Coordination of Human-Robot Teaming with Human Task Preferences

Matthew C. Gombolay, Cindy Huang, and **Julie A. Shah**

Massachusetts Institute of Technology

77 Massachusetts Avenue

Cambridge, Massachusetts 02139

gombolay@csail.mit.edu, cinhuang@mit.edu, julie_a_shah@csail.mit.edu

Abstract

Advanced robotic technology is opening up the possibility of integrating robots into the human workspace to improve productivity and decrease the strain of repetitive, arduous physical tasks currently performed by human workers. However, coordinating these teams is a challenging problem. We must understand how decision-making authority over scheduling decisions should be shared between team members and how the preferences of the team members should be included. We report the results of a human-subject experiment investigating how a robotic teammate should best incorporate the preferences of human teammates into the team's schedule. We find that humans would rather work with a robotic teammate that accounts for their preferences, but this desire might be mitigated if their preferences come at the expense of team efficiency.

Introduction

Human-robot teaming offers promise to increase the productivity of human labor and improve the ergonomics of manual tasks. However, the tight choreography required to safely and efficiently coordinate human-robot teams in time and space is a challenging computational problem. Task allocation and sequencing with upper and lowerbound temporal constraints is known to be NP-Hard (Bertsimas and Weismantel 2005). Fully autonomous solutions to the problem have been proposed by researchers in academia and by industry practitioners (Alsever 2011; Bertsimas and Weismantel 2005; Gombolay, Wilcox, and Shah 2013). However, these solutions require that a human has fully specified the relevant constraints and optimization criteria for the joint human-robot schedule. The interface between human and robotic agents has been long identified as the key bottleneck in the utilization of these advanced robotic systems (Casper and Murphy 2004). As a result, human factors researchers have sought to design supervisory control interfaces to bring the human into the decision-making loop to improve schedule quality and ease the burden of manually coding task specifications (Adams 2009; Barnes et al. 2011; Chen, Barnes, and Qu 2010; Cummings, Brzezinski, and Lee 2007; Goodrich et al. 2009; Jones et al. 2002; Hooten, Hayes, and Adams 2011). Some researchers have focused

on creating interfaces that can solicit feedback in the form of quantitative, qualitative, hard or soft constraints over various scheduling options that could be fed into a scheduling algorithm to generate candidate schedules (Ardissono et al. 2012; Clare et al. 2012; Hamasaki et al. 2004; Haynes et al. 1997; Macho, Torrens, and Faltings 2000; Zhang et al. 2012). Yet, the role of humans in collaborative teaming and joint-action is less well-studied from the point of view of control authority (Gombolay et al. 2014).

We have been conducting a series of experiments to determine how to best insert the human into the decision-making loop as a member of a human-robot team (Gombolay et al. 2013; 2014; 2015). We have discovered that human workers have a strong preference to give more control to robot teammates over scheduling decisions when that robotic teammate is able to generate more efficient schedules (Gombolay et al. 2014). We also found that human subjects tend to alter the way they schedule team activities depending on whether they were responsible for allocating work to the entire team or only themselves. Subjects tend to decouple their work from the rest of the team when they have control over only which tasks they will perform. On the other hand, when subjects can allocate tasks to the entire team, subjects were more willing to perform tasks that were dependent upon other team members completing prerequisite tasks (Gombolay et al. 2014).

Next, we sought to determine how the dynamics of decision-making authority would change as a function of team composition. Specifically, we sought to answer whether subjects would desire a different level of control over scheduling decisions depending on whether the human subject was a member of a mixed human-robot team or if the subject was a member of a strictly human team. We discovered a key difference for how subjects' perception of decision-making authority changes with team composition (Gombolay et al. 2015). Subjects inherently view human teammates as more intelligent and capable regardless of whether that teammate is responsible for scheduling decisions. However, the subjects' perception of the intelligence, value, and contribution of a robot teammate greatly increases if the robot is responsible for scheduling decisions (Gombolay et al. 2015).

