





# Demonstrating Performance Benefits of Human-Swarm Teaming


## Demonstration Track

William Hunt   
University of Southampton  
Southampton, United Kingdom  
w.hunt@soton.ac.uk

Jack Ryan   
University of Southampton  
Southampton, United Kingdom  
j.ryan@soton.ac.uk

Ayodeji O. Abioye   
University of Southampton  
Southampton, United Kingdom  
a.o.abioye@soton.ac.uk

Sarvapali D. Ramchurn   
University of Southampton  
Southampton, United Kingdom  
sdr1@soton.ac.uk

Mohammad D. Soorati   
University of Southampton  
Southampton, United Kingdom  
m.soorati@soton.ac.uk




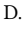

### ABSTRACT

Autonomous swarms of robots can bring robustness, scalability and adaptability to safety-critical tasks such as search and rescue but their application is still very limited. Using semi-autonomous swarms with human control can bring robot swarms to real-world applications. Human operators can define goals for the swarm, monitor their performance and interfere with, or overrule, the decisions and behaviour. We present the “Human And Robot Interactive Swarm” simulator (HARIS) that allows multi-user interaction with a robot swarm and facilitates qualitative and quantitative user studies through simulation of robot swarms completing tasks, from package delivery to search and rescue, with varying levels of human control. In this demonstration, we showcase the simulator by using it to study the performance gain offered by maintaining a “human-in-the-loop” over a fully autonomous system as an example. This is illustrated in the context of search and rescue, with an autonomous allocation of resources to those in need.

### KEYWORDS

Human-Swarm Teaming; Swarm Robotics; Human-robot Interaction; Simulation Environments; HARIS

#### ACM Reference Format:

William Hunt , Jack Ryan , Ayodeji O. Abioye , Sarvapali D. Ramchurn , and Mohammad D. Soorati . 2023. Demonstrating Performance Benefits of Human-Swarm Teaming: Demonstration Track. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023. IFAAMAS, 3 pages.

## 1 INTRODUCTION

Autonomous multi-agent systems are becoming increasingly ubiquitous in our everyday lives, from autonomous vehicles to home automation, and most pressingly in high-stakes scenarios such as disaster relief [1, 4]. In these settings, more so than any other, the cognitive burden on human operators to control a large collective of autonomous systems over a dispersed area is unmanageable [14]. A single human operator is unable to effectively control individual agents, designate tasks, and achieve overall mission objectives in a timely and precise manner. On the other hand, fully autonomous

systems are still lacking in their ability to handle a diverse range of real-world scenarios reliably, and as such their performance is not assured, nor deemed trustworthy, for safety-critical systems. [8, 9]

It is hypothesised that the optimal solution lies at the intersection between these two approaches – in human-swarm teaming [3] – and is therefore the primary goal of our proposed tool. An efficient human-swarm system requires a highly usable interface for the human operator(s) of the swarm to define tasks with priorities, allocate groups of agents to a task set, monitor the behaviour of the swarm, analyze the decisions of the system and intervene when needed [1].

The objective is to enable human-swarm teaming in various scenarios and allow researchers and practitioners to quickly and easily build and run user studies with a human-supervised swarm of robots. To achieve this objective, we extend a digital twin simulation platform, the Human Teaming Simulator (HutSim) [11], that facilitates the completion of a search and rescue task, in two different modes—fully autonomous and human-teaming modes. Existing multi-agent simulators such as Gazebo [5], Webots [6], and ARGoS [10] have limited utility for experiments on human-robot interaction. Features such as “loop functions” in ARGoS enable some waypoint and manual control of agents, but these simulators do not have (multi-)human interaction as a central concept. More recent projects such as AirSim focus primarily on physics-modelling for “Hardware-in-the-Loop” simulation [13] and are not intended for complex multi-robot human interaction. HARIS is a java-based Web application with low reliance on external libraries that can be easily deployed online and accessed in a web browser, making it ideal for large-scale user studies with multiple users that can access the same scenario instance simultaneously if desired. HARIS has built-in APIs for plugging in physical UAVs, allowing it to also act as a digital twin platform if desired. We use search and rescue as an exemplar of a safety-critical scenario that generalizes to other similar applications in disaster management, serving to illustrate the performance gain of human-swarm teaming in a realistic setting.

HARIS models all scenarios using a simple JSON format, meaning that all interface and algorithmic options can be customised on the fly and even the mode of operation and swarm behaviours can be changed without restarting the server. Users can also easily create their own JSON files to suit their experiment needs or use the built-in scenario editor to place targets, hazards, and agents on the realistic (google) map anywhere in the world; for example, within

*Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

5 minutes an operator could build a simulation of the actual site of a recent natural disaster for experimental and training purposes.

