

Human-Robot Swarm Interaction with Limited Situational Awareness

Gabriel Kapellmann-Zafra^(✉), Nicole Salomons, Andreas Kolling,
and Roderich Groß^(✉)

Natural Robotics Lab, Department of Automatic Control and Systems Engineering,
The University of Sheffield, Sheffield, UK
{gkapellmann,rgross}@sheffield.ac.uk

Abstract. This paper studies how an operator with limited situational awareness can collaborate with a swarm of simulated robots. The robots are distributed in an environment with wall obstructions. They aggregate autonomously but are unable to form a single cluster due to the obstructions. The operator lacks the bird's-eye perspective, but can interact with one robot at a time, and influence the behavior of other nearby robots. We conducted a series of experiments. They show that untrained participants had marginal influence on the performance of the swarm. Expert participants succeeded in aggregating 85 % of the robots while untrained participants, with bird's-eye view, succeeded in aggregating 90 %. This demonstrates that the controls are sufficient for operators to aid the autonomous robots in the completion of the task and that lack of situational awareness is the main difficulty. An analysis of behavioral differences reveals that trained operators learned to gain superior situational awareness.

1 Introduction

As multi-robot systems continue to assist humans in an increasing variety of roles, more and more humans will need to interact with them. Swarm robotic systems are a subset of multi-robot systems with characteristics, also observed in natural swarms, that seem to complicate such interactions. In particular, they use local sensing and communication capabilities, have no access to global information, and are governed by simple rules. Yet, complex behaviors may result from interactions among the robots and of robots with their environment.

Krause et al. [9] proposes that swarm-intelligent systems could become a useful tool for solving problems. Self-organization could lead to novel solutions to problems, for example, path finding in dynamic environments, exploration, or rescue and support. However, this does not mean that swarm intelligence will necessarily be the best solution for a particular situation. Interaction with humans can be beneficial, for example, to adapt and react to critical environment changes or make decisions in which human experience is important.

A problem with swarm systems is that the attractive features of their social structure also makes interactions with (external) users complex. There have been

several proposals suggested for the implementation of human interaction with robot swarms, surveyed in [8]. Some of these explored proposals are: controlling units as a leader [18], with haptic interactions [17] or with body gestures [1].

In this paper we consider further restrictions where operators can only receive local sensor information from a single robot and its local cluster and do not have access to global positions. This is more in line with the nature of distributed systems in which global state information can be difficult to obtain. In addition, it simulates better real-world scenarios such as search and rescue missions, where keeping visual contact with each member of the swarm would be infeasible.

The paper is organized as follows. In Sect. 2 related work is presented. Section 3 details the methodology used in this study. The results are presented in Sect. 4. Section 5 concludes the paper.

2 Related Work

Research on human-swarm interaction has produced a number of studies on the topic. A recent survey on human interaction with robotic swarms is available in [8]. One of the most frequently studied questions in human-swarm interaction is the design of appropriate control inputs for swarms. Four basic control approaches are distinguished in [8]: (1) algorithm switching; (2) parameter changing; (3) indirect control through environment influence; and (4) control through selected leaders. In our study we utilize algorithm switching and control through teleoperated leaders. These two approaches have been studied in a variety of scenarios.

In [3] a hybrid control approach allowed operators to teleoperate leaders and switch the swarm algorithm after teleoperation. In contrast to our work, the operator had access to global position information for all robots and algorithm switches did not propagate through the local swarm network. Teleoperation of leaders and their effect on the remaining swarm has been studied in [7] with an emphasis on determining what kind of flocking and motion behaviors can be generated from different human inputs.

The problem of obtaining and visualizing the information about the state of the swarm has been studied somewhat less than swarm controls. The studies that focused on controls usually assumed access to and present the position of all swarm robots in an interface. Work that considers limited access to state information due to bandwidth or latency restrictions has been presented in [14, 19]. An emphasis on how to display a limited amount of information while still allowing the human to detect patterns is found in [5]. In [11] a brief overview of potential display visualization for swarms is given.

We are not aware of any work in human-swarm interaction that studies the impact on and adaption of operators when removing access to global state information (such as position) and restricting the interaction to be strictly local and distributed. Such interaction schemes are pointed out as desirable in [2, 4]. It is not clear, however, what the cost of such an interaction scheme is with regard to the operator's ability to observe and control the swarm effectively. Our study is aimed to contribute towards this area.

