Towards Successful Coordination of Human and Robotic Work using Automated Scheduling Tools: An Initial Pilot Study

Matthew C. Gombolay†, Ronald J. Wilcox‡, Ana Diaz†, Fei Yu*†, and Julie A. Shah†
Department of Aeronautics and Astronautics†
Department of Mechanical Engineering‡
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
gombolay@csail.mit.edu, rjwilcox@mit.edu, anadiaz@mit.edu, fyu@stu.ust.hk, julie_a_shah@csail.mit.edu

Abstract— **With the latest advancements in robotic manufacturing technology, there is a desire to integrate robot workers into the labor force to increase productivity and efficiency. However, coordinating the efforts of humans and robots in close physical proximity and under tight temporal constraints poses challenges in planning and scheduling and the design of human-robot interaction. In prior work, we present a scheduling algorithm capable of performing the coordination of heterogeneous multi-agent teams. Given this capability, we now want to understand how best to implement this technology from a human-centered perspective. Humans derive purpose and identity in their roles at work, and requiring them to dynamically change roles at the direction of an automated scheduling algorithm may result in the human worker feeling devalued. Ultimately, overall productivity of the human-robot team may degrade as a result. In this paper, we report the results of a human-subject pilot study aimed at answering how best to implement such an automated scheduling system. Specifically, we test whether giving humans more control over the task allocation process improves worker satisfaction, and we empirically measure the trade-offs of giving this control in terms of overall process efficiency.**

I. Introduction

Robotic systems are increasingly entering domains previously occupied exclusively by humans. In manufacturing, there is an increasing desire to integrate robots into the workforce to leverage the heterogeneous strengths of both humans and robots. This integration requires a choreography of human and robotic work that meets upperbound and lowerbound temporal deadlines on task completion (e.g. assigned work must be completed within one shift) and spatial restrictions on agent proximity (e.g. robots must maintain four meter separation from other agents), to support safe and efficient human-robot co-work.

In our recent work [6], we present Tercio, a centralized task assignment and scheduling algorithm that scales to multi-agent, factory-size problems and supports on-the-fly replanning with temporal and spatial-proximity constraints. We demonstrate that this capability enables human and robotic agents to effectively work together in close proximity to perform manufacturing-relevant tasks. Tercio is made efficient through a fast, satisficing, incomplete multi-agent task sequencer inspired by real-time processor scheduling techniques. We demonstrate that Tercio generates near-optimal schedules for up to 10 agents and 500 tasks in less than 10 seconds.

In this paper, we will seek to understand how much control human workers should have over the assignment of roles and schedules when working in a team with robot partners. Successful integration of robot systems into human teams requires more than just algorithms capable of performing dynamic task allocation and scheduling. The mechanisms for coordination must be valued and appreciated by the human workers. Human workers often find identity and security in their roles or jobs in the factory and are used to some autonomy in decision-making. A human worker that is instead tasked by an automated scheduling algorithm may begin to feel that his or her value is importance. Even if the algorithm increases process efficiency at first, there is concern that taking control away from the human workers may alienate them and ultimately damage the productivity of the human-robot team. The study of human factors can inform the design of effective algorithms for collaborative tasking of humans and robots.

In this work, we present results from a pilot study (n=8) where participants collaborate with a robot on a construction task. In one condition, both the human and robot are tasked by Tercio, the automatic scheduling algorithm. In the second condition, the human worker is provided with a limited set of task allocations from which he or she can choose. We hypothesize that giving the human more control over the decision-making process will increase worker satisfaction, but that doing so will decrease system efficiency in terms of time to complete the task.

We begin in Section II with a brief review of related work in human factors studies of human-machine systems and task allocation. Section III briefly describes Tercio, the task allocation and scheduling algorithm we use in our experiment. In Section IV and Section V we describe our experimental method and we report the results from the human subject experiment. We then discuss the implications, limitations, and lessons learned from the results of our pilot

*This author collaborated on this work during a summer research position at MIT while he was a student at Hong Kong University of Science and Technology (HKUST). **This work was supported by Boeing Research and Technology and by the National Science Foundation (NSF) Graduate Research Fellowship Program (GRFP) under grant number 2388357.

study in Section VI. Lastly, we summarize our findings and discuss future work in Section VII.

