Human Factors and Simulation, Vol. 180, 2025, 1–11
https://doi.org/10.54941/ahfe1006383

A Human Multi-Agent Teaming Testbed: Escape Room Simulation

Lawrence Dale Perkins[1], Hakki Erhan Sevil[1], and Matthew Johnson[1,2]

[1]Department of Intelligent Systems and Robotics, University of West Florida, Pensacola, FL 32514, USA
[2]Florida Institute for Human and Machine Cognition, Pensacola, FL 32502, USA

ABSTRACT

Scalability of multi-agent systems is quickly becoming a crucial area of research as artificial agents become more capable and sophisticated. A key challenge to furthering research in human-multi-agent teams is having suitable test domains that are simple enough to allow thorough analysis, yet rich enough to foster appropriate types of human-agent interaction. In this paper, we present a set of new open-source test domains based on a popular team game; the escape room. The testbed simulations are designed and developed to explore the dynamics of human interactions with multi-agent systems and the factors that enable and inhibit scalability. Our simulation provides a valuable tool for researchers and practitioners in this evolving field.

Keywords: Human-agent teaming, Human-agent interactions, Simulation, Scalability

INTRODUCTION

Multi-Robot or more generally Multi-Agent Systems (MAS) have been studied extensively over the last few decades and promise the potential for parallel execution of distributed tasks, reducing mission timelines and human workloads (Wooldridge, 2009). Despite efforts such as the DARPA Offset program (Arnold et al., 2019), the promised benefits of such systems often fail to manifest (Johnson and Vera, 2019), and the anticipated scale is rarely achieved in operational systems, with the state of practice remaining at one human to one robot or less (Murphy and Burke, 2005). Therefore, we are interested in how to scale MAS that work closely with humans, in other words, Human Multi-Agent Teams (HMAT). This area is rapidly becoming a crucial topic of research as agents become more capable and sophisticated and is particularly relevant in the context of increasingly ubiquitous artificial intelligence and automation.

A key challenge to furthering research in HMAT is having suitable test domains that are simple enough to allow thorough analysis, yet rich enough to capture the important characteristics that foster appropriate types of human-agent interaction. We begin by laying out simulation requirements based on our perspective of essential elements to explore the dynamics of human interactions with MAS and the factors that enable and inhibit scalability. We then review related works and discuss how they did and did not support our requirements. Then, we introduce our simulation by describing the Escape Room domain and present a set of new open-source test domains designed and developed to meet our proposed HMAT testbed requirements. Finally, we discuss potential modifications for future testbeds that would further advance the ability to effectively study scalability in HMAT.

SIMULATION REQUIREMENTS

Our desire to understand and model the factors that influence scalability and related measures led us to search for a testbed simulation that could provide specific characteristics in the following areas:

- The Human Role
- Complexity
- Scalability
- Interactions
- Autonomy
- Availability

The Human Role: Today and for the foreseeable future, people remain the owner of the problems that HMATs are designed to help solve. As such, they should remain the primary stakeholder in the mission. Therefore, we desire a simulation domain in which the human owns the mission goals. The human should develop, or play a major role in developing, the strategy used to coordinate the efforts of the team toward the common goal. Interaction should not simply be handing out work like a task dispenser, nor should it be simply responding to requests like a help desk when the automation is stuck. When humans interact with other teammates, including agents, they bring mission-related information to the interaction that will help to further the team goals by influencing individual agent decisions and actions to better fit into the team strategy. People also interact with multiple team members to gain information from their teammates to establish awareness, help assess progress, and trigger adaptation as necessary. Interaction is more than task allocation and the human role is more than a task dispenser. They are the problem holder.

Complexity: The work domain should contain qualities found in more complex real-world domains. One such quality is nested hierarchical activity. Domains that do not require task decomposition into related subtasks are too simple and miss a key aspect important to HMAT. Another important quality is interrelatedness. The work of a team has natural interdependencies (Johnson et al., 2014; Johnson and Vera, 2019; Johnson and Bradshaw, 2021) and the coordination necessary for effective teamwork is due to managing these interdependencies (Malone and Crowston, 1994; Johnson et al., 2014). The simulation domain must capture the interdependent nature of HMAT. Lastly, to model most real-world team applications, the task should be time constrained. Time constraints are a typical driver for distributed work and efficient team decision-making and adaptation. These three qualities are important influences on team dynamics, whether the team members are people or agents, and a proper testbed domain needs to support all three.

