

The Effect of Display Type on Operator Prediction of Future Swarm States

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—Large teams of robots that operate collectively, whose behavior emerges from local interactions with neighbors, are known as swarms. While significant progress has been made improving the hardware, communication capabilities, and autonomous operation of these swarms, we still have much to learn about how human operators control and interact with them. This research is necessary if real world swarms are to be deployed in the future. The study presented here investigates different methods of displaying information about the swarm state to operators, and asks them to make predictions about the swarm’s future state. In the study, participants are shown swarms performing one of three different behaviors, and are asked to use the information available from the display to make their predictions. Results show that summarizing the swarm’s current state to just an average position and bounding ellipse allowed predictions as accurate as those made when full state information was shown. Furthermore, two leader-based methods were used, whereby the operators were shown only a small subset of the swarm. However, such display methods were inferior for prediction than either the summary center and ellipse or full information methods. With these results, and with participant feedback about the helpfulness of the four display types, we hope future studies can make more informed decision about interface design when it comes to the control of swarms.

I. INTRODUCTION

Robot swarms are large teams of robots that are distinct from multi-agent systems because each individual robot does not operate independently. Instead, swarms give rise to global, emergent behaviors due to the interactions and computation between neighboring robots. Swarms typically have remarkable robustness to the failure of individual members, and are scalable to large teams. Multi-robot systems, on the other hand, typically have explicitly-represented goals, have heterogeneous capabilities, and assume different roles [1], [2]. This means that the entire mission could suffer if an individual member fails. The main attraction of swarms are that they avoid this pitfall.

One of the difficulties currently facing researchers in robotics is how to properly control a large swarm of robots after it has begun some task. With a multi-robot system, the control laws or planned path can be individually updated on each robot; however this is both impractical and unhelpful when it comes to swarms. Addressing each member of a

swarm individually would take a large amount of time and, due to the interdependent nature of the agents, would have unpredictable results. Therefore, researchers in the field of human swarm interaction (HSI) typically control a swarm by either switching between pre-programmed behaviors [3], [4], or adjusting the parameters of a single behavior [5], [6].

Even with preprogrammed behaviors or parameter switching, there remains the problem of giving the operator the best possible picture of a swarm’s current state. This is essential if the operator is to make predictions about future states, which are key to determining the proper behavior or parameter to change. How the swarm is displayed to the human—be it on a computer screen, through a haptic device, or some other method—is the most important piece of the picture of a swarm’s current state. Furthermore, this display or device should highlight important properties of the current behavior in order to best represent what the operator needs to know. For instance, if the swarm is gathering together at a single point (rendezvous) the most important feature of the swarm is its footprint (i.e. the space the swarm covers). If instead, the swarm is moving through an environment to a goal point (as is the case with the flocking behavior), the distribution of headings and current speed are likely more important than the swarm’s footprint.

Therefore, we designed the study herein to investigate the idea that the display type has a significant impact on how an operator in an HSI system predicts the future state of a swarm. The study uses three different simple swarm behaviors: rendezvous, flocking, and dispersion (see Section III-B); and four different display types to test these ideas (see Section III-C). The display types used each highlight a different feature of the swarm and have different bandwidth requirements. First is *Full Information*, which shows the position and headings of every robot, serves as a control, and requires the most bandwidth between the swarm and human operator. Second is the *Centroid/Ellipse* display, which shows a bounding ellipse and the center point of the swarm and requires the least bandwidth. Finally are two leader-based methods, whereby the participants see only a small subset of the swarm—*MVEE Leaders*, which shows leaders at the edges of the swarm (low bandwidth), and *RCC Leaders*, which show leaders spaced evenly throughout the swarm (variable bandwidth, see Section

III-C).