II. Background

Developing human-machine systems that are able to leverage the strengths of both humans and their artificial counterparts has been the focus of much attention from both human factors engineers and researchers in artificial intelligence. Parasuraman has pioneered work examining adaptive automation to regulate operator workload [10]. When the operator is over-tasked, the automation can reduce the burden on the operator by assuming certain tasks. If the human-automation system experiences a period of low workload, the automation can shift more responsibility to the human operator to mitigate possible complacency or a reduction in manual skills [9].

The human-robot interface has long been identified as a major bottleneck in utilizing these robotic systems to their full potential [2]. As a result, we have seen significant research efforts aimed at easing the use of these systems in the field, including careful design and validation of supervisory and control interfaces [1], [4], [7], [8]. Other, related research efforts have focused on utilizing a human-in-the-loop to improve the quality of task plans and schedules [3], [4], [5]. This is particularly important when the search space is large or if it is not possible to represent all aspects of the environment in the problem formulation.

These prior efforts have focused primarily on utilizing a human-in-the-loop to improve plans and schedules for work performed by other agents. In this work, we are motivated by applications in the manufacturing domain where human workers will be performing physical tasks in coordination with robotic partners. In some cases, the human workers may also be responsible for tracking the progress and tasking the team. We seek to understand how much control human workers should have over the assignment of roles and schedules when working in teams with robots. While we necessarily want to maximize the efficiency of human-robot teams, we also desire for the human workers to value, accept, and cooperate with the new technology. With this in mind, the following sections of this paper will describe the pilot study we conducted to lend insight into trade-offs among flexibility in decision-making, overall task efficiency, and worker satisfaction.

III. Tercio: Task Allocation and Scheduling Algorithm

We use the Tercio algorithm [6] to perform task allocation and scheduling for our pilot study. In this section, we formulate the task assignment and scheduling problem for multiple robots moving and working in the same physical space as a mixed-integer linear program (MILP), as shown in Equations 1 - 9, and briefly discuss the benefits of the Tercio algorithm for this problem domain.

$$min\ f(A, A^{prev}, J, S, E, R, \tau) \tag{1}$$

subject to

$$\sum_{a \in Ag} A_{a,j} = 1, \forall j \in \tau \tag{2}$$

$$lb_i \leq t_m - t_n \leq ub_i, \forall i \in T, n, m \in \tau \tag{3}$$

$$t_k^E - t_k^S \geq lb_{a,k} - M(1 - A_{a,k}), \forall k \in \tau, a \in Ag \tag{4}$$

$$t_k^E - t_k^S \leq ub_{a,k} + M(1 - A_{a,k}), \forall k \in \tau, a \in Ag \tag{5}$$

$$t_j^S - t_i^E \geq M(1 - J_{i,j}), \forall i, j \in R \tag{6}$$

$$t_i^S - t_j^E \geq M J_{i,j}, \forall i, j \in R \tag{7}$$

$$t_j^S - t_i^E \geq M(1 - J_{i,j}) + M(2 - A_{a,i} - A_{a,j}) \quad \forall i, j \in \tau \tag{8}$$

$$t_i^S - t_j^E \geq M J_{i,j} + M(2 - A_{a,i} - A_{a,j}) \quad \forall i, j \in \tau \tag{9}$$

In this formulation, $A_{a,j} \in \{0, 1\}$ is a binary decision variable for the assignment of agent a to task j, and $J_{i,j}$ is a binary decision variable specifying the relative sequencing of two tasks i and j ($J_{i,j} = 1$ implies task i occurs before j). T is the set of all interval temporal constraints relating tasks, Ag is the set of all agents, τ is the set of all tasks, and t_i^S and t_i^E represent the start and end times of task i, respectively. ub_i and lb_i are the upper and lowerbound temporal constraints on the duration of task τ_i. R is the set of all task pairs (i, j) that are separated in space by less than the minimum separable distance. S and E are the set of task start and end times, respectively. Finally, $A_{a,i}^{prev}$ represents $A_{a,i}$ from the previous task allocation (if any). M is an artificial variable set to a large positive number, and is used to encode conditional constraints.