Scalability: An obvious need for a simulation that supports studying scalability is the ability to control and measure scalability. Control is straightforward, as it is typically easy to change the number of agents in a simulation. However, measurement is more challenging. The purpose of scaling is to ultimately improve speed and/or efficiency of goal achievement. While measuring such performance outcomes is trivial, it is important to be able to discern if the improvement is due to scalability or other factors, such as individual performance. Several predictive scalability measures around the concept of fanout have been proposed (Olsen and Goodrich, 2003; Crandall and Cummings, 2007; Cummings et al., 2007). Additionally, there is a need for more directly measured assessments such as number of active agents and performance plateau (Olsen, Wood and Turner, 2004). A scalability testbed should be instrumented to collect such measures. Two key measures in understanding scalability are interaction time and neglect tolerance, discussed next.

Interactions: Given the importance of interaction time on scalability, it is essential for the simulation to have the ability to measure and control the interaction between humans and agents. Additionally, it should be able to isolate and capture asynchronous interaction effort, such as when the human is observing an agent but not sending commands or requests. While improving and optimizing interfaces is important, it can confound scalability results. To understand scalability, the simulation needs to be able to isolate and measure when a human is interacting with an agent. This can then be used to inform interface design.

Autonomy: The other primary measure that influences scalability is neglect tolerance (Olsen and Goodrich, 2003) which reflects an agent's ability to conduct work independently. The simulation needs to be able to control and measure the time that an agent is self-sufficient in completing its assigned tasks. This can be achieved in many ways, such as changing the speed that an agent works, modifying the capabilities or intelligence of the agents, or adjusting the complexity of the agent's tasking. These are very different methods and introduce overlapping research areas, but to study HMAT and scalability, it is sufficient to understand that the simulation must have some control over what an agent can do and how long they can work independently.

Availability: Finally, the simulation needs to be available for use. If a simulation is available but does not meet all the needs, it may be possible to make modifications. However, this would require the simulation to be open source and include sufficient documentation to make the desired changes. Even when available and open source, implementation and modification of the simulation may prove to be too challenging to use. Ideally, multiple researchers working in similar spaces could reuse the same simulation for better validation and easier comparison of work. Unfortunately, most HMAT research uses a unique custom simulation and reuse is rare. Next, we review several studies using simulations and discuss how well they addressed the requirements we have laid out.

RELATED WORK

Table 1 lists related works that studied HMAT using various simulations and rates suitability in relation to the requirements presented in the previous section.

Human Role: We categorize the human role as either Stakeholder, Manager, or Help Desk in Table 1. By Stakeholder we mean the human is the problem owner and primary stakeholder. They must figure out a strategy and the plan of action necessary to move toward the goal. By Manager, we mean the human is not solving the domain problem, but instead the role is constrained to determining the task allocation necessary to support a predetermined plan of action. Lastly, by Help Desk we mean that the human operated as a help desk to assist agents when they got stuck.

Table 1: Selection of research with HMAT simulation environments.

Reference	Human Role	Complexity	Scalability	Interactions	Autonomy	Available
Adams, Hamell and Walker, 2024	Stakeholder	H, I, TC	yes	no	no	no
Rebensky et al., 2022	Help Desk	I	no	no	no	yes
Ruff et al., 2018	Manager	H, I, TC	no	no	no	no
Mercado et al., 2016	Manager	H, I, TC	no	no	no	yes
Lewis et al., 2014	Manager	H, TC	no	no	no	yes
Kidwell et al., 2012	Manager	H, I, TC	no	no	no	no
Goodrich et al., 2011	Stakeholder	none	no	no	no	no
Hardin and Goodrich, 2009	Stakeholder	H, I, TC	no	no	yes	no
Cummings et al., 2007	Manager	H, I, TC	yes	no	yes	no
Goodrich et al., 2007	Help Desk	none	no	yes	yes	no
Goodrich et al., 2007	Stakeholder	I	no	no	no	no
Goodrich et al., 2007	Help Desk	H, I	no	no	no	no
Goodrich et al., 2007	Stakeholder	H, I, TC	no	no	no	no
Crandall and Cummings, 2007	Stakeholder	H, I, TC	yes	no	yes	no
Lewis, Wang and Scerri, 2006	Stakeholder	H, I, TC	yes	no	no	no
Olsen, Wood and Turner, 2004	Manager	none	yes	no	yes	no
Zigoris et al., 2003	Stakeholder	H, I, TC	no	no	no	no