Neglect benevolence [13,19] is a concept that is concerned with the dynamic nature of emergent behaviors. Most swarm algorithms require time to converge to an emergent behavior and should their dynamics be disturbed, for example, by interacting with an operator, convergence may be delayed and the interaction may be detrimental. Hence, some swarms may benefit from a period of neglect. This stands in contrast to the concept of neglect tolerance studied in the context of multi-robot systems [15]. In these systems the performance deteriorates due to periods of neglect. In our study we observed positive effects and some learning of neglect benevolence dynamics by experienced operators, further supporting the evidence from [13] that human operators can learn to adapt the timing of their commands to the neglect benevolence of the swarm.

3 Methodology

3.1 Problem Formulation

We study a distributed interaction scheme between a human and a swarm of robots in the context of an aggregation task. The robots operate in a bounded environment with wall obstructions. They are equipped with motors, a communication device, a camera and proximity sensors. Initially, the robots are randomly distributed in the environment. Their goal is to aggregate into a single cluster in a given time period. By default, the robots execute the aggregation (clustering) behavior presented in [6]. Unlike [6], we consider environments with obstructions and robots which have limited range sensors, both of which can prevent aggregation.

The operator has access to a graphical interface that provides a connection to a single randomly selected robot. The robot, upon request, transmits either the readings from the proximity sensor or the camera. The operator can also issue motion commands to the selected robot and switch the behavior of it and all its neighbors to either clustering, following, or gossip, which can count the robots in the cluster. The operator does not have access to any other state information, but is shown a map of the environment prior to the experiment. Details on the robots, swarm behaviors and interface are provided in the following sections.

3.2 Robot and Simulation Platform

In our study, the operator interacts with a swarm of simulated robots. We use the open source physics simulator Enki [10], which treats the kinematics and dynamics of rigid objects in two dimensions.

We consider the e-puck miniature mobile robot [12]. Enki has a built-in model of the e-puck. The robot is represented as a disk with a diameter of 7.4 cm and a weight of 152 g. It is a differential wheeled robot. Each wheel can move backward and forward at different speeds with a maximum of 12.8 cm/s.

Each robot has a color camera, providing a horizontal field of view of 56 degree and a maximum range of 150 cm. We assume that the robot can use

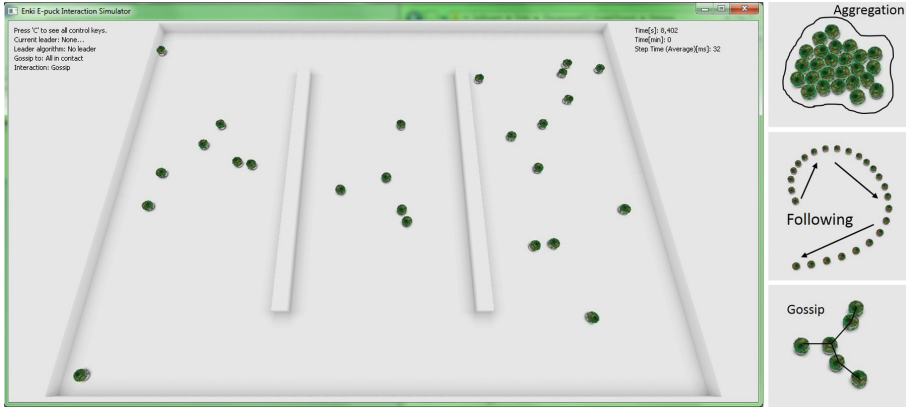


Fig. 1. Snapshot showing the simulation environment from a bird’s-eye perspective. The operator is not provided with this global state information (except for the control experiments). The robots can perform three swarm behaviors: aggregation (top right), following (middle right), and gossip (bottom right).

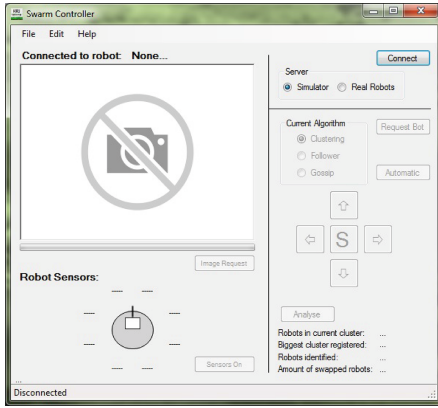
the camera to detect other robots in the direct line of sight. In addition, the robot has eight infra-red sensors distributed around its body and a simulated Bluetooth communication device. These sensors help the operator interact with the robots.