II. RELATED WORKS

A common metaphor both for thinking about and designing swarm systems is the biologically-inspired model. This model is used in numerous examples of HSI [5], [7] and serves as the metaphor guiding the design of the swarms used in the present study. Swarms following this model typically have control laws based on those seen in natural swarms, such as flocking birds, schooling fish, or ant colonies. We believe it is important to point out previous research investigating human perception of biological motion and common fate, as these demonstrate an innate ability for humans to recognize collective and biologically-inspired motion. In [8], the authors show that humans are quite good at recognizing dots on a screen with a common direction of motion, even with a significant amount of background noise and small deviation in the direction of motion. Similarly, in [9], the authors show that humans are well-suited to recognizing biological motion of other humans as well, even using as few as ten points.

Recognition of the current swarm behavior is also an important step to being able to predict and control a swarm. Indeed, it is a prerequisite for the predictions we will require of participants in the present study. While a companion paper to this one investigates recognition of behaviors, as opposed to prediction, previous work has also shed light on other methods of how recognition of swarm behaviors can occur. In [10], the authors show, through user studies, that human perception of biologically-inspired swarms is superior to non-biological unstructured motion. However, recognition was still inferior to rigid, structured motion, as is the case in perception of human motion (see [9]). Similarly, in [11], the authors demonstrate autonomous recognition of behaviors using a Bayesian classifier, giving positive results. They were able to demonstrate that near-perfect recognition could occur even with just two robots sampled from the swarm, although there were only two possible behaviors that each had the same base control laws (but different parameters).

Another important line of research that inspired this work can be found in the phenomenon of *neglect benevolence*, which is the idea that the optimal time for an operator to give input to a swarm may not be as soon as possible, due to the distributed and emergent nature of swarm behavior. In [12], this phenomenon was first discovered when results showed operators who gave quick commands back-to-back had worse performance on a swarm control task than those that allowed some time between subsequent commands. In [13], the authors show through simulations that this phenomenon does indeed exist, and propose an algorithm for computing optimal input times for certain cases. A follow up study in [14] demonstrated that human operators could also learn to recognize and adapt to neglect benevolence.

Ideally, an algorithm like the one used in [13] could be integrated into a display for an HSI system, so that the operator need not try to predict the optimal input time. However, using such an algorithm requires that the operator know what their

next input should be, and also requires some amount of prediction of the future state to determine if a second input is needed and what it should be. Therefore, in order to achieve the goal of integrating features that allow operators to account for neglect benevolence, or other as-yet-undiscovered phenomena, we must first understand how to fulfill the prerequisite task of accurately displaying the swarm to an operator.

III. STUDY METHODOLOGY

A. Overview and Hypotheses

The goal of this study is to determine if different methods of displaying a swarm change the ability of an operator to predict the swarm's future state. During a single trial, participants were shown a swarm performing one of three different behaviors: rendezvous, flocking, or dispersion (see Section III-B). Furthermore, the swarm is displayed on the screen using one of four display methods (see Section III-C). Each trial lasted 30 seconds, during which the participant was asked to draw on a touch screen what they predicted to be the final shape of the swarm at the end of the trial. Finally, each participant saw each behavior-display pairing twice, giving a total of 24 trials over the course of the study (3 behaviors \times 4 display types \times 2 viewings).

The swarm consists of 256 robots, and begins at random positions and orientations within a 24×24 meter box centered at a random point within 10 meters of the origin of the simulation environment. The environment extends from ± 100 meters in the x-dimension, and ± 35 meters in the y-dimension. These dimensions were chosen because they gave a close enough view of the swarm to make each detail of the GUI visible, while still allowing the entire swarm to be visible throughout each trial, regardless of the current behavior being performed.

