Using Haptic Feedback in Human-Swarm Interaction

Steven Nunnally, Phillip Walker, Michael Lewis
University of Pittsburgh, USA

Nilanjan Chakraborty, Katia Sycara
Carnegie Mellon University, USA

ABSTRACT

A swarm of robots is a large group of individual agents that autonomously coordinate via local control laws. Their emergent behavior allows simple robots to accomplish complex tasks. Since missions may have complex objectives that change dynamically due to environmental and mission changes, human control and influence over the swarm is needed. The field of Human Swarm Interaction (HSI) is young, with few user studies, and even fewer papers focusing on giving non-visual feedback to the operator. The authors will herein present a background of haptics in robotics and swarms and two studies that explore various conditions under which haptic feedback may be useful in HSI. The overall goal of the studies is to explore the effectiveness of haptic feedback in the presence of other visual stimuli about the swarm system. The findings show that giving feedback about nearby obstacles using a haptic device can improve performance, and that a combination of feedback from obstacle forces via the visual and haptic channels provide the best performance.

Keywords: Robotics, Experiments, Haptics, Swarms, Human-Robot Interaction, Human-Swarm Interaction

INTRODUCTION

A robotic swarm consists of a large collection of simple robots with limited sensing, communication, actuation, and computational capabilities. Individual robots act according to simple local rules and exhibit a wide range of behaviors, such as flocking (Reynolds, 1987; Couzin, Krause, James, Ruxton, & Franks, 2002; Spears & Spears, 2012; Bruemmer, 2002) without any centralized controller. However, for performing complex tasks like search and exploration in obstacle-filled environments, it is usually difficult to design local control laws for individual swarms that guarantee good performance of the overall system. To use swarm robotic systems in a complex mission, the presence of human operators are required to guide the behaviors of the swarm towards accomplishing mission goals. A key aspect of using a human to control a swarm is the transfer of information between the human and the swarm. The human has to obtain information about the state of the swarm in order to control it. In the extant literature, experimental studies in human swarm interaction have primarily explored the use of the visual channel of the human to transfer information about the swarm state. However, the use of the haptic channel has not been studied adequately in HSI, except for formation control tasks and with small multi-robot systems (fewer than 10 robots). Therefore, along with background literature, this chapter will present experiments to explore the benefits of using a haptic device to control swarm robots (in addition to the visual channel), using large numbers of robots to demonstrate scalability.

A key aspect of using a haptic device to feed information back about the robot state is to decide on the information (or cue) that should be fed back from the robots to the human. In many swarm robotic algorithms, potential field based methods (LaValle, 2006) are used for avoiding obstacles. Roughly speaking, when a robot is near an obstacle, the robot controller computes a virtual force from the obstacle that is inversely proportional to the distance (or some superlinear function of the distance) between the robot and the obstacle. Thus, the nearer a robot is to an obstacle, the greater the force it “experiences” that makes it move away from the obstacle. Therefore, one cue that can be fed back to the human is the net obstacle forces experienced by the robots.

In formation control tasks, since the robots have to usually maintain a rigid formation as they move, each robot has to track a path in order to maintain a formation. Thus, a natural cue in formation control is the tracking error of the robots in following their paths. The use of net forces from the obstacles along with tracking errors has been explored in the context of formation control (Son et al., 2011). In many applications of swarm systems, especially in obstacle-filled environments, it is not always desirable to move the robots in a formation. In these cases, there is no natural notion of tracking error, so the force cue from the obstacles can be fed back to the humans through the haptic device instead.

Although feeding back the obstacles forces is conceptually intuitive and computationally simple, it is not clear a priori whether such information is helpful in improving the performance of the human controller. This is because the force fed back is the (vector) sum of the forces acting on all the robots from all the obstacles. Thus, in some cases like moving from a room to another room through a doorway, one should keep pushing in the direction in which the resistance from the haptic device increases (since the resistance would increase initially and then drop as the robots start moving through the doorway). This is not the intuitive form in which haptics is usually used in human robot interaction. The usual intuition in using haptic feedback from obstacles is that one should *not* try to push in the direction in which the haptic resistance is increasing. Furthermore, since the force fed back to the human is an average of the forces from all the obstacles it is not clear whether it is necessary to use a haptic device to feed back the force or whether a visual representation of this force is enough. In other words, it is unclear whether having redundant information in two channels (haptic and visual) will help or hurt the performance since attending to redundant information may increase operator workload or cause confusion. Therefore, there are a few basic research questions that arise: (a) Does net obstacle force cue fed back through the haptic device help or hurt the performance of the human-swarm system? (b) Can only the visual channel be used to feed back the net obstacle force cue to the human or is a combination of both haptic and visual feedback needed for the obstacle force cue?

The latter half of this chapter will consist primarily of two user experiments conducted to answer the above questions in the context of using haptic feedback for tasks where the robots do not move in formation. These studies use a target searching task in which the human operators have to control a swarm of robots to find hidden targets in obstacle-filled environments. A between-subjects design is used, where there are four different conditions under which the participants perform the target finding task. There are two general findings from the studies: 1) obstacle forces felt by the operator can improve performance and 2) the combination of the feedback of the obstacle forces in the visual and haptic channels provide the best performance.

BACKGROUND

The topic of this chapter is at the intersection of research in human interaction with swarm robotic systems, and the use of haptic feedback in robotic systems. A key component in the human interaction with swarms is the distributed control laws used by the robots. The extant literature on designing distributed control laws for swarms as well as the use of haptic feedback in robotic systems is quite large. Therefore, the focus will be on the literature on human-swarm interaction, swarm control laws/mechanisms used in the context of human supervised control of swarms, and haptic interaction with multi-robot systems.

Human Swarm Interaction (HSI) is a relatively new field of study with a small but growing body of literature, and only a small portion of this literature explores using additional sensory channels to convey swarm state information to the operator. HSI has been studied for a variety of applications including discovering radiation sources (Bashyal & Venayagamoorthy, 2008), aiding firefighters in environments with low visibility (Naghsh, Gancet, Tanoto, & Roast, 2008), tracking a military convoy and surveying areas of interest (Fields, Haas, Hill, Stachowiak, & Barnes, 2009), and focusing the direction of a flock (Klarer, 1998). These works show the utility of HSI for these specific tasks, but only use the operator's vision to give feedback about the swarm.

Designing local control laws for individual robots that result in a global behavior for the whole swarm system has been extensively studied in the context of computer graphics (Reynolds, 1987) and robotic applications (Bullo, Cortés, & Martínez, 2009). In the context of HSI, potential field based methods (Olfati-Saber, 2006), physicomimetic methods (Spears & Spears, 2012), and bioinspired methods (Couzin et al., 2002) for controlling swarms have been used. Others have taken less direct approaches by using evolved controllers (Dorigo et al., 2004; Trianni et al., 2003), or amorphous computing (Bachrach, McLurkin, & Grue, 2008; Bachrach & Beal, 2006).

In (Kira & Potter, 2009), the authors use physicommimetics, and allowed operators to add forces to direct the swarm. In physicommimetics, individual agents in the swarm are viewed as physical particles, and the interactions between these agents are based on physical laws, and the final configuration of the particles is one that minimizes the overall potential energy of the system (Spears & Spears, 2012). Even though the forces might include additional knowledge about the swarm and the environment, the interface only showed each robot's position and heading.

The work on bio-inspired swarm control is influenced from work in collective animal behavior in nature (Couzin et al., 2002). In (Couzin et al., 2002), the authors use a flocking model similar to the original flocking proposed by (Reynolds, 1987) to achieve switching behaviors in schools of fish. This is followed by (Couzin, Krause, Franks, & Levin, 2005), whereby the authors introduce leaders into the swarm, and demonstrate their viability in swarm control. In (Goodrich, Pendleton, Sujit, & Pinto, 2011), the authors use the biologically-inspired swarm model in (Couzin et al., 2002) in the context of HSI, and show that switching the swarm between two topologies (i.e. a torus to a flock) was possible by adjusting the control parameters. These authors also make use of swarm "leaders" and "predators" which can pull and push other members of the swarm, respectively. Although the authors found the leaders more effective, there are still use cases where predators can be beneficial, such as scenarios where it is necessary to break up the swarm into separate smaller groups. Apart from adjusting control parameters, different control algorithm selection was also studied for controlling swarms by humans. Kolling et al. (2012) examines different selection techniques using virtual beacons for switching modes of swarm members between differing consensus-driven goals (i.e. rendezvous and deployment). Here again, the position and orientation information of all members of the swarm are available to the human.

Using neural networks and other learning algorithms to develop optimal control laws has also received attention from the HSI community. The evolved controller in (Dorigo, et al., 2004) is able to robustly exhibit both aggregation and coordinated motion—two common tasks in swarm robotics. Furthermore, their controller was also successfully ported to a real robots. Similarly, in (Trianni, et al., 2003), the researchers evolve controllers for swarm aggregation using a variety of parameters and initial conditions, showing that the evolutionary method can be robust to changes in these features. In terms of amorphous computing, in (Bachrach, McLurkin, & Grue, 2008; Bachrach & Beal, 2006), the authors introduce protoswarm—an amorphous computing language for controlling a swarm as a spatial computer. This work is extended in (Bachrach, Beal, & McLurkin, 2010) to show that such a framework could be used in real robotic swarms in a laboratory setting to achieve swarm-like behaviors, such as flocking, clustering, and tracking.

Although the different control laws mentioned above that have been used in the context of HSI aim to keep the swarm connected, there is no guarantee that the swarm will remain connected, especially in environments with obstacles. In (Giordano, Franchi, Secchi, & Bülthoff, 2013), the authors present a method of connectivity maintenance using the second smallest eigenvalue of the Laplacian matrix of the communication graph (the Fiedler number). This approach is robust and scalable, and allows for control of global objectives of the group (like formation control) with a human supervisor. The authors demonstrate their approach both in large numbers in simulation, and in groups of four in real-world case studies with a human controlling the robots through a 3-DoF haptic device.

There is also some work that does not assume that perfect state information of the robot swarm is available to the human. Communication bandwidth limitations or latency in data transmission from the swarm to the human along with imperfect localization capability of robots may be responsible for noisy state estimates in many situations. *Summary displays*, which display aggregate information about the swarm (e.g. a bounding ellipse) can be used in situations with limited bandwidth limitations, and *predictive displays*, which estimate future states of the swarm, can be used when latency is present. These have been used in Nunnally et al. (2012) and Walker et al. (2012) to maximize the operator's situational awareness despite limited information returned from the swarm. In (Nunnally et al., 2012), the authors performed experiments with bandwidth limitations on the information returned from the swarm to the operator. It was shown that a summary display (using the centroid of the swarm and the variance in positions) instead of individual state information about each robot, was enough for the subjects to successfully control the swarm in spite of the bandwidth limitations. In (Walker et al., 2012), the authors similarly placed significant latency in the human-swarm communication channel, and then used a predictive display to show that this could help overcome the lack of current information available to the operator. These studies either focus on the various features of the control to make the swarms goal-directed, or focus on the various visual feedback techniques to give the operator more information about the swarm’s state. There are few similar studies which implement a second sensory channel, namely haptic feedback, in HSI. These will be reviewed now.

