Human swarm interaction using plays, audibles, and a virtual spokesperson

Chaffey, Patricia, Artstein, Ron, Georgila, Kallirroi, Pollard, Kimberly, Nasihati Gilani, Setareh, et al.

Patricia Chaffey, Ron Artstein, Kallirroi Georgila, Kimberly A. Pollard, Setareh Nasihati Gilani, David M. Krum, David Nelson, Kevin Huynh, Alesia Gainer, Seyed Hossein Alavi, Rhys Yahata, Anton Leuski, Volodymyr Yanov, David Traum, "Human swarm interaction using plays, audibles, and a virtual spokesperson," Proc. SPIE 11413, Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II, 114130V (22 April 2020); doi: 10.1117/12.2557573

SPIE.

Event: SPIE Defense + Commercial Sensing, 2020, Online Only, California, United States

Human swarm interaction using plays, audibles, and a virtual spokesperson

Patricia Chaffey[a], Ron Artstein[a], Kallirroi Georgila[a], Kimberly A. Pollard[b], Setareh Nasihati Gilani[a], David M. Krum[a], David Nelson[a], Kevin Huynh[a], Alesia Gainer[a], Seyed Hossein Alavi[a], Rhys Yahata[a], Anton Leuski[a], Volodymyr Yanov[a], and David Traum[a]

[a]Institute for Creative Technologies, University of Southern California, Los Angeles, CA, 90094, USA

[b]U.S. Army Research Laboratory, Los Angeles, CA, 90094, USA

ABSTRACT

This study explores two hypotheses about human-agent teaming: 1. Real-time coordination among a large set of autonomous robots can be achieved using predefined "plays" which define how to execute a task, and "audibles" which modify the play on the fly; 2. A spokesperson agent can serve as a representative for a group of robots, relaying information between the robots and human teammates. These hypotheses are tested in a simulated game environment: a human participant leads a search-and-rescue operation to evacuate a town threatened by an approaching wildfire, with the object of saving as many lives as possible. The participant communicates verbally with a virtual agent controlling a team of ten aerial robots and one ground vehicle, while observing a live map display with real-time location of the fire and identified survivors. Since full automation is not currently possible, two human controllers control the agent's speech and actions, and input parameters to the robots, which then operate autonomously until the parameters are changed. Designated plays include monitoring the spread of fire, searching for survivors, broadcasting warnings, guiding residents to safety, and sending the rescue vehicle. A successful evacuation of all the residents requires personal intervention in some cases (e.g., stubborn residents) while delegating other responsibilities to the spokesperson agent and robots, all in a rapidly changing scene. The study records the participants' verbal and nonverbal behavior in order to identify strategies people use when communicating with robotic swarms, and to collect data for eventual automation.

1. INTRODUCTION

Robotic assets and other autonomous systems are becoming increasingly important in a variety of industry, civilian, and military contexts. As these systems become more complex and capable, their use is dynamically evolving from one where a single human controls a single robot to one where multiple robots (or swarms) are controlled by a single human. Human-multirobot systems have been used in applications such as terrain mapping,[1] transport,[2] construction,[3] and search and rescue operations.[4] The use of multiple robots can allow for benefits such as greater surveillance ability, additional physical power available for work, more robustness in the overall system, and an opportunity to conduct diverse and coordinated activities. A significant challenge in implementing swarms of robots is controlling the group in a manner that is efficient for one human to oversee:[5] with an increasing number of robots comes additional strain on the human operator to monitor, instruct, and communicate with the robots. This situation can lead to increased stress, workload, and fatigue for a human operator,[6,7] which in turn may lead to operator errors[4,8,9] or loss of situation awareness.[10] Human factors can place limits on the number of robots that can successfully be controlled.[11]

The workload demands placed on the human operator vary by how they control or communicate with the robotic assets. While human involvement can increase swarm performance in complex and cluttered environments,[12] the human can also be a bottleneck to performance if they must control robots individually or approve all decisions.[4,13] Controlling or communicating with each robot on an individual level can result in high workload, especially if the operator must also coordinate the interplay of the robots.[14] In contrast, workload can be reduced if the operator can use high-level commands[14] or interact via a swarm "spokesperson".[15]

Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II, edited by Tien Pham, Latasha Solomon, Katie Rainey, Proc. of SPIE Vol. 11413, 114130V

It is important, then, to develop methods of communication and control that can potentially reduce human workload by allowing them to issue high-level commands or to query and interact with a swarm via an autonomous intermediary acting as a spokesperson.

In this paper, we explore the idea of controlling a swarm of robots using human language alone, a method which has already seen use with teams of virtual agents[16] and individual robots.[17] From a human controller's perspective, verbal communication is very natural, as it does not require learning an unfamiliar communication system or control interface. To avoid the need for individual communication with each robot, we utilize two models of communication intended to treat the swarm as a unit:[15] (1) a sports-team inspired "playbook" capability that assigns specific roles to members of the swarm that can be coordinated with other units, and (2) a virtual swarm "spokesperson" who communicates with humans and acts as a human-like front-end interface to the complex swarm and imparts information from the collective knowledge of the swarm.