The robots, control laws, and environment were developed in a simulation using CUDA C and OpenGL. Participants used a 23-inch Dell touch screen display with a resolution of $1920 \times 1080p$ to both view the swarm and give input indicating their prediction. During a trial, the participant is asked to use the touch screen to draw a shape that they believe would bound the final position of the swarm at the end of the 30-second trial period. A progress bar is shown at the top of the screen, and each robot was represented with a 5-pixel diameter circle centered at its position, and a 25-pixel long, 2-pixel thick line indicating that robot's heading (see Figure 1). Participants could give multiple predictions over the course of the experiment—allowing them to update previous predictions if they deemed them inaccurate. Each prediction input was converted into the convex hull of the user-drawn points and, along with the convex hull of the final positions of the swarm members, is used for data analysis purposes to determine the accuracy of the prediction input.

Before completing the 24 trials in a random order, the participants were allowed to complete three training trials to get accustomed to the interface and touch screen. These training trials consisted of one each of the three behaviors, all with a *Full Information* display. Upon completing the 24

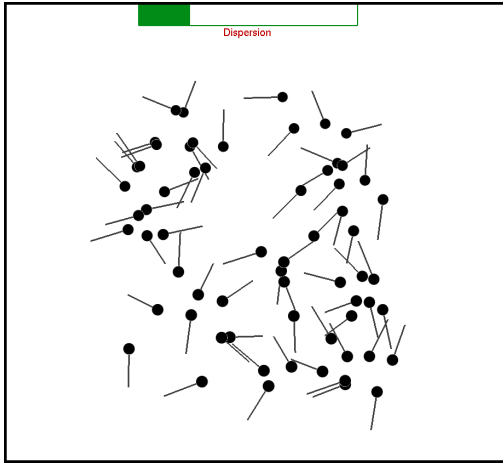


Fig. 1. Illustration of the graphical user interface for the study, showing the progress bar and current behavior (top) and the swarm of robots (middle).

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

Variable	Value	Description
d_1	1.0	Close range (meters)
d_2	2.0	Close-mid range (meters)
d_3	3.0	Mid range (meters)
r	5.0	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

trials of the main experiment, participants finished with a short survey, asking them to rank the four displays in order of their perceived helpfulness, both overall and for each of the three behaviors. Because each participant viewed all trials in every condition, this study uses a within-subjects design.

We hypothesize that the *Full Information* display type will give the best prediction accuracy across all behaviors, with the *Centroid/Ellipse* display giving either equal or slightly lower accuracy. If the results presented in [15] hold here, the two display types should give equal performance. Furthermore, we believe that the leader-based displays will both give lower predictions than either the *Full Information* or *Centroid/Ellipse* displays overall, however for individual behaviors that may not hold true. Specifically, for the flocking behavior, where the current heading of the swarm members is necessary for determining their future position, we hypothesize that one or both of the leader-based displays will outperform the *Centroid/Ellipse* display, as the latter does not display any heading information.

B. Behaviors

Several common parameters are used in the following behavior algorithms. They are defined in Table I.

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 $R = (x_{min}, y_{min}, x_{max}, y_{max})$, where x_{min} and y_{min} are the minimum x and y coordinate values of the robots in the neighbor set N within r , respectively, and x_{max} and y_{max} are the maximum x and y coordinate values of the robots in N . The computation of this cohesion vector is adapted from the parallel circumcenter algorithm in [16]. The final goal vector is then computed using the following equation:

$$\vec{g}_r = w_r \vec{r}_r + w_c \vec{c}_r \quad (1)$$

2) *Dispersion*: The dispersion behavior was computed using two component vectors. First, the repel vector, \vec{r}_d , is computed by taking the sum of the vectors from each robot in N within d_3 . A cohere vector is also used to keep the swarm from dispersing indefinitely. This vector, \vec{c}_d , is computed by taking the sum of the vectors from this robot's to each neighbor robot's (x, y) position, as long as the neighbor was within the range $[d_2, r]$. The final goal vector is then computed using the following equation:

$$\vec{g}_d = w_r \vec{r}_d + w_c \vec{c}_d \quad (2)$$

3) *Flocking*: For the flocking algorithm, leaders were selected according to the distributed minimum volume enclosing ellipse (MVEE) algorithm [17], which determines the robots which together form the minimum bounding ellipsoid of the swarm. These leaders were given an identical random heading, g_h , from $[0, 2\pi)$. The remaining robots each computed a repulsion vector \vec{r}_f identically to \vec{r}_r the rendezvous algorithm (Section III-B1), except using d_2 as the maximum neighbor range instead of d_1 . The cohesion vector, \vec{c}_f , is computed by taking the sum of the vectors from this robot's to each neighbor robot's (x, y) position, as long as the neighbor was within the range $[d_2, r]$. 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, $\vec{a} = \langle \cos(g_h), \sin(g_h) \rangle$. 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 .

Both leaders and non-leaders computed their goal vector according to the following equation:

$$\vec{g} = w_r \vec{r}_f + w_a \vec{a} + w_c \vec{c}_f \quad (3)$$

4) *Movement Towards Goal*: In the above behaviors, both velocity $v = \|\vec{g}\|$ and angular velocity, α , were capped at the values given in Table I. 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 0.108 radians. Once rotated, the robot would then move forward at the maximum velocity v_{max} . Again, for each step this would be $v_{max}/60s$, or approximately 0.017 meters.

C. Display Types

Four different display types were used to show the current state of the swarm to the participant. The first is called *Full Information*, and shows each robot's (x, y) position and heading, updated at 60Hz. This display requires a high amount of bandwidth from the swarm to operator, as it must transmit three variables (x coordinate, y coordinate, and heading) for each of the 256 robots at 60Hz.

The second display type, called *Centroid/Ellipse* showed only three pieces of information. First, an ellipse bounding the swarm was displayed in red, using the distributed MVEE algorithm presented in Section III-B3 to select a set of robots that defined the ellipse. A second smaller ellipse, in yellow, was displayed within the bounding ellipse. This smaller ellipse bounded the middle 50% of the swarm, which allowed the participants to see some measure of spatial distribution throughout the swarm. For instance, if the inner ellipse was off center relative to the bounding ellipse, the participant can infer the swarm is denser on one side than the other. The final piece of information shown to the participants is a green cross, with the intersection at the centroid of the swarm. The centroid is defined as midpoint of both the minimum and maximum coordinate (x, y) in each dimension. This display requires significantly less information, as the swarm need only transmit the five points which define each ellipse (20 variables), plus two variables for the coordinates of the centroid.

The third and fourth display types showed participants a small subset of the swarm, called the *leaders*. The first of these display types showed only the leader set computed by the MVEE algorithm, this is referred to as the *MVEE Leaders* display. This requires the lowest amount of bandwidth amongst all display types, as the algorithm gives a maximum of eight leaders, which can be defined by 16 variables. The second display type showed the leader set as computed by the modified random competition clustering algorithm, first presented in [18] and based on the RCC algorithm in [19]. This is referred to as the *RCC Leaders* display. The amount of information to be transmitted using this display type depends on the number of leaders selected, which in turn depends on the structure of the communication graph of the swarm. The number of leaders can range from one (if the graph is fully connected) to $N/2$ (if the communication graph forms a chain). All display types are shown visually in Figure 2.

D. Participant details

Participants were recruited from the University of Pittsburgh and surrounding area. There were 26 participants in total, with an age range of 18-65 years old. Of the 26 participants, 13 were female and 13 were male.

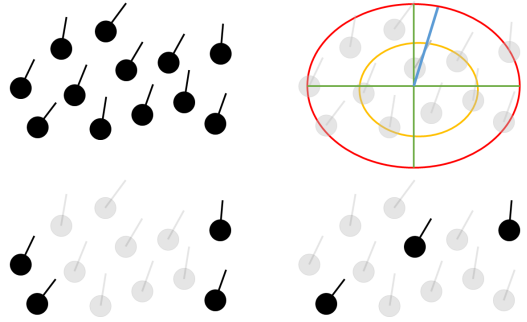


Fig. 2. The four different display types used in the experiment: *Full Information* (top left), *Centroid/Ellipse* (top right), *MVEE leaders* (bottom left), and *RCC leaders* (bottom right). Greyed-out robots are still present in the simulation, but not visible to the participants. Note that the four robots defining the enclosing ellipse in the *Centroid/Ellipse* display are the same ones selected as leaders for the *MVEE Leaders* display.