Haptics in Human-Robot Interaction

In the extant literature, haptic feedback in multirobot systems has been explored mostly in the context of formation control. Researchers in Franchi, Giordano, Secchi, Son, & Bülthoff (2011) and P. Giordano, Franchi, Secchi, & Bülthoff (2011) have developed a control interface giving operators haptic feedback about obstacle avoidance of a small group of robots in path following tasks. The group of robots maintain a rigid formation using a decentralized spring force control, while the operator controls the centroid of the agents to follow a predefined path around obstacles. Later work (Franchi, Secchi, Son, Bülthoff, & Giordano, 2012; P. R. Giordano et al., 2013) removed the rigid formation assumption and allowed for a looser formation control with links between robots allowed to be broken or established. This formation control was coupled with autonomous decisions to split and join around obstacles in the path of the robots. Further publications showed that this control technique was successful with latency, and results were demonstrated both in simulation and with real hardware in three dimensions (Lee, Franchi, Giordano, Son, & Bülthoff, 2011; Secchi, Franchi, Bülthoff, & Giordano, 2012; Franchi, Secchi, Son, et al., 2012; P. R. Giordano et al., 2013; Franchi, Secchi, Ryll, Bülthoff, & Giordano, 2012; Franchi, Masone, et al., 2012; Lee, Franchi, Son, Bülthoff, & Giordano, 2013).

In (Franchi, Secchi, Son, et al., 2012), the authors propose a bilateral teleoperation system with haptic feedback to allow a human operator to control a heterogeneous team of robots in a variety of settings. The authors use a leader-follower paradigm and demonstrate their results with a single leader controlling up to eight robots. In (Lee et al., 2013), the authors give a general architecture for multi-robot teleoperation with multiple leaders. While the use of a leader is not necessarily a detriment, the approach of the studies in this chapter use a broadcast method to communicate with all robots, and there are no designated leader robots. This choice was made, since (arguably) not having special leader robots is in spirit with the swarm robotics literature of being robust to failure of individual swarm members.

Another key distinction of the studies herein from the extant literature is in terms of the metric used for evaluating the performance of the swarm. The basic metric for evaluation in human supervisory haptic control of multi-robot systems is the position tracking error (of the centroid of the swarm to a desired path) which is a sensible metric for formation control tasks where paths have to be followed while maintaining connectivity of the robots. Other metrics of human performance that are more relevant to swarm robotic systems: namely, maneuverability and perceptual sensitivity, have been studied in (Son et al., 2013). Their experimental results show that haptic feedback is not always beneficial in all performance measures—specifically, the perceptual sensitivity measure used therein. This means that haptic feedback might be better suited toward maneuverability and control of the swarm overall, whereas visual feedback may be better for giving the operator awareness of the environment surrounding the swarm.

In (Son et al., 2011; Son et al., 2013), the authors also compared the following different sources for haptic feedback as an operator directs a group of robots to follow a path: 1) repelling forces from obstacles, 2) forces matching the robots’ inertia, 3) a combination of the previous two. The results showed that haptic feedback should be closely tied to environmental information which relates to the given task and not internal swarm state information. The studies herein replace the path following task with an environment exploration task, but the results of Son et al

(2011) still apply, since haptic feedback based on obstacle avoidance is environmental information that can improve performance of the task. Other researchers have shown similar findings with small groups of robots and were also able to prove no collisions and stability with control theory (Rodríguez-Seda et al., 2010). However, as stated before all the studies above were done in the context of path following tasks, an important class of tasks in robotics, but by no means the only one. Furthermore, much of the focus in the previous literature was also on designing haptic feedback laws with passivity property so that theoretical guarantees on stability of tracking are possible. Although these contributions are quite novel and useful for HSI, the following studies are geared toward environment exploration applications where the robots are not necessarily following a planned path, but are being driven by the human to explore different regions of the environment. Therefore, the metrics of performance are quite different from those existing in the literature, and the main focus is on the usefulness of haptic feedback when there is no task-performance related error (e.g., position tracking error) to be fed back to the human via the haptic channel. Thus, the belief is that the contributions of these studies is complementary to the existing literature on haptics in robotic control, and will supplement the existing discussion.

The specific hypothesis for each user study will be discussed in their corresponding sections; however, in general our hypothesis is that adding haptic feedback cues about environment obstacles, in addition to visual feedback of the swarm's position within the environment, will increase the performance of operators in terms of exploration related coverage metrics (like number of targets found, area covered etc.) in obstacle-filled environments more so than either one alone.

USER STUDIES

There are two studies presented that explore the utility of haptic feedback in HSI, see Table 1A. Because the two studies are similar, this section will present the overall design and characteristics of both, with specific differences discussed in each corresponding study section.

The participants in each study are asked to apply virtual forces to the robots in order to influence the swarm towards unexplored areas of environments in search of targets. Each study compares the differences of conditions via a between-subject design. The forces from the obstacles on the robots are fed back to the operator in two different ways: 1) via a haptic device (haptic feedback) and 2) via visual feedback on the interface in a side panel (Force Feedback panel), shown in Table 1B as the Haptic and Visual rows, respectively. The forces felt in the haptic feedback and shown in the Force Feedback panel are exactly the same, but the conditions differ in whether or not they are felt in the input device or shown on the display.

There are two groups in Study I. The first group sees the Force Feedback panel and feels the haptic feedback (labeled as haptics condition in Table 1B). The second group see the Force Feedback panel but does not feel the haptic feedback (labeled as control condition in Table 1B).

In Study II, the participants are divided into four groups. One group feels haptic feedback and sees the Force Feedback panel (labeled as **HV** condition in Table 1). The second group receives haptic feedback without the Force Feedback panel (labeled as **HO** condition in Table 1). Another group sees the Force Feedback panel but does not have haptic feedback (labeled as **CV** condition in Table 1). The final group does not see the Force Feedback panel or feel the haptic feedback (labeled as **CO** in Table 1). Note that the haptics condition from Study I corresponds to the HV condition of Study II and the control condition from Study I corresponds to the CV condition of Study II. The comparisons between conditions in both studies are based on three metrics: 1) number of targets found, 2) environment coverage percentage, and 3) swarm cohesion. These three metrics help explore the effects of haptic feedback with and without a visual representation of the force.

The participants are presented with the same environments in random order within each study. The environments have four variables: 1) environment type, 2) clutter, 3) hidden obstacles, and 4) distractions. Both studies have one environment with hidden obstacles. The only difference between the environments with hidden obstacles and all others is that the obstacles are not displayed to the operator in the main map panel. Study I's hidden obstacle environment is a cluttered hallway (labeled as C in Table 1) while Study II's hidden obstacle environment is a cluttered office (labeled as D in Table 1). Study II has one environment with distractions where single digit addition problems block the main map panel in order to distract the operator (labeled as A in Table 1). This environment is a cluttered hallway. The only difference between all other environments in the studies is the features with in the environment, like doorways, hallways, and clutter. The following subsections will present the setup and characteristics common to both studies.

Robot Control

These experimental studies use a control algorithm for the robots based on the autonomous potential field swarm control algorithm, which uses repulsive forces from obstacles and both repulsive and attractive forces from other robots in order to deploy robots and cover as much area as possible (LaValle, 2006). Human input is allowed in order to explore other parts of the map. The virtual force, which can be mapped to the robot's goal velocity and heading is given by the following:

$$F = F_o + F_r + F_h$$

F is the force applied to each robot, F_o is the force vector due to obstacle interactions, F_r is the force vector due to neighboring robot interactions, and F_h is the force vector due to human interactions. F determines the heading and speed, up to a maximum possible speed, for each robot using its own sensed data from the environment and the broadcast information from the operator.

More precisely, for a given robot located at q_i, let $O = \{o_0, o_1, \dots, o_n\}$ be the set of obstacles in range r_i of the robot. Thus, let $d_i = |q_o - q_i|$ is the distance between the obstacle and the robot, and $\boldsymbol{f}_o = q_o - q_i$ be the vector from the obstacle to the robot, and k_o be the gain parameter for F_o, then:

$$F_o = k_o \sum_{o \in O} \frac{\boldsymbol{f}_o}{d_i^2}$$

F_o forces the robots to spread around the environment, avoiding obstacle collisions and overly redundant coverage of the area around obstacles. The average F_o across all robots is shown in the feedback panel on the right side of Figure 1 and felt in the haptic device if the participant is in the haptics condition.

Similarly, let $N = \{n_0, n_1, \dots, n_n\}$ be all robots in range r_i of the robot located at q_i (the set of neighbors); $d_i = |q_n - q_i|$ be the Euclidean distance between the neighboring robot and the robot at q_i; $\boldsymbol{f}_n = q_n - q_i$ = be the vector from the neighbor to the robot; and z_b and z_e be the inner and outer radii of the neutral zone, respectively. Finally, let k_r be the gain parameter for F_r, then:

$$F_r = \begin{cases} k_r \sum_{r_n \in R} \left(\frac{\boldsymbol{f}_n}{d_i}\right)^2 \ if\ d_i < z_b \\ -k_r \sum_{r_n \in R} \left(\frac{\boldsymbol{f}_n - z_e}{d_i - z_e}\right)^2 \ if\ d_n > z_e \end{cases}$$

The competing forces are used to disperse the robots in a controlled manner for coverage, while the attraction force keeps the swarm from breaking apart. This attractive force is an addition to the deployment algorithm of Howard et. al. (2002) to overcome the limited sensing range, r_i which is 4 meters in these studies, see Figure 3. While the attraction does not guarantee that the swarm remains one cohesive unit, it gives the operator the opportunity keep most of the swarm in one large group that has the ability to sense target with the redundancy threshold required to mark targets.

Finally, to describe the human input component, let $\boldsymbol{h}$ be the input force vector from the haptic device; $r_h = |h|$ be the magnitude of this vector; h_{max} be the maximum allowable value of r_h; and k_h be the gain parameter for F_h, then:

$$F_h = k_h \left(\frac{\boldsymbol{h}}{h_{max}}\right)^2$$

This force moves the robot in the general direction that the operator applies. Each robot determines its own F, then uses this vector as its motion vector. If the magnitude of F is greater than the robot's maximum speed, which is different for each study, then the magnitude of this vector is reduced to that maximum.