(H = hierarchical, I = interdependent, TC = time constraint)

Complexity: The simulations reviewed have either none or some combination of the qualities listed in the prior section. Table 1 indicates which quality each of these simulations have. If the task domain includes hierarchical work, the table contains a 'H', if the work was interdependent, 'I', and if there were time constraints, 'TC'.

Scalability: Most of the simulations reviewed do not specifically discuss if they can control or measure scale. While control may be constrained by processing speed of the simulator, adding more agents is typically not the issue. Measuring the number of agents active at a given point in the simulation or throughout the scenario should be easy to do. However, the number of agents that are effectively working to complete the task can be harder to distinguish. Table 1 indicates simulations that specifically discuss or present a measure of scalability.

Interaction: Previous studies concentrated on interface design to optimize the presentation of all applicable information to human users. While this may be important to maintain operator awareness, it confounds understanding interaction time by making it impossible to isolate the times the human is focusing on an individual agent. This is practical in a real-world interface, but the goal of our testbed was not to optimize for a specific scenario, but to quantify what enables scalability. Table 1 indicates if the simulation attempts to isolate both explicit interactions as well as interaction effort.

Autonomy: The autonomy of the agents varied significantly. Some agents did not operate over time but rather provided decision support to the human. To measure scalability, the goal was to implement agents that work over a period of time, and for the simulation to provide the ability to exert influence and measure their effective time. Similar to scalability, this time measure of autonomy is not discussed in many papers. In fairness, some HMAT work was not focused on scalability, but adjustable autonomy or other related concepts. However, recording when agents are actively operating in an environment should be feasible. Table 1 indicates if the study specifically tried to control and measure the time an agent acted independently.

Availability: Finally, many simulations are described in the literature but are difficult to find and may not be available or open source. Table 1 indicates which simulations appear to be available and open source.

After researching available testbeds, we did not find many simulations that met the desired criteria, were readily available, open source, easy to use, and most importantly, provided the desired human interaction dynamics. As such, we developed a suitable simulation testbed to explore measures of scalability and how both the characteristics of the technology and the human users affect the scalability of HMAT.

HUMAN MULTI-AGENT TEAMING TESTBED DESCRIPTION

Escape Room Domain

We desired a test domain that had complexity characteristics described above yet was natural and easy for untrained people to learn quickly. Escape rooms are immersive team games that have become popular in the last few decades. The typical escape room consists of a small team of two to ten players in a themed facility with a few rooms. To win the game, the players need to work together looking for a series of clues that will reveal keys to open locks and doors to escape the room within a given time limit. This time limit drives the team needs to spread out and explore the room searching for information. This natural distributed work drives a need to share information discovered. When clues are found, players communicate crucial pieces of information to teammates. The escape room models a complex domain and the aspects of teamwork we were looking for: exploration, partial knowledge, distributed execution, interdependent decision making, and layered hierarchical puzzle solving in time-critical scenarios. We developed two versions of the escape room discussed next.

Escape Room Simulation, Version 1

Human Role: In our initial escape room simulation, the team consists of a human player and a set of simulated homogeneous robots. The human directs the robots individually and collects information to solve the puzzle. To escape the room, the human participants must guide a team of robots to find, pick up, and deliver keys to open a series of locks on a two-dimensional map (see Figure 1). To add motivation, scores are calculated and displayed for picking up keys and unlocking locks. In this simulation, the human is the stakeholder responsible for devising a strategy and action plan to achieve the goal.