Based on this prior work, we posit that robotic teammates can improve the productivity of human-robot teams and gar-
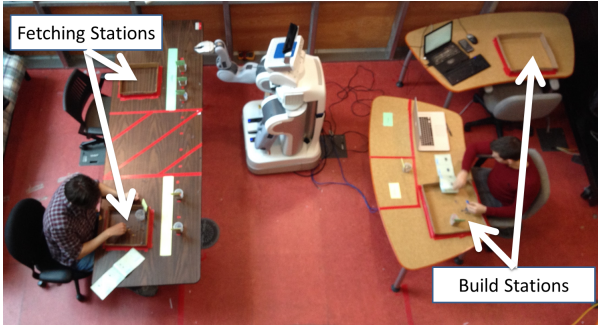
Figure 1: This figure depicts a diagram of the laboratory room where the experiment took place. There are two locations where the human and robot workers can inspect part kits during a fetching task, and two locations where the human workers can build the part kits.

ner the appreciation of human-workers. However, the robot must also understand the human factors aspects of how humans wish to perform their tasks. While human teammates may naturally develop an understanding of how each member of his or her team prefers to do individual tasks, a robot operator would need to manually encode such preferences into a scheduling algorithm (or rely on more sophisticated machine learning techniques) for a robot to consider these preferences. If the robot can produce a schedule that has a minimum makespan (i.e., shortest overall duration) yet does not consider the preferences of individual workers (e.g., a given worker may prefer one type of task over another), workers may be resistant to working with the robotic teammate.

In this paper, we report the results of an initial human-subject experiment ($n = 17$) investigating the effect of a robotic teammate including or ignoring the preferences of human subjects for which types of tasks they would prefer to perform. Specifically, we hypothesized that subjects would more highly rate their subjective experience working on a human-robot team with a robot that allocated tasks to the team members based on the preferences of the subjects for which tasks they would prefer to complete. At the conclusion of our pilot study, we found that the incorporation of preferences for task types into the scheduling processes is important. However, the robotic teammate must also balance these task-based preferences with human idle time and team efficiency. To better ascertain the role of preferences in the scheduling of human-robot teams relative to human idle time and team efficiency, we propose a follow-on experiment in future work.

We begin in with a formal definition of the scheduling problem of interest. Next, we describe our experimental design to study the role of scheduling with and without the preferences of the human team members. We then present the results of statistical testing for both objective and subjective measures of team fluency. Finally, we discuss the implications of our findings and propose a follow-on experiment for future work.

Formal Problem Definition

We begin with a formal definition of the problem of scheduling a team of heterogeneous agents to complete a set of tasks with upper and lowerbound temporal constraints and shared resources (e.g., spatial locations), as shown in Equations 1-13.

$$\min z, z = g\left(\{A^a_{\tau^j_i}|\tau_i \in \boldsymbol{\tau}, a \in A\}, \right.$$
$$\left. \{J_{\langle \tau^j_i, \tau^y_x\rangle}|\tau^j_i, \tau^y_x \in \boldsymbol{\tau}\}, \{s_{\tau^j_i}, f_{\tau^j_i}|\tau^j_i \in \boldsymbol{\tau}\}\right) \quad (1)$$

subject to

$$\sum_{a \in A} A^a_{\tau^j_i} = 1, \forall \tau^j_i \in \boldsymbol{\tau} \quad (2)$$

$$ub_{\tau^j_i} \geq f_{\tau^j_i} - s_{\tau^j_i} \geq lb_{\tau^j_i}, \forall \tau^j_i \in \boldsymbol{\tau} \quad (3)$$

$$f_{\tau^j_i} - s_{\tau^j_i} \geq lb^a_{\tau^j_i} - M\left(1 - A^a_{\tau^j_i}\right), \forall \tau^j_i \in \boldsymbol{\tau}, a \in A \quad (4)$$