IV. RESULTS

The main method for evaluation of participants' predictions of the final state of the swarm was accuracy as compared to the actual final state (the positions of the robots at the end of the 30 seconds of the trial). Each time the user gave a prediction, their chain of drawn points was used to compute a convex hull around those points. Similarly, the final robot state was represented by the convex hull of the positions of the robots at the end of 30 seconds. Given two convex hulls, it is therefore easy to compute the accuracy according to the following equation:

$$A = H_u \cap H_r / H_u \cup H_r \quad (4)$$

where A is the accuracy variable, H_u is the convex hull of the user-drawn estimate, and H_r is the convex hull of the final robot positions. We decided on this measure of accuracy because it penalizes inaccuracy in the position of the swarm, as well as both overestimating the shape (i.e. drawing a large shape covering much more area than the swarm takes up in reality) and underestimating the shape (drawing a small shape in the middle of the swarm).

Using an analysis of variance test (ANOVA), results presented in Figure 3 show that there are significant differences for accuracy between both behavior ($F(2, 1651) = 414.4, p < .001$) and display type ($F(3, 1650) = 8.68, p < .001$). The accuracy of predictions for the dispersion (65.8%) behavior were significantly higher than both rendezvous ($A = 43.1\%$, $t = 26.97, p < .001$) and flocking ($A = 40.0\%$, $t = 27.63, p < .001$). Accuracy for rendezvous was also significantly higher than flocking ($t = 3.35, p < .001$).

For display types, there was no significant difference between the accuracies for the *Full Information* (50.5%) and *Centroid/Ellipse* (51.2%) displays, however both were significantly more accurate overall than the *MVEE* (45.5%, $p < .001$ for each) or *RCC* (46.9%, $p = .008$ and $p = .002$, respectively) *Leader* displays. Accuracy between the two leader-based displays was not significantly different.

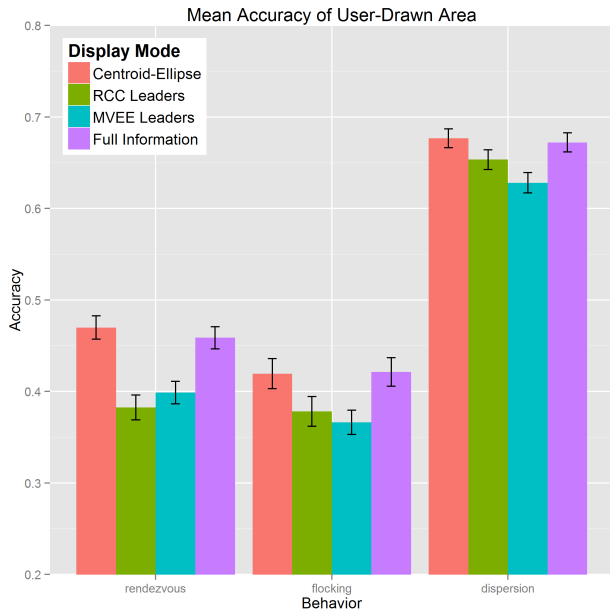


Fig. 3. Average accuracy of all participants across both behavior and display type. Error bars show the standard error.