Figure 1 provides an overview of the simulation environment. The robots operate in a rectangular arena of size 400×300 cm that contains two walls, which are symmetrically arranged. Their lengths are $2/3$ of the corresponding side length of the arena and divide it in three equally sized areas joined only at the extremes. The walls are sufficiently tall to prevent robots at opposite sites from perceiving each other.

3.3 Swarm Behaviors

Each robot can execute three swarm algorithms (the corresponding behaviors are shown in Fig. 1):

- The *aggregation* algorithm is identical to the one reported in [6]. By default, this algorithm is executed. Each robot measures whether another robot is in its direct line of sight or not. It maps this binary sensor input onto a pair of constant wheel velocities. For simplicity we state the velocity values after scaling them from -1 to 1 . If another robot is perceived, the velocity pair is $(1, -1)$; the robot thus turns clockwise on the spot. Otherwise, the scaled velocity pair is $(-0.7, -1)$; the robot thus moves backward, following a clockwise circular trajectory. As shown in [6], this simple algorithm leads to the overall aggregation of the swarm, provided the sensing range is sufficiently large and no obstacles are present in the environment.



(a) GUI



(b) Epuck Frame

Fig. 2. (a) Graphical user interface that the participants used in the human-robot swarm interaction study. (b) Image taken by a simulated robot and provided to the operator via the graphical user interface.

- The *follower* algorithm uses the same line-of-sight sensor and reactive control architecture as the aggregation algorithm. The wheel velocity constants are however different. If another robot is perceived, the robot moves straight forward $(1, 1)$, attempting to approach the detected robot; otherwise, the robot rotates anti-clockwise on the spot $(-1, 1)$.
- The *gossip* algorithm prevents the selected robot from changing its position (yet, the operator has control over its orientation). The robot requests all other robots in its neighborhood to stop. These requests get relayed, so that all ‘connected’ robots finally stop. Only in this mode the operator can obtain a count of the connected robots. The counting algorithm is explained in [16].

The robots do not use their IR sensors for obstacle avoidance. Nevertheless, the user can detect any obstacle by monitoring a robot’s sensors (IR or camera).

3.4 User Interface

The interaction between the operator and the robot swarm occurs through the graphical user interface (GUI) shown in Fig. 2a. The operator can connect with one random robot at a time (“Request Bot” button). The operator is shown the robot’s (unique) identification number and which of the three algorithms is currently being executed.

Once connected to a robot, the operator has two options to obtain information from its sensors. To simulate bandwidth limitations of the hardware, only one of these options can be selected at a time:

- Requesting an image of the camera: By clicking the “Image Request” button, the user is shown a 80×60 pixels snapshot of the robot’s camera as shown

in Fig. 2b. Between requesting and displaying the image, a 1 s delay occurs, emulating the time the Bluetooth protocol would take to transfer the data.

- Monitoring the robot’s other sensors: By activating the “Sensors On/Off” button, the operator can either observe the status of the binary line-of-sight sensor, indicating whether another robot is perceived, or, they can see the raw values of the proximity sensors. Unlike the camera image, the sensor data is updated periodically.

The operator has two options to influence the robots:

- The operator issues basic motion commands to the currently selected robot. These are forward, backward, rotate left, rotate right and stop. When in the gossip mode, the forward and backward buttons are disabled.
- The operator changes the algorithm that is being executed on the selected robot to either aggregation, follower or gossip. The change is broadcast from the selected robot to the entire local network of robots connected via IR, and all robots in the network change their algorithm as well. When disconnecting from a robot, the algorithm which it is currently executing remains active. However, it is not possible to disconnect from a robot while in gossip mode. This is to avoid robots from being left in a static position.

3.5 Experimental Setup

A series of human-robot swarm interaction experiments were conducted. The study received ethical approval by The University of Sheffield. All participants were students of the university and their age ranged between 18 and 39.