## 2 METHODOLOGY

To demonstrate the flexibility of HARIS, we present an example use case for human-swarm teaming in disaster relief which is to identify a number of unknown casualties in an area using a fixed number of UAVs. The methodology employed combines HARIS with a novel experimental design, to demonstrate the potential for performance gain by employing human-swarm teams in response to a real-world search and rescue scenario.

In this scenario, the swarm autonomously performs dynamic allocation using the embedded algorithms (i.e., max-sum decentralised coordination algorithm [2, 12]) where the swarm optimises for the total energy cost of an allocation. Human operator preferences can also be modelled as constraints used to modify the original max-sum algorithm when computing task allocation [11]. The design of HARIS is based on our previous work including a user study with 100 online participants [1], interviews with 6 experienced drone pilots [9], and a survey with 26 swarm experts from academia [7]. We used the result of these studies to design and test the usability of the interface elements and functionalities such as custom interface lenses which can be switched on-the-fly to meet the operator’s needs.

Figure 1 shows an experiment where four UAVs are autonomously assigned to unknown targets on a map in HARIS (these represent areas with potential casualties). The agents move towards their targets and find out whether or not they are actual casualties. The operator can monitor the views of the agents and their planned schedule on the right-hand panel, or can manually reassign them using the task view.

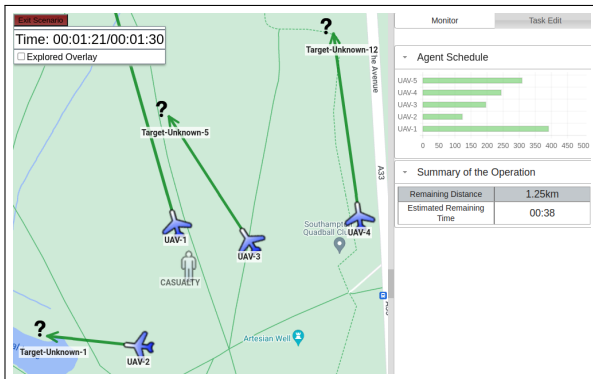


Figure 1: The HARIS interface uses Google Maps to provide a realistic command centre for the swarm

## 3 DEMONSTRATION

*Design:* Users are presented with a Google Maps API view showing  $N$  UAVs (5 by default) and  $M$  unknown targets (12 by default). The swarm will autonomously use the max-sum algorithm to efficiently allocate UAVs to these targets. As the UAVs approach their targets, a popup shows a simulated camera feed of the area, which initially

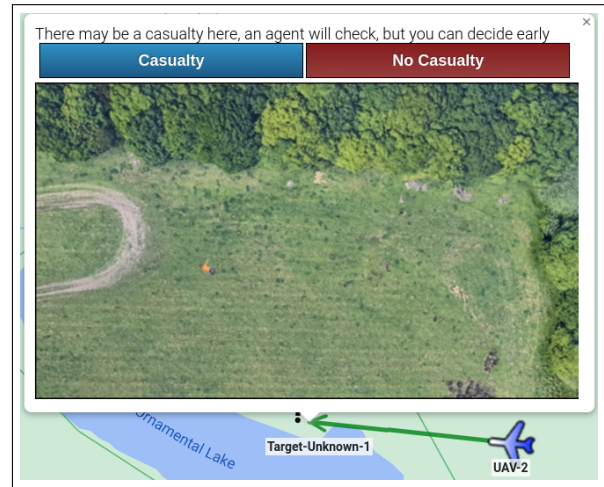


Figure 2: Users are shown a simulated image of the target which gets clearer as the agent gets closer to the target. The user can classify the target early, saving the operation time.

has a low resolution, but the resolution incrementally increases as the UAV approaches the target.

The human user can support the swarm by viewing these images and providing an earlier classification of “Casualty” or “No Casualty” once the image becomes sufficiently clear. This allows a human to augment the swarm with their superior reasoning skills, and empower the existing strategy by reducing the number of wasted journeys, where UAVs travel to the target only to find that it is not a real casualty. Figure 2 shows the popup with corresponding buttons for human interaction.

*Evaluation:* Performance is assessed by a combined metric of classifications per minute  $\times$  accuracy of classifications. We expect a human-swarm team to classify more targets within the time limit, leading to a higher score. We ran the scenario using an entirely autonomous swarm (i.e., no human assistance), and provide this as a benchmark with which users can compare their performance against. This allows us to clearly demonstrate that even untrained users can empower the swarm, increasing its performance.

HARIS is easy to deploy online, and so we have made the scenario accessible, along with a video demonstrating the use case in action <sup>1</sup>.