Experimental Setup

The following studies use the same user interface. The environment, targets, and swarm consisting of 30 differential drive P2AT robots is simulated in Stage v. 3.2.2 (Gerkey, Vaughan, & Howard, 2003). The graphical interface and robot control laws are implemented using the Robot Operating System (ROS) (Quigley et al., 2009), see Figure 1. The main map panel is on the left of Figure 1. The robots' positions within the environment are given by a gray circle with a line pointing in the direction of their heading. There are different colored targets that are hidden from the operator and randomly distributed around the environments. When a robot senses a target, it will turn that color in the interface so the operator knows the robot is reporting a target. The sensors exhibit false alarms and misses, so a target is only marked on the interface and considered found when a threshold of five robots sense the target simultaneously. This feature was implemented to represent the error inherent in real-world sensors due to environmental or hardware conditions. The necessary redundant coverage creates a secondary goal for the users to maintain a cohesive swarm to accomplish the task of marking targets. Once a target is marked, a square of that target's color will appear over the position of the target in the interface.

The robots' target sensors are faulty and miss at a rate of p_n in all conditions, where:

$$p_n = \left(1 - (r - d)\right)^{\alpha}$$

Here, d is the distance to the target, r is the maximum range of the sensor, and α is the decay rate set to 4.0 for these studies. The sensors also generated false positives, where an imaginary target was reported at a randomly chosen position at the edge of the target sensor's range. The false positive interval t_r is set to some randomly sampled value between 6 and 10 seconds for each of the 30 robots in the studies presented. The automated target marking system only marks a target when a threshold of five robots could sense the same target, using redundancy to eliminate the faulty sensor error.

Traversable areas are represented by white space and obstacles are represented by black, except for the hidden environments where the area of operation is all white and the operator does not see the obstacles. Operators can scroll and zoom the main map panel using a mouse as needed to complete the task. The keyboard is used in one environment in Study II to respond to a mathematical question, which provides the operator a distraction.

A Phantom Omni device is used as the input and haptic feedback output device as shown in the operator's left hand in Figure 2. The device's coordinate frame is centered on a circle on the desktop. The operator moves the end effector within this circle to create a persistent force vector F_h that is broadcast to the members of the swarm. On the screen, the input force panel shows F_h next to the main map panel so that operators do not have to look down at the input device to determine their influence, as shown on the right side of Figure 1. The force feedback panel, above the Input Force panel, shows the average force from the obstacle on the robots (average of F_o, described in the previous section). Note that for both the force feedback panel and the input force panel, the magnitude of the force is given by the length of the line within the circle. In Figure 1, the input force and the force feedback is identical, because it shows a scenario where the robots are being pushed against the obstacles and they cannot move in the direction of input force. The same force displayed in the Force Feedback panel is mapped to the haptic device and felt by the operators in a haptics condition. The input forces, robot positioning data, and marked targets were logged every second for data analysis purposes.

Experimental Design

These user studies have a between-subjects design. Each study divided the participants into different groups for the different conditions. Half of the participants in each study received physical force feedback in the Phantom device as described above, hereafter referred to as the haptics condition. The other half received only visual feedback, hereafter called the control condition. All other variables remained constant.

Participants were first given explanations of the robot control, the interface, the importance of a cohesive swarm due to the coverage redundancy required to mark targets. They were then given ten minutes of practice to gain experience and ask questions. In each condition, the environments were presented in random order.

Study I

The first study is considered a pilot to explore the utility of haptic feedback in HSI in searching targets with different features.

Study Parameters

This study had 20 participants, with 10 users in the control condition (visual feedback only) and 10 in the haptics condition (haptic and visual feedback). Preliminary experimentation showed that the following parameters were acceptable to accomplish the task: $k_o = 4$, $k_r = 2$, $k_h = 5$, $n_b = 2.0$, $n_e = 2.5$, $h_{max} = 4$. These gains make F_h stronger than F_o, which in turn is twice as strong as F_r. The swarm is able to maintain enough cohesion to mark targets and follow the operator's influence to explore the environments without collisions with the obstacles. Each robot's maximum speed is 0.5 m/s. The robot senses neighboring robots and obstacles with a 360 degree field of view and 4 degree resolution. Participants were given ten minutes to explore each environment described below to find as many targets as possible.

Environments

The participants attempted to find as many objects as possible in four different obstacle filled environments. The order of the environments are randomized and the starting position of the robots was in the bottom right corner. All environments contained 50 targets.

The first environment is a corridor maze called the Hallway environment, see Figure 4(A). The width of the halls was such that the robots could mark all targets in the corridor if the swarm traveled down the center. Drift of the swarm's centroid from the hallway center, however, might cause the participant to miss targets along the edges.

The second environment is a structured environment with rooms and narrow doorways between them, called the Doors environment, see Figure 4(B). The open rooms were intended to be simple to explore, while the doorways could cause problems, allowing only a few robots through at a time. These choke points cause the exploration to slow and possibly break robots away from the swarm, creating smaller groups less likely to mark targets.

The third environment is a corridor maze type environment with funnel points and traps, see Figure 4(C). The participants did not see the environments but had to navigate the robotic swarm using only the haptic feedback, the Force Feedback panel, and the behavior of the robots. The shape and hazard types of the environment were explained to the participants before each corresponding condition.

The final environment is structured with obstacles that have numerous edges and concave corners requiring exploration, see Figure 4(D). This environment explores the ability for the operator to traverse doorways and explore interior corners of various obstacle to find targets.

Hypotheses

The hypotheses are that the haptic condition will perform better than the control condition in each of the environments:

- Hallway environment (hypothesis A)
- Doors environment (B)
- Hidden Halls environment (C)
- Complex environment (D)

Support for hypothesis A could show that haptic feedback aided the operators in traversing the center of the hallway. Support for hypothesis B could show that haptic feedback helped guide the operators through the center of the doorways to allow more robots through at a time, increasing exploration efficiency and swarm cohesion. The Hidden Halls environment is difficult, but the haptic condition may receive intuitive obstacle force information through haptic feedback to allow for better environment understanding, hence hypothesis C. Support for hypothesis D could show that haptic feedback traverses doorways and complex obstacles more efficiently by inducing the best strategy.

Results

All ANOVA results where $p < .05$ are considered significant for this study. Any p-values between .05 and .15 gives potential trends indicating possible areas for future research to clarify. All mean comparisons between conditions are shown as box plots and the dots show outliers in the data as defined by Tukey tests, while relations between variables are shown with linear regression on scatter plots. Below, three performance metrics are presented: 1) number of targets found, 2) environment coverage percentage, and 3) swarm cohesion.

Targets Found

There are no significant benefits observed from haptic feedback in this performance metric. However, haptic feedback shows marginally significant improvements over the control condition in the Doors ($F = 3.25$, $p = .088$) and Complex ($F = 2.38$, $p = .140$) environments, see Figure 5. Therefore, Hypotheses B and D only have marginally statistically significant support and will be more closely examined in the second study.

Environment Coverage

It is possible that comparing marked targets does not best represent performance due to the effects of random distribution of the targets. Instead, the percentage of the environment covered by five robots simultaneously throughout the trial may better show the operator's ability to explore the environment given the task parameters. Results show significant improvement for haptics in the Complex environment ($F = 6.24$, $p = .022$), which supports hypothesis D. Furthermore, Figure 6 shows the correlation between this coverage percentage and the number of targets found ($p < .001$), indicating that the two measures may bear more similarity in longer or more in depth studies.

Swarm Cohesion

Swarm cohesion is also an important metric depending on the task. In this instance, operators are required to overcome faulty sensors and accomplish the goal of marking targets using redundant coverage. In other tasks, it might be better to break up the swarm to cover more area. In either case, it is always important to understand the effects on swarm cohesion, so the average number of connected components throughout trials is used to compare the user's ability to maintain a cohesive swarm.

The haptics condition shows significant increases in the number of connected components in the Complex environment ($F = 4.91$, $p = .040$). Hypothesis D, therefore, is not supported by investigating swarm cohesion because operators in the haptics condition broke up the swarm into more groups than the control condition.

Study II

The results of Study I showed promise, but the data did not support conclusions that haptic feedback would help users improve the operator's performance in all of environments. This study sought to outline the strengths of the haptics condition by overloading the operator's visual channel, which can loosely model situations in which swarm operators must deal with multiple streams of information at once, sometimes even for different tasks. The environments explained below were built to distract, increase workloads, and extend environmental features used in Study I.

Study Parameters

This study had 64 participants, so there were 16 users in the control condition with visual force feedback cues named the control-visual (**CV**) condition, 16 in the haptics condition with visual force feedback cues named the haptic-visual (**HV**) condition, 16 in the control condition without visual force feedback cues named the control-only (**CO**) condition, and 16 in the haptics condition without visual force feedback cues named the haptic-only (**HO**) condition. There were fifteen repeat participants from the first study who were randomly placed in a condition and environment order, without regards to their previous participation in Study I. Nine were placed in the HV (of which six were in the haptics condition in Study I) and six in the CV condition (of which one was in the control condition in Study I). The participants in Study II had a similar setup of hardware as in Study I, see Figure 2, except the Phantom was moved to be used by the right hand and the keyboard was moved within reach to respond to math problems in one of the environments. For this study, the values assigned to them were changed from Study I to increase the robot speed, which requires the operator to increase focus on the main map panel. The parameters were adjusted based on this speed increase to maintain similar controllability as in Study I: $k_o = 5$, $k_r = 3$, $k_h = 7$, $n_b = 2.5$, $n_e = 3$, h_{max}. The robots' max speed is 0.6 m/s. The robot senses other robots and obstacles with a 360 degree field of view and 4 degree resolution. All environments share these constraints except for the Speed environment, which is explained in the next section. Participants were given 15 minutes to explore each environment described below to find as many targets as possible.

Environments

The participants worked to find as many targets as possible in four different obstacle filled environments. The starting position of the robots was again in the bottom right corner, and all environments contained 60 targets.

The Math environment is a corridor maze, see Figure 7(A). The width of the halls was such that the robots could mark targets along both of the walls if the swarm traveled down the center, but there were choke points and traps which slowed the participant down. Each participant was instructed that the optimal strategy is to avoid the traps. Single digit addition problems blocked the main map panel randomly every ten to fifteen seconds until a correct answer was given, leaving only the side panel with the Force Feedback and Input vectors visible. The participant received math problems in the 10 minute training period as well, at a random interval between 30 and 60 seconds, in order to get familiar with the math interface. The robots continue to receive input and the interface still gives both visual and haptic representations of the force feedback, if the condition allows, to the operator while the math problem is blocking the main map panel. This condition corresponds to a scenario where navigation is a secondary task and an infrequent primary task requires full visual attention (i.e. checking video surveillance).

The Speed environment and parameters are structured to increase focus on the main map panel, see Figure 7(B). The dead ends and intersections require focus and decision-making for operators to determine the best path to explore as much area as possible in the allotted time. The distinctive feature of this environment is greater speed to increase the operator's focus on the main map panel, which required changes in the parameters to avoid obstacle collisions. The speed was set to 1.0 m/s and the force gains were set to $k_o = 10$, $k_r = 4$, $k_h = 5$. All other constants remained the same. This increase in speed created a volatile swarm requiring more focus which was more likely to break up around obstacles. The participants were warned of this.

