2014 Human Factors *Prize Finalist*

Analyzing the Effects of Human-Aware Motion Planning on Close-Proximity Human–Robot Collaboration

Przemyslaw A. Lasota and **Julie A. Shah**, Massachusetts Institute of Technology, Cambridge, Massachusetts

Objective: The objective of this work was to examine human response to motion-level robot adaptation to determine its effect on team fluency, human satisfaction, and perceived safety and comfort.

Background: The evaluation of human response to adaptive robotic assistants has been limited, particularly in the realm of motion-level adaptation. The lack of true human-in-the-loop evaluation has made it impossible to determine whether such adaptation would lead to efficient and satisfying human–robot interaction.

Method: We conducted an experiment in which participants worked with a robot to perform a collaborative task. Participants worked with an adaptive robot incorporating human-aware motion planning and with a baseline robot using shortest-path motions. Team fluency was evaluated through a set of quantitative metrics, and human satisfaction and perceived safety and comfort were evaluated through questionnaires.

Results: When working with the adaptive robot, participants completed the task 5.57% faster, with 19.9% more concurrent motion, 2.96% less human idle time, 17.3% less robot idle time, and a 15.1% greater separation distance. Questionnaire responses indicated that participants felt safer and more comfortable when working with an adaptive robot and were more satisfied with it as a teammate than with the standard robot.

Conclusion: People respond well to motion-level robot adaptation, and significant benefits can be achieved from its use in terms of both human–robot team fluency and human worker satisfaction.

Application: Our conclusion supports the development of technologies that could be used to implement human-aware motion planning in collaborative robots and the use of this technique for close-proximity human–robot collaboration.

Keywords: human–robot interaction, motion-level adaptation, team fluency, human satisfaction

Address correspondence to Przemyslaw A. Lasota, Massachusetts Institute of Technology, 77 Massachusetts Ave., Cambridge, MA 02139, USA; e-mail: plasota@mit.edu.

HUMAN FACTORS
Vol. 57, No. 1, February 2015, pp. 21–33
DOI: 10.1177/0018720814565188
Copyright © 2015, Human Factors and Ergonomics Society.

INTRODUCTION

In many domains today, robots are often deployed in complete isolation from humans. Although the physical separation of people and robots can be an effective strategy for some applications, a lack of human–robot integration prevents robots from being used in domains that stand to benefit from robotic assistance. The final assembly of aircraft and automobiles, for example, is still mostly a manual operation, involving minimal use of automation (D. Amirehteshami, Boeing Research and Technology, personal communication, July 2011; S. Bartscher, Innovation Product Division, BMW Group, personal communication, February 2012). Although the capabilities of robotic systems continue to expand, many tasks in domains such as these require a level of judgment, dexterity, and flexible decision making that surpasses the abilities of current robots, causing tasks to become human dominated. However, there are also many non-value-added tasks within these domains that could be performed by robotic assistants. Allowing robots to collaborate with people in shared workspaces and to perform such tasks has the potential to increase productivity and efficiency, providing a strong incentive for the development of technologies that foster this collaboration.

A significant amount of research has been conducted in recent years in support of this goal, across a variety of complementary domains. The first step toward creating robots that can successfully collaborate with people is allowing the robots to navigate a shared workspace, which requires the development of specialized path-planning algorithms and frameworks designed with the human element in mind (Bellotto, 2012; Chung & Huang; 2010; Kitade, Satake, Kanda, & Imai, 2013; Kruse, Kirsch, Sisbot, & Alami, 2010; Trautman & Krause, 2010; Ziebart et al.,

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

2009). Sisbot and colleagues developed a framework that uses parameters including human–robot separation distance and the field of vision and stance of the human to generate safe and socially acceptable robot paths (Sisbot, Marin-Urias, Alami, & Simeon, 2007). An extension of this framework also reasons on task constraints and human kinematics (Sisbot, Marin-Urias, Broquere, Sidobre, & Alami, 2010). To further enhance human–robot co-navigation, Chung and Huang (2011) investigated the development of a predictive path-planning framework based on pedestrian motion models trained via statistical learning methods. Bennewitz, Burgard, Cielniak, and Thrun (2005) showed that a combination of clustering and hidden Markov models could be used to learn and predict human motion patterns, allowing for improved co-navigation.

Prediction of human actions and trajectories has been studied with a variety of approaches (Hawkins, Vo, Bansal, & Bobick, 2013; Hoffman & Breazeal, 2007a; Koppula & Saxena, 2013; Luber, Stork, Tipaldi, & Arras, 2010; Mori et al., 2006; Thompson, Horiuchi, & Kagami, 2009). When it comes to prediction on the task level, results from prior work have indicated that the observation of changes to the entropy rate of a Markov chain produced from a task description encoded as a Markov decision process could be used to encode the uncertainty of a robot about what action a human will perform next (Nikolaidis et al., 2013). Lenz et al. (2008) used a high-level architecture for joint human–robot action that could predict human tasks based on knowledge databases and decision processes. Alami et al. (2006) approached this problem by encoding discrete sets of human and robot actions, allowing for the incorporation of task-specific rules and preferences that could then be used to predict likely sequences of human actions. Kwon and Suh (2012) showed how Bayesian networks could be used to perform simultaneous inference on temporal and causal information, allowing for the determination of what task a robot should take and when it should perform that task. Hoffman and Breazeal (2007b), on the other hand, used a formulation based on a first-order Markov process and showed that it could be successfully used for anticipatory action selection.

Prior work has indicated that predictions about human actions are possible without the use of specific task models, through observation of motion features of the human worker. Mainprice and Berenson (2013), for example, showed that early stages of human motion can be assessed to predict what action a human is taking in a reaching task and that these predictions can subsequently be used to select tasks that would avoid motion conflict. Also, work by Doshi and Trivedi (2009) indicated that head motion can be a useful factor for predicting human actions in the domain of advanced driver assistance systems, supporting the results found earlier by Pentland and Liu (1999).