We test our model using a simulated search and rescue mission. In the simulation, a small town is under threat of an approaching natural disaster (simulated wildfire), and a single-human/multirobot team has been tasked to find residents and help them evacuate safely. A human participant assumes the role of a swarm operator, tasked with the swift and safe evacuation of the town residents before the fire arrives at their location. The operator handles the high-level assignment of duties to members of the swarm, while low-level decisions are taken by the swarm itself. The operator also handles communication with residents who need to be evacuated. This paper describes the simulation environment we built for this purpose, and a data collection effort where 31 participants interacted with the system assuming the role of operator. At present, our simulation environment relies on two human controllers to interpret the participant's verbal instructions and input them into the simulation; the collected data will enable the training of an automated language understanding component to make the simulated environment fully autonomous.

2. BACKGROUND AND RELATED WORK

In emergency disaster situations, there is extreme time pressure to conduct searches and locate survivors within the first 48 hours, as the mortality rate increases substantially after this time.[18] Emergency responders are limited to whatever manpower can be mustered, as well as what technology is available to them to aid in their tasks.

The idea of using robots for search and rescue is well established,[19] with disasters being a driving force for advancement.[20] Natural catastrophes (e.g., earthquakes) or man-made disasters (e.g., terrorist attacks) can kill thousands of people in a very short amount of time, and the critical period (48–72 hours) after the disaster is when first responders want to utilize all the available assets.[21] After this period, victim mortality drastically increases owing to exposure and lack of food, water, and medical treatment.

Historically, ease of use has been an impediment in using robots during disaster events. In 2001, the World Trade Center rescue response teams had access to various robots, but only utilized a few. Some reasoning for the exclusion included physical constraints, but in one case, the robot was not used because the user interface provided with it was still in its infancy (suitable for developers, but not emergency personnel).[20] This highlights the importance of developing an interface that is intuitive to responders with various levels of training with robots, as the rate at which rescuers can be taught to use these tools correlates to the accessibility of training.[22]

A robotic swarm can be defined by the coordination between robotic units, which is reliant on shared information and distributed algorithms. Units within the swarm will be able to traverse areas dangerous to human life, mitigating some of the risk that typically befalls emergency responders working in the field. In addition, the swarm will need to be able to communicate effectively with survivors, as well as impart knowledge to emergency personnel responding to or controlling the swarm's movements.

There are various research groups worldwide who focus on developing robots for Urban Search And Rescue (USAR). The Center for Robot-Assisted Search And Rescue[23] is the oldest institution devoted to promoting the use of ground, marine, and aerial unmanned systems for public safety. Their mission is fostering unmanned systems used by formal emergency management agencies through voluntary national and international activities.

One of the first attempts for designing collaborative systems of mobile robots introduced an algorithm for a distributed team of autonomous mobile robots to search for an object.[24] Upon finding the object, the robots would gather around it and "rescue" it in a collaborative manner.

Other exploration of issues that arise when using robots in search and rescue include a collapsed building scenario[25] and technology for a swarm of robots assisting fire-fighters, which discusses the swarming algorithms by which the robots react to and follow humans.[26] There is also work on different metrics for performance evaluation of urban search and rescue tasks.[27, 28]

Rescue robotics is a "team of teams" process with a 2:1 human to robot ratio.[18] Previous research has developed two theoretical models: a workflow model identifying the major tasks, actions, and roles in robot-assisted search and a general information flow model capturing when, how, and why information is communicated between robots and other team members.[18] Our goal is to enable human operators to handle multiple robots at the same time and move beyond theory to computational models that can be implemented and used in practice.

As the number of robots increases, the need for coordination among robots as well as for assistance to the human operator exceeds current state-of-the-art capabilities.[14] Thus, in order to achieve a wider deployment of robots for practical tasks in various areas we need to expand human span of control over teams of robots.[14] Earlier research has shown that direct involvement of the human operator in the coordination of multiple robots can be very difficult and that performance deteriorates when a human operator is asked to control and coordinate the decisions of more than 8–12 robots.[4] This result motivates our use of a virtual spokesperson as a means to reduce the cognitive load on the human operator, and in turn allow for the simultaneous control and coordination of a large number of robots.

3. THE SIMULATION ENVIRONMENT

The primary objective of this project is to study how humans interact and communicate with swarms of robots using predefined plays and a virtual spokesperson. The concept was developed through several rounds of paper prototyping, with team members acting out all the roles in a wildfire scenario (operator, virtual spokesperson, robots, and town residents). We then implemented a simulation environment that allows a participant to take on the role of a control center operator in such a situation. The environment is extendible, so in the future, participants may take on other roles such as town residents or rescue workers.

3.1 Scenario

The environment simulates a wildfire as it approaches (and eventually consumes) a small town. The rescue task involves locating residents, convincing them to evacuate, and in some cases helping them to reach safety. The task is time-sensitive because each resident must evacuate before the fire reaches their location. Some of the challenges include efficiently locating residents; stubborn residents who need to be convinced to leave their homes; residents in need of physical assistance for evacuation; and a scarcity of resources such as search drones, evacuation vehicles, and most importantly time. All of these conspire to put a cognitive load on the operator as they work to manage the rapidly developing situation.

The resources available to the operator in the simulation are ten drones and one transport vehicle. The drones can perform a variety of tasks such as monitoring the fire's spread, searching for residents, playing recorded warnings, opening an audio communication channel with residents, following a resident's movement, or leading a resident to safety. The transport vehicle can drive and evacuate one resident household on each trip. To keep track of the situation, the operator can view a map of the town which is updated live with the location of the drones and vehicle, the areas that have been searched, and any information that was discovered such as the location of the fire and any found residents (Figure 1).