The objective function $f(A, A^{prev}, J, S, E, R, \tau)$, shown in Equation 1, is application specific. In our pilot study, the objective function served to minimize the makespan, or total process duration. Equation 2 ensures that each task is assigned to one agent. Equation 3 ensures that the temporal constraints relating tasks are met. Equations 4 and 5 ensure that agents are not required to complete tasks faster or slower than they are capable. Equations 6 and 7 sequence actions to ensure that agents performing tasks maintain safe buffer distances from one another. These equations enforce spatial constraints by requiring tasks that overlap in physical space do not also overlap in time. Equations 8 and 9 ensure that each agent only performs one task at a time. Note Equations 6 and 7 couple the variables relating sequencing constraints, spatial locations, and task start and end times, resulting in tight dependencies among agents' schedules.

Solving the task allocation and scheduling problem as a MILP for large-scale manufacturing applications (i.e., tens of agents and hundreds of tasks) is computationally intractable; the key bottleneck is the binary decision variables for the sequencing of tasks, which is exponential with the number of tasks. To alleviate this bottleneck we developed Tercio, a hybrid algorithm that uses a MILP to solve the task allocation problem in Equations 1-2 and a heuristic scheduler to provide near-optimal schedules that satisfy Equations 3-9.

By leveraging problem structure and utilizing a polynomial time approach for the sequencing of tasks, we are able to dynamically perform task allocation and scheduling for ten agents and hundreds of tasks in less than twenty seconds on average [6], in contrast to current methods that timeout at approximately five agents and fifty tasks. Next, we discuss how we use Tercio to perform task allocation and scheduling for our pilot study.

IV. Methods

The purpose of this pilot study is to gain insight into how to best integrate multi-agent task allocation and scheduling algorithms to improve the efficiency of coordinated human and robotic work. We hypothesize that keeping the human worker in the decision making process by letting him/her decide the task assignments will decrease performance of the system (i.e., increase time to complete the task objective) but will increase the human appreciation of the overall system. We conducted a pilot study to assess the tradeoffs in task efficiency and user satisfaction, depending on whether or not the worker is allowed control over task allocation.

For our pilot study, we consider two experimental conditions:

- Condition 1: A task allocation and scheduling algorithm (i.e., Tercio) specifies the plan.
- Condition 2: The human worker is given a limited set of task allocations from which he or she can choose.

A. Experimental Setup

The task objective given to the human-robot team is to complete two Lego kits, each consisting of seven steps. The parts required to build each of the seven steps for both Lego kits are placed into bins away from the build area. There are two classes of roles for the agents: *builder* or a *part fetcher*. If an agent is a builder, then the agent is responsible for building either one or both of the Lego kits. If an agent is a fetcher, that agent retrieves part bins and transports them to a specified builder. We enforce that the fetcher can only fetch one part bin at a time.

If a robot is assigned to fetch parts, we must have some mechanism of informing the robot when and to whom to fetch the part bin. *A priori* we tuned the temporal constraints of the task network so that the robot would fetch the next part bin for a human builder just before the human finished the previous step. In a future experiment, we want to incorporate a closed-loop feedback mechanism so that the timing of the fetching adapts to when the human is ready for the parts. We discuss this further in Section VI.

We use two KUKA Youbots (See Figure 1), which are mobile manipulator robots (height 65cm, weight 31.4 kg). To control the robots movement through the space, we implemented a simple closed-loop control system. A PhaseSpace motion capture system, which tracks the locations of LEDs that pulse with unique frequencies, providing real-time feedback of the state of the robots (i.e., $\vec{x} = [x, y, z]$ for each robot).