Complexity: The simulation contains the layered tasks of collecting distributed keys to open locks that reveal more keys for more locks. Although a player can simply interact with just one robot the whole time, they will not be able to complete the puzzle in the allotted time. To escape the room, the player will have to leave the robot and allow it to operate independently while interacting with other robots. Keys take the form of colored letters. The simulation retains the desired complexity modelled in the escape room scenario.

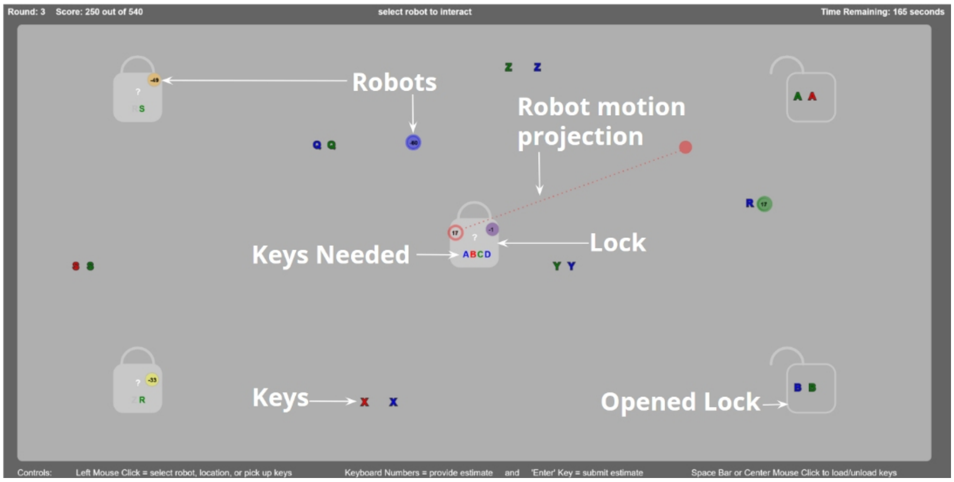

Figure 1: Human multi robot teaming escape room simulation: robots are circles identified by their colors. The main lock to escape is in the middle, with additional locks concealing information. Locks need a specific set of keys. Keys are differentiated by both color and letter.

Scalability: The number of active robots per second is recorded. Any robot activity was determined to be above the performance threshold. The number of active robots serves as the measure of scalability. This measure allows for comparisons against different models of fan-out which in turn, inform what factors play a critical role in modelling and increasing HMAT scalability.

Interaction: Interactions are constrained to measure time the player spends with each robot. During these interactions, players estimate how long each robot will work independently. Continuous tracking of all robots is unavailable; instead, while in a global view, participants see projected paths of robot movements based on their estimates. This prevents direct observation and helps isolate the interaction. To maintain awareness, information about keys and locks remains visible, and countdown timers based on player estimates help track time and reduce memory burden.

Autonomy: To control the time a robot acts independently, the speed of the robot's movement and key handling is predetermined based on the directed task. For motion, this is a function of the distance divided by the robot's speed and is easily adjustable in the simulation. The automation in this simulation is elementary. The robots can travel to the specified location and, once there, can be instructed to load or unload a key. This very simple form of automation allowed us to easily calculate how long the robot would act autonomously but required greatly simplifying typical real-world challenges.

Availability: The simulation was developed using the p5 JavaScript library, and source code are publicly available at https://git.ihmc.us/dperkins/human-multirobot-teaming.

Escape Room Simulation, Version 2

Human Role: In our second escape room simulation, we aimed to increase the capabilities (autonomy) of the agents, as well as allow the human more capability and flexibility in the level of tasking (interaction). The team structure and goals did not change, and the human remains the stakeholder. However, now a player can give the agents multiple tasks and the tasks are more sophisticated (e.g. go to a room, search for a certain color key, pick up a key, take it to a door, put down a key, and move inside the room). The player has two views; global view and interaction view. In both views, players can observe the timer, score, agent icons, and map as in Figure 2. In the global view, players can also see the queues of all agents on the map. In agent view, the player has agent specific detailed information; location, any other agents, traffic, or keys in the field of view, the dynamic path plan, and a status bar when loading/unloading a key.