$$s_{\tau^y_x} - f_{\tau^j_i} \geq W_{\langle \tau^j_i, \tau^y_x\rangle}, \forall \tau^j_i, \tau^y_x \in \boldsymbol{\tau}|, \forall W_{\langle \tau^j_i, \tau^y_x\rangle} \in \boldsymbol{TC} \quad (5)$$

$$f_{\tau^y_x} - s_{\tau^j_i} \leq D^{rel}_{\langle \tau_i, \tau_j\rangle}, \forall \tau^j_i, \tau^y_x \in \boldsymbol{\tau}|\exists D^{rel}_{\langle \tau^j_i, \tau^y_x\rangle} \in \boldsymbol{TC} \quad (6)$$

$$f_{\tau^j_i} \leq D^{abs}_{\tau_i}, \forall \tau_i \in \boldsymbol{\tau}|\exists D^{abs}_{\tau_i} \in \boldsymbol{TC} \quad (7)$$

$$s_{\tau^y_x} - f_{\tau^j_i} \geq M\left(A^a_{\tau^j_i} + A^a_{\tau^y_x} - 2\right)$$
$$+ M\left(J_{\langle \tau^j_i, \tau^y_x\rangle} - 1\right), \forall \tau_i, \tau_j \in \boldsymbol{\tau}, \forall a \in A \quad (8)$$

$$s_{\tau^j_i} - f_{\tau^y_x} \geq M\left(A^a_{\tau^j_i} + A^a_{\tau^y_x} - 2\right)$$
$$- M\left(J_{\langle \tau^j_i, \tau^y_x\rangle}\right), \forall \tau_i, \tau_j \in \boldsymbol{\tau}, \forall a \in A \quad (9)$$

$$s_{\tau^y_x} - f_{\tau^j_i} \geq M\left(J_{\langle \tau^j_i, \tau^y_x\rangle} - 1\right), \quad (10)$$
$$\forall \tau_i, \tau_j \in \boldsymbol{\tau}|R_{\tau^j_i} = R_{\tau^y_x} \quad (11)$$

$$s_{\tau^j_i} - f_{\tau^y_x} \geq -M\left(J_{\langle \tau^j_i, \tau^y_x\rangle}\right), \quad (12)$$
$$\forall \tau_i, \tau_j \in \boldsymbol{\tau}|R_{\tau^j_i} = R_{\tau^y_x} \quad (13)$$

In this formulation, $A^a_{\tau^j_i} \in \{0, 1\}$ is a binary decision variable for the assignment of agent a to task τ_i, $J_{\langle \tau^j_i, \tau^y_x\rangle} \in \{0, 1\}$ is a binary decision variable specifying whether τ_i comes after or before τ_j, and $s_{\tau^j_i}, f_{\tau^j_i} \in [0, \infty)$ are the start and finish times of τ_i. $\boldsymbol{TC}$ is the set of simple temporal constraints relating task events. Equation 1 is a general objective that is a function of the decision variables $\{A^a_{\tau^j_i}|\tau^j_i \in \boldsymbol{\tau}, a \in A\}$, $\{J_{\langle \tau^j_i, \tau^y_x\rangle}|\tau_i, \tau_j \in \boldsymbol{\tau}\}$, and $\{s_{\tau^j_i}, f_{\tau^j_i}|\tau^j_i \in \boldsymbol{\tau}\}$. Equation 2 ensures that each task is assigned to a single agent. Equation 3 ensures that the duration of each $\tau_i \in \boldsymbol{\tau}$ does not exceed its upper and lowerbound durations. Equation 4 requires that the duration of task τ^j_i, $f_{\tau^j_i} - s_{\tau^j_i}$, is no less than the time required for agent a to complete task τ_i. Equation 5

requires that τ_x^y occurs at least $W_{\langle \tau_i^j, \tau_x^y \rangle}$ units of time after τ_i^j. Equation 6 requires that the duration between the start of τ_i^j and the finish of τ_x^y is less than $D_{\langle \tau_i^j, \tau_x^y \rangle}^{rel}$. Equation 7 requires that τ_i^j finishes before $D_{\tau_i^j}^{abs}$ units of time have expired since the start of the schedule. Equations 8-9 enforce that agents can only execute one task at a time. Equations 11-13 enforce that each resource R_i can only be accessed one agent at a time.