Fig. 1. A picture of a KUKA Youbot. Image Courtesy of KUKA Robotics.

When the experiment begins, the initial assignment of the roles is as follows:

1) The human subject is responsible for completing one of the Lego kits.
2) One Youbot, called "Burra", is responsible for completing the second Lego kit
3) A second Youbot, called "Kooka", is responsible for fetching the Lego parts for both the subject and the robot Burra.

After "Kooka" fetches the first three part bins for the human subject and "Burra", we simulate a failure of "Kooka" and inform the human subject of the malfunction. Because "Kooka" was responsible for fetching the remaining four part bins for the human subject and "Burra", the assignment of roles must change to complete the build process. We recall that we want to observe the effects of giving the human subject control over his or her task assignment versus using an autonomous task allocation scheduling algorithm (i.e., Tercio).

For participants assigned to Condition 1, the Tercio algorithm is used to assign roles to the human and robot workers in an optimal manner. In the optimal solution, the participant and the robot Burra work independently, fetching their own part bins and completing their own Lego kits. For participants assigned to Condition 2, the human worker decides the roles of both himself or herself and the robot Burra. The two options proposed to the human worker are:

1) The human worker and "Burra" work independently
2) "Burra" fetches all remaining part sets and the human worker completes the two Lego kits.

After task allocation has been performed, either by Tercio or the human participant, the human and "Burra" complete their respective tasks. The completion of both Lego construction tasks marks the end of the trial.

In manufacturing domains, the initial task allocation and scheduling is performed offline, well in advance of the actual work. The challenge for researchers is performing dynamic recomputation of these plans in response to disturbances during runtime. Currently, this recomputation is performed by a human supervisor and the process can be costly to the manufacturer. For this reason, we want to understand the tradeoffs of using an automated or manual mechanism for recomputing the task allocation online *after* a significant online disturbance.

B. Data Collection

All experiments were performed at the Interactive Robotics Group (IRG) laboratory at MIT. The experiment was approved by the Massachusetts Institute of Technology's Committee on the Use of Humans as Experimental Subjects (COUHES) and informed consent was obtained from each subject prior to each experimental session.

We tested eight subjects in total; four subjects were randomly assigned to each of the two conditions. The distribution of gender was 1 male and 3 females for Condition 1, and 2 males and 2 females for Condition 2. Subjects' age ranged between 23-27 years old with a mean and standard deviation of 25 years, 3 months $\pm$ 1 year, 9 months of age. All subjects were recruited from the MIT graduate student population. Because the target population consists of workers in a manufacturing environment, our pilot study suffers from a selection bias. In a future experiment, we will attempt to collaborate with industry to test the use of automated task allocation and scheduling algorithms on the target population.

Time to complete the task objective (i.e. finish building both Lego sets) was measured using a stopwatch. This time includes the time that the human spent deciding how to reallocate him/herself and "Burra" to the remaining fetching and building tasks after "Kooka" malfunctions. We include this decision time to better simulate the extra time the decision-making process would take versus the computer-assigned decision. At the end of the trial, each participant completed a questionnaire asking them to rate their level of agreement with the seven statements using a five-point Likert scale from 1 (strongly disagree) to 5 (strongly agree):

1) I was satisfied by the robot system's performance.
2) I would use this robot system next time this set of tasks was to be repeated.
3) The robots collaborated well with me.
4) The robots and I performed the tasks in the least possible time.
5) I was content with the task allocation after the robot malfunctioned.
6) I felt safe and comfortable with the robots.
7) The robots were necessary to the successful completion of the tasks.

C. Statistical Analysis

Performance data (time to complete the task objective in seconds) were tested for normality using the Kolmogorov-Smirnov test (Condition 1: $p = 0.994$; Condition 2: $p = 0.49$). A one tail t-test for two independent samples with equal variances was used to compare the two conditions. Prior to that, the samples were tested for equal variances using the F test ($p = 0.08$).