Participants were given a 10 min presentation explaining the mission, the three swarm behaviors, and the user interface. They were also shown a snapshot of the simulation environment (see Fig. 1).

The default group of participants, referred to as *untrained participants*, were not provided with the opportunity to test the system in advance of the experiment. Overall, data for 38 untrained participants were collected. The data for three participants were excluded as they did not complete all three trials.

Six further participants received training on the system prior to conducting trials. Three of these received 60 min training (five to six trials), these are referred to as *trained participants*. Three further participants, chosen from the developer team, received several hours of training and are considered as *experts*.

All participants conducted three trials with 25 robots and lasting 10 min (600s) each. The untrained participants were further assigned to one of two conditions at random:

- Blind-Blind-Blind (BBB): Participants of this group had no access to global state information (i.e., the bird’s-eye perspective) during any of their trials. There were 19 participants in this group.
- Visual-Blind-Blind (VBB): Participants of this group had access to global state information for the entire duration of their first trial (referred to as VBB_V), but had no access to that information during the second and third trials (referred to as VBB_B). There were 16 participants in this group.

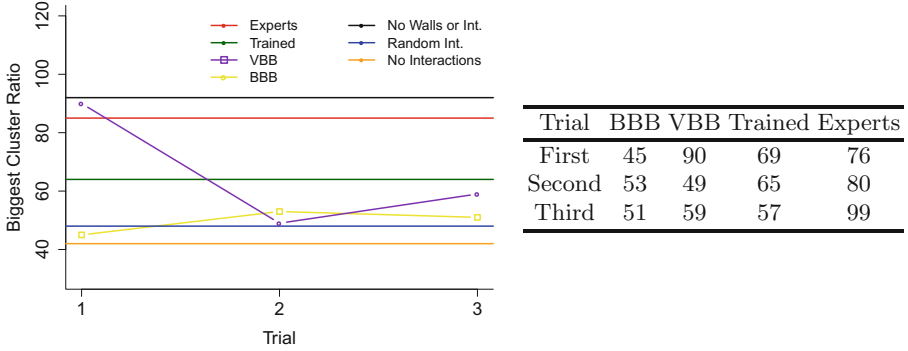


Fig. 3. The graph presents the percentage of robots in the biggest cluster at the end of the trial for each group of untrained participants (BBB and VBB). The performance of all baselines and the average performance in the last three trials of trained and expert participants are plotted as lines to provide a reference performance. The table presents the percentage of robots in the biggest cluster at the end of the trials with human participants.

Throughout all trials, the robots’ positions and the participant’s interactions through the interface were recorded.

4 Results

4.1 Performance Metrics and Baseline Performance

The main performance metric is the number of robots in the largest cluster. A pair of robots is considered in close proximity if the distance between their centers is less than 15 cm. We consider two robots that are in close proximity to belong to the same cluster. Moreover, if $\{a, b\}$ belong to the same cluster and $\{b, c\}$ belong to the same cluster, then the same holds true for $\{a, c\}$.

We establish the following baselines for comparison:

- *No Interaction*: This is the performance of the swarm in the absence of any interaction with an operator. In other words, each robot of the swarm executes the aggregation algorithm for the entire duration of the trial.
- *No Walls or Interactions*: This is the performance of the swarm when aggregating in the absence of wall obstructions and interactions with an operator. These represent the ideal conditions for the algorithm as presented in [6].
- *Random Interactions*: This is the performance of the swarm when interacting with a virtual operator agent choosing random commands drawn from a distribution that models the average participant across all trials.

For each of the baseline performance measures, 10 trials of 600s were conducted. The table in Fig. 3 shows the average size of the biggest cluster at the end of the trial. Random commands resulted in slightly better performance than no interactions but with a larger standard deviation.

4.2 Operator Performance

We compared the performance of untrained, trained and expert operators to the baseline performance. Figure 3 presents this comparison where untrained operators with access to real-time global state information of the position of all robots (trial 1 for group VBB) aggregate 90 % of robots and perform as well as the ‘no walls or interactions’ baseline. This validates the efficacy of the swarm controls available to the operator. Operators were able to use the available controls to mitigate the shortcomings of the aggregation algorithm in the presence of obstacles.