The Doors2 environment is the Doors environment from Study I rotated. This will help compare the results of the new swarm parameters to those found in Study I across all three metrics. The participants were informed that the best strategy was to explore each room thoroughly before moving through the doorway to the next room because the swarm took more time when traversing the doorways.

The Hidden Complex environment is structured with obstacles that have numerous edges and concave corners requiring exploration, see Figure 7(D). The obstacles in this environment are hidden from the interface so the operators had to blindly explore the area using the swarm's behavior, side panel information (if given visual force feedback cues), and haptic feedback (if in a haptic condition). The participants were told it was an office structure with obstacles in the room, and that the best strategy was to sweep the rooms avoiding a lot of force from the walls and then finding the exit, similar to the strategy in the Doors2 environment. They were also instructed to use marked targets as landmarks when doorways were thought to be found in case the new room had no other exits.

Hypotheses

To discuss the hypothesis in this study with the hypotheses in Study I, hypotheses from Study I will have a 1 followed by the character (i.e. 1A, 1B, etc) and these hypothesis will have a 2 followed by the character (i.e. 2A, 2B, etc). It is expected that haptic feedback (conditions HV and HO) will improve performance over the corresponding conditions without haptic feedback (conditions CV and CO) respectively in all environments:

- Math environment (hypothesis 2A)
- Speed environment (2B)
- Doors2 environment (2C)
- Hidden Complex environment (2D)

Support for hypothesis 2A will show that operator's with haptic feedback have better control when distracted with math problems. The operators with haptic feedback may also have better control and keep the robots away from obstacle forces that might break up the volatile swarm in the Speed environment, hence hypothesis 2B. Because there was support for hypothesis 1B, hypothesis 2C should hold since the Doors2 environment is similar to the Doors environment. Support for hypothesis 2D would show that haptic feedback can help operators interpret the hidden obstacles in the Hidden Complex environment and explore more of the environment.

This study further explores the need of the Force Feedback panel displayed visually. For hypothesis 2E, it is expected that the performance between conditions HV and HO will not be different throughout the environments. Support of this hypothesis would show that the visual force cues do not increase the operator's performance and if hypotheses 2A - 2D hold, then haptic feedback is the only contributing factor. The CO condition should obviously perform worse in all environments because of the lack of information the user receives (hypothesis 2F).

Results

The greater number of participants in each condition allows the results to have increased statistical significance, so only ANOVA results where $p < .010$ are considered, where $p < .005$ is considered significant and $p < .010$ marginally significant, and an avenue for future investigation. Below, three metrics are presented: 1) number of targets found, 2) environment coverage percentage, and 3) swarm cohesion.

Targets Found

Preliminary analysis of the number of targets marked shows marginally significant improvement of condition HV over conditions HO ($F = 3.32$, $p = .079$) and CO ($F = 3.61$, $p = .067$) in the Doors2 environment and condition CO over condition HV ($F = 3.31$, $p = .079$) in the Hidden Complex environment (see Figure 8). Hypothesis 2C is supported by this result, but hypotheses 2D and 2E are not supported.

Environment Coverage

Figure 9 shows the correlation between map coverage of five robots simultaneously and the number of targets found ($p < .001$). This confirms that there is a natural correlation between the coverage metric and the task, but this metric does not include the effects of the random distribution of targets. Table 2 shows the hypothesis testing, for environments with significant results, of the percentage of environment coverage by at least 5 robots simultaneously comparing all conditions against each other over all environments. As confirmed by the results of the first study, and Figure 8, the hidden environments are very difficult and create a different set of challenges for the operators. If the Hidden Complex environment is removed and all conditions compared, a significant improvement in performance is evident for condition HV over the two control conditions ($p < .001$ for CV, $p < .005$ for CO). This supports hypotheses 2A-2C and 2F since haptics can improve the operator's ability to effectively explore different environments.

The results do not support hypothesis 2E, however, as operators in condition HV shows significant improvement of environmental coverage when compared to condition HO ($p < .002$ for HO). Hypotheses 2C and 2F are supported while hypothesis 2E is not supported, where condition HV shows significant improvement over conditions HO and CO and condition HV shows a marginally significant improvement over condition CV. The Speed environment shows marginal improvement of condition HV over CV, supporting hypothesis 2B.

Swarm Cohesion

The average number of connected components in each trial for all conditions is used as a measure of swarm cohesion, see Table 3. Figure 10 shows a significant decrease in connected components of conditions HV and CV and condition CO and a marginally significant decrease between conditions CV and HO. This result supports hypotheses 2A-2D and 2F, but does not support hypothesis 2E. Similar results are shown when the environments are split up where conditions HV and CV have significantly fewer connected components than conditions CO and HO for the Math environment, see Figure 11. The Hidden Complex environment also shows a significant decrease from conditions HV and CV and condition CO as well as marginal decreases from conditions HV and CV and condition HO.

DISCUSSION

The results support the hypotheses that haptic feedback can increase performance of HSI with the environment exploration task. Operators were better able to explore a variety of environments when haptic feedback was present than when it was not. Operators could traverse doorways and bottlenecks better and explore more area and mark more targets. Haptic feedback also allowed operators better performance when distracted and under stress of volatile swarms. Finally, haptic feedback allowed operators to find more targets and cover more area in most environments.

Results from Study II did not, however, support the hypothesis that showing the Force Feedback panel was unnecessary for the operators who felt haptic feedback. The operators in the haptics-visual condition outperformed operators in the haptic-only condition showing the necessity for the Force Feedback panel. While haptic feedback can help operators explore environments better, it seems to only do so when there is a visual representation of that feedback. This visual representation could help aid the operator in interpreting the force they feel through the input device. Without the visual representation of the force, operators do not mark as many targets or explore as much of the environments. Evidence even exists that the visual force feedback may be even more important, since operators in the control-visual condition found more targets than the haptics-only condition in the Hidden Complex environment.

The results of Study I conflict with findings from Study II concerning swarm cohesion. Study I found evidence that haptic feedback would split up the swarm more so than the control condition in the Complex environment, even though operators in the haptic condition covered more area in this environment. The best strategy was not as described in the directions, where a more cohesive swarm will cover more area and, therefore, mark more targets, but the haptic condition led operators to a better strategy by splitting up the swarm around the complex obstacles in the environment. Study II found that visual force feedback, with or without haptic feedback, could help operators maintain a more cohesive swarm. Unlike in Study I, this strategy often lead to better performance since the haptics-visual condition explored more area than any other condition in all environments except for the Hidden Complex environment.

CONCLUSIONS AND FUTURE WORK

The overall goal of the above studies was to explore the effectiveness of haptic feedback in the presence of other visual stimuli about the swarm system. It was found that providing haptic feedback of obstacles to the operator, with accompanying visual force information, was the most beneficial for the performance of the system. Operators with haptic feedback repeatedly outperformed operators without haptic feedback in numerous environments by discovering targets in unexplored areas of the environment. Some assumptions must be overcome (i.e. perfect localization of robots in the environment), but the results show that operators can quickly understand the control structure and convey their intent to accomplish their goal. The hidden environments seem to create difficulty for the haptic feedback, however. No evidence supports conclusions that haptic feedback would hinder the operator from exploring more of the environment if the placement of the obstacles are not known, as was the case in the hidden environments.

These findings show the importance for continuing work with integrating haptic feedback to HSI. There are many other control parameters to consider for the operators, such as changing the radii of the attraction and repelling zones of the robots so that operators can increase or decrease their coverage. This would allow operators to expand the swarm in instances where larger coverage is needed, or control the swarm in instances involving tight spaces. The results also require further exploration of environments where the structural layout is not known, because the operator does not have previous information about the environment.

Future work in using haptics in HSI should explore the utility of removing erroneous robot positions to see if haptic feedback can provide information to improve performance. This should be helpful because robots commonly have localization errors. Therefore, an experimental setup where only the static map of the environment is shown in the main map panel (see the left side of Figure 1) and starting area is known and the robots' (wrong) positions are not shown so operators must "feel" their way around the environment to discover the number of targets in different rooms. This could relate to a real-world environment where robots can explore all rooms blindly and the operator will have to direct resources to rooms of more importance afterwards. Finally, haptic feedback techniques should be explored for different tasks (i.e. a swarm collectively moving a large object). The mission given to the operators in these studies lends itself to obstacle avoidance forces, but if the robots should be pushing an object, they would require a different control algorithm. A list of tasks must be examined to determine when haptic feedback techniques will need to be different or if they can even be helpful at all. Overall, the work presented in this chapter should demonstrate that the integration of haptic devices in HSI systems is promising, and further research should be directed toward this problem.

ACKNOWLEDGEMENTS

Some results presented in this article were also reported in Nunnally, Chakraborty, Walker, Lewis, & Sycara (2013a, 2013b). This work was funded by ONR Science of Autonomy Grant N0001409-10680.

REFERENCES

Bachrach, J., & Beal, J. (2006). Programming a sensor network as an amorphous medium. *Distributed computing in sensor systems.*

Bachrach, J., Beal, J., & McLurkin, J. (2010). Composable continuous-space programs for robotic swarms. *Neural Computing and Applications*, 19(6), 825–847.

Bachrach, J., McLurkin, J., & Grue, A. (2008). Protoswarm: a language for programming multi-robot systems using the amorphous medium abstraction. In *Proceedings of the 7th international joint conference on autonomous agents and multiagent systems* (pp. 1175–1178).

Bashyal, S., & Venayagamoorthy, G. (2008). Human swarm interaction for radiation source search and localization. *Swarm intelligence symposium* (pp. 1–8).

Blackwell, T. M., & Bentley, P. (2002). Improvised music with swarms. In *Computational intelligence, proceedings of the world on congress on* (Vol. 2, pp. 1462–1467).

Bruemmer, D. (2002). A robotic swarm for spill finding and perimeter formation (Tech. Rep.). *DTIC Document*.

Bullo, F., Cortés, J., & Martínez, S. (2009). *Distributed control of robotic networks*. Princeton University Press.

Couzin, I. D., Krause, J., Franks, N. R., & Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025), 513–516.

Couzin, I. D., Krause, J., James, R., Ruxton, G. D., & Franks, N. R. (2002). Collective memory and spatial sorting in animal groups. In *Journal of theoretical biology*, 218(1), 1–11.

Dorigo, M., Trianni, V., Şahin, E., Groß, R., Labella, T. H., Baldassarre, G., ... & Gambardella, L. M. (2004). Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, *17*(2-3), 223-245.

Fields, M., Haas, E., Hill, S., Stachowiak, C., & Barnes, L. (2009). Effective robot team control methodologies for battlefield applications. In *International conference on intelligent robots and systems* (pp. 5862–5867).

Franchi, A., Giordano, P., Secchi, C., Son, H. I., & Bülthoff, H. (2011). A passivity-based decentralized approach for the bilateral teleoperation of a group of UAVs with switching topology. In *International conference on robotics and automation* (p. 898-905).