Managing all the resources individually is beyond the capabilities of a single operator. The operator therefore has at their disposal a set of "plays", or high-level commands, which instruct a group of drones to perform a certain action: the operator only needs to issue a command with parameters (for example, instruct a specified set of drones to search a certain area), and the drones self-organize to carry out the instruction. Available plays are listed in Figure 2. To communicate with the various resources, the operator utilizes a virtual agent that acts as a "spokesperson" for the swarm. The spokesperson can relay verbal messages to and from the swarm,

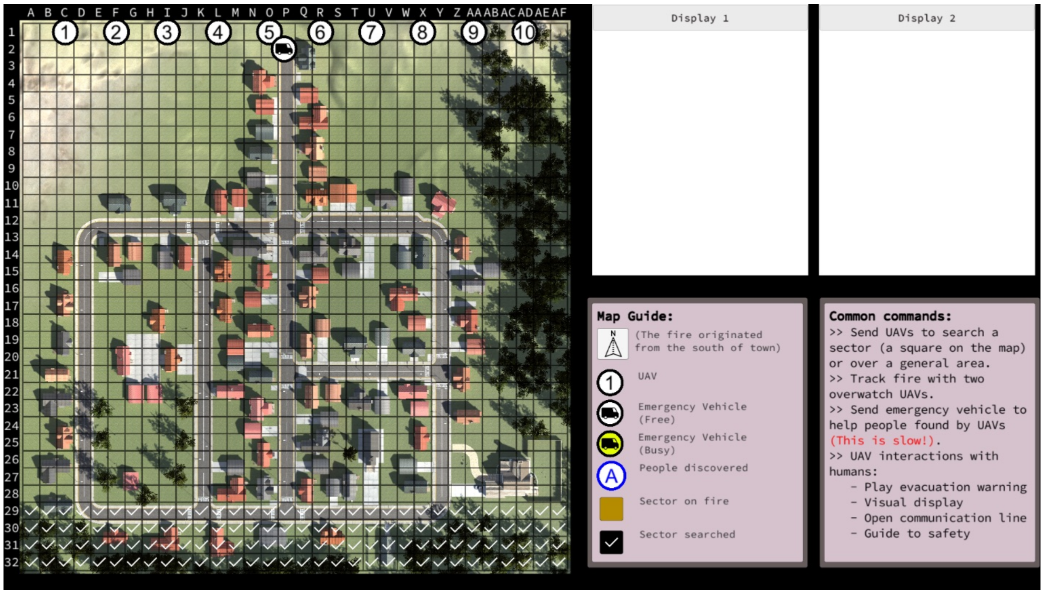

Figure 1. The operator's interface at the beginning of the simulation, showing the location of the ten drones and the rescue vehicle at the top of the map. Check marks at the bottom of the screen show areas that are known to be without residents. The map does not show the location of fire or any residents, because the drones have not yet been instructed to look for those.

Overwatch Send one or more drones to monitor the fire's spread. Drones will position themselves in the best spots to observe and report. Two drones are sufficient to cover the entire town.

Search Send one or more drones to search for residents in a designated area. Drones will self-organize to conduct a systematic search within the area. Drones search at a constant speed, so a search will conclude faster if it uses more drones or covers a smaller area.

Warning Instruct all drones to inform nearby residents of the current danger and the requirement to evacuate. Only residents that are close to a drone will hear the warning, and there is no guarantee they will heed it.

Monitor Instruct a drone to follow the location of an identified resident.

Guide Instruct a drone to guide a resident to safety. The drone will find the optimal route by itself. It is up to the resident to follow the drone, so this instruction is best given after the resident has agreed to comply.

Figure 2. Available "plays" which guide the behavior of drone swarms.

Cooperative Couple A couple who are capable and willing to evacuate, but need to be made aware of how dangerous the fire truly is. They will leave after a warning from either the operator or the virtual spokesperson.

Babysitter An individual accompanied by two children who is capable and willing to evacuate, but is anxious about leaving without the parents' approval. She can be persuaded to leave by either the operator or the virtual spokesperson.

Stubborn Man A resident who resists evacuation orders due to reluctance to leave behind his work. He will only leave if successfully persuaded by the operator.

Stubborn Couple A couple under the notion that their preparations to their house are enough to withstand the encroaching wildfire. They will only leave if successfully persuaded by the operator.

Van Driver A man tasked with transporting a small group of elderly patients who has become stranded with his patients following an accident on the road. This group requires the transport vehicle to successfully evacuate.

Figure 3. Town residents.

relieving the operator from the need to communicate individually with swarm members. The spokesperson can also perform some tasks on its own, such as communicating with residents or giving advice to the operator. Thus, the spokesperson agent acts to reduce the cognitive load on the operator, and an important factor to a successful evacuation is finding the right tasks to delegate to the spokesperson. The interaction between operator and spokesperson (and by extension, all other elements of the simulation) is done via spoken language.

The scenario includes five residents, each of whom requires different actions in order to successfully evacuate (Figure 3). The location of each resident is determined at random at the beginning of each simulation. Talking to the residents is key in getting them to agree to evacuate; the operator can talk to a resident directly through an audio channel enabled by a drone, or they can delegate the task to the spokesperson. The operator does not know the residents' characteristics ahead of time, nor do they know which resident they encounter at the time that resident is discovered. The operator therefore needs to figure out the necessary resources and actions through interaction with the resident. When multiple residents are discovered in short order, the operator needs to decide with the spokesperson who will handle each case.