Untrained participants in their final trial aggregated 51 % (BBB) or 59 % (VBB) of robots into a single cluster, an improvement over the no interactions baseline that aggregated 42 % (two-sided Mann–Whitney test, p -values = 0.049 and 0.029). The blind trials of both groups of untrained operators (trials 1, 2 and 3 for BBB and trials 2 and 3 for VBB) did not perform significantly better than the random interaction baseline (two-sided Mann–Whitney test, p -values = 0.985 and 0.481). Note that the proportion of the types of instructions is identical but the random interactions baseline does obviously not exploit any sensory information. This suggests that operators have similar difficulties in exploiting local sensory information.

A comparison between the blind trials of BBB and VBB shows no significant differences in performance (two-sided Mann–Whitney test, p -values = 0.215). This suggests a minimal learning effect of the initial trial with global state information. It further supports the conclusion that operator performance in blind trials was diminished due to a lack of situational awareness rather than lack of planning. If it were due to a lack of planning the trial with global state information would be expected to have facilitated the learning of plans.

Trained and expert operators were able to obtain significantly improved performance in their three trials over the random interactions baseline (two-sided Mann–Whitney test, p -values = 0.029 and 0.001). They aggregated 57–69 % and 76–99 % of robots respectively.

In summary, the results show a dramatic drop in performance of untrained operators when removing access to global state information, with performance on par with a random agent. The recovery of performance for trained and expert operators shows that learning does occur and warrants a closer look, in the following section, at the actions and strategies that are being learned.

4.3 Interaction Analysis

A detailed history of the operators’ actions was recorded throughout all trials. The data is grouped into three categories: (i) the operator moves the robot, (ii) the operator uses the robot’s sensors and (iii) the operator switches between algorithms. Figure 4a shows the distribution of time spent on these activities for the last three trials for untrained, trained and expert operators.

As expected, untrained operators with access to global state information (trial 1 in group VBB) rarely request local sensory information and instead

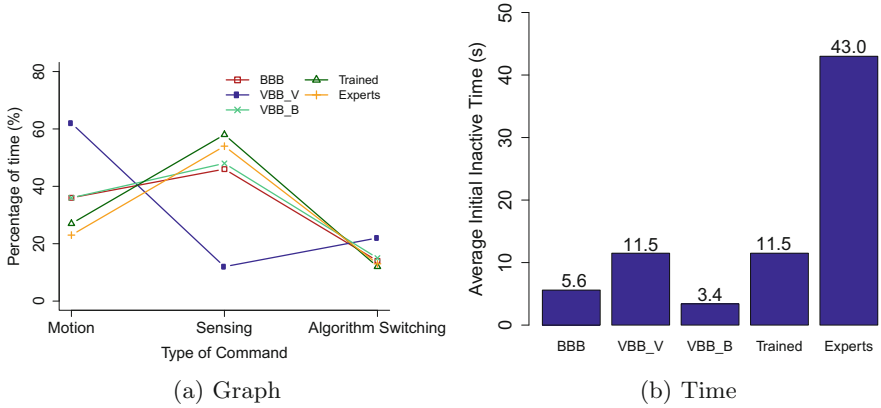


Fig. 4. (a) Percentage of time spent with type of command. (b) Average initial inactive time.

move the robots and switch algorithms more frequently. Untrained operators in the blind trials, however, spend a larger proportion on obtaining sensor information to recover some situational awareness. The key observation is found by comparing trained and expert operators to untrained operators.

The improved performance of the later seems to rely on more requests for sensor observations while reducing the amount of time spent moving the robots. The time spent switching algorithms is identical between all groups. Given that the time spent on motion commands is significantly less for trained and expert operators than for untrained operators with global state information, the efficiency of the motion commands for the former group was higher. This is likely where the training effect materializes.

In addition to varying the time spent on certain activities we observed a difference in the initial interactions with the swarm, that is, the time the operators waited at the beginning of the trials before performing the first interaction. Figure 4b show the average time that operators waited at the beginning of the trials. This period of inactivity allows the swarm to exhibit local aggregation behavior and form small clusters within parts of the environment. These clusters may be controlled more effectively than dispersed robots. Operators that interact with the swarm too early may disturb this process and have hence less effective subsequent interactions. This suggests evidence for the concept of neglect benevolence in these experiments that is being learned and exploited by trained and expert operators. It is worth noting that untrained participants with access to global state information also increase their initial period of inactivity while observing the global dynamics, yet they do not repeat this in subsequent blind trials.