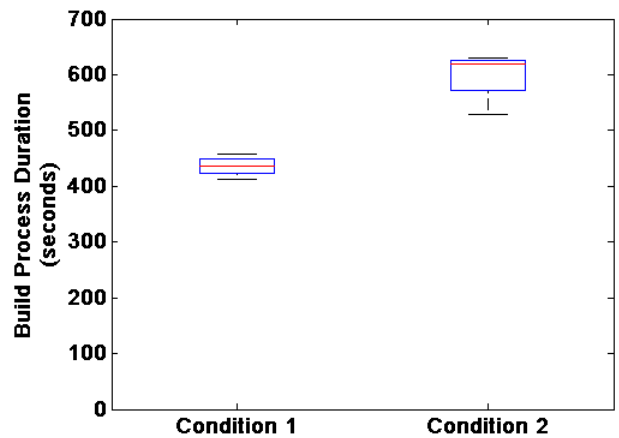

Fig. 2. Boxplot showing the median, quartile and standard deviations of the performance of the human subjects in both conditions.

Human appreciation of the system (or human satisfaction) data were calculated for each subject as the average of all seven questions in the questionnaire. Thus, every subject had an ordinal score between 1 and 5. The Mann-Whitney U test (a.k.a Wilcoxon test) were used to compare the two conditions. In all cases, significance was taken at the $\alpha = 0.05$ level. Data is presented as the average $\pm$ standard deviation.

V. Results

A. Performance

The average time for the four participants in Condition 1 was found to be 436 ± 19.1 s. Similarly, the average time for the four participants in Condition 2 was found to be 598.8 ± 47.5 s. Figure 2 shows the boxplot of the performance results (time to complete the two Lego kits) across the two conditions (in Condition 1, the algorithm decides assignments; in Condition 2, the subject decides assignments). Time to complete the task was significantly higher in Condition 2 than in Condition 1 ($p < 0.001$).

B. Human Appreciation of the System

We measure the subject's appreciation of the system as the mean of his/her scores across the seven questions of our post-test questionnaire. The average questionnaire rating for the four participants in Condition 1 is 3.54 ± 0.18. Similarly, the average questionnaire rating for the four subjects in Condition 2 was found to be 3.22 ± 0.60. Figure 3 shows the boxplot of the rating results across the two conditions (in Condition 1, the algorithm decides assignments; in Condition 2, the subject decides assignments) The non-parametric Mann-Whitney U test did not find any significant difference between the two conditions ($U = 5 > U_{crit} = 1$). Furthermore, the average rating in Condition 1 is higher than the average rating in Condition 2, which is in disagreement with the initial hypothesis.

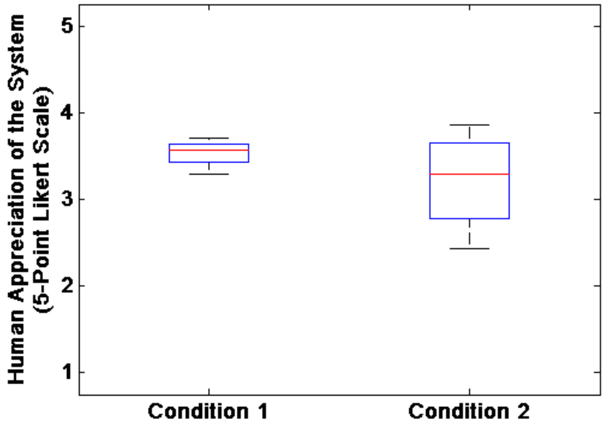

Fig. 3. Boxplot showing the median, quartile and standard deviations of our measure of human appreciation of the autonomous system based on a five-point Likert scale.