The worst-case time complexity of a complete solution technique for this problem is dominated by the binary decision variables for allocating tasks to agents ($A_{\tau_i^j}^a$) and sequencing ($J_{\langle \tau_i^j, \tau_x^y \rangle}$), and the complexity is given by $O\left(2^{|A||\tau|^3}\right)$, where $|A|$ is the number of agents and $|\tau|$ is the number of tasks. Agent allocation contributes $O\left(2^{|A||\tau|}\right)$, and sequencing contributes $O\left(2^{|\tau|^2}\right)$.

In our work, we focus our investigation on the effect of incorporating the preferences of human team members when generating the team's schedule. Preferences can come in a variety of forms. For example, humans may have preferences for the duration of events (e.g., how long it takes to complete a task) or the duration between events (e.g., the lowerbound or upperbound on the time between two tasks) (Wilcox, Nikolaidis, and Shah 2012). In our investigation, we consider preferences for types of tasks. For example, a worker may prefer to complete a drilling task rather than a painting task. These preferences can be included in the mathematical formulation in Equations 1-13 as an objective function term where one seeks to maximize the number of preferred-tasks assigned to the subject as shown in Equation 14). Alternatively, one could incorporate preferences as a set of constraints to enforcing a minimum or maximum level of preferred work assigned to the subject as shown in Equations 15-16). In these equations, k_{lb} is a lowerbound on the preferred task time allocated to the subject and k_{ub} is an upperbound on the cumulative duration of non-preferred tasks assigned to the subject.

We chose to model the inclusion of preferences as a set of constraints guaranteeing that subjects perform at most one task of the type they do not prefer (Equation 16). For the purpose of human subject experimentation where one must control for confounds, this approach offered greater control over schedule content as opposed to including a preference term in the objective function. The challenge with using an objective function model is that one must tune one or more coefficients (e.g., α in Equation 14) in the objective function to trade off the contribution of the schedule efficiency (i.e., makespan) with the importance of adhering to preferences. In practice, we found this tuning to be difficult across a variety of subjects with differing task completion times, $lb_{\tau_i^j}$.

$$\max z, z = \alpha \times g\left(\{A_{\tau_i^j}^a | \tau_i \in \boldsymbol{\tau}, a \in A\},\right.$$

$$\{J_{\langle \tau_i^j, \tau_x^y \rangle} | \tau_i^j, \tau_x^y \in \boldsymbol{\tau}\}, \{s_{\tau_i^j}, f_{\tau_i^j} | \tau_i^j \in \boldsymbol{\tau}\}\right)$$

$$+ (1 - \alpha) \times \left(\sum_{\tau_i^j \in \boldsymbol{\tau}_{preferred}} A_{\tau_i^j}^{subject} \times lb_{\tau_i^j}\right) \quad (14)$$

$$k_{lb} \leq \sum_{\tau_i^j \in \boldsymbol{\tau}_{preferred}} A_{\tau_i^j}^{subject} \times lb_{\tau_i^j}^a \quad (15)$$

$$k_{ub} \geq \sum_{\tau_i^j \in \boldsymbol{\tau}_{preferred}^c} A_{\tau_i^j}^{subject} \times lb_{\tau_i^j}^a \quad (16)$$

Experimental Design

We conducted a human-subject experiment ($n = 17$) to study how a robotic teammate's inclusion or ignorance of the preferences of human teammates over scheduling decisions affects the dynamics of the team. Our human-robot manufacturing team consisted of the human subject, a robotic assistant, and a human assistant. The human subject was capable of both fetching and building, and the robot assistant was only capable of fetching. One of the experimenters played the role of a third teammate for all subjects and was capable of both fetching and building. The third human teammate was included to more realistically represent the composition of a human-robot team in a manufacturing setting. We used a Willow Garage PR2 platform, as shown in Figure 1, as the robotic assistant for our human-robot team. The robot used Adaptive Monte Carlo Localization (AMCL) (Fox 2003) and the standard *Gmapping* package in the Robot Operating System (ROS) for navigation.