3.2 Implementation

The simulation is implemented in the Unity game engine (https://unity.com), using the Residential Buildings Pack from Gabro Media* with trees from SpeedTree.† We added an overlay grid of 32×32 squares over the map, a set of rules for governing the behavior of the various simulation elements, and an interface for inputting parameters to the simulation. The following rules govern the behavior of simulation elements:

- The fire always spreads from the south, but its precise location and the speed at which it travels are determined by parameters chosen at random for each simulation.
- The locations of the five residents are chosen at random for each simulation.
- When a swarm of drones searches an area, they divide it into sections and each drone searches a section systematically following a line pattern. If the swarm changes mid-search (for example, one drone leaves to guide a resident), then the remaining members adjust their patterns to cover the area.
- When a drone guides a resident to safety, it follows the most efficient route out of town.

*https://assetstore.unity.com/packages/3d/environments/urban/residential-buildings-pack-39272
†https://assetstore.unity.com/packages/3d/vegetation/speedtree/free-speedtrees-package-29170

Figure 4. Simulation control room: the "simulation wizard" on the right controls events in the simulated environment; the "agent wizard" on the left monitors the video feed from the participant, and controls the virtual spokesperson agent and town residents.

- When the transport vehicle evacuates a resident to safety, it follows the most efficient route to the resident and out of town.

The spokesperson agent is implemented separately from the main application using the ICT Virtual Human Toolkit,[29] which provides components for the agent's appearance, non-verbal behavior, and synthesized speech.

At present we do not have trained language processing components that can interpret the participants' speech, so all verbal communication is done using the Wizard of Oz methodology, with human controllers standing in for the language understanding components, listening to the user speech and issuing commands that the system can act on.[30,31] Our system currently employs two controllers: a "simulation wizard" responsible for interpreting instructions to the swarm and inputting them into the simulation environment, and an "agent wizard" responsible for the verbal communication by the spokesperson agent and the town residents (Figure 4). The simulation wizard uses a control panel inside the Unity application to input system parameters such as drone search patterns and evacuation instructions (the panel is visible in Figure 4). The agent wizard triggers individual utterances by the spokesperson and town residents using a separate wizard interface on a laptop (Figure 5).[32] The interface has a large array of buttons that trigger predefined utterances, including acknowledgements, confirmation of actions, status updates, and suggestions to the operator. Buttons on the interface are color-coded in order to help wizards quickly find the desired buttons, maintaining a natural conversational pace. The wizard also has the option of typing a novel character utterance in real time, in case the existing utterances are insufficient. Utterances are realized through speech synthesis. Our intention is to use the data collected from the interactions to train language understanding components that will enable the future replacement of the wizards with automated components, following the trajectory of other human-robot communication projects.[33]

4. DATA COLLECTION

The system was used for data collection in October 2019, for the purpose of analyzing verbal and nonverbal behavior strategies when communicating with a robotic swarm, and in order to provide training data for the eventual automation and extension of the system.

4.1 Materials and setup

Data were collected using the simulation environment described in section 3. Participants were seated in a room in front of a Samsung 50-inch KU6300F Flat Smart 4K UHD TV showing the live game environment (Figure 6).

Figure 5. Control interface for the agent wizard: the top row of buttons changes the screen to quickly switch between agents; the current visible screen triggers utterances for the virtual spokesperson. Buttons are color coded as follows: **Red** Urgent responses for time-sensitive situations; **Orange** Clarification questions about directions; **Yellow** Information about available resources; **Green** Confirmation of executed actions; **White** Utterances that require the wizard to input an extra parameter (e.g., drone number or row identifier); **Gray** Other utterances that do not require additional input.

Between the participant and the TV was a small table where the participant could put a sheet with written instructions. The participant's face and upper body were recorded by a Logitech web camera on top of the TV, while their voice was recorded by a Sennheiser HSP-4 headworn microphone, connected to the computer through a Focusrite Scarlett 2i2 USB recording audio interface. To the right of the main display was a smaller Dell monitor showing the virtual spokesperson. Two Dell speakers for playing system alerts and the voices of the spokesperson and residents were located just beneath the main display. A Zoom handy video recorder behind the participant's right shoulder transmitted a live audio and video feed of the interaction to the control room (this feed was not recorded). A MacBook Pro laptop computer, used for inputting questionnaire data, was placed next to the small monitor, on a table with a separate chair.

The two wizards sat in a separate control room (Figure 4, above). The room contained a large Dell monitor for the simulation environment, a second large monitor and speakers providing a live feed from the experiment room, and a MacBook laptop computer for the agent wizard interface.

4.2 Participants

A total of 31 participants were recruited for the study through Craigslist (https://craigslist.org). Participants were all adults, English speakers, with normal or corrected vision and hearing. We did not collect any demographic data on the participants. Participants were paid $30 for their effort.

4.3 Procedure

Participants started by giving informed consent, then filled out pre-interaction questionnaires on the laptop computer. Participants then moved to sit in front of the main display, were given oral instructions and watched a short tutorial video, and were fitted with the microphone. After clarifying any remaining questions regarding the simulation, the participants then engaged in the simulation of trying to save residents from an oncoming

Figure 6. Participant in the experiment room: the main display shows the participant view of the simulation environment; the monitor to the right shows the virtual spokesperson.

wildfire. Interaction was verbal, mostly between the participant and the virtual spokesperson, and occasionally with town residents. When all the residents had either been saved or died, the experimenter entered the room, stopped the simulation, and the participant moved back to the side table and filled a questionnaire about the interaction. The process repeated itself for a second trial, followed by a second post-interaction questionnaire. After this final questionnaire, the experimenter came in for a debrief and answered any questions participants had about the experiment. Participants were thanked, paid, and escorted out.