Conversely, enabling a human to predict robot actions is also important. The legibility of robot trajectories—meaning their successful portrayal of the intent of the robot—is often cited as of key importance for fluid human–robot interaction (Dragan, Lee, & Srinivasa, 2013; Dragan & Srinivasa, 2013; Sisbot et al., 2010, 2007). However, Dragan et al. (2013) observed that legible robot trajectories are often not the same as what a person would predict, making legibility and predictability contradictory properties of robot motion. Indeed, there must be a careful balance between these two properties, as the benefits of legibility suffer when robots move beyond a "trust region" of expectation (Dragan & Srinivasa, 2013). Legibility has also been studied in the context of human–robot co-navigation (Kruse, Basili, Glasauer, & Kirsch, 2012).

The ability to predict what action a robot will take is just one of many aspects to consider when measuring the effects robotic assistants have on humans and the implications of these effects on the quality of interaction. Arai, Kato, and Fujita (2010) indicated that several parameters, including separation distance, end effector speed, and advance notice of robot motion, have a significant effect on mental strain among human operators. Results from an experiment by Meisner, Isler, and Trinkle (2008) evaluating a robot controller designed for human-friendly trajectories suggested that the prevention of collisions between humans and robots during a co-navigation task is not sufficient to maintain human comfort. The effects of robotic assistants on people were also assessed in work by Unhelkar, Siu, and Shah (2014), who showed that there are significant differences between human–robot and human–human team dynamics, as measured by interaction

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

and idle times and subjective evaluations of team fluency, situational awareness, comfort, and safety. Hoffman and Breazeal (2007b) observed that human–robot teams in which participants worked with robots that anticipated their actions were able to perform a task more efficiently and had more favorable perceptions of the contribution of the robot to their team's success than did those who worked with a nonadaptive robot.

Objective

The majority of prior work aimed at bringing humans and robots closer together and allowing for close-proximity interaction described previously focuses on the development of frameworks or algorithms without evaluating human response to these technologies. The developed systems were, for the most part, evaluated either only in simulation or with very limited experimentation with actual human participants. However, without fully evaluating human response to adaptive robotic assistants through human subject experimentation, it is impossible to predict whether these technologies would improve team efficiency or human satisfaction. It is possible, for example, that the decreased predictability of an adaptive robot could cause human workers to trust it less than a preprogrammed robot, leading to decreased team efficiency.

The works mentioned earlier that evaluated human response to robotic assistants either did not deal with adaptive systems at all or only considered adaptation at the task level. Motion-level robot adaptation, however, is critical for true close-proximity interaction; and, to the authors' knowledge, there has been no work thus far aimed at evaluating the human response to this type of robot adaptation. The objective of this work, therefore, was to develop a robotic assistant capable of motion-level adaptation and to evaluate through human subject experimentation whether this type of adaptation leads to more efficient teamwork and a more satisfied human coworker.

Human-Aware Motion Planning

The specific motion-planning adaptation technique we implemented to investigate these questions is a method we call *human-aware motion planning*. In this technique, the system attempts to predict the next action of a human worker and approximate the portion of the shared workspace the human will be using during this action, based on an appropriate motion model. The system then uses this prediction to modify the robot's motions to avoid this portion of the shared workspace, in an attempt to eliminate motion conflicts.

To illustrate this technique, consider the shared workspace depicted in Figure 1. The left side of the figure indicates a shared workspace in which a human and robot place screws and apply a sealant, respectively. If we can accurately predict that the human worker will place a screw at the third hole from the left, beside the two screws already placed, we can then approximate the portion of the shared workspace that the human worker will use within the next several moments. A simple and effective motion model for this particular task is the approximation of the predicted workspace occupancy of the human via a cylinder that encompasses the worker's arm as he or she completes the placing task, as this is the only part of the human the robot can reach. This cylinder is indicated in the virtual representation of the workspace depicted on the right side of Figure 1. Once the robot has made a prediction of the workspace occupancy of the human, it can adapt its motion planning by selecting a path to its own goal—in this case, the second screw from the left—that avoids the area that the model predicts the human will occupy. This human-aware motion is depicted in the figure as a solid arrow, and the simple, shortest-path motion the robot would have otherwise taken is shown as a dashed arrow.

METHOD

To investigate human response to robot adaptation at the motion-planning level, we devised a human subject experiment in which participants worked cooperatively with a robot to perform a collaborative task within a shared workspace. The particular robot used in our experiment was the ABB IRB-120, shown in Figure 1, which is a standard 6-degrees-of-freedom industrial robot. A PhaseSpace motion capture system was used to track the human within the workspace and to detect the actions he or she performed. A real-time safety system was deployed on the robot to adjust speed as necessary according to the separation distance between the human and robot,

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

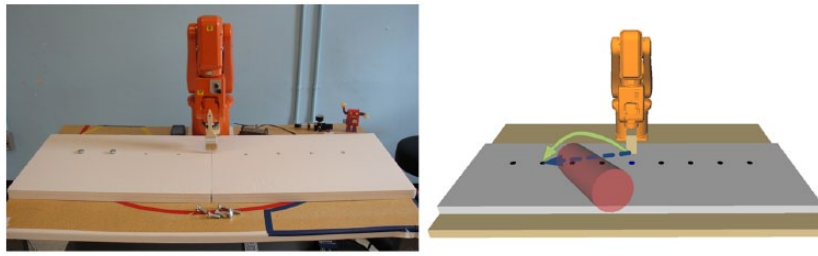

Figure 1. Illustration of human-aware motion planning. The left panel depicts a shared workspace in which a human and robot are placing and sealing screws, respectively. The right panel depicts both the standard, shortest-path motion (dashed arrow) and a human-aware motion (solid arrow) that the robot could take given the expected human workspace occupancy, represented by the cylinder.

with the ability to gradually decrease the robot's speed to a complete stop (see Lasota, Rossano, & Shah, 2014, for details).