Franchi, A., Masone, C., Grabe, V., Ryll, M., Bülthoff, H. H., & Giordano, P. R. (2012). Modeling and control of UAV bearing-formations with bilateral high-level steering. In *International Journal of Robotics Research: Special Issue on 3D Exploration, Mapping, and Surveillance*, 31(12), 1504–1525.

Franchi, A., Secchi, C., Ryll, M., Bülthoff, H. H., & Giordano, P. R. (2012). Shared control: Balancing autonomy and human assistance with a group of quadrotor UAVs. In *IEEE Robotics and Automation Magazine, Special Issue on Aerial Robotics and the Quadrotor Platform*, 19(3), 57–68.

Franchi, A., Secchi, C., Son, H. I., Bülthoff, H. H., & Giordano, P. R. (2012). Bilateral teleoperation of groups of mobile robots with time-varying topology. In *Robotics, IEEE Transactions on*, 28(5), 1019–1033.

Gerkey, B., Vaughan, R., & Howard, A. (2003). The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proceedings of the 11th international conference on advanced robotics* (pp. 317–323).

Giordano, P., Franchi, A., Secchi, C., & Bülthoff, H. (2011). Experiments of passivity-based bilateral aerial teleoperation of a group of UAVs with decentralized velocity synchronization. In *IEEE/RSJ international conference on intelligent robots and systems* (pp. 163–170).

Giordano, P. R., Franchi, A., Secchi, C., & Bülthoff, H. H. (2013). A passivity-based decentralized strategy for generalized connectivity maintenance. In *The International Journal of Robotics Research*, 32(3), 299–323.

Goodrich, M., Pendleton, B., Sujit, P., & Pinto, J. (2011). Toward human interaction with bio-inspired robot teams. In *International conference on systems, man, and cybernetics* (pp. 2859–2864).

Howard, A., Matarić, M., & Sukhatme, G. (2002). An incremental self-deployment algorithm for mobile sensor networks. In *Autonomous Robots*, 13(2), 113–126.

Kira, Z., & Potter, M. (2009). Exerting human control over decentralized robot swarms. In *4th international conference on autonomous robots and agents* (pp. 566–571).
Klarer, P. (1998). Flocking small smart machines: An experiment in cooperative, multi-machine control (*Tech. Rep*.). Sandia National Labs, Albuquerque, NM (United States).
Kolling, A., Nunnally, S., & Lewis, M. (2012). Towards human control of robot swarms. In *Proceedings of the 7th international conference on human-robot interaction* (pp. 89–96).
LaValle, S. (2006). *Planning algorithms*. Cambridge University Press.
Lee, D., Franchi, A., Giordano, P., Son, H., & Bülthoff, H. (2011). Haptic teleoperation of multiple unmanned aerial vehicles over the internet. In *International conference on robotics and automation* (pp. 1341–1347).
Lee, D., Franchi, A., Son, H. I., Bülthoff, H. H., & Giordano, P. R. (2013). Semi-autonomous haptic teleoperation architecture of multiple unmanned aerial vehicles. In *IEEE/ASME Transactions on Mechatronics: Focused Section on Aerospace Mechatronics*, 18(4), 1334–1345.
Naghsh, A., Gancet, J., Tanoto, A., & Roast, C. (2008). Analysis and design of human-robot swarm interaction in firefighting. In *The international symposium on robot and human interactive communication* (pp. 255–260).
Nunnally, S., Walker, P., Chakraborty, N., Lewis, M., & Sycara, K. (2013, October). Using Coverage for Measuring the Effect of Haptic Feedback in Human Robotic Swarm Interaction. In *Systems, Man, and Cybernetics, 2013 IEEE International Conference on* (pp. 516-521).
S., Chakraborty, N., Walker, P., Lewis, M., & Sycara, K. (2013b). Using haptics in human robotic swarm interaction. In *International annual meeting of the human factors and ergonomics society*.
Nunnally, S., Walker, P., Kolling, A., Chakraborty, N., Lewis, M., Sycara, K., & Goodrich, M. (2012, October). Human influence of robotic swarms with bandwidth and localization issues. In *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on* (pp. 333-338).
Olfati-Saber, R. (2006). Flocking for multi-agent dynamic systems: Algorithms and theory. In *IEEE Transactions on Automatic Control*, 51(3), 401–420.
Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., Ng, A. (2009). ROS: an open-source robot operating system. In *ICRA workshop on open source software* (Vol. 3).
Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioral model. In *ACM SIGGRAPH computer graphics* (Vol. 21, pp. 25–34).
Rodríguez-Seda, E., Troy, J., Erignac, C., Murray, P., Stipanovic, D., & Spong, M. (2010). Bilateral teleoperation of multiple mobile agents: Coordinated motion and collision avoidance. In *Transactions on Control Systems Technology*, 18(4), 984-992.
Secchi, C., Franchi, A., Bulthoff, H. H., & Giordano, P. R. (2012, May). Bilateral teleoperation of a group of UAVs with communication delays and switching topology. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on* (pp. 4307-4314). IEEE.
Son, H., Chuang, L., Franchi, A., Kim, J., Lee, D., Lee, S., Bülthoff, H., Robuffo Giordano, P. (2011). Measuring an operator's maneuverability performance in the haptic teleoperation of multiple robots. In *IEEE/RSJ international conference on intelligent robots and systems* (pp. 3039–3046).

Son, H. I., Franchi, A., Chuang, L. L., Kim, J., Bülthoff, H. H., & Giordano, P. R. (2013). Human-centered design and evaluation of haptic cueing for teleoperation of multiple mobile robots. In *Cybernetics, IEEE Transactions on*, 43(2), 597–609.
Spears, W., & Spears, D. (2012). Physicomimetics: *Physics-based swarm intelligence*. Springer-Verlag New York Inc.
Trianni, V., Groß, R., Labella, T. H., Şahin, E., & Dorigo, M. (2003). Evolving aggregation behaviors in a swarm of robots. In *Advances in artificial life* (pp. 865-874). Springer Berlin Heidelberg.
Walker, P., Nunnally, S., Lewis, M., Kolling, A., Chakraborty, N., & Sycara, K. (2012, October). Neglect benevolence in human control of swarms in the presence of latency. In *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on* (pp. 3009-3014).

KEY TERMS AND DEFINITIONS

Haptic Feedback: Information fed back from a system to a user that is tactile in nature.

Human-Swarm Interaction: A system including a human operator in control of a semi-autonomous swarm. In such a system, the human uses the information returned by the swarm about the environment and swarm itself to give new commands and inputs to the swarm.

Flocking: A swarm behavior whereby all agents align their velocities and headings while maintaining some minimum distance from one another.

Performance Metric: A quantifiable measurement pertaining the swarm's operation that can be used to investigate the effects of some independent variable in the swarm system.

Predictive Display: A graphical user interface that displays a prediction about a swarm's future state, in addition to current information.

Summary Display: A graphical user interface displaying summary information about the swarm state, such as a bounding centroid or average heading.

Swarm: A robust, scalable system composed of numerous robots, which coordinate through local interaction only.

FIGURES AND TABLES

Table 1A

	Participants	
	Study I	Study II
Number of participants	20	64
Number of conditions	2	4
Participants per condition	10	16
Number of environments	4	4

Minutes per environment	10	15
Targets in environments	50	60
Number of robots	30	30

Table 1B

	Conditions					
	Study I		Study II			
	Haptics	Control	HV	CV	HO	CO
Haptics	yes	no	yes	no	yes	no
Visual	yes	yes	yes	yes	no	no

Table 1C

	Environments							
	Study I				Study II			
	A	B	C	D	A	B	C	D
Environment type	halls	office	halls	office	halls	halls	office	office
Clutter	no	no	yes	yes	yes	no	no	yes
Hidden obstacles	no	no	yes	no	no	no	no	yes
Distraction	no	no	no	no	yes	no	no	no

Table 1: A comparison of the studies. The top portion shows general differences between the two studies, the middle portion shows differences between the conditions, and the bottom portion shows differences between the environments. Here, HV and HO are the conditions where haptic feedback from obstacles is given either with or without corresponding visual feedback, respectively. CV and CO represent the conditions where operators received no haptic feedback, with either with or without visual feedback from obstacles. Note that participants always received visual feedback of the positions and movement of the swarm in the main viewport. The conditions and environments for each study are described further in section 7 for Study I and section 8 for Study II.

Results for Environment Coverage

	Over All Environments		
	CV	HO	CO
HV	3.92 (**.050**)	3.49 (.064)	1.35 (.247)
CO	0.72 (.399)	0.52 (.471)	
HO	0.02 (.896)		

	All Environments Except Hidden Complex		
	CV	HO	CO
HV	7.51 (**.007**)	6.24 (**.014**)	4.32 (**.040**)
CO	0.43 (.513)	0.08 (.777)	
HO	0.17 (.685)		

	Doors2 Environment		
	CV	HO	CO
HV	3.31 (.079)	4.71 (**.038**)	4.75 (**.037**)

CO	0.00 (.987)	0.00 (.959)	
HO	0.00 (.955)		

Table 2: Hypothesis testing on the percentage of the environment covered by 5 robots simultaneously comparing all conditions against each other. Only results for environments with significant results are shown. The format of each cell is the F-value followed by the p-value.

Results for Swarm Cohesion

	Over All Environments		
	CV	HO	CO
HV	0.03 (.873)	2.62 (.108)	4.23 (**.042**)
CO	4.81 (**.030**)	0.23 (.635)	
HO	3.11 (.080)		

	Math Environment		
	CV	HO	CO
HV	0.03 (.869)	6.68 (**.015**)	13.32 (**<.001**)
CO	9.22 (**.005**)	0.06 (.813)	
HO	5.04 (**.032**)		

	Hidden Complex Environment		
	CV	HO	CO
HV	0.01 (.918)	3.37 (.077)	5.05 (**.032**)
CO	5.26 (**.029**)	0.35 (.557)	
HO	3.58 (.068)		

Table 3: Hypothesis testing on the average number of connected components comparing all conditions against each other. Only environments with significant differences are shown. The format of each cell is the F-value followed by the p-value.

