

Human-Multi-Robot Team Collaboration using Advising Agents

(Doctoral Consortium)

Ariel Rosenfeld
Bar-Ilan University
Ramat-Gan, Israel
arielros1@gmail.com

ABSTRACT

The number of multi-robot systems deployed in field applications has risen dramatically over the years. Nevertheless, the part of the human operator in these systems has been mostly overlooked. In this work we propose a novel approach for utilizing automated advising agents when assisting an operator to better manage a team of multiple robots in complex environments. We introduce an advice provision methodology and exemplify its implementation using automated advising agents in two real-world human-multi-robot team collaboration tasks: the Search And Rescue (SAR) task and the warehouse operation task. Our intelligent advising agents were evaluated through extensive field trials, with over 100 human operators and both simulated and physical mobile robots, and showed a significant improvement in the team's performance.

Keywords

Human-Robot Interaction; Human-Agent Interaction; Advice Provision

1. INTRODUCTION

In recent years multi-robot systems have been applied to complex tasks which used to be performed by humans alone. Two prominent examples are the Search And Rescue (SAR) [3] and the warehouse operation [2] tasks. In such multi-robot systems the part of the human operator is often assumed to be marginal. Two assumptions are made in this case: the first is that the robots perform relatively smoothly, with the infrequent need for human intervention. Second is that the human operator is only required to perform a single task at any given moment. Reality, however, can be more complicated on both accounts.

The deployment of robots in real-world environments has shown that robots usually face difficulties in completing their tasks. Specifically, failures are common. In such situations, a human operator must get involved in order to solve the problem. Furthermore, human operators may be occupied by numerous tasks at a time. For example, in a warehouse operation setting, the human worker may be employed in

packing merchandise, refiling inventory and handling robot malfunctions, all at the same time. The prioritization of the operator's tasks has mostly been overlooked in system design, leaving the prioritization in each operator's hands. Sub-optimal prioritization of the operator's tasks may result in sub-optimal performance by the robot-team and in a high cognitive workload for the operator [8].

Improving the performance of human-multi-robot systems can be done using one of the following approaches: Either (1) Improving the robot's hardware and software, thus relying less on human supervision (making the robots more autonomous): or (2) Improving the efficiency of the Human-Robot Interaction (HRI). Assuming we are given a team of robots and we cannot control the reliability of their hardware or software, our work deals with improving the HRI in order to allow a person to control a team of many (possibly unreliable) robots at once while executing different domain-specific tasks.

In this thesis, we present a novel methodology that enhances operators' performance by using intelligent advising agents. The agents provide advice to the operator regarding which actions she should take and acts as a smart filter between the robots' requests and the human operator. Our methodology is not restricted to any certain hardware or algorithm used by the robots, and we consider these factors constants. Note that intelligent advising agents have been successfully implemented in different domains, for example [6, 1, 5]. However, to the best of our knowledge, this is the first work on utilizing advising agent technology in the HRI field.

We present an optimization heuristic which models the maximization of the human-multi-robot team's performance by selecting when and which advice to provide the operator in a greedy fashion, using a k -steps-lookahead search. We have extensively tested the suggested heuristic in designing advising agents deployed in both simulated environments (using the Gazebo robot simulation toolbox¹) and physical deployment with more than 100 human operators. Experimental results show that our advising agents were able to significantly enhance the operators' performance when managing a large team of mobile robots in both the SAR and warehouse operation tasks.

The main contribution of this work is in showing, for the first time, that an intelligent agent which supports the human operator can lead to better performance of the human-multi-robot team.

Appears in: *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016)*, J. Thangarajah, K. Tuyls, C. Jonker, S. Marsella (eds.), May 9–13, 2016, Singapore.

Copyright © 2016, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

¹<http://www.gazebosim.org/>

2. THE SAR TASK

In our SAR task, the operator remotely supervises and controls 10 unmanned mobile ground robots using a computer program. The operator is required to find and correctly classify green objects in a clustered terrain². Our SAR task is conducted in two distinct environments: a real-world office building floor that was investigated in both simulation and physical deployment, and a medium-sized warehouse yard that was investigated in a simulation.

In our SAR task, robots can malfunction. We account for three malfunction types that our physical robots can experience: **Sensing** – where the robot’s camera/laser stops working. In such problems, the operator can send the robot back to home base, where the experiment operator can fix it. **Stuck** – where the robot cannot move (for example, due to an obstacle), and needs the operator to help it get loose. In some cases the Stuck malfunction is terminal. **WiFi** – where the robot stops sending and receiving data. In such cases the operator can mark the area as a “no entrance” zone. Upon losing the WiFi signal, the robot is programmed to return to its last point of communication. Furthermore, the operator could perform the following operations: she could **tele-operate** a robot from one place to another, she could order a robot to autonomously **navigate** to a desired point of interest and, once an object is detected by a robot, she could **classify** the object as a ball or a box.

