100% noise, for a total of 24 videos. The




Fig. 2. Illustration showing an initial random state (top left) and each of the three behaviors used in the study: rendezvous (top right), flocking (bottom left) and dispersion (bottom right). Lines away from each dot indicate the heading of that robot. Note that due to size constraints in this paper, the full viewport bounds are not shown, nor are the full 2048 robots used.

videos with 100% noise served as a baseline, to ensure that participants could also discriminate between an organized behavior and no behavior at all.

C. Robots and Simulation

The videos were generated via a simulation of 2,048 robots programmed in CUDA C and OpenGL run on an Nvidia GTX 980 GPU, which allows for such a large number to be used. All the robots began at random positions in a 2D plane, bounded by $[-15,15]$ meters in both the x - and y -axis. The camera showed as far as $[-20,20]$ meters in each direction, to allow the participants to see the entire swarm at the beginning, as well as to allow for some expansion in the overall swarm area before leaving the viewport. The videos were each 20 seconds long, to give the participants plenty of time to view the behaviors and distinguish between them, although participants could select a response at any time. Total viewing time before selecting a response was recorded along with their response for data analysis. A simulation step was performed 60 times a second, whereby each robot sensed its neighbors and moved according to the algorithms described in Section III-D, giving a total of 1,200 steps per video. We will now introduce the algorithms used to generate each of the three behaviors, as shown in Figure 2.

D. Behavior Algorithms

In the below algorithms, several common parameters are used for the behaviors. They are defined in Table I.

TABLE I
PARAMETERS USED IN THE RENDEZVOUS, FLOCKING, AND DISPERSION ALGORITHMS OF THIS STUDY.

Variable	Value	Description
d_1	0.5	Close range (meters)
d_2	1.5	Close-mid range (meters)
r	2.0	Maximum range (meters)
v_{max}	1.0	Maximum velocity (m/s)
α_{max}	6π	Maximum angular velocity (rad/s)
w_a	1.0	Align vector weight
w_c	0.9	Cohere vector weight
w_r	1.0	Repel vector weight

1) *Rendezvous*: The rendezvous behavior was determined by the following algorithm for each robot. Two vectors were computed: the repel vector, \vec{r}_r and the cohere vector, \vec{c}_r . Vector \vec{r}_r is the sum of vectors from each neighbor robot within d_1 to this robot's (x, y) position. Vector \vec{c}_r is computed by taking the midpoint of the rectangle $R = (x_{max}, y_{max}, x_{min}, y_{min})$, where x_{min} and y_{min} are the minimum x and y coordinates of the robots in the neighbor set N within r , respectively, and x_{max} and y_{max} are the maximum x and y coordinates of the robots in N . The computation of this cohesion vector is adapted from the parallel circumcenter algorithm in [18]. The final goal vector is then computed using the following equation:

$$\vec{g} = w_r \vec{r}_r + w_c \vec{c}_c \quad (3)$$

2) *Dispersion*: The dispersion behavior was computed using only one component vector, the repel vector, \vec{r}_d , which is computed by taking the sum of the vectors from each robot in N within the maximum range r . The goal vector for dispersion is equal to \vec{r}_d .

3) *Flocking*: For the flocking algorithm, leaders were selected according to the distributed MVEE algorithm [19], which determines the robots which together form the minimum bounding ellipsoid of the swarm. These leaders were given an identical random goal point outside the bounding box of the swarm. The remaining robots each computed a repulsion vector \vec{r}_f identically to the rendezvous algorithm (Section III-D1), except using d_2 as the maximum range instead of d_1 . The cohesion vector, \vec{c}_f , is computed by taking the average of the vectors from this robot's to each neighbor robot's (x, y) position within r . An alignment vector, \vec{a} , is also used only by the flocking behavior, but was computed differently, depending on whether the robot was a leader or not. If the robot is a leader, \vec{a} is set to the vector from its current position to the goal point. If the robot is not a leader, but it is within range r of one, this robot will set \vec{a} to match the closest leader. If the robot is not a leader nor in range of one, $\vec{a} = \sum_{n=1}^N a_n$, where a_n represents the alignment vector of the n -th neighbor in N , the set of neighbors of this robot within range r .