4.4 Measures taken

The primary data are audio and video recordings of the participants' communication with the characters in the environment: audio and video were recorded continuously by the microphone and web camera. Microsoft Platform for Situated Intelligence (PSI)[34,35] was used as a platform for data collection. Using PSI allows automatic synchronization across multiple input streams with different latency and frequencies. Video and audio recordings as well as the experiment log files from the wizard control system and from the game environment were all stored in a PSI store format.

A measure of task success was the number of residents saved from the fire, ranging from 0 to 5 for each trial. Task success was logged manually on a spreadsheet immediately after each trial, together with experimenter notes about the trial.

Participants filled out several questionnaires. Three questionnaires were taken before the interaction: Guilford-Zimmerman spatial orientation,[36] mini-IPIP,[37] and the Need for Cognition;[38] another three questionnaires were taken after each interaction: NASA TLX,[39] one-item mental effort,[40] and an internally developed questionnaire about rapport and character perception.

4.5 Results

As of this writing the recordings have not yet been transcribed, so our analysis is based on the task success and experimenter notes. A total of 29 participants completed both trials and are included in the analysis. The mean number of lives saved in the first trial was 2.1 (median 2), and in the second trial the mean was 3.0 (median 3). The increase in lives saved from the first trial to the second was significant (paired t-test, $t = 3.08$, $df = 28$, $p < 0.005$ two-tailed), showing that participants were getting better at interacting with the simulation: only 7 participants lost more lives in the second trial, 6 participants saved the same number of lives, and 16 participants saved more lives in the second trial. There was also a strong positive correlation between the number of lives

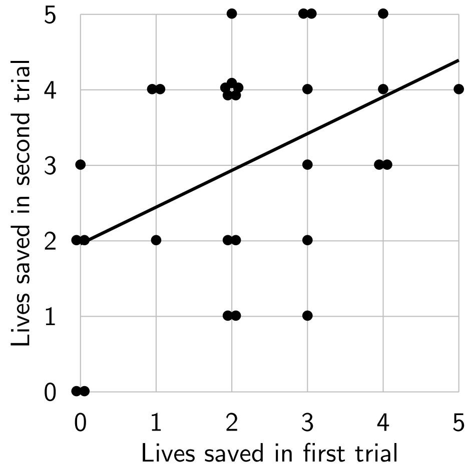

Figure 7. Positive correlation between lives saved in first and second trials ($r = 0.45$): each dot represents one participant.

saved in the first and second trials ($r = 0.45$, $df = 27$, $p < 0.05$), showing that some participants were better than others at saving lives (Figure 7).

We observed several strategies used by the participants. Generally, participants who engaged the spokesperson for advice and suggestions had better outcomes. Some participants tried to control drones individually, but this method was too slow and not many lives were saved; giving very detailed search instructions was also less efficient than relying on the swarm's self-organizing search capabilities. Five participants relied heavily on the rescue vehicle, sending it to all residents whether it was needed or not; this turned out to be effective for two participants, who used this strategy while engaging with the spokesperson, but produced poor results for the other three, who did not engage the spokesperson while using the same strategy. Some participants had difficulty understanding the various elements in the simulation: for example, one participant asked residents to follow the rescue vehicle rather than ride in it, and another kept trying to contact residents whose coordinates were engulfed in flames, not realizing that they had died. One participant engaged heavily in social conversation with the spokesperson but never followed the spokesperson's suggestions or advice; this participant did not save any lives on either trial.

The conversations with the spokesperson were quite varied. Participant utterances included, for example, instructions to individual drones or groups of drones; questions and answers about the current situation on the ground; consultations about possible courses of action; and questions and answers about capabilities. The spokesperson provided verbal alerts about events detected by the drones, such as finding survivors; requests for approval to take action; and advice. Most participants used the grid coordinates, though some preferred descriptive terms such as "the four corner intersection", a pattern consistent with other studies on human-robot navigation.[33]

Conversations with town residents included a variety of persuasion techniques to convince the residents to evacuate, from calm reasoning to threats. Some participants attempted to speak individually to each resident, while others relied more on the spokesperson. Several speakers gave up on trying to convince stubborn residents when they could not think of other persuasion strategies.

5. FUTURE STEPS

The next steps for using the collected data will be to transcribe participant utterances from the audio recordings, extract participant non-verbal behaviors such as eye gaze and head position from the video recordings, time-align the transcripts with agent utterances and simulation events recovered from system logs, and annotate participant and agent utterances with communicative intent (speech acts). We intend to develop a taxonomy of intents, based on speech act theory,[41,42] by adapting existing taxonomies of human-robot dialogue[43,44] to capture distinctions that are relevant to the domain of human-swarm communication for wildfire rescue.

The annotated data will be useful for several purposes. We can use the annotations to develop machine learning models for predicting the communicative intent from the utterance texts. Such prediction can help speed up the annotation of future data, and is a necessary step in developing a language understanding component to enable full automation of the agents (both the spokesperson and the town residents).