The robot was programmed such that it could operate in two motion-planning modes: standard and human aware. A database of robot motions was computed offline for every possible combination of human and robot actions using the constrained bidirectional rapidly exploring random tree (CBiRRT) algorithm (Berenson, Srinivasa, & Kuffner, 2011). The decision to generate robot motion plans offline was made to ensure that the robot motions were consistent throughout the experiment, as the CBiRRT algorithm is based on the rapidly exploring random tree algorithm, which inherently may produce different motions each time. Both human-aware and standard motions were planned with the CBiRRT algorithm.

Task

The task used in the experiment is depicted in Figure 2. During this task, human participants placed eight screws at designated locations on a table, while the robot simulated the application of a sealant to each screw by dipping a brush in a centrally positioned container and then moving the brush over the screw. Each participant was instructed to twist the screw one full rotation before moving on to the next one. The task, from the placement of the first screw to the sealant of the last, took 52 s on average.

The screws were placed by the participants in a predefined order. This, in effect, simulated perfect human action prediction and allowed us to focus on measuring the effects of the motion

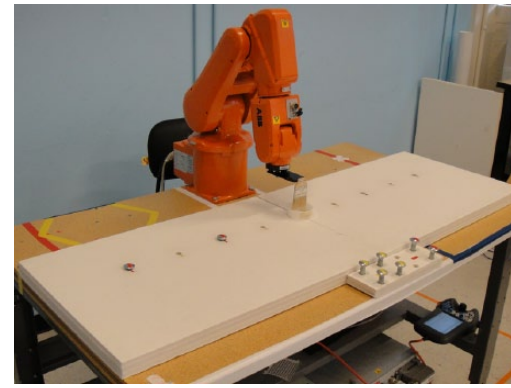

Figure 2. Photograph of the task setup used in the human subject experiments.

adaptation independent of the accuracy of action prediction. The screws were colored either yellow or red. The participants first placed the four red screws and were then instructed to wait for a sound cue before proceeding to the yellow set of screws. Splitting the screw placement in this manner allowed us to control the types of motion conflicts the subject would experience by preventing the participant from working too quickly and losing synchronization with the robot.

Numbering the screws from left to right, the screws were placed according to the following sequence: 1, 3, 8, 6, 2, 4, 7, 5. This sequence was selected in order to balance between conflicting and nonconflicting motion.

Participants

Our participant pool consisted of 20 Massachusetts Institute of Technology affiliates,

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

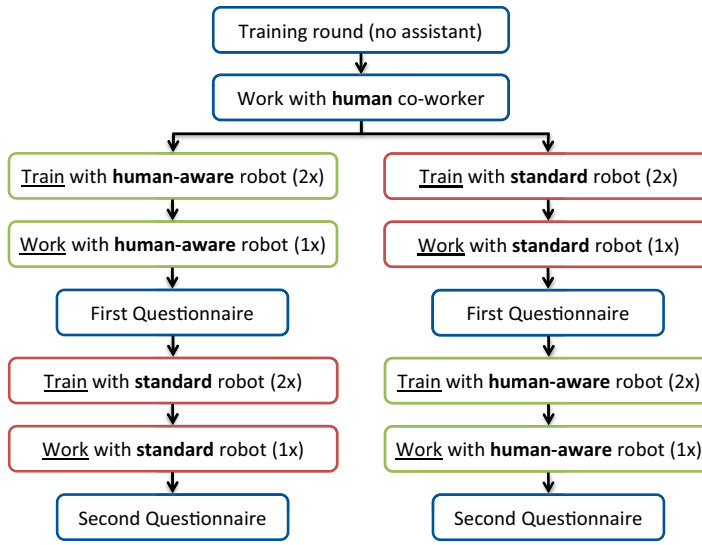

Figure 3. Diagram depicting the experimental procedure. The group on the left represents the "human-aware-first" condition, whereas the group on the right represents the "standard-first" condition.

including undergraduate and graduate students, postdoctoral associates, and visiting students. Of the 20 participants, seven were female and 13 were male, with ages ranging from 19 to 39 years ($M = 27.1$, $SD = 6.24$).

Procedure

Our experiment had a repeated-measure design that consisted of two randomly selected groups of subjects: those who first worked with a human-aware robot ($n = 11$) and those who initially worked with a robot using standard motion planning ($n = 9$). The subjects were not informed prior to completion of the experiment what these two conditions were or what was being measured.

The experiment procedure is depicted in Figure 3. First, both groups executed a training round, during which the participants placed the screws without an assistant applying the sealant, to familiarize themselves with the task. Next, all participants performed the task with a human assistant to implicitly prime the subjects to work with a robot in a manner similar to working with a person during subsequent task executions. To prevent any unintentional bias from the experimenter, who acted as the human coworker, this first task execution was conducted in a double-blind fashion, with the experimenter unaware of which of the two conditions each participant had been assigned to.

After performing the task with a human coworker, the participants then performed the same task with a robotic assistant. Participants under the "human-aware-first" condition worked with a robot in the human-aware motion-planning mode, whereas those under the "standard-first" condition worked with the robot in the standard motion planning mode. Each participant performed two training task executions with the robot to practice working with a robotic assistant and to build a mental model of its behavior. After two training executions, the task was performed one more time, followed by the administration of the first questionnaire.

In the next phase of the experiment, the mode of the robot was switched to the opposite of what each participant had experienced thus far. Prior to continuing, participants were informed that the robot would move differently during the second phase, allowing for anticipation of the change in motion so that the new robot behavior would not startle them. To avoid revealing that the modes were linked to whether the robot adapted its motions or not, the participants were simply told that a few robot motion parameters were changed. After learning of the change, the

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

TABLE 1: Questionnaire Items

Satisfaction with the robot as a teammate
1. I trusted the robot to do the right thing at the right time.
2. The robot did not understand how I wanted to do the task.
3. The robot kept getting in my way.
4. The robot and I worked well together.
Perceived safety and comfort
5. I felt safe when working with the robot.
6. The robot moved too fast for my comfort.
7. The robot came too close to me for my comfort.
8. I trusted the robot would not harm me.

participants performed another set of training rounds with the robot and a final task execution. Upon completion of the final task, participants responded to a second questionnaire that directly compared the robot behavior in the second mode with that of the first mode. A video showing sample task executions for both conditions can be viewed at the following link: http://youtu.be/Dk5XVQBDJpU.