As one of our main hypotheses is that different behaviors will require different displays to maximize the accuracy of prediction, we investigated the data within each behavior to determine if this hypothesis is true. Results show that, while accuracies in the two leader-based displays were lower than either the *Full Information* or *Centroid/Ellipse* displays overall across all behaviors, when investigating the dispersion behavior significantly, prediction accuracy using *RCC Leaders* (65.3%) were not significantly different than the *Full Information* or *Centroid/Ellipse* displays, while the *MVEE Leaders* display gave significantly lower prediction accuracies (62.8%) than both the *Full Information* ($t(234.11) = 2.88, p = .004$) and *Centroid/Ellipse* displays ($t(238.98) = 3.21, p = .002$) displays.

A. Timing of Predictions

Accuracy is only one of the important qualities of a prediction; the other is the time a prediction is made. Earlier predictions are more beneficial than later predictions, assuming equal accuracy, because they allow information to be gained earlier for whoever or whatever (operator or HSI system) uses them. Therefore, we repeated the previous analysis using only predictions made in the first half (first 15 seconds) of a trial. These were termed *early predictions*. Results show that the differences in accuracy between each behavior for early predictions remain the same as when comparing between behaviors using all predictions ($F(2, 720) = 522.00, p < .001$), with dispersion giving significantly higher accuracy (63.1%) than both rendezvous (34.4%, $t(485) = 25.33, p < .001$) and flocking (25.4%, $t(454.14) = 30.62, p < .001$). Similarly, early predictions for the rendezvous trials were on average




Fig. 4. Participant rankings of each display type across all behaviors.

more accurate than those in flocking trials ($t(472.06) = 7.12, p < .001$).

When restricting predictions to only the first half of a trial, the differences between the display types overall are not present ($F(3, 719) = 0.91, p = .436$). However, when looking within behaviors, the differences are evident. For rendezvous, early predictions using the *Centroid/Ellipse* display were significantly more accurate (37.5%) than the *RCC Leader* display ($A = 31.7\%$, $t(120.86) = 2.40, p = .018$), and marginally significantly more accurate than the *MVEE Leader* display ($A = 33.0\%$, $t(119.89) = 1.85, p = .067$). The *Full Information* display, surprisingly, was not significantly different from either of the two leader-based displays or the *Centroid/Ellipse* display.

Similar results are found for dispersion, except with the *Full Information* display being superior, instead of the *Centroid/Ellipse* display. Early predictions using the *Full Information* display (66.3%) were significantly more accurate than both the *MVEE leader* display ($A = 61.6\%$, $t(90.84) = 2.10, p = .038$) and *RCC leader* display ($A = 61.2\%$, $t(114.27) = 2.52, p = .013$). There were no significant differences between the *Centroid/Ellipse* display and either leader-based method, nor was there a significant difference in early prediction accuracies between the two leader displays.

B. Participant Feedback

Feedback from participants, while subjective, is important for interface designers of any system merging both autonomy and human intelligence—HSI systems are no different. To get qualitative feedback from participants, they were asked to complete a short survey ranking their preferred display types after the experiment completed. The responses here clearly show that the *Full Information* condition was widely preferred by participants, with the *Centroid/Ellipse* display coming in a distant second. Neither of the leader-based display methods received the top ranking overall by any participant; however, the *MVEE Leader* display does seem to be preferred slightly more than the *RCC Leader* display (see Figure 4).

V. DISCUSSION AND CONCLUSIONS

The results of this study suggest that there are differences in the predictive abilities of operators depending on what

behavior a swarm is exhibiting, as well as what display mode is being used. While results show that the dispersion behavior allowed for significantly more accurate predictions overall, we believe this is due to the fact that small differences between the boundaries of the user-drawn and final robot convex hulls make up a smaller percentage of the overall area of their union than in the flocking or rendezvous behaviors. Specifically, due to the fact the participants could not be perfectly precise down to the pixel level, there is always going to be some inaccuracy in drawing the estimated final swarm shape; however, because this inaccuracy does not increase with swarm size, it makes up a larger percentage of the overall swarm area when the space covered by the swarm is small—as it is in flocking and rendezvous—thus giving those behaviors higher errors.