Another potential use of the annotated corpus is the modeling of team performance (in this case, a team of one person and several agents) based on the individual behaviors observed in the interactions.[45] In addition to the number of town residents saved, we can also extract from the logs the number of residents discovered, giving us two measures of team success. While some variation is due to initialization parameters (location of survivors and the fire), there are many individual behaviors of participants and agents that can affect the final team outcome. For example:

1. To what extent should the participant give instructions to drones individually, as opposed to allowing groups of drones to self-organize?
2. When should the participant step in to talk to a resident, as opposed to letting the drone or spokesperson communicate with them?
3. Under what conditions should the participant ask for a scarce rescue vehicle, as opposed to instructing a drone to guide a resident to safety?
4. When is it a good idea for the participant to ask for advice, and when should they make a decision themselves?
5. When is it useful for an agent (spokesperson or drone) to initiate an action, and what should they ask permission for?
6. Which events should the spokesperson communicate verbally to the participant?
7. When is it appropriate for the spokesperson to interrupt the participant from performing a task?

We can use the corpus of annotated interactions to learn models that predict team outcomes from individual and team communicative behaviors, for example by finding correlations between communicative intent and other behavior indicators, and good/bad team outcomes, and applying linear regression to determine which indicators contribute the most to team outcomes.[46]

We can also learn models of optimal behavior of individuals and the team as a whole by using available data and/or simulations and reinforcement learning. In a pilot experiment performed in simulation[47] we addressed the issues (2) and (7) mentioned above, i.e., when it is appropriate for the participant to step in and talk to a resident directly and when it is appropriate for the spokesperson to interrupt the participant and ask for help. We used reinforcement learning to automatically learn a policy to be followed when a drone has located residents in order to help them evacuate. The policy learned to choose the most appropriate action given the situational context, i.e., warn the residents through the drone or the spokesperson, ask the participant to intervene if needed, guide the residents to safety, or just wait for more information. This is an example of a policy that optimizes the use of the available assets (drones, spokesperson, and participant) to optimize team efficiency. The cost of having the spokesperson engage in conversation with the residents is higher than the cost of just warning them through the drone. This is because the spokesperson is responsible for multiple tasks whereas the drone is already occupied with monitoring the residents. Similarly, the cost of having the participant engage in dialogue with the residents is higher than the cost of having the spokesperson or the drone try to persuade the residents to evacuate, especially when the participant is busy handling more urgent situations. At the same time, the proximity of the fire is an unpredictable factor that should be taken into account. If the fire is rapidly approaching and the residents refuse to evacuate then there may not be enough time for warnings and the participant should be urgently asked to intervene. In this pilot experiment, we varied the distance of the fire (as well as the rate of spread of fire), the level of the cooperativeness of the residents, and the level of busyness of the participant. Our results were highly promising in terms of the percentage of residents saved especially given the fact that we take into account potential randomness in the aforementioned parameters.

Another use of reinforcement learning may be to improve the autonomous behavior of the spokesperson and the swarm entities, and their ability to communicate and adapt to interaction with humans. Reinforcement learning is well suited for this task as it is designed for maximizing long-term goals rather than focusing on immediate results that are not necessarily optimal in the long run. For example, a potential reward to be maximized could be the number of residents saved; reinforcement learning can then optimize the behavior of the spokesperson to achieve that goal. In addition, reinforcement learning could be used offline for designing new plays to populate our playbook. That is, we can provide the swarm entities with rewards (based on individual and common goals), and let them train against one another in order to come up with the optimal play for tackling a particular problem. Because in this case multiple agents will be learning at the same time, we will have to use multi-agent reinforcement learning, which does not assume that the learning environment is stationary.[48–50]

Other possible extensions of the present work are to expand the simulation, either by adding new roles for human participants (e.g., as town residents or rescue workers), or by developing new search and rescue scenarios (earthquake or flooding) or other swarm management scenarios.

6. CONCLUSION

We developed a platform for investigating the use of verbal communication to interact with robotic swarms, and used it two explore two models of interaction: predefined "plays" that can allow succinct instructions to guide a group of self-organizing robots, and a virtual "spokesperson" to serve as the verbal interface to the swarm. Our initial data collection suggests that both of these components are helpful for interaction with the swarm in a wildfire rescue scenario. Future work will include deeper analysis of the interaction, automation of the remaining system components, and extension of the platform to allow additional exploration of human-swarm verbal interaction.

Acknowledgements

The work depicted here was sponsored by the U.S. Army Research Laboratory (ARL) under contract number W911NF-14-D-0005. Statements and opinions expressed and content included do not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

REFERENCES

[1] Simmons, R., Apfelbaum, D., Burgard, W., Fox, D., Moors, M., Thrun, S., , and Younes, H., "Coordination for multi-robot exploration and mapping," in [*AAAI/IAAI*], 852–858 (2000).

[2] Adams, J. A., "Multiple robot/single human interaction: Effects on perceived workload," *Behaviour & Information Technology* **28**(2), 183–198 (2009).

[3] Gauci, M., Nagpal, R., and Rubenstein, M., "Programmable self-disassembly for shape formation in large-scale robot collectives," in [*Distributed Autonomous Robotic Systems*], Gross, R., Kolling, A., Berman, S., Frazzoli, E., Martinoli, A., Matsuno, F., and Gauci, M., eds., *Springer Proceedings in Advanced Robotics* **6**, 573–586, Springer (2018).

[4] Wang, H., Lewis, M., Velagapudi, P., Scerri, P., and Sycara, K., "How search and its subtasks scale in n robots," in [*ACM/IEEE International Conference on Human-Robot Interaction (HRI)*], 141–148 (2009).

[5] Kolling, A., Walker, P., Chakraborty, N., Sycara, K., and Lewis, M., "Human interaction with robot swarms: A survey," *IEEE Transactions on Human-Machine Systems* **46**, 9–26 (2016).