Dependent Measures

The dependent measures we considered when evaluating human response to human-aware motion planning and the possible derived benefits were split into two main groups: quantitative metrics of team fluency and subjective evaluation by the participants. Our first group of metrics was based on those proposed by Hoffman and Breazeal (2007b), expanded to consider additional measures, including task execution time, the amount of concurrent motion, human and robot idle time, and human–robot separation distance.

The second group of metrics was based on subjective evaluation of the robotic assistant according to the questionnaire responses. The questions, shown in Table 1, were intended to determine each participant's satisfaction with the robot as a teammate as well as his or her perceived safety and comfort.

Participants responded to the questions using a 5-point Likert scale ranging from *strongly disagree* to *strongly agree* for the first questionnaire and from *much less* to *much more* for the second questionnaire.

Hypotheses

Based on the dependent measures described in the previous section, the two main hypotheses in this experiment were as follows:

Hypothesis 1: Using human-aware motion planning will lead to more fluent human–robot teamwork, including shorter task execution time, more concurrent motion, shorter human and robot idle times, and a greater separation distance between the human and robot, compared with standard motion planning.

Hypothesis 2: Participants will be more satisfied with the human-aware robot's performance as a coworker and will feel more comfortable and safe when working with it compared with a robot that uses standard motion planning.

RESULTS

Quantitative Team Fluency Metrics

When we compared the participants' performance while working with a human-aware robot to their performance working with a robot using standard motion planning, significant differences were found for all quantitative team fluency metrics. The Shapiro-Wilk test (at $\alpha = .05$) was used to assess the normality of the data. With the exception of human idle time, $W(20) = 0.817$, $p = .002$, the distributions of the differences between the human-aware and standard task executions were not significantly different from normal, minimum $W(20) = 0.943$,

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

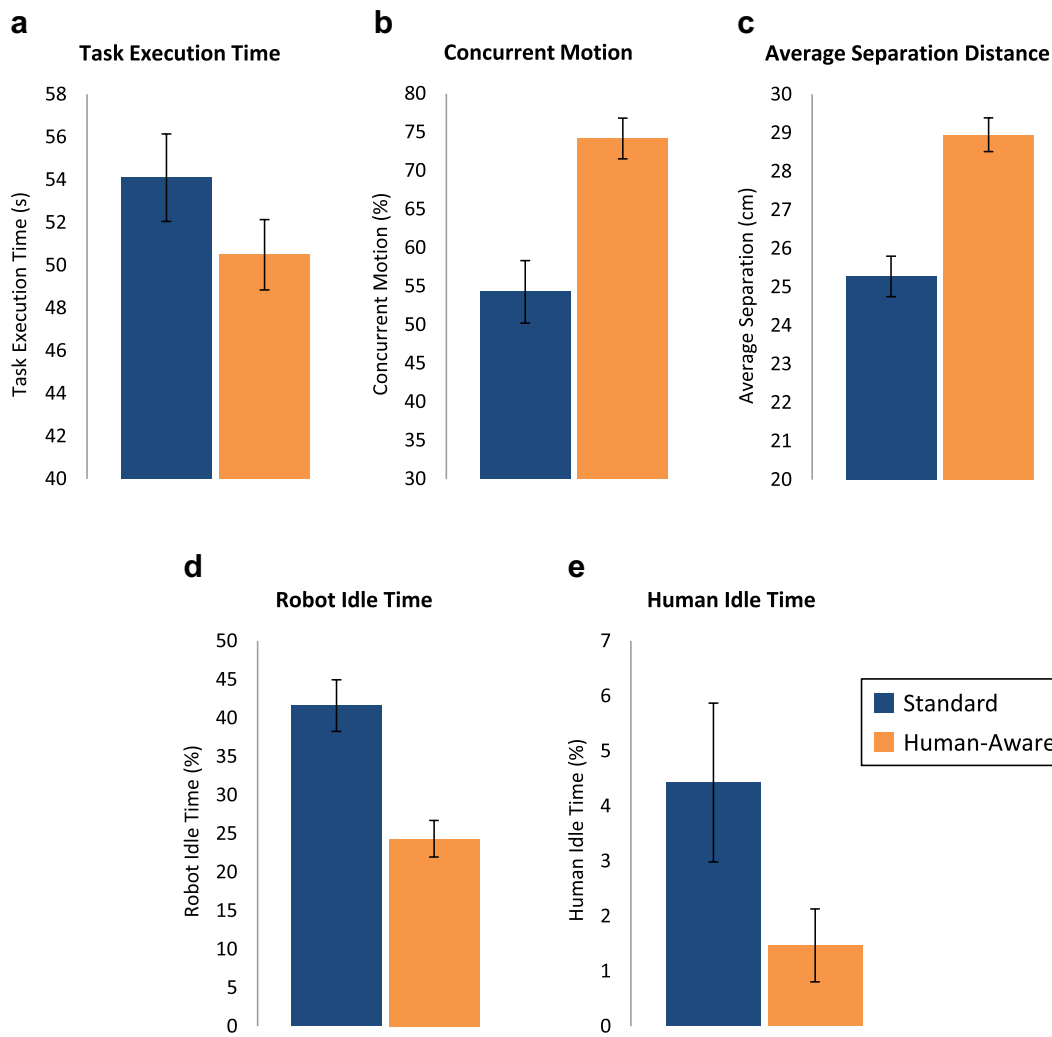

Figure 4. Mean values, with error bars indicating standard error of the mean (SEM), of (a) task execution time, (b) percentage of concurrent motion, (c) average separation distance between the human and robot, (d) robot idle time, and (e) human idle time for the standard and human-aware robot executions.

p = .274. One-way repeated-measures ANOVAs were used where appropriate. Human idle time was instead analyzed with the Wilcoxon signed-rank test. The results of the statistical tests revealed that when working with a human-aware robot, participants completed the task 5.57% faster, $F(1, 19) = 4.95$, $p = .038$, with 19.9% more concurrent motion, $F(1, 19) = 53.82$, $p < .001$; 2.96% less human idle time, $Z = -2.48$, $p = .013$; 17.3% less robot idle time, $F(1, 19) = 54.79$, $p < .001$; and a 15.1% larger separation distance, $F(1, 19) = 61.18$, $p < .001$. The mean values for each of these metrics, along with error bars depicting standard error of the mean, are depicted in Figure 4 for both robot modes.