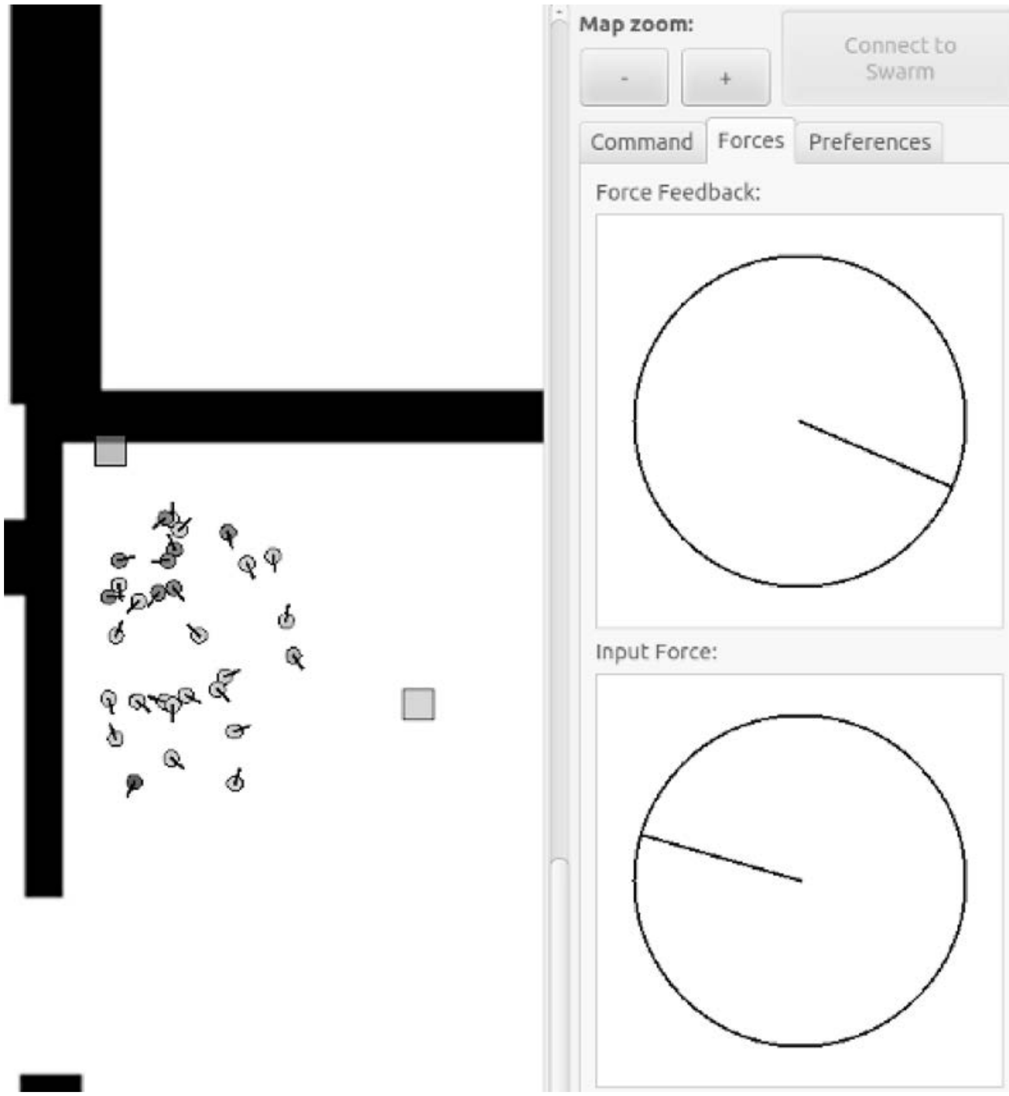

Figure 1. The GUI used for every condition of the study. The left side shows the robots' estimated positions, obstacles, and marked targets. The right side shows the Force Feedback panel calculated in equation 2, and the Input Force panel calculated in equation 4.

Figure 2. Participants used an Omni Phantom device (left) and mouse to influence the swarm and control the interface

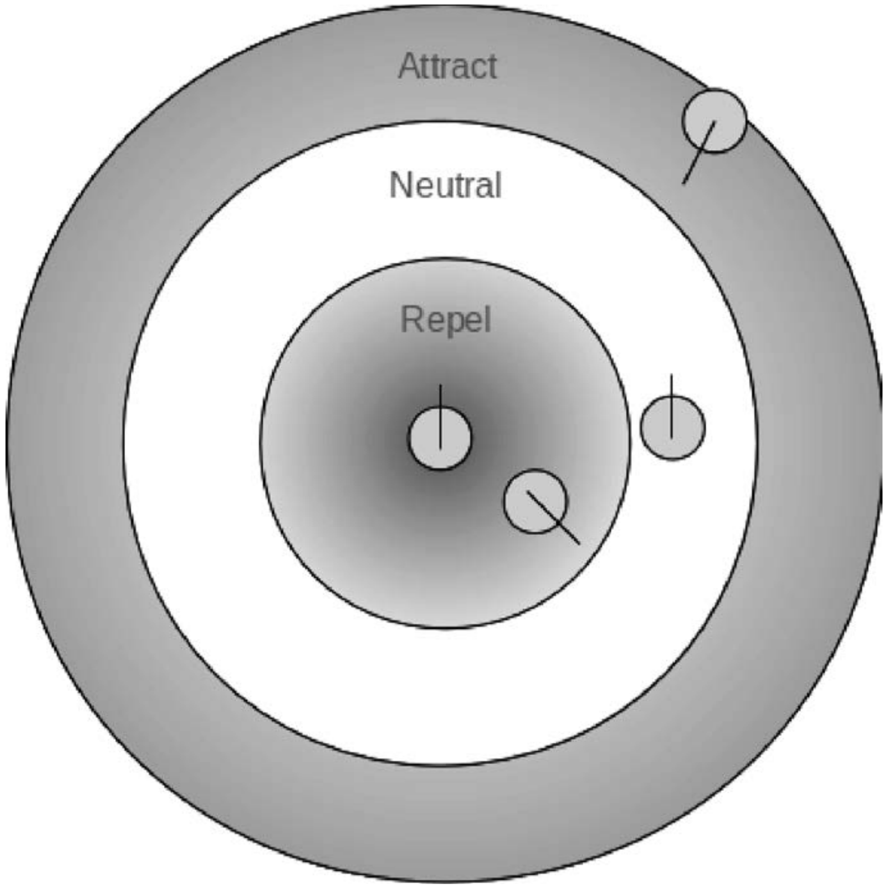

Figure 3. Figure visualizing the different zones for the robot control algorithm. The closer the robots are, the stronger the repulsive force, the further the robots are, the strong the attractive force. The robots will naturally stabilize to the neutral zone without F_o and F_h, the forces given by obstacles and human input, respectively.

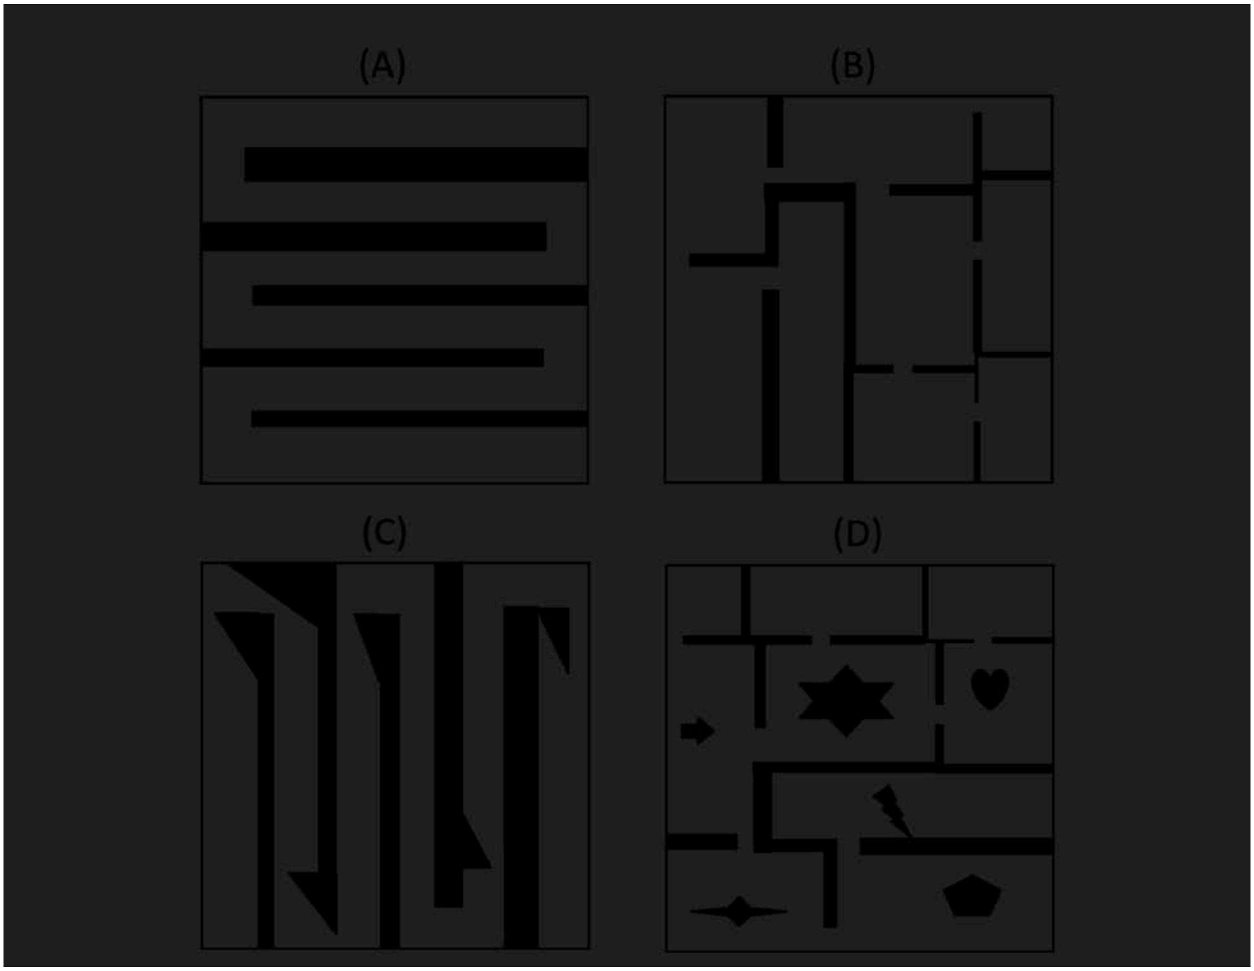

Figure 4. *Four environments used in the study: (A) Hallway, (B) Doors, (C) Hidden Halls, and (D) Complex. The robots always began the condition in the bottom right corner.*