If the robot is a leader, the goal vector is represented by the following (notice that leaders do not have a repel vector):

$$\vec{g} = w_a \vec{a} + w_c \vec{c}_f \quad (4)$$

If the robot is not a leader, the goal vector is represented by the following:

$$\vec{g} = w_r \vec{r}_f + w_a \vec{a} + w_c \vec{c}_f \quad (5)$$

In the above behaviors, both velocity $v = \|\vec{g}\|$ and angular velocity, α , were capped at the values given in Table I. α was the difference in heading between the goal vector at the previous time step, \vec{g}_{t-1} , and the goal vector at the current time step, \vec{g}_t .

4) *Movement Towards Goal*: Once a robot has computed its goal vector using the relevant component vectors for the current behavior (repulsion, flocking, cohesion), the next state of the robot is computed by first turning the robot toward the heading of the goal vector, up to a maximum change of α_{max} , the maximum angular velocity. Because there were 60 simulation steps per second, for each step this maximum angular velocity would be $\alpha_{max}/60s$, or approximately 0.314 radians. Once rotated, the robot would then move forward at the maximum velocity $v_{max} = 1m/s$. Again, for each step this would be $v_{max}/60s$, or approximately 0.017 meters.

E. Experiment Details

Participants were selected from the Amazon Mechanical Turk user-base, and were given a short questionnaire to complete, asking their age, gender, average weekly computer use, and average weekly spent playing computer games. After the videos were finished, each participant was asked to describe their strategy for recognizing the behaviors, if any. A total of 72 participants were collected. Of these, 32 participants were female and 40 were male, and ages ranged from 20 to 72 years old (median of 32).

IV. RESULTS

Due to how the noise was adjusted for each behavior-type (see Section III-B), the final correct answer given by the participant for each behavior is also the maximum noise percentage for that behavior where recognition was still successful. Therefore, for each participant, we can easily find the maximum noise percentage with correct recognition for each behavior. Using a Welch's t-test, rendezvous was found to be significantly easier to recognize than either flocking ($t(137.86) = 3.521, p < .001$) or dispersion ($t(140.16) = 2.619, p = .01$). Flocking and dispersion were not significantly different ($t(141.49) = 0.889, p = .375$). Figure 3 shows these results graphically.

A similar measure of recognizability is the number of correct answers for each behavior, which should roughly correlate with the maximal noise percentage presented previously. Indeed, the same results were found here. Using a Welch's t-test, rendezvous allowed for more correct answers by participants than either flocking ($t(133.25) = 3.153, p = .002$) or dispersion ($t(138.44) = 4.146, p < .001$), with no significant difference between flocking and dispersion ($t(140.6) = 1.278, p = .204$). Interestingly, we found that participants correctly recognized a lack of behavior (100% noise) significantly less often than any of the three behaviors ($t = 5.310, df = 114.92,$




Fig. 3. The average maximum noise percentage for each behavior where participants still answered correctly.

$p < .001$ for rendezvous; $t = 3.480, df = 98.95, p < .001$ for flocking; and $t = 2.645, df = 104.49, p = .009$ for dispersion).

Results also show that, on average, the longer a participant views a video before submitting a response the more likely they were to be correct, but only for the flocking behavior. Taking longer to view a video of flocking behavior before giving a response corresponded to correct responses at higher noise percentages for flocking ($F = 14.94, df = 70, p < .001, r^2 = 0.164$). This effect was marginally significant for rendezvous ($p = .053$), and there was no such effect for dispersion ($p = .919$). These results were again mirrored when considering total correct answers instead of maximum noise.