In regard to data collected up to the administration of the first survey (see Figure 3), each participant was exposed only to a single robot type. Consequently, we can treat the team fluency metric data from the two groups of participants, those who worked with a standard robot and those who worked with a human-aware robot, as independent. Significant differences

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

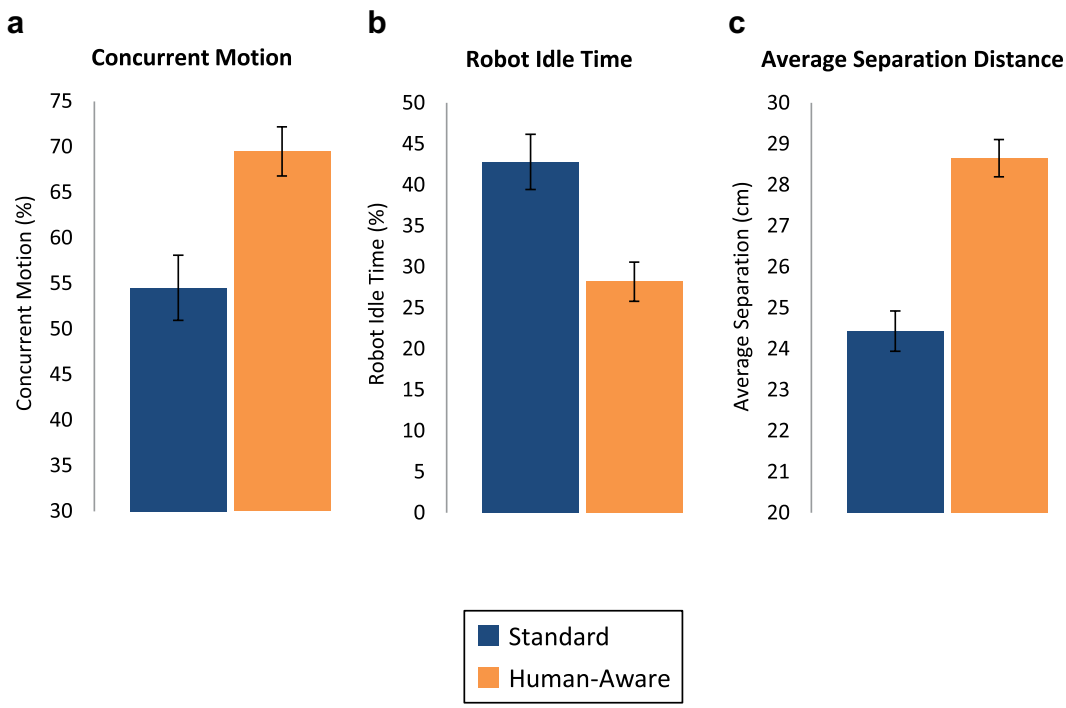

Figure 5. Mean values, with error bars indicating standard error of the mean (SEM), of (a) percentage of concurrent motion, (b) robot idle time, and (c) average separation distance between the human and robot for groups of participants working with the standard and human-aware robots prior to exposure to the second robot type.

between these populations emerged for three of the five team fluency metrics considered when analyzed with one-way ANOVAs. Once again, the normality of the data was assessed with the Shapiro-Wilk test, and none of the distributions was significantly different from normal, minimum $W(11) = 0.862$, $p = .061$. The group of participants that worked with the human-aware robot completed the task with 15.0% more concurrent motion, $F(1, 18) = 5.68$, $p = .028$; 14.6% less robot idle time, $F(1, 18) = 6.41$, $p = .021$; and a 17.3% larger separation distance, $F(1, 18) = 19.83$, $p < .001$. The mean values of both groups of participants for each of these metrics along with error bars depicting standard error of the mean are shown in Figure 5.

Subjective Evaluation

As indicated in Figure 3, the first questionnaire for the subjective assessment of the robotic assistant was administered prior to each participant's exposure to the second robot mode, regardless of whether he or she was in the human-aware-first or standard-first condition. As a result, we can treat the responses of these questionnaires as having come from two independent populations: one that worked with the human-aware robot and one that worked with a standard robot. We used the Mann-Whitney-Wilcoxon test to determine whether these two groups provided significantly different responses to the questions listed in Table 1. Significant differences (at $\alpha = .05$) were found for three of the questions, with participants exposed to the human-aware robot disagreeing more strongly with "The robot did not understand how I wanted to do the task" ($p = .012$), "The robot kept getting in my way" ($p < .001$), and "The robot came too close to me for my comfort" ($p = .05$).

After being exposed to both conditions, the participants were asked to directly compare the behavior of the robot in the first mode with its behavior during the second mode. Since "first mode" and "second mode" signified different modes depending on whether a participant was

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

in the human-aware-first or standard-first condition, we could once again treat the two groups as independent and test whether the responses were significantly different between groups using the Mann-Whitney-Wilcoxon test. All of the questions yielded significantly different responses, with the exception of "The robot moved too fast for my comfort." The group of participants that experienced the human-aware robot second (standard-first condition), when compared with the group that experienced the standard robot second (human-aware-first condition), more strongly sided toward *much more* when responding to the following statements regarding the second robot mode: "I trusted the robot to do the right thing at the right time" ($p < .001$), "The robot and I worked well together" ($p < .001$), "I felt safe when working with the robot" ($p < .001$), and "I trusted the robot would not harm me" ($p = .008$) and more strongly sided toward *much less* for the statements "The robot did not understand how I wanted to do the task" ($p = .046$), "The robot kept getting in my way" ($p < .001$), and "The robot came too close to me for my comfort" ($p < .001$).