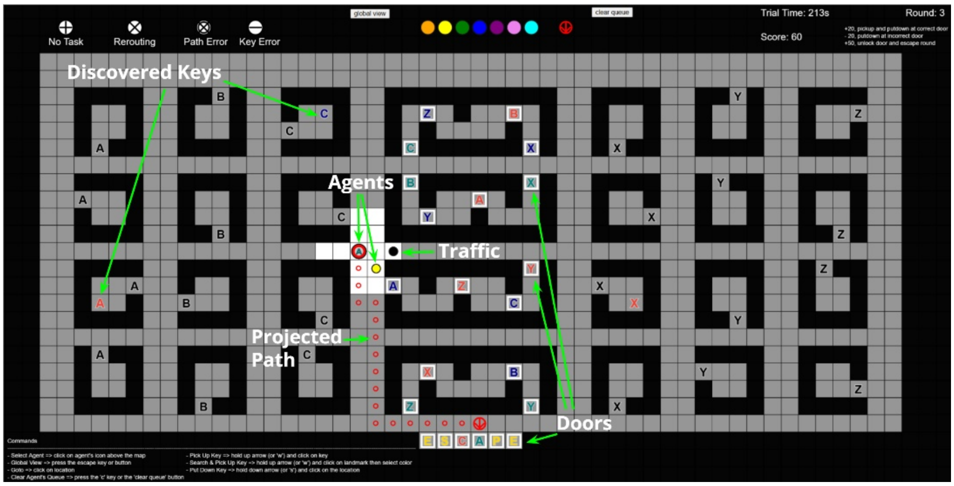

Figure 2: HMAT escape room simulation: agents are circles identified by their colors. The doors that spell 'ESCAPE' at bottom center need to be opened to escape the round. The rooms in the center of the map have the keys needed but are locked behind other doors. Keys are differentiated by both color and letter.

Complexity: In this version, more realistic challenges are introduced. The map, now a graph with open and closed (walls) spaces, required agents to navigate around walls and other agents in the open unoccupied spaces. Traffic agents were also added to increase task complexity. The locks from the original simulation were replaced with doors (framed in white). To escape the round, the doors at the bottom of the map that spell out *ESCAPE* need to be opened. The four yellow keys E, S, P, and E are located inside the rooms in the middle of the map and blocked by closed doors. The outer rooms are open to explore and contain all the A, B, C, X, Y, or Z keys needed to open the other doors. Labels identify what letter to expect inside, but the colors are hidden. Hence, multiple agents need to search different rooms to find the keys needed to open the doors. The complexity of the task domain retains the qualities of the first simulation but adds more hierarchical depth to the work and more uncertainty.

Scalability: Again, the number of active agents is recorded. Any agent activity is deemed to be above the performance threshold. The number of active agents serves as the measure of scalability. This measure allows for comparisons against different models of fan-out which in turn, inform what factors play a critical role in modelling and increasing HMAT scalability.

Interaction: Version 2 allowed the human to queue commands, building up a more complex sequence of activity. The simulation was still designed to isolate and measure the interactions between humans and agents. Additionally, we wanted to explore the effects of communication on teaming. Therefore, the simulation has trials with visual and audible messages from the agents to the human. The agents communicate four different messages; no task, rerouting, path error, and key error. The messages appear as symbols over the agent's icons located above the map, with a key to the messages displayed on the left. Flags also appear below the icon to increase the salience of the messages. While the four described messages are reactive, another visual message, the agent's projected path is a proactive form of communication.