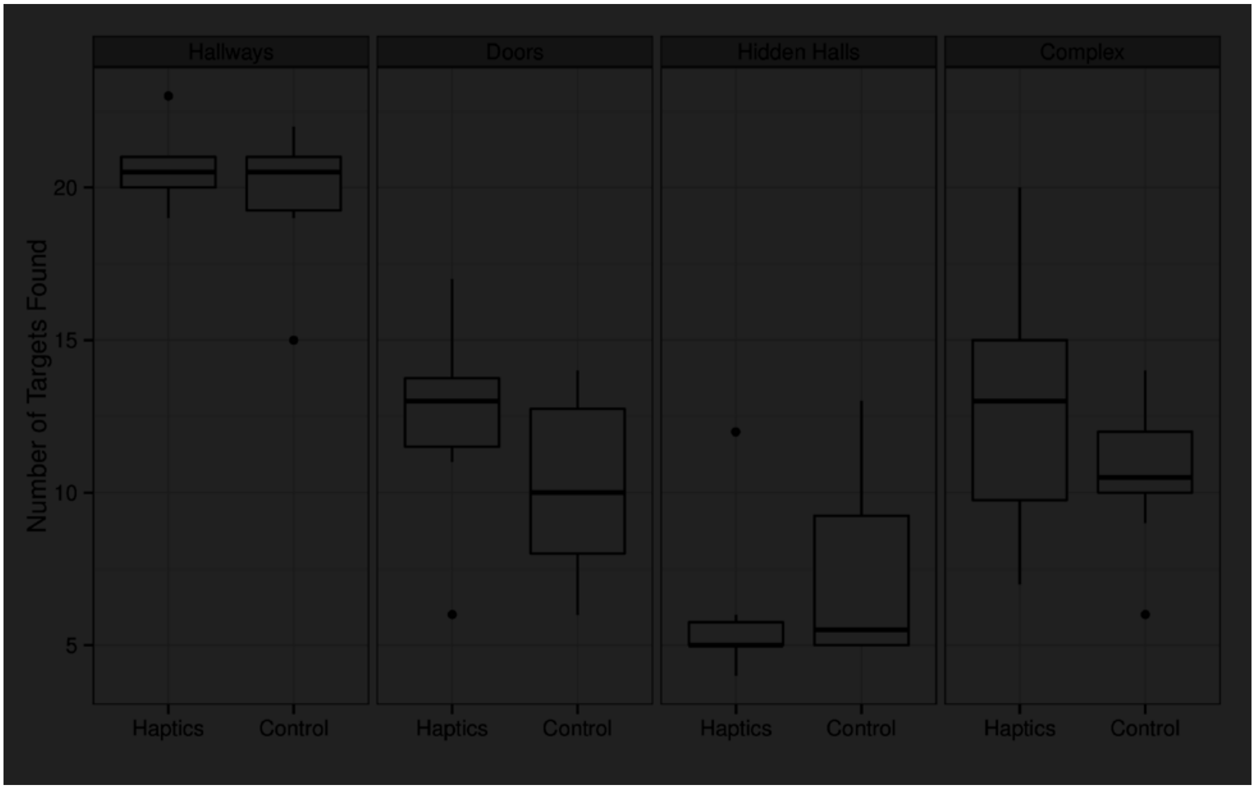

Figure 5. A box plot around the median number of targets found in each environment between each condition. Each box represents 10 trials. Dots outside the box plot represent outliers.

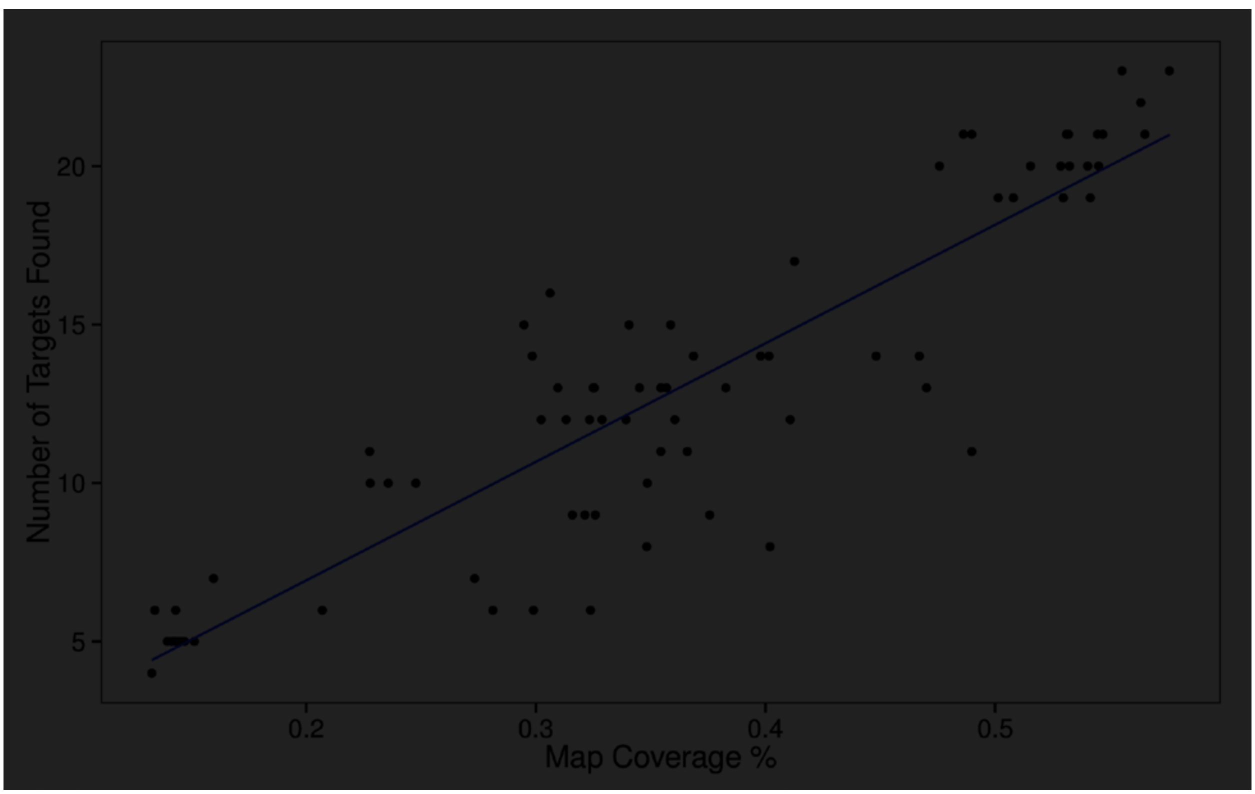

Figure 6. The relationship between the number of targets found and the percentage of the traversable area of the map covered by at least five robots simultaneously in Study I.

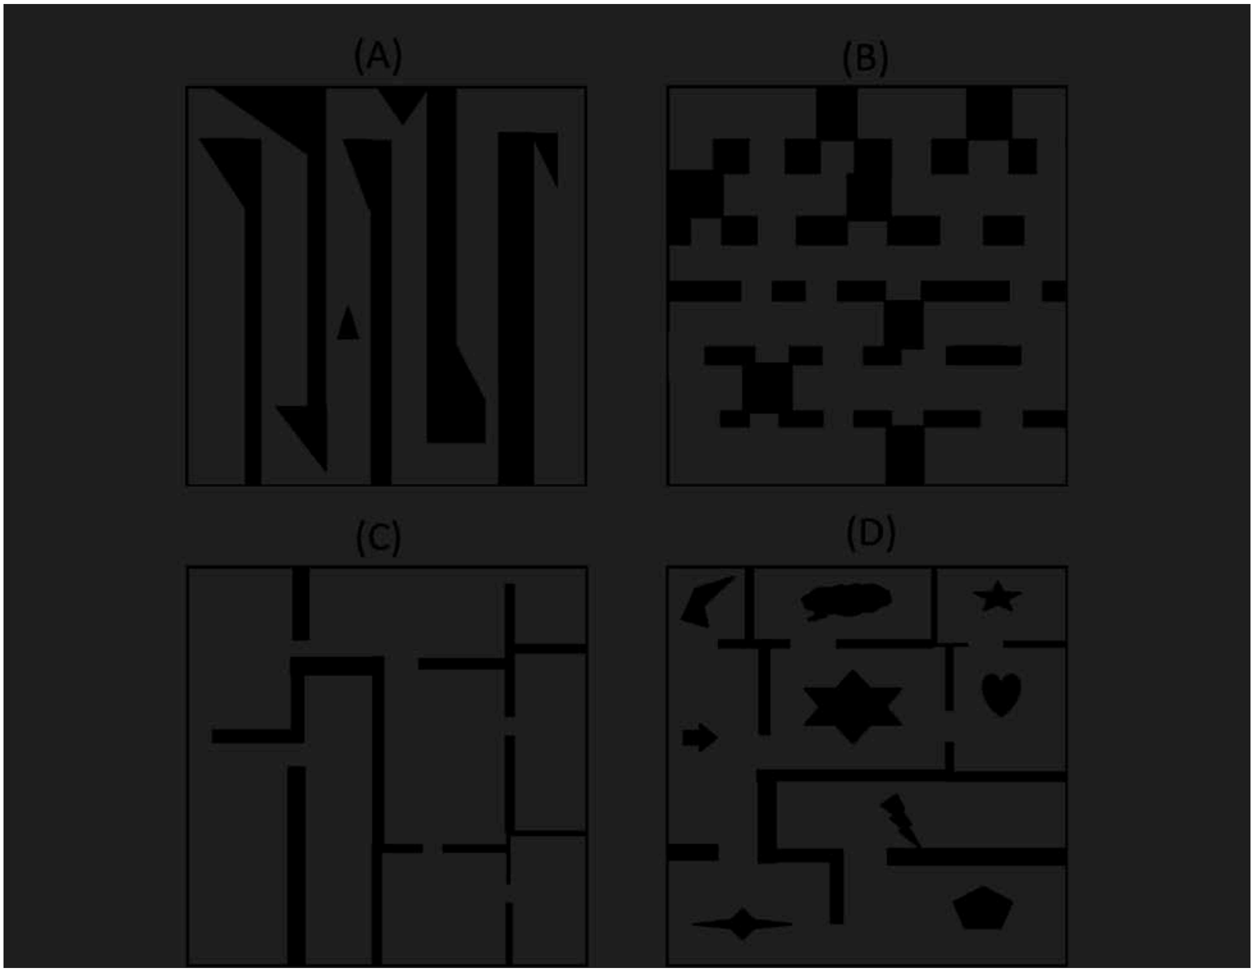

Figure 7. Four environments used in the study, the robots always started in the lower right corner: (A) Math, (B) Speed, (C) Doors2, and (D) Hidden Complex