DISCUSSION

Differences in Team Fluency

The results presented in the previous section provide strong support for both of our hypotheses. The significant differences in favor of human-aware motion planning observed in all quantitative metrics of team fluency indicate that this type of motion planning leads to more effective human–robot teamwork (Hypothesis 1). Further, the fact that significant differences emerged between the two independent groups of participants before exposure to the second robot mode for most of these metrics provides very strong evidence in support of human-aware motion planning (Hypothesis 1).

The major significance of this result is that these improvements can be achieved among participants who were never previously exposed to a robot capable of human-aware motion planning. After only two practice task executions, and without any explanation of the system's capability to adapt its motion planning, participants were able to take advantage of the adaptation and form a more effective team with the robot. This finding suggests that given minimal demonstration of the robot's ability to avoid motion conflicts, human workers inherently begin to exploit this capability, just as they would when working with a human assistant. In light of previously mentioned research, which indicated that some participants expect a robot to adapt at the task level as a human would (Hoffman & Breazeal, 2007b), it is possible that a similar expectation is placed on motion-level adaptation and that a lack of adaptation leads not only to inefficient teamwork, as shown by the quantitative metrics, but also to an unsatisfied human worker, as discussed in the next section.

Differences in Human Satisfaction and Perceived Safety and Comfort

Even before exposure to both robot modes, as indicated by the results of the first questionnaire, we already observed significant differences between the groups with regard to satisfaction with the robot as a teammate: Participants working with the human-aware robot disagreed more strongly with statements like "The robot did not understand how I wanted to do the task" and "The robot kept getting in my way." Additionally, we found that participants were less comfortable with the robotic assistant that used standard motion planning, agreeing more strongly with the statement "The robot came too close to me for my comfort."

Once the participants finished working with both the human-aware and standard robots and filled out the comparison questionnaire, these results were even more pronounced: In addition to the three items that yielded significantly different responses during the first questionnaire, when directly comparing the two robot modes, participants agreed more strongly with the phrases "I trusted the robot to do the right thing at the right time," "The robot and I worked well together," "I felt safe when working with the robot," and "I trusted the robot would not harm me" when describing the human-aware robot. The first two of these questionnaire items indicate that our participants experienced more satisfying interaction with the robot in human-aware mode, and the last two items suggest that the participants also felt more comfortable and safe when working with the human-aware robot.

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

Collectively, these results provide strong support for our second hypothesis (Hypothesis 2).

The only item on the questionnaire that did not yield significantly different results between human-aware and standard robots was "The robot moved too fast for my comfort." This is, in a way, an expected result, as the base speed of the robot was the same under both conditions. The safety system mentioned in the Method section would slow and eventually stop the robot if the distance between the human and robot workers fell below certain thresholds, but this behavior was identical under both conditions.

Since the safety system was running identically under both conditions, the higher perceived safety and comfort ratings when working with a human-aware robot are interesting phenomena. In terms of physical safety, participants in both conditions were at an equally low risk of unwanted contact with the robot due to the use of the safety system. However, as a result of the robot taking evasive maneuvers under the human-aware condition, participants were exposed to fewer sudden robot stops while under this condition; we hypothesize that participants felt safer when working with the human-aware robot for this reason. Although physical safety was the same in both conditions, the observed higher degree of perceived safety with the human-aware robot is an important result, as low perceived safety can be a high-stress situation for a human worker, and continuous exposure to stress has been shown to have negative long-term effects on health (McEwen & Stellar, 1993).

Human and Robot Idle Time

Another interesting observation follows from a comparison of the percentages of human and robot idle time for the two modes. As one can see from Figure 4, the average human idle time was far shorter than robot idle time in both conditions. Analysis with a paired t test indicated statistical significance for the difference between human idle time and robot idle time when participants worked with both the standard motion planning robot, $t(19) = 11.8$, $p < .001$, and the human-aware robot, $t(19) = 10.1$, $p < .001$. This finding suggests that in a human–robot team when motion conflict prevents both agents from performing their tasks simultaneously, people prefer to perform their own task and make the robot wait, rather than waiting for the robot to perform its task first. This result is similar to that observed by Unhelkar et al. (2014), whose experiment indicated that human workers are more likely to make a robotic assistant wait for them rather than a human assistant. In the present experiment, we showed that human workers prefer to make the robot wait rather than wait themselves; on the basis of these results and those observed by Unhelkar et al., it appears that people consider human time more valuable than robot time, whether the time in question is their own or that of another human.

Learning Effects

Several precautions were taken to minimize confounding due to learning effects. Recall that the experiment included two training rounds preceding the test tasks. This training allowed participants to grow accustomed to the task; mistakes, such as placing screws in the wrong sequence, were common during training but were far less common during final task execution. The training also contributed to a learning effect confound in that it provided participants with experience in performing the task in the first condition that carried over to subsequent task executions in the second condition. We mitigated this effect by randomizing the order in which the two robot modes were presented to our participants to counterbalance the potential impact of the experience gained from these training rounds.

We also tested whether the order in which our participants were exposed to the two conditions had a significant effect on the experiment metrics across robot modes. We evaluated whether the groups that worked with the human-aware (or standard) robot in the first and second conditions performed significantly differently in terms of human idle time and other metrics. Analysis with one-way ANOVAs found that with the exception of percentage concurrent motion in the human-aware mode, $F(1, 18) = 4.53$, $p = .047$, none of the metrics in either mode differed significantly based on the order in which it was measured, minimum $F(1, 18) = 3.71$, $p = .07$.

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

Limitations

Our experiment results lend support that human-aware motion planning can provide benefits in terms of perceived safety and team fluency. Our experiment was designed as a first investigation, and there are a number of ways in which follow-on experiments can be improved. First, these results were found using a tabletop task with a fixed-base robot. Although the task was simple, it was designed to be representative in that a large set of manual assembly operations are performed at workbenches, and most current industrial robots systems are fixed based and not yet freely mobile. As was mentioned in the Method section, the task was also carefully designed to provide a balance between conflicting and nonconflicting motions. This design makes our task, and the presented results, a good representation of similar tasks—namely, those in which a human and a stationary robot work at small distances of separation, such as assembly of consumer electronics or collaborative tasks on assembly lines. It is an open question, however, whether these results will generalize to other types of human–robot collaboration, such as applications in which the robot is mobile or not visible to the human worker at all times.