A. Individual Differences

Effects of demographic data were also analyzed to determine if age, gender, computer use, or video gaming frequency impacting the performance of operators. While there was no correlation between age and the average maximum noise across all behaviors and responses, age did have a small positive correlation with total average viewing time of the videos ($F(1, 70) = 5.697, p = .020, r^2 = 0.062$), although this did not translate to higher maximum noise values for correct responses in the flocking behavior, as might be expected. Surprisingly, there was a difference in performance between genders, with females recognizing behaviors at a higher noise level than males ($F(1, 70) = 5.26, p = .025, 0.057$), although like the correlation between age and viewing time, this effect was small. This could be due to the fact that females, on average, viewed the videos for a longer period ($F(1, 70) = 2.975, p = .089, r^2 = 0.027$), although this effect was marginal. Computer usage was assessed by asking participants to estimate how often they used a computer in a week, at 10-hour intervals (i.e. 0-10 hours, 10-20 hours, etc.). Higher computer use correlated with better recognition at higher noise rates, but for rendezvous only, and the results

TABLE II
EXAMPLE RESPONSES CHARACTERISTIC OF THE COMMON THEME OF
GLOBAL FOCUS AND RECOGNIZING PATTERNS.

ID	Description of Strategy
p8	Tried to look at everything as a whole and pick out certain patterns
p14	I liked to let my eyes unfocus. That really helped to see the picture as a whole.
p15	I tried to unfocus my eyes and recognize the general movement pattern of the dots.
p18	I watched for the density to change in the picture and then tried to discern some sort of pattern from that.
p21	I tried to let my focus widen and not stare too hard.
p30	The strategy I used for recognizing behaviors was to unfocus my eyes from any particular spot and try to notice if there seemed to be a pattern within the large group.
p37	I tried to analyze a specific area in the picture and look for patterns and then I would look at the picture overall.
p40	I forced myself to stare at the video without blinking to try and catch any distinct patterns.

were only marginally significant ($F(1, 70) = 3.00$, $p = .088$, $r^2 = 0.027$).

Perhaps the most telling results for individual differences come from the subjective qualitative responses of the participants. The final question of the survey asked if they could describe any strategies they used for recognizing behaviors, and there were many common themes across the responses, primarily supporting the idea that humans are good at recognizing patterns and collective motion. A common strategy seemed to be to unfocus and view the bigger picture to recognize global patterns, instead of focusing on individual robots. For instance, many participants mentioned “unfocusing [their] eyes” or “watching for patterns to emerge”. Some characteristic responses are reported in Table II.

V. DISCUSSION AND CONCLUSIONS

The results of this study and feedback from the participants clearly show that some behaviors are easier to recognize than others, and that humans use the Gestalt properties of swarm behaviors such as common alignment, common velocity, proximity, etc. to recognize different behaviors. Rendezvous, which involves an easily visible aggregation to a common point, was the easiest to recognize from background noise than either flocking or dispersion. Furthermore, flocking benefited from longer viewing times, as this gave the user more time to pick out the common fate of the flocking (non-noise) swarm members, as we hypothesized. The average highest level of noise where recognition was still successful is 85.38%, which while high, is not as high as the signal-to-noise ratio for common fate reported in [17]. This is likely because the swarm behaviors we tested were slightly more complex than simple common motion—primarily due to interactions with neighbors.

The responses by participants further reinforce the idea that operators take a holistic approach to viewing the collective motion inherent in emergent swarm behaviors. This could mean that the underlying metaphor used for the design of swarm algorithms may be less important, as long as the end

result is recognizable via base perceptual mechanisms, such as collective motion and common fate as studied in [17]. Furthermore, this also helps explain why flocking was the only behavior to benefit from longer viewing times, as it requires significantly longer for the robots to reach a consensus on direction than it does for the robots to begin rendezvous or dispersion. These results lay the groundwork for research to develop intelligibility metrics that will allow estimation of the intelligibility of swarm behaviors based on Gestalt characterization of swarm dynamics. This capability coupled with the long sought ability to design control laws to produce desired emergent behaviors could provide the grounding needed to make HSI a practical technology.