In our scenario, there are two types of tasks: fetching part kits and assembling part kits. Fetching a part kit required walking to one of two inspection stations where the kits were located, inspecting the part kit and carrying it to the build area. The architecture of our fetching task is analogous to what is required in many manufacturing domains: to adhere to strict quality assurance standards, fetching a part kit requires verification from one to two people that all off the correct parts are in the kit, and certification from another person that the kit has been verified.

We imposed a set of additional constraints to mimic an assembly manufacturing environment. A part kit must have been fetched before it could be built, and no two agents were able to occupy the same fetching or build station at the same time. As shown in Figure 1, there were two fetching and two build stations. Four part kits were located at one fetching station, and four kits were located at the second fetching station.

Agents were required to take turns using the fetching stations. Allowing workers to sort through parts from multiple kits at the same location risked mixing the wrong part with the wrong kit. Furthermore, in manufacturing, if a part or part kit is missing from an expected location for too long, work in that area of the factory will temporarily halt until

the missing pieces are found. As such, we imposed a 10-minute deadline from the time that the fetching of a part kit began until that part kit had been built, for similar reasons.

Assembly of the Lego model involved eight tasks $\tau = \{\tau_1, \tau_2, \ldots, \tau_8\}$, each of which was composed of a *fetch* and *build* subtask $\tau_i = \{\tau_i^{fetch}, \tau_i^{build}\}$. The time each subject took to complete each subtask $C_i^{subject-fetch}$ and $C_i^{subject-build}$ was measured during an experiment training round. The timings for the robot $C_i^{robot-fetch}$ and human assistant $C_i^{assist-fetch}$ and $C_i^{assist-build}$ (performed by an experimenter) were collected prior to the experiments.

At the beginning of the experiment, subjects were told that the robot wanted to know which tasks subjects preferred to complete: fetch tasks or build tasks. Subjects were then treated to three experimental conditions in a within-subjects experimental design:

- *Positive* - The robot would generate a schedule incorporating the preferences of the subject.

- *Neutral* - The robot would ignore the preferences of the subject.

- *Negative* - The robot would schedule the team as if the preferences of the subject were opposite (e.g., subjects preferring to build would be scheduled as if they preferred fetching).

Subjects were not informed *a priori* of the different conditions. As such, subjective evaluations of the team dynamics in each condition would not be biased by an expectation of the robot catering to the subjects' preferences or not. We established the following hypothesis:

Hypothesis 1: Subjects would rather work with a robotic teammate that includes their scheduling preferences than one that is unaware of their preferences, and subjects would rather work with a robotic teammate ignorant to their preferences than one that actively schedules against their preferences.

To schedule the human-robot team, we adapted a dynamic scheduling algorithm called Tercio (Gombolay, Wilcox, and Shah 2013). To handle the preferences of the subject, we added a constraint into the Tercio task allocation formulation as shown in Equation 16. In the positive condition, subjects could be assigned only one task that did not align with their preferences. For example, subjects preferring to build could be assigned at most one fetching task (and vice versa). In the negative condition, subjects could be assigned a maximum of one task that aligned with their preferences. For example, subjects preferring to build could be assigned at most one build task (and vice versa). In the neutral condition, Tercio's task allocation subroutine would run without alteration.