To illustrate the above more qualitatively, Fig. 5(a-f) shows a sequence of snapshots taken from an example trial. The initial positions of the robots are randomly distributed through the arena (a). Because of the aggregation

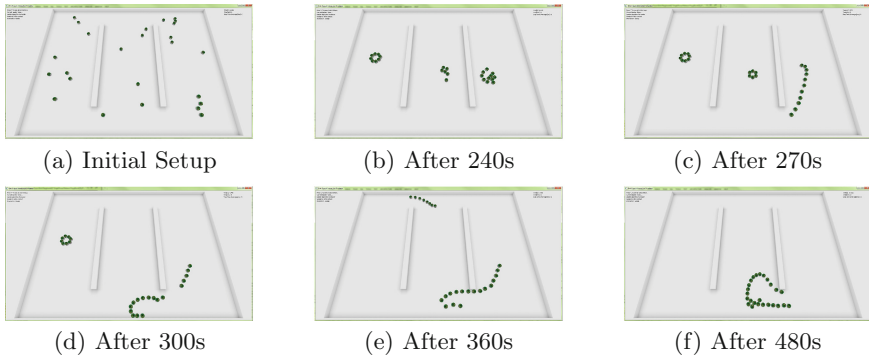


Fig. 5. Sequence of snapshots taken during a trial with an expert participant. The expert was not provided with the bird’s-eye view of the scene, which is depicted here.

algorithm, the robots start grouping and form three clusters (b). The operator then starts moving the right cluster to the center area (c). The operator finds the third cluster and guides it to the center area (d). Again, when the robots are in visual range, they attempt to group together (e). Finally, the operator is monitoring the process until the swarm reports a complete aggregation of the swarm (f).

5 Conclusions and Future Work

This study investigated a distributed human-swarm interaction scheme in which operators have access to only local information when aiding a swarm in an aggregation task. Operators had access to swarm controls with which they were able to complete the aggregation task successfully when given global state information. When given only local information, however, untrained operators did not perform significantly better than random interactions. Nor did they exhibit a significant learning effect within three trials. Furthermore, operators that once were given global state information did not demonstrate improved performance on subsequent trials when being restricted to local information. This suggests no learning benefit from having observed the global dynamics once.

Trained and expert operators, with at least one hour of training, showed significantly improved performance suggesting the task is solvable. These operators compensated the lack in global situational awareness with increased requests for local sensory information while reducing the number of motion commands. Expert operators performed nearly as well as the baseline performance of the autonomous algorithm under ideal conditions, that is, without obstructions.

In addition, we observed evidence for neglect benevolence for trained and expert operators. These operators waited at the beginning of the trial for the swarm to converge to the emergent local clusters. From this configuration, interactions with the swarm were more beneficial as emerging clusters could be

changed into leader-follower formations more easily. Untrained operators disturbed and interacted with the swarm prior to it settling into local clusters.

Overall, our findings suggest that exposure to global swarm dynamics does not necessarily accelerate learning, neither for improving situational awareness nor for understanding swarm dynamics to accommodate for neglect benevolence. In addition, learning to interact with a swarm through a distributed interaction scheme that relies on local information requires training times even for simple tasks and interfaces. This should inform future research on human-swarm interaction.

Simulation experiments, however, only offer a limited potential to validate human-swarm interaction schemes as they simulate simplified dynamics. Future work is planned to investigate the interaction scheme with physical robots and a wider range of tasks. This will answer whether the presented findings generalize to other scenarios and whether the dynamics of physical robots interfere with successful human-swarm interaction.

Acknowledgments. The first author acknowledges scholarship support by Consejo Nacional de Ciencia y Tecnologia (CONACYT). All authors thank the students of The University of Sheffield who volunteered to participate in the experiment.