In light of the recent results presented here, in [10], and the results of work on neglect benevolence, we believe the next logical step for this line of research is to investigate switching between behaviors to achieve a goal more complex than merely recognizing the current behavior. For instance, can we determine the optimal switching time between two behaviors to minimize the time to a final goal state? Similarly, how can a display aid the operator in recognizing when a switch should occur? In future work, we plan on addressing these questions in user studies that build off the results and user feedback presented here.

ACKNOWLEDGMENTS

This research has been sponsored in part by AFOSR grant FA9550-15-1-0442 and ONR award N0001409-10680.

REFERENCES

- [1] L. E. Parker, “Multiple mobile robot systems,” *Springer Handbook of Robotics*, pp. 921–941, 2008.
- [2] M. Lewis, “Human interaction with multiple remote robots,” *Reviews of Human Factors and Ergonomics*, vol. 9, no. 1, pp. 131–174, 2013.
- [3] J. C. Barca and Y. A. Sekercioglu, “Swarm robotics reviewed,” *Robotica*, vol. 31, no. 03, pp. 345–359, 2013.
- [4] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, “Swarm robotics: a review from the swarm engineering perspective,” *Swarm Intelligence*, vol. 7, no. 1, pp. 1–41, 2013.
- [5] A. Kolling, S. Nunnally, and M. Lewis, “Towards human control of robot swarms,” in *Proceedings of the 7th international conference on Human-robot interaction*. ACM, 2012, pp. 89–96.
- [6] G. Coppin and F. Legras, “Autonomy spectrum and performance perception issues in swarm supervisory control,” *Proceedings of the IEEE*, no. 99, pp. 590–603, 2012.
- [7] Z. Kira and M. Potter, “Exerting human control over decentralized robot swarms,” in *Autonomous Robots and Agents, 2009. ICARA 2009. 4th International Conference on*. IEEE, 2009, pp. 566–571.
- [8] D. S. Brown, S. C. Kerman, and M. A. Goodrich, “Human-swarm interactions based on managing attractors,” in *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*. ACM, 2014, pp. 90–97.
- [9] J.-P. de la Croix and M. Egerstedt, “Controllability characterizations of leader-based swarm interactions,” in *AAAI Fall Symposium: Human Control of Bioinspired Swarms*, 2012.
- [10] A. Seiffert, S. Hayes, C. Harriott, and J. Adams, “Motion perception of biological swarms,” in *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*. Cognitive Science Society, 2015, pp. 2128–2133.
- [11] G. Johansson, “Visual perception of biological motion and a model for its analysis,” *Attention, Perception, & Psychophysics*, vol. 14, no. 2, pp. 201–211, 1973.

- [12] D. S. Brown and M. A. Goodrich, "Limited bandwidth recognition of collective behaviors in bio-inspired swarms," in *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2014, pp. 405–412.
- [13] G. Wagner and H. Choset, "Gaussian reconstruction of swarm behavior from partial data," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 5864–5870.
- [14] P. Walker, S. Nunnally, M. Lewis, A. Kolling, N. Chakraborty, and K. Sycara, "Neglect benevolence in human control of swarms in the presence of latency," in *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on*. IEEE, 2012, pp. 3009–3014.
- [15] S. Nagavalli, L. Luo, N. Chakraborty, and K. Sycara, "Neglect benevolence in human control of robotic swarms," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 6047–6053.
- [16] S. Nagavalli, S.-Y. Chien, M. Lewis, N. Chakraborty, and K. Sycara, "Bounds of neglect benevolence in input timing for human interaction with robotic swarms," in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 2015, pp. 197–204.
- [17] F. Stürzel and L. Spillmann, "Perceptual limits of common fate," *Vision research*, vol. 44, no. 13, pp. 1565–1573, 2004.
- [18] F. Bullo, J. Cortés, and S. Martínez, *Distributed Control of Robotic Networks*, ser. Applied Mathematics Series. Princeton University Press, 2009, to appear. Electronically available at <http://coordinationbook.info>.
- [19] W. Luo, S. S. Khatib, S. Nagavalli, N. Chakraborty, and K. Sycara, "Asynchronous distributed information leader selection in robotic swarms," in *IEEE International Conference on Automation Science and Engineering*, August 2015.