Based on our previous studies showing the importance of team efficiency (Gombolay et al. 2014; 2015), we sought to control for how schedule duration would affect how subjects perceived the team dynamics. As such, we ran Tercio until an approximately equal makespan schedule would be generated for all three conditions.

Table 1: Subjective Measures - Post-Trial Questionnaire

Robot Teammate Traits
1. The robot was intelligent.
2. The robot was trustworthy.
3. The robot was committed to the task.
Working Alliance for Human-Robot Teams
4. I feel uncomfortable with the robot. (reverse scale)
5. The robot and I understand each other.
6. I believe the robot likes me.
7. The robot and I respect each other.
8. I feel that the robot worker appreciates me.
9. The robot worker and I trust each other.
10. The robot worker perceives accurately what my goals are.
11. The robot worker does not understand what I am trying to accomplish.
12. The robot worker and I are working towards mutually agreed upon goals.
13. I find what I am doing with the robot worker confusing. (reverse scale)
Additional Measures of Team Fluency
14. I was satisfied by the team's performance.
15. I would work with the robot the next time the tasks were to be completed.
16. The robot increased the productivity of the team.
17. The team collaborated well together.
18. The team performed the tasks in the least time possible.
19. The robot worker was necessary to the successful completion of the tasks.
20. The human worker was necessary to the successful completion of the tasks.
21. I was necessary to the successful completion of the tasks.

Results

Subjects received post-trial questionnaires after each trial, consisting of 21 Likert-scale questions, as shown in Table 1. Our questionnaire was inspired by the work of Hoffman (Hoffman 2013) and the adaptation of the "Working Alliance Index" for human-robot teams. We added questions 14-21 based on our own insight. Subjects also received a post-test questionnaire after completing the three trials. This questionnaire gathered demographic information, and included three additional Likert-scale questions summarizing the experience of the subjects, as well as two open-ended questions.

We found statistically significant evidence that human subjects would prefer working with a robot that included the subjects' preferences when making scheduling decisions based on questions 22-24 in Table 2 ($p < 0.001$). Subjects would rather work with a robotic teammate that included his or her preferences rather than if the robot was unaware of the subjects' preferences ($p < 0.001$). Furthermore, subjects would rather work with a robot that was unaware of his or her preferences rather than working with a robot that sched-

Table 2: Subjective Measures - Post-Test Questionnaire

| **Overall Preference** |
| --- |
| 22. If the robot scheduled me to do the tasks I preferred, I would want to work with the robot again. |
| 23. If the robot did not know my preferences when scheduling, I would want to work with the robot again. |
| 24. If the robot scheduled me to do different tasks than what I preferred, I would want to work with the robot again. |
| **Open Response Questions** |
| 25. Which of the three scenarios did you prefer and why? |
| 26. If you were going to add a robotic assistant to a manufacturing team, to whom would you give the job of rescheduling the work and why? |

Figure 2: This figure shows the mean and standard error of the amount of work in seconds assigned to the subject by the robotic teammate. Horizontal bars with an asterisk denote statistical significance ($p < 0.01$).

uled opposite of his or her preferences ($p = 0.001$). We also found that subjects felt that the robot liked them more (Question 6) in the neutral condition when the robot was unaware of the subjects' preferences than when the robot scheduled opposite of their the negative condition ($p = 0.048$). These results supports our hypothesis that the preferences of human workers are important for a robotic teammate to include when making scheduling decisions.