Autonomy: To handle the extra environment difficulties, agents were given new capabilities: path planning, task queuing, search functions, and optional communications. Like version 1, the neglect time was controllable through speed, but version 2 added more capabilities and task queuing that also impact neglect time. When agents plan a path, they use local information from their field of view (2 squares in each direction) and the global map of open and closed spaces. If an agent's path is blocked, the agents will replan if possible and continue to the destination. Occasionally, the agents or traffic will block the only path available, in which case the human player will have to clear the task queue and deconflict the agents. Traffic agents are not controlled by the players, and if blocked just stop and wait until their path clears before resuming. Like the first version, the time an agent acts independently is measured and recorded. Furthermore, the goal was to project how long they could have acted independently if the human had not interacted. In the first simulation, the agents travelled in a straight line and did not have to avoid obstacles, which simplified this calculation. For the second simulation, this was not possible due to traffic uncertainty causing path rerouting. Therefore, the agent's self-sufficient time is measured and calculated as the time the agent acted independently before the human interaction plus the time it would have taken to complete the projected path, assuming no replanning. The last half of this calculation is an estimate, therefore could require calibration if heavy traffic causes frequent rerouting.

Availability: The simulation was developed using the p5 JavaScript library, and source code are publicly available at https://git.ihmc.us/dperkins/human-multi-agent-teaming.

Future Simulations

These two versions of the Escape Room simulation have proven to be useful in establishing factors that drive scaling limitations in Human Multi-Agent Teams. However, this is only the foundation; the objective is to research the factors that enable and inhibit effective scaling of HMAT. The results should inform design practices that enable more effective and scalable HMAT. To advance these tools to support HMAT research, we envision several areas to extend the simulations:

1. Increase the abilities of the agents
2. Increase agent to agent teaming
3. Implement hierarchical sub teams (e.g., swarms)
4. Implement the escape room in a 3-dimensional virtual simulation.

Each addition introduces real-world complexities that advance the understanding of critical elements of teaming. They also present experimental challenges (e.g., isolating and controlling variables). Implementing all four simultaneously may be difficult; however, adding each one thoughtfully and incrementally can lead to a suite of simulations that enhance HMAT research.

CONCLUSION

With the increase of artificial intelligence and agent capabilities and the desire to employ ever larger multi-agent teams, it is important to have a simulation to test out theories to better understand what influences, constrains, and enables effective HMATs. The Escape Room simulation provides a suitably complex domain where humans are stakeholders in mission goals, the interaction and observability of the agents are controllable elements, and the autonomy of the agents is easily adjustable. Our simulation is open source, easy to use, and provides the desired human interaction dynamics, making it a valuable tool for researchers in this evolving field. We encourage others to explore their questions on human multi-agent teaming using this scenario and the provided testbeds.

REFERENCES

Adams, J. A., Hamell, J. and Walker, P. (2024). Can a Single Human Supervise a Swarm of 100 Heterogeneous Robots? IEEE Transactions on Field Robotics, 2, pp. 46–80. doi: 10.1109/tfr.2024.3502316.

Arnold, R., Carey, K., Abruzzo, B. and Korpela, C. (2019). What is a robot swarm: A definition for swarming robotics. In: 2019 IEEE 10th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON). IEEE, pp. 0074–0081. doi: 10.1109/UEMCON47517.2019.8993024.

Crandall, J. W. and Cummings, M. L. (2007). Developing performance metrics for the supervisory control of multiple robots. In: 2007 2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE, pp. 33–40. doi: 10.1145/1228716.1228722.

Cummings, M. L., Nehme, C. E., Crandall, J. and Mitchell, P. (2007). Predicting Operator Capacity for Supervisory Control of Multiple UAVs. In: C. J. Singh, L. C. Jain, A. Mizutani and M. Sato-Ilic, eds., Innovations in Intelligent Machines - 1. [online] Springer Berlin Heidelberg, pp. 11–37. doi: 10.1007/978-3-540-72696-8_2.

Goodrich, M. A., McLain, T. W., Anderson, J. D., Sun, J. and Crandall, J. W. (2007). Managing autonomy in robot teams: Observations from four experiments. In: 2007 2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI). IEEE, pp. 25–32. doi: 10.1145/1228716.1228721.

Goodrich, M. A., Pendleton, B., B, S. P. and Pinto, J. (2011). Toward human interaction with bio-inspired robot teams. In: 2011 IEEE International Conference on Systems, Man, and Cybernetics. IEEE, pp. 2859–2864. doi: 10.1109/ICSMC.2011.6084115.