Instead, we believe the results within each behavior are more telling. Because the *Full Information* and *Centroid/Ellipse* displays gave equal accuracy regardless of the display type or behavior, we can confidently say that, in most bandwidth-limited situations, displaying a bounding ellipse and centroid is preferable to individual information about each robot. This confirms the results of the swarm control study conducted in [15]. This does not mean that leader-based displays are useless, however. Because the task in this study involved estimating the spatial qualities of a swarm, the *Full Information* and *Centroid/Ellipse* displays may have been superior because they give better spatial information. The *Centroid/Ellipse* display in particular gives *only* spatial information. The leader-based displays, however, may be better suited for tasks where other characteristics need to be estimated—such as consensus of the swarm on a goal direction, or overall connectivity of the sensing or communication graphs. Therefore, we believe future studies should extend this research to predictions where non-spatial information is needed instead.

Another possible explanation for the high accuracy of predictions in the *Centroid/Ellipse* display is that this display preserves the global, or Gestalt-type properties of a swarm. Namely, with a centroid and ellipse operators get a clear picture of the overall shape and position of the swarm in a single glance. With leader-based displays, operators must infer this information from the positioning of the leaders. Doing so while the swarm is moving may be difficult enough to give the lower performance seen here.

A clear direction for future studies is to investigate how display types like the ones presented here improve not only an operator's ability to predict and recognize swarm behaviors, but also to control them. Specifically, when switching between behaviors, future studies should investigate how the display should be changed to adapt to the new behavior. Should the display characteristics remain the same to preserve the situational awareness of the operator, or should they be adapted to best show the properties of the swarm necessary for the current behavior? Using the methods presented here, we believe future work can and should address this next step.

ACKNOWLEDGMENTS

This research is 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] A. Kolling, K. Sycara, S. Nunnally, and M. Lewis, "Human swarm interaction: An experimental study of two types of interaction with foraging swarms," *Journal of Human-Robot Interaction*, vol. 2, no. 2, pp. 103–128, 2013.
- [4] P. Walker, S. Nunnally, M. Lewis, N. Chakraborty, and K. Sycara, "Levels of automation for human influence of robot swarms," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 57, no. 1. SAGE Publications, 2013, pp. 429–433.
- [5] I. Couzin, J. Krause, R. James, G. Ruxton, and N. Franks, "Collective memory and spatial sorting in animal groups," *Journal of theoretical biology*, vol. 218, no. 1, pp. 1–11, 2002.
- [6] 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.
- [7] M. A. Goodrich, B. Pendleton, S. Kerman, and P. Sujit, "What types of interactions do bio-inspired robot swarms and flocks afford a human?" in *Robotics: Science and Systems VIII*. MIT Press, 2013, pp. 105–112.
- [8] F. Stürzel and L. Spillmann, "Perceptual limits of common fate," *Vision research*, vol. 44, no. 13, pp. 1565–1573, 2004.
- [9] G. Johansson, "Visual perception of biological motion and a model for its analysis," *Attention, Perception, & Psychophysics*, vol. 14, no. 2, pp. 201–211, 1973.
- [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] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] S. Nunnally, P. Walker, A. Kolling, N. Chakraborty, M. Lewis, K. Sycara, and M. Goodrich, "Human influence of robotic swarms with bandwidth and localization issues," *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on*, pp. 333–338, 2012.
- [16] 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>.
- [17] 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.
- [18] P. Walker, S. Amirpour Amraii, M. Lewis, N. Chakraborty, and K. Sycara, "Control of swarms with multiple leader agents," in *Systems, Man and Cybernetics (SMC), 2014 IEEE International Conference on*. IEEE, 2014, pp. 3567–3572.
- [19] K. Xu and M. Gerla, "A heterogeneous routing protocol based on a new stable clustering scheme," in *MILCOM 2002. Proceedings*, vol. 2. IEEE, 2002, pp. 838–843.