Second, our participant pool consisted of Massachusetts Institute of Technology affiliates with a rather narrow age distribution centered around 27 years of age. It is possible that, on average, the members of our participant pool are more likely to be accepting of autonomy and robots than is an average worker in a factory environment. Participants from relevant occupations that are more representative of the target population would strengthen future experiments.

A third limitation is related to the length of the work session experienced by our participants. In the experiment, each participant worked with the robot for 5 min in each of the two conditions. This period is relatively short compared to work session lengths that would be experienced by workers at a factory. Over time, a worker may become more efficient working with a robot that uses standard motion planning by learning how to time his or her actions to minimize motion conflicts. Conversely, a longer work session with a standard robot may lead to greater fatigue than that caused by working with a human-aware robot, since the worker has to be more focused on avoiding motion conflicts. This increased fatigue may affect team fluency and efficiency. The potential effects of longer work sessions on team fluency of human–robot teams with and without motion-level adaptation is an open question.

CONCLUSION

In this paper, we presented an adaptive motion-planning technique that we call human-aware motion planning, which uses the prediction of human actions and workspace occupancy to actively avoid potential motion conflicts that could arise in close-proximity human–robot interaction. We then described a human subject experiment aimed at studying human response to this motion-planning technique and quantifying its potential benefits through quantitative team fluency metrics and subjective evaluation of satisfaction, safety, and comfort. The results of this experiment indicated that people learn to take advantage of human-aware motion planning even when performing novel tasks with very limited training and with no indication that the robot's motion planning is adaptive. It was also shown that participants working with a human-aware robot form a more effective human–robot team and are able to perform a collaborative task in less time, with more concurrent motion and less human and robot idle time, and while maintaining a greater separation distance, compared to working with a standard robot. Furthermore, subjective evaluations indicated that human workers were more satisfied with the human-aware robot as a teammate and perceived it to be more comfortable and safe to work with. This finding signifies that human-aware motion planning leads to more satisfying human–robot interaction.

That human-aware motion planning leads to a higher degree of perceived safety—and thus reduced potential for stress-related health problems—while simultaneously improving team fluency is a very important result. The ability to show simultaneous improvements to efficiency and human worker satisfaction and well-being makes human-aware motion planning a highly desirable tool for close-proximity human–robot

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

interaction and brings us one step closer toward the successful introduction of robotic assistants into previously human-only domains. Furthermore, by showing that human-aware motion planning is effective through actual human subject experiments, rather than through simulation, we can strongly motivate future research in all the facets that would make a real-time human-aware system possible, including action prediction algorithms, the development of rapid motion–planning techniques, human motion model development, and many others.

One important point to consider, however, is that the improvements presented in this paper were obtained with the robot having perfect knowledge of what action the human would take next, since the sequences of actions were preset. Consequently, these results should be viewed as an upper bound of possible increases to team fluency and human worker satisfaction. As one might imagine, the improvements would depend very highly on prediction accuracy and would not be as pronounced with imperfect action prediction. Nonetheless, the results shown in this paper are substantial, leading us to believe that significant improvements can be derived from human-aware motion planning. Evaluating how the performance of a team using a human-aware robot can change according to varying levels of prediction accuracy is very important and is a planned future avenue of research.

KEY POINTS

- To determine whether motion-level robot adaptation leads to efficient and satisfying human–robot interaction, human-in-the-loop evaluation is necessary.
- We conducted a human subject experiment in which participants performed a collaborative task with an adaptive robot incorporating human-aware motion planning and with a baseline robot using shortest-path motions.
- When working with the adaptive robot, participants completed the task faster, with more concurrent motion, with less human and robot idle time, and with a larger human–robot separation distance.
- Participants also indicated they were more satisfied with the adaptive robot as a teammate and felt safer and more comfortable working with it.
- The results of the experiment indicate that humans respond well to motion-level adaptation, even with very limited training and without being informed the robot is adaptive.
- The positive response to human-aware motion planning motivates the use of this technique in close-proximity human–robot interaction and the development of supporting technologies.

REFERENCES

Alami, R., Chatila, R., Clodic, A., Fleury, S., Herrb, M., Montreuil, V., & Sisbot, E. A. (2006). Towards human-aware cognitive robots. In *Proceedings of the Fifth International Cognitive Robotics Workshop (the AAAI-06 Workshop on Cognitive Robotics)* (pp. 39–46). Palo Alto, CA: Association for the Advancement of Artificial Intelligence.

Arai, T., Kato, R., & Fujita, M. (2010). Assessment of operator stress induced by robot collaboration in assembly. *CIRP Annals–Manufacturing Technology*, *59*, 5–8.

Bellotto, N. (2012, March). *Robot control based on qualitative representation of human trajectories*. Paper presented at AAAI Symposium on Designing Intelligent Robots: Reintegrating AI, Stanford, CA.

Bennewitz, M., Burgard, W., Cielniak, G., & Thrun, S. (2005). Learning motion patterns of people for compliant robot motion. *International Journal of Robotics Research*, *24*, 31–48.

Berenson, D., Srinivasa, S., & Kuffner, J. (2011, October). Task space regions: A framework for pose-constrained manipulation planning. *International Journal of Robotics Research*, *30*, 1435–1460.

Chung, S.-Y., & Huang, H.-P. (2010). A mobile robot that understands pedestrian spatial behaviors. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 5861–5866). Piscataway, NJ: IEEE.

Chung, S.-Y., & Huang, H.-P. (2011). Predictive navigation by understanding human motion patterns. *International Journal of Advanced Robotic Systems*, *8*, 52–64.

Doshi, A., & Trivedi, M. M. (2009). On the roles of eye gaze and head dynamics in predicting driver's intent to change lanes. *IEEE Transactions on Intelligent Transportation Systems*, *10*, 453–462.