Hardin, B. and Goodrich, M. A. (2009). On using mixed-initiative control: a perspective for managing large-scale robotic teams. In: HRI '09: Proceedings of the 4th ACM/IEEE international conference on Human robot interaction. [online] New York, NY, USA: Association for Computing Machinery, pp. 165–172. doi: 10.1145/1514095.1514126.

Johnson, M. and Bradshaw, J. M. (2021). How Interdependence Explains the World of Teamwork. In: W. F. Lawless, J. Llinas, D. A. Sofge and R. Mittu, eds., Engineering Artificially Intelligent Systems: A Systems Engineering Approach to Realizing Synergistic Capabilities. [online] Springer International Publishing, pp. 122–146. doi: 10.1007/978-3-030-89385-9_8.

Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., Riemsdijk, M. B. van and Sierhuis, M. (2014). Coactive Design: Designing Support for Interdependence in Joint Activity. Journal of Human-Robot Interaction, [online] 3(1), pp. 43–69. doi: 10.5898/JHRI.3.1. Johnson.

Johnson, M. and Vera, A. (2019). No AI Is an Island: The Case for Teaming Intelligence. AI Magazine, [online] 40(1), pp. 16–28. doi: 10.1609/aimag.v40i1.2842.

Kidwell, B., Calhoun, G. L., Ruff, H. A. and Raja Parasuraman (2012). Adaptable and Adaptive Automation for Supervisory Control of Multiple Autonomous Vehicles. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, [online] 56(1), pp. 428–432. doi: 10.1177/1071181312561096.

Lewis, M., Chien, S.-Y., Siddarth M. Chakraborty, N. and Sycara, K. (2014). Task Switching and Single vs. Multiple Alarms for Supervisory Control of Multiple Robots. In: Harris, D. (eds) Engineering Psychology and Cognitive Ergonomics. EPCE 2014. Lecture Notes in Computer Science, 8532, pp. 499–510. doi: 10.1007/978-3-319-07515-0_50.

Malone, T. W. and Crowston, K. (1994). The Interdisciplinary Study of Coordination. ACM Comput. Surv., [online] 26(1), pp. 87–119. doi: 10.1145/174666.174668.

Mercado, J. E., Rupp, M. A., Jessie, Barnes, M. J., Barber, D. and Procci, K. (2016). Intelligent Agent Transparency in Human–Agent Teaming for multi-uxv Management. Human Factors, [online] 58(3), pp. 401–415. doi: 10.1177/0018720815621206.

Murphy, R. R. and Burke, J. L. (2005). Up from the rubble: Lessons Learned about HRI from Search and Rescue. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, [online] 49(3), pp. 437–441. doi: 10.1177/154193120504900347.

Olsen, D. and Goodrich, M. (2003). Metrics for Evaluating Human-Robot Interactions. In: Proceedings of PERMIS. Gaithersburg: Faculty Publications, p. 4.

Olsen, D. R., Wood, S. B. and Turner, J. (2004). Metrics for Human Driving of Multiple Robots. In: IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. IEEE, pp. 2315–2320, Vol. 3. doi: 10.1109/ROBOT.2004.1307407.

Rebensky, S., Carmody, K., Ficke, C., Carroll, M. and Bennett, W. (2022). Teammates Instead of Tools: The Impacts of Level of Autonomy on Mission Performance and Human–Agent Teaming Dynamics in Multi-Agent Distributed Teams. Frontiers in Robotics and AI, 9. doi: 10.3389/frobt.2022.782134.

Ruff, H., Calhoun, G., Frost, E., Behymer, K. and Bartik, J. (2018). Comparison of Adaptive, Adaptable, and Hybrid Automation for Surveillance Task Completion in a Multi-Task Environment. Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 62(1), pp. 155–159. doi: 10.1177/1541931218621036.

Wooldridge, M. (2009). An introduction to MultiAgent systems. 2nd ed. Wiley Publishing.

Zigoris, P., Siu, J., Wang, O. and Hayes, A. T. (2003). Balancing Automated Behavior and Human Control in multi-agent systems: A Case Study in RoboFlag. In: Proceedings of the 2003 American Control Conference, 2003. IEEE, pp. 667–671, vol. 1. doi: 10.1109/ACC.2003.1239096.