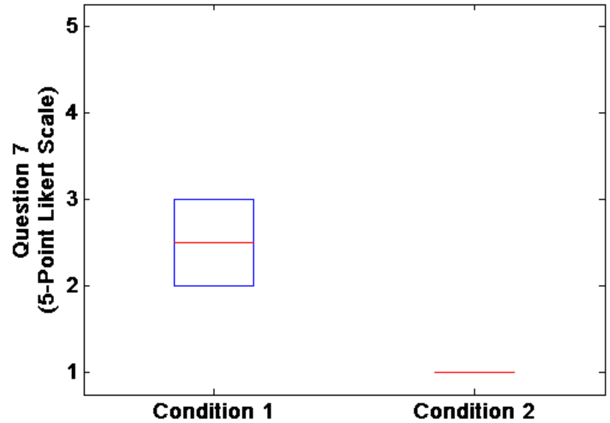

Fig. 4. Boxplot showing the median, quartile and standard deviations of our measure of human appreciation of the autonomous system based on a five-point Likert scale.

Question	Condition 1	Condition 2	P-Value
1	4.00 ± 0.00	3.25 ± 0.96	0.13
2	3.25 ± 0.96	3.00 ± 1.41	0.65
3	4.25 ± 0.96	4.00 ± 1.41	0.88
4	1.75 ± 0.50	2.00 ± 1.41	0.87
5	4.00 ± 0.82	4.50 ± 0.58	0.34
6	5.00 ± 0.00	4.75 ± 0.50	0.32
7	2.50 ± 0.58	1.00 ± 0.00	$< \mathbf{0.01}$

Table 1: Table reporting the mean and standard deviation of the ratings for each experimental condition and question on our post-test questionnaire. This table also shows the results of statistical testing for these ratings; statistically significant results ($\alpha = 0.05$) are shown in bold.

We also evaluate the ratings across subjects for the individual questions. The mean and standard deviations for each experimental condition as well as the p-vales are shown for each question in Table 1. While the ratings for Questions 1-6 are not statistically significant, there is a statistically significant difference between the subjects in each experimental condition for Question 7 ($p < 0.01$) shown in Figure 4. All four subjects in Condition 2 reported that they "strongly disagree" with the statement that "The robots were necessary to the successful completion of the tasks." While we cannot establish the cause for the subjects in Condition 2 believing that the robots were not necessary for the completion of the task, this devaluing of the roles of the robots does correspond to the humans having more direct supervisory control over the system. In a future experiment, we plan to revise the post-test questionnaire to better ascertain the underlying sentiment of the subjects in regards to their views of the overall system, the value of the team members, and their personal metric for evaluating system optimality.

VI. Discussion

A. Evaluation of Time to Complete the Task

Time to complete the task objective was significantly higher in Condition 2 than Condition 1. All four human-robot teams in Condition 2 needed more time than any of the human-robot teams in Condition 1. Surprisingly, three of the four participants in Condition 2 chose the non-optimal solution (human completes both Legos and the robot "Burra fetches the parts). Only one participant chose the optimal solution and his/her time needed to complete the task was the least in his/her group, although still higher than all four teams in Condition 1. These results indicate that making decisions takes time. Even in the case when workers chose the optimal solution, the time needed to complete the tasks was higher, possibly due to additional time for decision making.

B. Evaluation Human Appreciation of the System

No statistically significant differences were observed in worker appreciation of the system contrary to what was hypothesized. The three participants from Condition 2 that chose the non-optimal solution present the worst ratings among all participants in the experiment. Interestingly, the participant from Condition 2 that chose the optimal-condition presents the highest rating, even above the ratings from participants in Condition 1. These results lend support to the hypothesis that human satisfaction is most affected by an inherent sense of efficiency while freedom of choice plays a secondary role. In our pilot study, high efficiency with freedom of choice yields the highest satisfaction, while high efficiency without freedom was second. When the participant is given freedom of choice, resulting in low efficiency, human satisfaction appears to decrease drastically.

The lower average scores in Condition 2 could alternatively be attributed to the disruption and difficulty the robot malfunction caused the human participant. From the human's perspective, deciding under time pressure amongst a set of options for how to resolve the problem may be less preferable than having an automated algorithm resolve the problem. To understand the true factors that affect both performance and human appreciation of the system, we will conduct a full experiment based on the observations from this pilot study.