Dragan, A., Lee, K. C., & Srinivasa, S. S. (2013). Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human–Robot Interaction (HRI)* (pp. 301–308). Piscataway, NJ: IEEE.

Dragan, A., & Srinivasa, S. (2013, June). *Generating legible motion*. Paper presented at Robotics: Science and Systems, Berlin, Germany

Hawkins, K. P., Vo, N., Bansal, S., & Bobick, A. (2013). Probabilistic human action prediction and wait-sensitive planning for responsive human–robot collaboration. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots (HUMANOIDS 2013)* (pp. 15–17). Piscataway, NJ: IEEE.

Hoffman, G., & Breazeal, C. (2007a). Cost-based anticipatory action selection for human–robot fluency. *IEEE Transactions on Robotics*, *23*, 952–961.

Hoffman, G., & Breazeal, C. (2007b). Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. In *Proceedings of the ACM/IEEE International Conference on Human–Robot Interaction* (pp. 1–8). New York, NY: ACM.

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015

Kitade, T., Satake, S., Kanda, T., & Imai, M. (2013). Understanding suitable locations for waiting. In *Proceedings of the 8th ACM/IEEE International Conference on Human–Robot Interaction* (pp. 57–64). Piscataway, NJ: IEEE.

Koppula, H. S., & Saxena, A. (2013, June). *Anticipating human activities using object affordances for reactive robotic response*. Paper presented at Robotics: Science and Systems, Berlin, Germany.

Kruse, T., Basili, P., Glasauer, S., & Kirsch, A. (2012). Legible robot navigation in the proximity of moving humans. In *2012 IEEE Workshop on Advanced Robotics and Its Social Impacts (ARSO)* (pp. 83–88). Piscataway, NJ: IEEE.

Kruse, T., Kirsch, A., Sisbot, E. A., & Alami, R. (2010). Exploiting human cooperation in human-centered robot navigation. In *RO-MAN 2010* (pp. 192–197). Piscataway, NJ: IEEE.

Kwon, W. Y., & Suh, I. H. (2012). A temporal Bayesian network with application to design of a proactive robotic assistant. In *2012 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3685–3690). Piscataway, NJ: IEEE.

Lasota, P., Rossano, G., & Shah, J. A. (2014, August). *Toward safe close-proximity human-robot interaction with standard industrial robots*. Paper presented at the 10th IEEE International Conference on Automation Science and Engineering (CASE), Taipei, Taiwan.

Lenz, C., Nair, S., Rickert, M., Knoll, A., Rosel, W., Gast, J., . . . Wallhoff, F. (2008). Joint-action for humans and industrial robots for assembly tasks. In *The 17th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN 2008)* (pp. 130–135). Piscataway, NJ: IEEE.

Luber, M., Stork, J. A., Tipaldi, G. D., & Arras, K. O. (2010). People tracking with human motion predictions from social forces. In *2010 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 464–469). Piscataway, NJ: IEEE.

Mainprice, J., & Berenson, D. (2013). Human–robot collaborative manipulation planning using prediction of human motion. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)* (pp. 299–306). Piscataway, NJ: IEEE.

McEwen, B. S., & Stellar, E. (1993). Stress and the individual: Mechanisms leading to disease. *Archives of Internal Medicine*, *153*, 2093–2101.

Meisner, E., Isler, V., & Trinkle, J. (2008). Controller design for human–robot interaction. *Autonomous Robots*, *24*, 123–134.

Mori, A., Uchida, S., Kurazume, R., Taniguchi, R.-i., Hasegawa, T., & Sakoe, H. (2006). Early recognition and prediction of gestures. In *18th International Conference on Pattern Recognition, (ICPR 2006)* (Vol. 3, pp. 560–563). Piscataway, NJ: IEEE.

Nikolaidis, S., Lasota, P., Rossano, G., Martinez, C., Fuhlbrigge, T., & Shah, J. (2013). Human–robot collaboration in manufacturing: Quantitative evaluation of predictable, convergent joint action. In *Proceedings of the International Symposium on Robotics (ISR)* (pp. 1–6). Piscataway, NJ: IEEE.

Pentland, A., & Liu, A. (1999). Modeling and prediction of human behavior. *Neural computation*, *11* (1):229–242.

Sisbot, E. A., Marin-Urias, L. F., Alami, R., & Simeon, T. (2007). A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, *23*, 874–883.

Sisbot, E. A., Marin-Urias, L. F., Broquere, X., Sidobre, D., & Alami, R. (2010). Synthesizing robot motions adapted to human presence. *International Journal of Social Robotics*, *2*, 329–343.

Thompson, S., Horiuchi, T., & Kagami, S. (2009). A probabilistic model of human motion and navigation intent for mobile robot path planning. In *4th International Conference on Autonomous Robots and Agents (ICARA 2009)* (pp. 663–668). Piscataway, NJ: IEEE.

Trautman, P., & Krause, A. (2010). Unfreezing the robot: Navigation in dense, interacting crowds. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 797–803). Piscataway, NJ: IEEE.

Unhelkar, V. V., Siu, H. C., & Shah, J. A. (2014). Comparative performance of human and mobile robotic assistants in collaborative fetch-and-deliver tasks. In *Proceedings of the 2014 ACM/IEEE International Conference on Human–Robot Interaction* (pp. 82–89). New York, NY: ACM.

Ziebart, B. D., Ratliff, N., Gallagher, G., Mertz, C., Peterson, K., Bagnell, J. A., . . . Srinivasa, S. (2009). Planning-based prediction for pedestrians. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2009)* (pp. 3931–3936). Piscataway, NJ: IEEE.

Przemyslaw A. Lasota is a PhD candidate in the Department of Aeronautics and Astronautics at Massachusetts Institute of Technology, where he received his master's degree in aeronautics and astronautics in 2014.

Julie A. Shah is an assistant professor in the Department of Aeronautics and Astronautics at Massachusetts Institute of Technology, where she received her PhD in aeronautics and astronautics in 2011.

Date received: June 3, 2014
Date accepted: October 31, 2014

Downloaded from hfs.sagepub.com at Massachusetts Institute of Technology on February 19, 2015