[6] Cummings, M. L. and Guerlain, S., "Developing operator capacity estimates for supervisory control of autonomous vehicles," *Human Factors* **49**(1), 1–15 (2007).

[7] Peters, J. R., Srivastava, V., Taylor, G. S., Surana, A., and amd Francesco Bullo, M. P. E., "Human supervisory control of robotic teams: Integrating cognitive modeling with engineering design," *IEEE Control Systems Magazine* **35**(6), 57–80 (2015).

[8] Guo, Z., Chen, R., Zhang, K., Pan, Y., and Wu, J., "The impairing effect of mental fatigue on visual sustained attention under monotonous multi-object visual attention task in long durations: An event-related potential based study," *PloS One* **11**(9), e0163360 (2016).

[9] Crandall, J. W. and Cummings, M. L., "Identifying predictive metrics for supervisory control of multiple robots," *IEEE Transactions on Robotics* **23**(5), 942–951 (2007).

[10] Valero-Gomez, A., de la Puente, P., and Hernando, M., "Impact of two adjustable-autonomy models on the scalability of single-human/multiple-robot teams for exploration missions," *Human Factors* **53**(6), 703–716 (2011).

[11] Humann, J. and Pollard, K. A., "Human factors in the scalability of multirobot operation: A review and simulation," in [*IEEE International Conference on Systems, Man, and Cybernetics (SMC)*], IEEE, Bari, Italy (2019).

[12] Kolling, A., Sycara, K., Nunnally, S., and Lewis, M., "Human swarm interaction: An experimental study of two types of interaction with foraging swarms," *Journal of Human-Robot Interaction* **2**, 103–129 (2013).

[13] Velagapudi, P., Scerri, P., Sycara, K., Lee, W.-H., and Wang, J., "Scaling effects in multi-robot control," in [*IEEE/RSJ International Conference on Intelligent Robots and Systems*], 2121–2126 (2008).

[14] Sycara, K. and Lewis, M., "Human control strategies for multi-robot teams," in [*WSEAS International Conference on Computers*], 149–154 (2012).

[15] Chaffey, P., Artstein, R., Georgila, K., Pollard, K. A., Gilani, S. N., Krum, D. M., Nelson, D., Huynh, K., Gainer, A., Alavi, S. H., Yahata, R., and Traum, D., "Developing a virtual reality wildfire simulation to analyze human communication and interaction with a robotic swarm during emergencies," in [*LTC '19: 9th Language and Technology Conference*], (May 2019).

[16] Traum, D., Rickel, J., Marsella, S., and Gratch, J., "Negotiation over tasks in hybrid human-agent teams for simulation-based training," in [*Proceedings of AAMAS 2003: Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*], 441–448 (July 2003).

[17] Lukin, S. M., Gervits, F., Hayes, C. J., Moolchandani, P., Leuski, A., Rogers III, J. G., Sanchez Amaro, C., Marge, M., Voss, C. R., and Traum, D., "ScoutBot: A dialogue system for collaborative navigation," in [*Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL), System Demonstrations*], 93–98 (July 2018).

[18] Murphy, R. R., "Human-robot interaction in rescue robotics," *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews* **34**, 138–153 (2004).

[19] Shah, B. and Choset, H., "Survey on urban search and rescue robots," *Journal of the Robotics Society of Japan* **22**(5), 582–586 (2004).

[20] Casper, J. and Murphy, R. R., "Human-robot interactions during the robot-assisted urban search and rescue response at the World Trade Center," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* **33**, 367–385 (2003).

[21] Stein, M., "The critical period of disasters: Insights from sense-making and psychoanalytic theory," *Human Relations* **57**(10), 1243–1261 (2004).

[22] USFA, "Technical rescue program development manual," Report FA-159, 2/96, United States Fire Administration, Federal Emergency Management Agency (1995).

[23] "The Center for Robot-Assisted Search And Rescue." http://crasar.org/ (2019).

[24] Jennings, J. S., Whelan, G., and Evans, W. F., "Cooperative search and rescue with a team of mobile robots," in [*International Conference on Advanced Robotics (ICAR)*], 193–200 (1997).

[25] Casper, J. L., Micire, M., and Murphy, R. R., "Issues in intelligent robots for search and rescue," in [*Unmanned Ground Vehicle Technology II*], **4024**, 292–303, International Society for Optics and Photonics (2000).

[26] Penders, J., Alboul, L., Witkowski, U., Naghsh, A., Saez-Pons, J., Herbrechtsmeier, S., and El-Habbal, M., "A robot swarm assisting a human fire-fighter," *Advanced Robotics* **25**(1-2), 93–117 (2011).

[27] Jacoff, A., Messina, E., Weiss, B. A., Tadokoro, S., and Nakagawa, Y., "Test arenas and performance metrics for urban search and rescue robots," in [*IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(Cat. No. 03CH37453)*], **4**, 3396–3403 (2003).

[28] Messina, E. and Jacoff, A., "Performance standards for urban search and rescue robots," in [*Unmanned Systems Technology VIII*], **6230**, 62301V, International Society for Optics and Photonics (2006).

[29] Hartholt, A., Traum, D., Marsella, S. C., Shapiro, A., Stratou, G., Leuski, A., Morency, L.-P., and Gratch, J., "All together now: Introducing the Virtual Human Toolkit," in [*Intelligent Virtual Agents: 13th International Conference, IVA 2013, Edinburgh, UK, August 29–31, 2013 Proceedings*], Aylett, R., Krenn, B., Pelachaud, C., and Shimodaira, H., eds., *Lecture Notes in Computer Science* **8108**, 368–381, Springer, Heidelberg (August 2013).