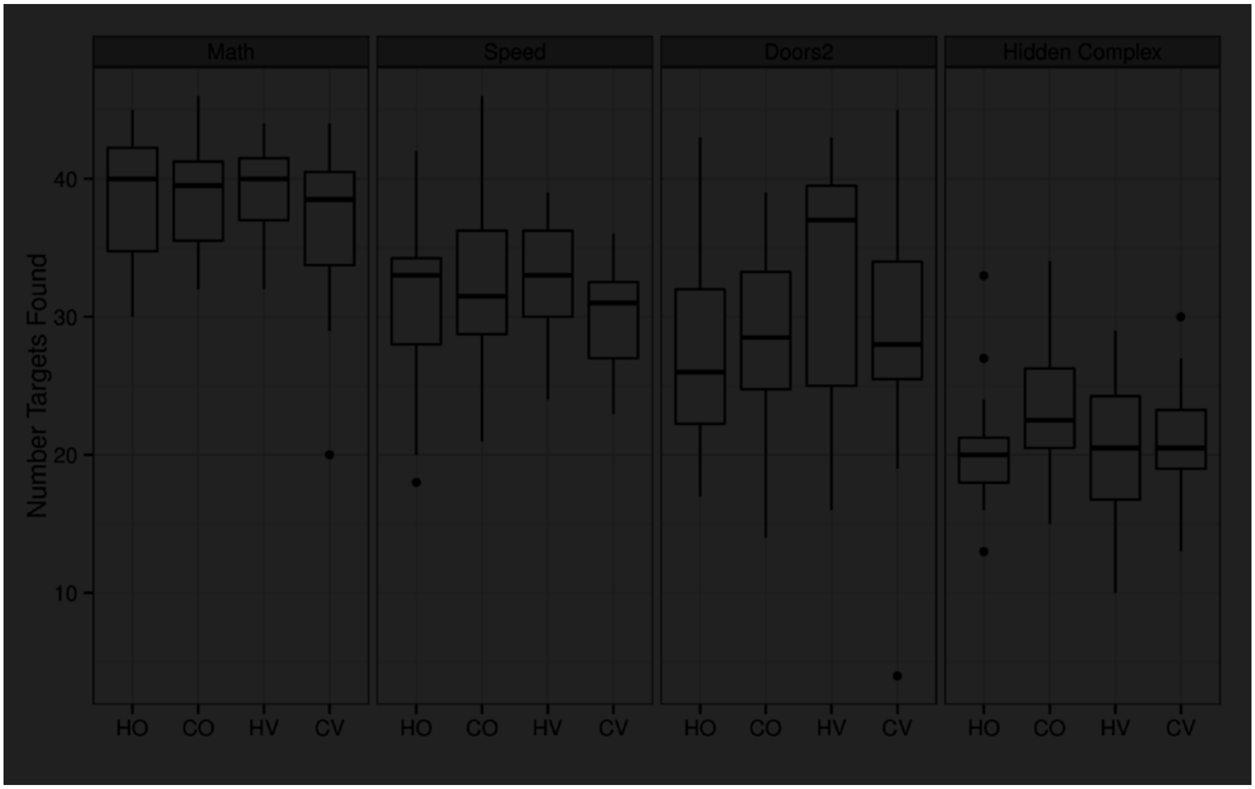

Figure 8. A box plot around the median number of targets found in each environment between each condition (HV is haptics-visual, CV is control-visual, HO is haptics-only, and CO is control-only). Each box represents 16 trials.

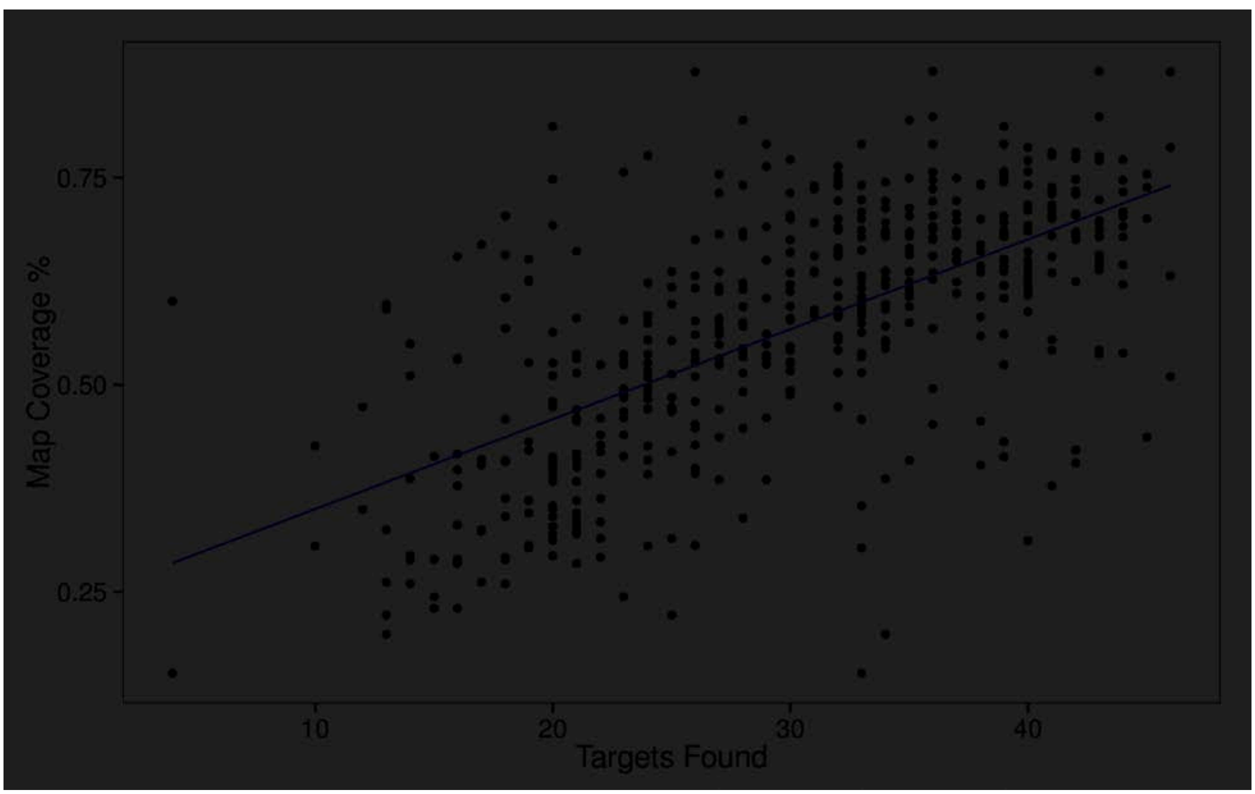

Figure 9. The relationship between the number of targets found and the percentage of the traversable area of the map covered by at least five robots simultaneously in Study II

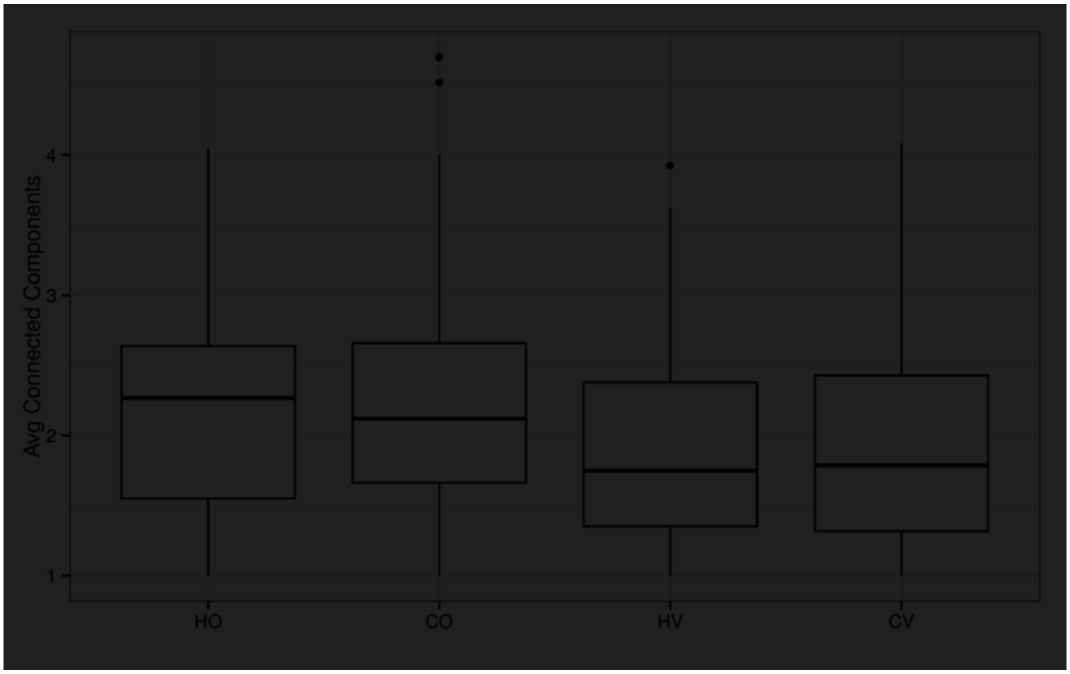

Figure 10. Box plots around the median average connected components for all environments between conditions (HV is haptics-visual, CV is control-visual, HO is haptics-only, and CO is control-only). Each box represents 64 trials.

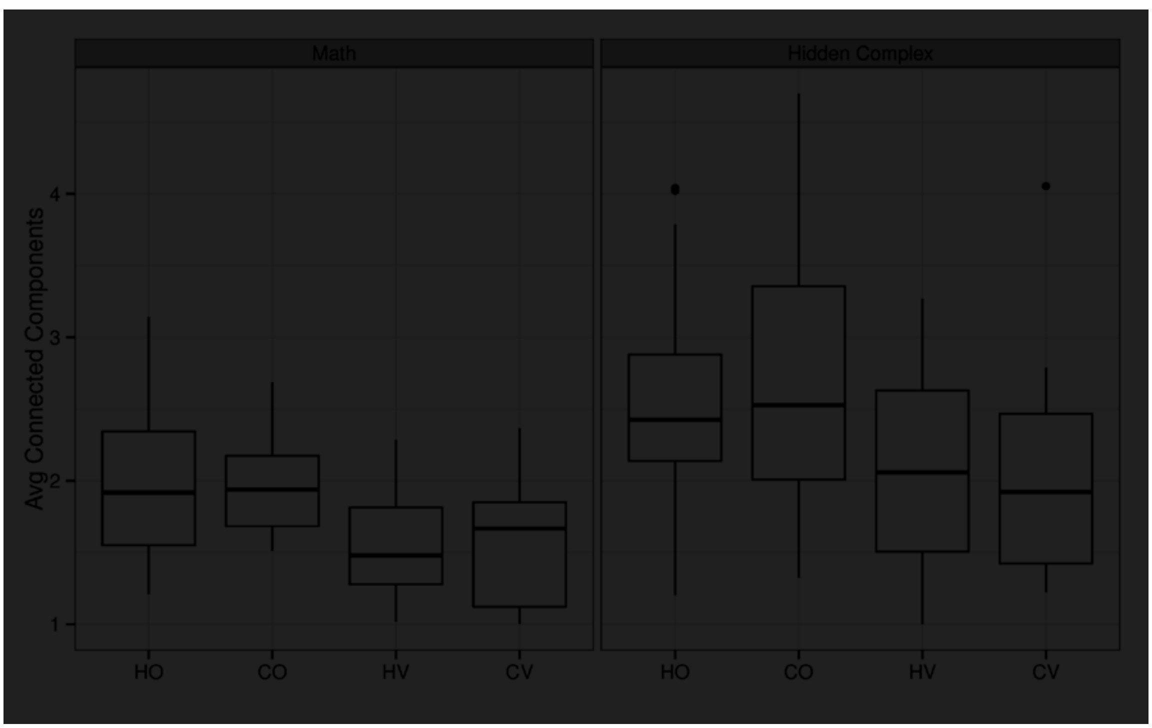

Figure 11. *Box plots around the median average connected components for the Math and Hidden Complex environments between conditions (HV is haptics-visual, CV is control-visual, HO is haptics-only, and CO is control-only). Each box represents 16 trials.*