## 4 CONCLUSION

This demonstration showcases our human-in-the-loop simulator, HARIS, by building an example use case which enables a direct comparison of untrained human users forming human-swarm teams, with fully autonomous multi-agent swarms in a safety-critical context. Further, the performance gain is measured quantitatively through a combined speed-accuracy metric for comparison between both individuals, and the autonomous swarm. HARIS is accessible online and can be used in a variety of search and rescue scenarios and will facilitate future research into the optimal strategies for human-swarm teaming.

<sup>1</sup>Link to HARIS Demo website: <https://uos-hutsim.cloud/demo/aamas2023/>

## REFERENCES

- [1] Mohammad D Soorati, Jediah Clark, Javad Ghofrani, Danesh Tarapore, and Sarvapali D Ramchurn. 2021. Designing a user-centered interaction interface for human–swarm teaming. *Drones* 5, 4 (2021), 131.
- [2] Francesco Maria Delle Fave, Alex Rogers, Zhe Xu, Salah Sukkarieh, and Nicholas R Jennings. 2012. Deploying the max-sum algorithm for decentralised coordination and task allocation of unmanned aerial vehicles for live aerial imagery collection. In *2012 IEEE International Conference on Robotics and Automation*. IEEE, 469–476.
- [3] Aya Hussein and Hussein Abbass. 2018. Mixed initiative systems for human-swarm interaction: Opportunities and challenges. In *2018 2nd Annual Systems Modelling Conference (SMC)*. IEEE, 1–8.
- [4] Mauro S Innocente and Paolo Grasso. 2019. Self-organising swarms of firefighting drones: Harnessing the power of collective intelligence in decentralised multi-robot systems. *Journal of Computational Science* 34 (2019), 80–101.
- [5] Nathan Koenig and Andrew Howard. 2004. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*(IEEE Cat. No. 04CH37566), Vol. 3. IEEE, 2149–2154.
- [6] Olivier Michel. 2004. Cyberbotics Ltd. Webots™: professional mobile robot simulation. *International Journal of Advanced Robotic Systems* 1, 1 (2004), 5.
- [7] Mohammad Naiseh, Sarvapali D. Ramchurn, and Mohammad D. Soorati. 2022. Outlining the design space of eXplainable swarm: experts' perspective. *International Symposium on Distributed Autonomous Robotic Systems (DARS)* (2022).
- [8] Changjoo Nam, Phillip Walker, Huao Li, Michael Lewis, and Katia Sycara. 2019. Models of trust in human control of swarms with varied levels of autonomy. *IEEE Transactions on Human-Machine Systems* 50, 3 (2019), 194–204.
- [9] Katie J Parnell, Joel E Fischer, Jediah R Clark, Adrian Bodenmann, Maria Jose Galvez Trigo, Mario P Brito, Mohammad D Soorati, Katherine L Plant, and Sarvapali D Ramchurn. 2022. Trustworthy UAV relationships: Applying the Schema Action World taxonomy to UAVs and UAV swarm operations. *International Journal of Human–Computer Interaction* (2022), 1–17.
- [10] Carlo Pinciroli, Vito Trianni, Rehan O'Grady, Giovanni Pini, Arne Brutschy, Manuele Brambilla, Nithin Mathews, Eliseo Ferrante, Gianni Di Caro, Frederick Ducatelle, et al. 2012. ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm intelligence* 6, 4 (2012), 271–295.
- [11] Sarvapali D Ramchurn, Joel E Fischer, Yuki Ikuno, Feng Wu, Jack Flann, and Antony Waldock. 2015. A study of human-agent collaboration for multi-UAV task allocation in dynamic environments. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- [12] Alex Rogers, Alessandro Farinelli, Ruben Stranders, and Nicholas R Jennings. 2011. Bounded approximate decentralised coordination via the max-sum algorithm. *Artificial Intelligence* 175, 2 (2011), 730–759.
- [13] Shital Shah, Debadepta Dey, Chris Lovett, and Ashish Kapoor. 2017. AirSim: High-Fidelity Visual and Physical Simulation for Autonomous Vehicles. In *Field and Service Robotics*. arXiv:arXiv:1705.05065 <https://arxiv.org/abs/1705.05065>
- [14] David St-Onge, Marcel Kaufmann, Jacopo Panerati, Benjamin Ramtoula, Yanjun Cao, Emily BJ Coffey, and Giovanni Beltrame. 2019. Planetary exploration with robot teams: Implementing higher autonomy with swarm intelligence. *IEEE Robotics & Automation Magazine* 27, 2 (2019), 159–168.

## ACKNOWLEDGEMENTS

This work was conducted in the "Open-Source Interaction Interface for Human-Swarm Teaming" pump priming project, as part of the Trustworthy Autonomous Systems Hub, supported by the UK Engineering and Physical Sciences Research Council (EP/V00784X/1). The project also received funding from the UK Research and Innovation Centre for Doctoral Training in Machine Intelligence for Nano-electronic Devices and Systems (EP/S024298/1).