However, we also found that the amount of work allocated to subjects had a strong impact on the subjective perception of team dynamics. Subjects felt more strongly that the robot did not understand what the participant was trying to accomplish in the positive condition when the robot included the subjects' preferences than in the negative condition when the robot scheduled opposite of the subjects' preferences ($p = 0.048$). To better understand this result, we report the amount of work the robotic teammate allocated to the subject in each of the conditions in Figure 2. We conducted an analysis of variance to determine that the robot allocated a statistically significantly different amount of work to the subject as a function of how the robot included the subjects' preferences when scheduling the team (ANOVA $F(2, 48) = 5.16, p = 0.009$). Interestingly, we found that subjects were allocated statistically significantly more work as measured in seconds in the negative condition (M= 448, SD= 113) when the robot scheduled the subjects' work opposite of their preferences as opposed to the positive (M= 373, SD= 92), $t(16) = 1.86$, $p = 0.04$, or neutral conditions (M= 345, SD= 82), $t(17) = 2.14$, $p = 0.03$, as measured in seconds. In collecting subjects' preferences for which types of tasks they would rather complete, we found that the vast majority of subjects reported they would rather build Lego part kits than fetch Lego part kits. In the positive condition, subjects would be given a maximum of one fetching task, and, in the negative condition, subjects would be given a maximum of one building task. The third teammate, the human assistant, was typically more proficient at building than the average subject. As such, the optimal allocation of work typically would have the human assistant performing most of the building tasks and the subject supporting
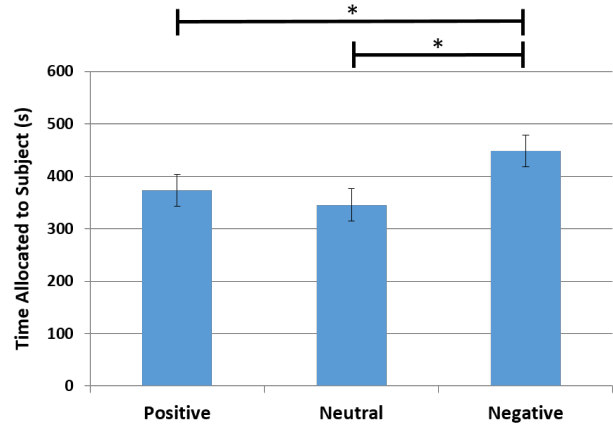
with more fetching tasks. we recall that the robot teammate could only fetch parts kits. As such, the negative condition actually afforded the teams better overall efficiency and would have subjects performing more work. Based on this result, we propose that subjects' preferences for task types need to be balanced with an innate desire of human workers be better utilized.

Contribution and Future Work

We conducted an initial human-subject experiment to better understand the role of the workflow preferences of human workers on a human-robot team. In our pilot study, we found statistically significant results that subjects would rather work with a robotic teammate that included the preferences of the subjects for which tasks they would prefer to complete as opposed to a robot that was unaware of or scheduled opposite of the preferences of the human teammate. However, we also found that a robot prioritizing the preferences of human workers may decrease team efficiency and decrease the human workers' belief that the robot understands the team's objectives. Specifically, subjects' preference for completing certain tasks decreased the team's efficiency and led the subjects to rated their experience more negatively.

To better understand the role of preferences for robotic teammates' scheduling of human-robot co-work, we propose a follow-on study. In this study, we will consider two key variables, team efficiency and the degree to which subjects' workflow preferences are included. By controlling for how team efficiency might decrease if subjects' preferences are counterproductive, we can isolate the individual effects of team efficiency and the inclusion of human workflow preferences for human-robot teaming.

Acknowledgements

References

Adams, J. A. 2009. Multiple robot-single human interaction: effects on perceived workload and performance. *Behavior and Information Technology* 28(2):183–298.

Alsever, J. 2011. Robot workers take over warehouses. *CNN Money*. http://money.cnn.com/2011/11/09/smallbusiness/kiva_robots/.

Ardissono, L.; Petrone, G.; Torta, G.; and Segnan, M. 2012. Mixed-initiative scheduling of tasks in user collaboration. In *Proceedings of the Eighth International Conference on Web Information Systems and Technologies*, 342–351.

Barnes, M. J.; Chen, J. Y. C.; Jentsch, F.; and Redden, E. S. 2011. Designing effective soldier-robot teams in complex environments: training, interfaces, and individual differences. In *Proceedings of the International Conference on Engineering Psychology and Cognitive Ergonomics (EPCE)*, 484–493. Springer.