VII. Conclusions and Future Work

We aim to understand how best to integrate robots to work in coordination with a human workforce. In the manufacturing domain, human workers are often provided some flexibility in decision-making for how to execute their work. As a means of integrating robots into human teams, we have developed an algorithm that takes a centralized approach to producing agent-task assignments and schedules. However, concerns exist that humans may not appreciate or may even reject a system that does not allow them enough flexibility in how to do their work. The pilot study we have conducted is a first step towards understanding how much control over decision-making a human worker should be provided.

We varied the amount of control that our participants had over the task assignments in the human-robot team. Results from the study supported our first hypothesis; giving the human workers more control decreased temporal efficiency. Our second hypothesis stated that worker appreciation of the technology would benefit from providing the human with some control over the decision-making. This hypothesis was not supported with statistical testing. However, we did find a trend where workers with more control who chose the optimal solution were the most satisfied, and workers with more control who chose the suboptimal solution were the least satisfied.

Pilot studies are designed to provide guidance on how to design a follow-on large scale experiment. While our pilot provides initial results and trends, these results were obtained through a small scale experiment and are not sufficient to provide recommendations on algorithm and interface design. Based on this pilot study, we plan on running a full human subject experiment with a number of changes: 1) the roles of the robots will be redesigned so that they are more highly valued by the human participants (e.g., having the fetching task be more realistic) 2) experimenter interference will be reduced by removing tethered power, and 3) the experiment will be redesigned to better isolate the dependent variable of worker satisfaction from confounding factors (e.g., uncoordinated timing between the fetching robot and the human).

References

[1] Michael J. Barnes, Jessie Y.C. Chen, Florian Jentsch, and Elizabeth S. Redden. Designing effective soldier-robot teams in complex environments: training, interfaces, and individual differences. In *EPCE*, pages 484–493. Springer, 2011.

[2] Jennifer Casper and Robin Roberson Murphy. Human-robot interaction in rescue robotics. *IEEE Transaction on Systems, Man, and Cybernetics (SMCS)*, 34(2):138–153, 2004.

[3] Andrew S Clare, Mary (Missy) L. Cummings, Jonathan P. How, Andrew K. Whitten, and Olivier Toupet. Operator objective function guidance for a real-time unmanned vehicle scheduling algorithm. 9:161–173, 2012.

[4] Mary (Missy) L. Cummings, Amy S. Brzezinski, and John D. Lee. Operator performance and intelligent aiding in unmanned aerial vehicle scheduling. *IEEE Intelligent Systems*, 22(2):52–59, March 2007.

[5] Edmund H. Durfee, James C. Boerkoel Jr., and Jason Sleight. Using hybrid scheduling for the semi-autonomous formation of expert teams. *Future Generation Computer Systems*, July 2013.

[6] Matthew C. Gombolay, Ronald J. Wilcox, and Julie A. Shah. Fast scheduling of multi-robot teams with temporospatial constrints. In *Proc. Robots: Science and Systems (RSS)*, 2013.

[7] Michael A. Goodrich, Bryan S. Morse, Cameron Engh, Joseph L. Cooper, and Julie A. Adams. Towards using UAVs in wilderness search and rescue: Lessons from field trials. *Interaction Studies, Special Issue on Robots in the Wild: Exploring Human-Robot Interaction in Naturalistic Environments*, 10(3):453–478, 2009.

[8] Henry L. Jones, Stephen M. Rock, Dennis Burns, and Steve Morris. Autonomous robots in SWAT applications: Research, design, and operations challenges. *AUVSI*, 2002.

[9] Raja Parasuraman and Dietrich H. Manzey. Complacency and bias in human use of automation: An attentional integration. 52:381–410, 2010.

[10] Raja Parasuraman, Mouloua Mustapha, and Brian Hilburn. Adaptive aiding and adpative task allocation enchance human-machine systems.