References

1. Alonso-Mora, J., Haegeli Lohaus, S., Leemann, P., Siegwart, R., Beardsley, P.: Gesture based human-multi-robot swarm interaction and its application to an interactive display. In: Proceedings of 2015 IEEE International Conference on Robotics and Automation (ICRA 2015), pp. 5948–5953. IEEE (2015)
2. Barca, J.C., Sekercioglu, Y.A.: Swarm robotics reviewed. *Robotica* **31**(03), 345–359 (2013)
3. Bashyal, S., Venayagamoorthy, G.K.: Human swarm interaction for radiation source search and localization. In: Proceedings of 2008 IEEE Swarm Intelligence Symposium (SIS 2008), pp. 1–8. IEEE (2008)
4. Brambilla, M., Ferrante, E., Birattari, M., Dorigo, M.: Swarm robotics: a review from the swarm engineering perspective. *Swarm Intell.* **7**(1), 1–41 (2013)
5. Brown, D.S., Kerman, S.C., Goodrich, M.A.: Limited bandwidth recognition of collective behaviors in bio-inspired swarms. In: Proceedings of 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), pp. 405–412 (2014)
6. Gauci, M., Chen, J., Li, W., Dodd, T.J., Groß, R.: Self-organized aggregation without computation. *Int. J. Robot. Res.* **33**(8), 1145–1161 (2014)
7. Goodrich, M.A., Pendleton, B., Kerman, S., Sujit, P.: What types of interactions do bio-inspired robot swarms and flocks afford a human? In: Proceedings of 2013 Robotics: Science and Systems (RSS VIII), pp. 105–112. MIT Press (2013)
8. Kolling, A., Walker, P., Chakraborty, N., Sycara, K., Lewis, M.: Human interaction with robot swarms: a survey. *IEEE Trans. Hum. Mach. Syst.* **PP**(99), 1–18 (2015)
9. Krause, J., Ruxton, G.D., Krause, S.: Swarm intelligence in animals and humans. *Trends Ecol. Evol.* **25**(1), 28–34 (2010)
10. Magnenat, S.: Laboratory of Intelligent Systems, EPFL, Lausanne: Enki reference documentation (2005). <http://lis2.epfl.ch/resources/download/doc1.0/libenki/>. Accessed 25 Aug 2015

11. Manning, M.D., Harriott, C.E., Hayes, S.T., Adams, J.A., Seiffert, A.E.: Heuristic evaluation of swarm metrics' effectiveness. In: Proceedings of Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI 2015), pp. 17–18. ACM (2015)
12. Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., Magnenat, S., Zufferey, J.C., Floreano, D., Martinoli, A.: The e-puck, a robot designed for education in engineering. In: Proceedings of 9th Conference on Autonomous Robot Systems and Competitions, vol. 1, pp. 59–65 (2009)
13. Nagavalli, S., Chien, S.Y., Lewis, M., Chakraborty, N., Sycara, K.: Bounds of neglect benevolence in input timing for human interaction with robotic swarms. In: Proceedings of 10th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI 2015), pp. 197–204. ACM (2015)
14. Nunnally, S., Walker, P., Kolling, A., Chakraborty, N., Lewis, M., Sycara, K., Goodrich, M.: Human influence of robotic swarms with bandwidth and localization issues. In: Proceedings of 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC 2012), pp. 333–338. IEEE (2012)
15. Olsen, D.R., Goodrich, M.A.: Metrics for evaluating human-robot interactions. In: Proceedings of 2003 Workshop on Performance Metrics for Intelligent Systems (PerMIS 2013), vol. 2003, p. 5 (2003)
16. Salomons, N., Kapellmann-Zafra, G., Groß, R.: Human management of a robotic swarm. In: Alboul, L., Damian, D., Aitken, J.M. (eds.) TAROS 2016. LNCS, vol. 9716, pp. 282–287. Springer, Heidelberg (2016). doi:[10.1007/978-3-319-40379-3_29](https://doi.org/10.1007/978-3-319-40379-3_29)
17. Setter, T., Fouraker, A., Kawashima, H., Egerstedt, M.: Haptic interactions with multi-robot swarms using manipulability. *J. Hum. Robot Inter.* 4(1), 60–74 (2015)
18. Walker, P., Amirpour Amraii, S., Chakraborty, N., Lewis, M., Sycara, K.: Human control of robot swarms with dynamic leaders. In: Proceedings of 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2014), pp. 1108–1113. IEEE (2014)
19. Walker, P., Nunnally, S., Lewis, M., Kolling, A., Chakraborty, N., Sycara, K.: Neglect benevolence in human control of swarms in the presence of latency. In: Proceedings of 2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC 2012), pp. 3009–3014. IEEE (2012)