Bertsimas, D., and Weismantel, R. 2005. *Optimization over Integers*. Belmont: Dynamic Ideas.

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

Chen, J. Y.; Barnes, M. J.; and Qu, Z. 2010. Roboleader: an agent for supervisory control of mobile robots. In *Proceedings of the International Conference on Human-Robot Interaction (HRI)*.

Clare, A. S.; Cummings, M. L.; How, J. P.; Whitten, A. K.; and Toupet, O. 2012. Operator objective function guidance for a real-time unmanned vehicle scheduling algorithm. *AIAA Journal of Aerospace Computing, Information and Communication* 9:161–173.

Cummings, M. L.; Brzezinski, A. S.; and Lee, J. D. 2007. Operator performance and intelligent aiding in unmanned aerial vehicle scheduling. *IEEE Intelligent Systems* 22(2):52–59.

Fox, D. 2003. Adapting the sample size in particle filters through kld-sampling. *International Journal of Robotics Research (IJRR)* 22:985–1003.

Gombolay, M. C.; Wilcox, R. J.; Artiles, A. D.; Yu, F.; and Shah, J. A. 2013. Towards successful coordination of human and robotic work using automated scheduling tools: An initial pilot study. In *Proceedings of Robotics: Science and Systems (RSS) Human-Robot Collaboration Workshop (HRC)*.

Gombolay, M. C.; Gutierrez, R. A.; Sturla, G. F.; and Shah, J. A. 2014. Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. In *Proceedings of the Robots: Science and Systems (RSS)*.

Gombolay, M. C.; Gutierrez, R. A.; Clarke, S. G.; Sturla, G. F.; and Shah, J. A. 2015. Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. *Autonomous Robots* 1–20.

Gombolay, M. C.; Wilcox, R. J.; and Shah, J. A. 2013. Fast scheduling of multi-robot teams with temporospatial constrints. In *Proceedings of the Robots: Science and Systems (RSS)*.

Goodrich, M. A.; Morse, B. S.; Engh, C.; Cooper, J. L.; and Adams, J. A. 2009. 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.

Hamasaki, M.; Takeda, H.; Ohmukai, I.; and Ichise, R. 2004. Scheduling support system for academic conferences based on interpersonal networks. In *Proceedings of ACM Hypertext*.

Haynes, T.; Sen, S.; Arora, N.; and Nadella, R. 1997. An automated meeting scheduling system that utilizes user preferences. In *Proceedings of the First International Conference on Autonomous Agents*, AGENTS '97, 308–315. New York, NY, USA: ACM.

Hoffman, G. 2013. Evaluating fluency in human-robot collaboration. In *International Conference on Human-Robot Interaction (HRI), Workshop on Human Robot Collaboration*.

Hooten, E. R.; Hayes, S. T.; and Adams, J. A. 2011. A comparison of communicative modes for map-based tasking. In *IEEE International Conference on Systems, Man, and Cybernetics*.

Jones, H. L.; Rock, S. M.; Burns, D.; and Morris, S. 2002. Autonomous robots in SWAT applications: Research, design, and operations challenges. *Association for Unmanned Vehicle Systems International*.

Macho, S.; Torrens, M.; and Faltings, B. 2000. A multi-agent recommender system for planning meetings. In *Proceedings of the ACM Conference on Autonomous Agents, Workshop on Agent-based Recommender Systems*.

Wilcox, R. J.; Nikolaidis, S.; and Shah, J. A. 2012. Optimization of temporal dynamics for adaptive human-robot interaction in assembly manufacturing. In *Proceedings of Robotics: Science and Systems (RSS)*.

Zhang, H.; Law, E.; Miller, R.; Gajos, K.; Parkes, D.; and Horvitz, E. 2012. Human computation tasks with global constraints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, 217–226. New York, NY, USA: ACM.