Due to the exponential size of the robots’ state space and the relative independence of the operator’s possible operations (each operation can be performed regardless of others), we only examined the 1-step-lookahead heuristic.

We have extensively tested the performance of our advising agent, which deployed the 1-step-lookahead heuristic, in a within-subjects experimental design, in both simulation and physical deployment with 44 human subjects.

In our physical deployment, a major **100%** increase in the average number of detected objects was recorded (14.1 vs. 7 objects, $p < 0.001$) when subjects were equipped with our advising agent. An increased average of covered terrain (462 vs. 305 square meters, $p < 0.05$) and a significant decrease in the average time that robots stay idle (2720 vs. 3244 seconds, $p < 0.05$) were also recorded.

Our full results are reported in [4] and a short video³ summarizing this work is available at <http://vimeo.com/119434903>.

3. WAREHOUSE OPERATION TASK

We focus on a small warehouse acting as a fulfilment center (also known as a packing center) where there is one human worker and 10 mobile ground robots capable of transporting shelves from one place to another over the warehouse floor. At the packing stations, it is the human worker’s job to unload products from the arrived shelves. The task is conducted in a simulated warehouse which we built using the Gazebo robot simulation toolbox.

We account for two malfunction types that our simulated robots can experience: **Stuck** – where the robot cannot move (for example, due to a fallen product), and needs the operator to help it get loose. In some cases the Stuck malfunction is terminal. **Path** – where the robot cannot find a feasible path to reach its goal. In such cases, the opera-

tor can remove fallen products and tele-operate the different robots to clear the blocked path. Furthermore, the operator could perform the following operations: she could **tele-operate** a robot from one place to another, she could **unload** an item from an arrived shelf, she could direct a robot to **move** a shelf from one place to another and, once all requested items in a customer’s order have been unloaded, she could **pack** the order and send it.

Note that the warehouse operator’s possible operations are strongly connected and dependent; for example, a worker cannot pack a customer’s order before attaining all the requested merchandise. This property necessitates the identification of these connections and their consideration in the advice provision process.

We designed and evaluated 2 advising agents, deploying the 1-step-lookahead and 2-steps-lookahead heuristics. Through an extensive empirical study with 60 human operators, we show that subjects equipped with an advising agent filled significantly more orders compared to subjects using a non-advising agent (Silent) by up to 13% (on average), $p < 0.05$. We further show that our agent that uses the 2-steps-lookahead heuristic, significantly outperforms the agent that uses the 1-step-lookahead heuristic, $p < 0.05$. Surprisingly, subjects equipped with the 1-step-lookahead heuristic agent accepted significantly more pieces of advice compared to subjects equipped with the 2-steps-lookahead heuristic agent. This phenomenon will be investigated in future work.

Our full results are reported in [7] and a short video summarizing this work is available at <http://www.youtube.com/watch?v=rC1a4c6Voco>.

REFERENCES

- [1] A. Azaria, A. Rosenfeld, S. Kraus, C. V. Goldman, and O. Tsimhoni. Advice provision for energy saving in automobile climate-control system. *AI Magazine*, 36(3):61–72, 2015.
- [2] E. Guizzo. Three engineers, hundreds of robots, one warehouse. *Spectrum*, 45(7):26–34, 2008.
- [3] Y. Liu and G. Nejat. Robotic urban search and rescue: A survey from the control perspective. *Journal of Intelligent & Robotic Systems*, 72(2):147–165, 2013.
- [4] A. Rosenfeld, N. Agmon, A. A. Maksimov, Oleg, and S. Kraus. Intelligent agent supporting human-multi-robot team collaboration. In *IJCAI*, 2015.
- [5] A. Rosenfeld, A. Azaria, S. Kraus, C. V. Goldman, and O. Tsimhoni. Adaptive advice in automobile climate control systems. In *AAMAS*, pages 543–551, 2015.
- [6] A. Rosenfeld and S. Kraus. Providing arguments in discussions based on the prediction of human argumentative behavior. In *AAAI*, pages 1320–1327, 2015.
- [7] A. Rosenfeld, A. Noa, O. Maksimov, and S. Kraus. Human-multi-robot team collaboration for efficient warehouse operation. In *Autonomous Robots and Multirobot Systems (ARMS)*, 2016.
- [8] P. Squire and R. Parasuraman. Effects of automation and task load on task switching during human supervision of multiple semi-autonomous robots in a dynamic environment. *Ergonomics*, 53(8):951–961, 2010.

²The objects can indicate victims, flames, weapons, etc.

³Best Robotics Video - IJCAI 2015.