[30] Dahlbäck, N., Jönsson, A., and Ahrenberg, L., "Wizard of Oz studies – why and how," *Knowledge-Based Systems* **6**, 258–266 (1993).

[31] Green, A., Huttenrauch, H., and Eklundh, K. S., "Applying the Wizard-of-Oz framework to cooperative service discovery and configuration," in [*RO-MAN IEEE International Workshop on Robot and Human Interactive Communication*], 575–580 (2004).

[32] Artstein, R., Leuski, A., Maio, H., Mor-Barak, T., Gordon, C., and Traum, D., "How many utterances are needed to support time-offset interaction?," in [*Proceedings of the Twenty-Eighth International Florida Artificial Intelligence Research Society Conference*], 144–149, AAAI Press, Hollywood, Florida (May 2015).

[33] Marge, M., Bonial, C., Foots, A., Hayes, C., Henry, C., Pollard, K. A., Artstein, R., Voss, C. R., and Traum, D., "Exploring variation of natural human commands to a robot in a collaborative navigation task," in [*Proceedings of the First Workshop on Language Grounding for Robotics*], 58–66, Association for Computational Linguistics, Vancouver, British Columbia, Canada (August 2017).

[34] Bohus, D., Andrist, S., and Jalobeanu, M., "Rapid development of multimodal interactive systems: A demonstration of platform for situated intelligence," in [*ICMI '17: Proceedings of the 19th ACM International Conference on Multimodal Interaction*], 493–494, ACM, Glasgow, United Kingdom (November 2017).

[35] Andrist, S., Bohus, D., and Feniello, A., "Demonstrating a framework for rapid development of physically situated interactive systems," in [*14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*], 668, IEEE, Daegu, South Korea (March 2019).

[36] Guilford, J. P. and Zimmerman, W. S., "The Guilford-Zimmerman aptitude survey," *Journal of Applied Psychology* **32**(1), 24–34 (1948).

[37] Donnellan, M. B., Oswald, F. L., Baird, B. M., and Lucas, R. E., "The Mini-IPIP scales: Tiny-yet-effective measures of the Big Five factors of personality," *Psychological Assessment* **18**(2), 192–203 (2006).

[38] Cacioppo, J. T., Petty, R. E., and Kao, C. F., "The efficient assessment of need for cognition," *Journal of Personality Assessment* **48**(3), 306–307 (1984).

[39] NASA Ames Research Center, *NASA Task Load Index (TLX)* (1986).

[40] Paas, F. G. W. C., "Training strategies for attaining transfer of problem-solving skill in statistics: A cognitive-load approach," *Journal of Educational Psychology* **84**(4), 429–434 (1992).

[41] Searle, J. R., [*Speech Acts: An Essay in the Philosophy of Language*], Cambridge University Press (1969).

[42] Bunt, H., Alexandersson, J., Choe, J.-W., Fang, A. C., Hasida, K., Petukhova, V., Popescu-Belis, A., and Traum, D., "ISO 24617-2: A semantically-based standard for dialogue annotation," in [*Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC)*], 430–437 (May 2012).

[43] Bonial, C., Lukin, S. M., Foots, A., Henry, C., Marge, M., Pollard, K. A., Artstein, R., Traum, D., and Voss, C. R., "Human-robot dialogue and collaboration in search and navigation," in [*Proceedings of the AREA Workshop: Annotation, Recognition, and Evaluation of Actions*], (May 2018).

[44] Traum, D., Henry, C., Lukin, S., Artstein, R., Gervitz, F., Pollard, K., Bonial, C., Lei, S., Voss, C. R., Marge, M., Hayes, C. J., and Hill, S. G., "Dialogue structure annotation for multi-floor interaction," in [*Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC)*], 104–111 (May 2018).

[45] Kalia, A. K., Buchler, N., DeCostanza, A., and Singh, M. P., "Computing team process measures from the structure and content of broadcast collaborative communications," *IEEE Transactions on Computational Social Systems* **4**, 26–39 (June 2017).

[46] Georgila, K., Gordon, C., Choi, H., Boberg, J., Jeon, H., and Traum, D., "Toward low-cost automated evaluation metrics for Internet of Things dialogues," in [*Proceedings of the 9th International Workshop on Spoken Dialogue Systems Technology (IWSDS)*], (2018).

[47] Vaidyanath, S., Georgila, K., and Traum, D., "Using reinforcement learning to manage communications between humans and artificial agents in an evacuation scenario," in [*Proceedings of the Thirty-Third International Florida Artificial Intelligence Research Society Conference*], AAAI Press (2020).

[48] Georgila, K., Nelson, C., and Traum, D., "Single-agent vs. multi-agent techniques for concurrent reinforcement learning of negotiation dialogue policies," in [*Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*], 500–510 (2014).

[49] Xiao, G. and Georgila, K., "A comparison of reinforcement learning methodologies in two-party and three-party negotiation dialogue," in [*Proceedings of the Thirty-First International Florida Artificial Intelligence Research Society Conference*], 217–220, AAAI Press, Melbourne, Florida (May 2018).

[50] Bowling, M. and Veloso, M., "Multi-agent learning using a variable learning rate," *Artificial Intelligence* **136(2)**, 215–250 (2002).