

Initiative in Robot Assistance during Collaborative Task Execution

Jimmy Baraglia*, Maya Cakmak†, Yukie Nagai*, Rajesh Rao† and Minoru Asada*

*Department of Adaptive Machine Systems, Osaka University, Japan

†Computer Science of Engineering Department, University of Washington, USA

Email: jimmy.baraglia@gmail.com

Abstract—Collaborative robots are quickly gaining momentum in real-world settings. This has motivated many new research questions in human-robot collaboration. In this paper, we address the questions of whether and when a robot should take *initiative* during joint human-robot task execution. We develop a system capable of autonomously tracking and performing table-top object manipulation tasks with humans and we implement three different initiative models to trigger robot actions. *Human-initiated help* gives control of robot action timing to the user; *robot-initiated reactive help* triggers robot assistance when it detects that the user needs help; and *robot-initiated proactive help* makes the robot help whenever it can. We performed a user study (N=18) to compare these trigger mechanisms in terms of task performance, usage characteristics, and subjective preference. We found that people collaborate best with a proactive robot, yielding better team fluency and high subjective ratings. However, they prefer having control of when the robot should help, rather than working with a reactive robot that only helps when it is needed.

I. INTRODUCTION

Task-oriented robots in homes, factories, and small businesses have the potential to improve productivity and quality of everyday tasks while reducing the workload of humans. However, many tasks in these environments involve difficult-to-automate steps because they require dexterous manipulation or a “human touch.” As a result, these tasks are best suited for *collaborative* or *joint* execution with humans, taking best advantage of the strengths of robots and humans. Although such joint task executions come naturally to human-human teams, achieving similar fluency and comfort in human-robot teams poses many challenges.

Previous work has tackled many of these challenges. Some of the main research threads have investigated ways to compute robot action plans that improve joint task performance while reducing the load on the human [1], tracking and anticipating human motion to enable execution of such task plans [2], [3], [4], and designing robot behaviors to improve team effectiveness and fluency [5], [6]. While past work provides useful insights into *how* a robot should help as part of joint human-robot tasks, in this paper we focus on the question of *when* a robot should help. In particular we investigate the factor of **initiative** in robot assistance **during task execution**. We ask two questions:

- 1) Should the robot take initiative or let the human control the robot’s participation in the task?
- 2) When should the robot take initiative?



Fig. 1. Different initiative models for robot assistance during collaborative task executions: human-initiated help (left), robot-initiated reactive help (middle), and robot-initiated proactive help (right).

To address these questions, we investigate different mechanisms for triggering robot assistance in the context of joint table-top manipulation tasks. We develop a joint task execution system capable of autonomously performing a number of object manipulation tasks as well as monitoring end-to-end human task executions. We implement three trigger mechanisms (Fig. 1): (i) *human-initiated help* which gives control of robot action timing to the user, (ii) *robot-initiated reactive help* in which assistance is triggered when the robot detects that the user needs help, and (iii) *robot-initiated proactive help* in which the robot helps whenever it can. We present findings from a user study (N=18) in which participants performed different tray preparation tasks in three conditions involving the different assistance trigger mechanisms. The study demonstrated that people collaborate best with a proactive robot in terms of team fluency metrics and prefer the proactive help over other conditions. They prefer the human-initiated help over the reactive help, even though it results in higher human idle times and slower task completion. Given control of when the robot should help in this condition, they tend to divide the tasks equally between themselves and the robot, similar to the division that emerges with the proactive robot.

II. RELATED WORK

A. Human-Robot Collaboration

In recent years, collaborative robots designed to work side-by-side with humans have gained momentum in real-world settings. This has fueled a large body of research on human-robot collaboration. One of the core research threads tackles the problem of *task planning* for joint human-robot tasks. Among others, Shah *et al.* generates a robot action plan so as to minimize human-idle time [7]. Hayes and Scasclatti

developed a collaboration planner that reduces cognitive and physical load on the human [1]. Another vein of research focuses on low-level motion planning for the robot within a collaborative context [5], [8], [9], [10], with an eye towards improving team fluency and the user’s sense of safety.

Researchers have studied other low-level behaviors, besides robot motion, that impact collaboration and enable coordination of actions during task execution [11]. For example, St.Clair and Mataric demonstrated that robot verbal feedback improves team performance [12]. Awais *et al.* proposed mechanisms to mitigate breakdowns in the joint tasks [13]. Others focus on the coordination of micro-interactions that occur during collaboration, such as object hand-overs, using gaze [14] or adapting timing of motions to the human’s state [15]. Chao and Thomaz developed mechanisms to coordinate sharing of common resources during collaboration, such as the speaking floor or part of the workspace that both the human and the robot need to access [6].

Besides generation of robot behaviors, another key problem in human-robot collaboration is perception of the human. Preliminary work by Hoffman and Breazeal suggests that anticipatory perceptual simulation improves efficiency and fluency in teamwork [2], [16]. With the help of new sensing and human tracking technologies, many others followed with models of action or motion anticipation in the context of human-robot collaboration [17], [13], [18], [19].

As mentioned in Sec. I, this paper focuses on the question of *initiative* about when a robot should help. We note that joint human-robot task planning implicitly addresses the question of *when* a robot should help by producing a plan that specifies the order and timing of human and robot actions. The key difference of the scenario considered in our work is that we do not assume pre-planning of the task prior to execution. Rather, the allocation of task components occur during task execution depending on both the human’s and the robot’s behaviors.

B. Robot Assistance and Help

Given our emphasis on in-situ, ad-hoc collaboration, rather than planned collaboration, previous work on robot help is highly relevant. In fact, many of these relate to one or more of the different help behaviors studied in this paper. For example, Kwon *et al.*’s work [20] is akin to our robot-initiated proactive help. Cuntoor *et al.* consider human instruction as part of the collaboration [21], similar to our human-initiated help condition. Najmaei and Kermani’s prediction-based reactive control model for collaboration [22] is akin to our robot-initiated reactive help. Sakita *et al.* design different robot assistance behaviors triggered in different conditions; such as taking over when the human’s both hands are occupied or providing verbal disambiguation when the user’s hesitation is detected [23]. Similarly, Baraglia and Nagai proposed a developmentally motivated behavior in which the robot intervenes to help when it detects that effects of a human’s action were not as predicted, *i.e.*, the human action failed [24].

C. Initiative in Human-Robot Interaction

In the context of human-robot collaboration, one study by Gombolay *et al.* is particularly relevant. They investigate decision-making authority in the planning process and find that people are willing to give control to the robot for the efficiency benefits [25]. While our results are consistent with theirs, our study differs in its focus on authority over assistance timing *during task execution*, as opposed to authority over assistance allocation *during task planning*. Groten *et al.* looked at shared decision making in the context of haptic collaborations [26]. Cakmak *et al.* investigated initiative in robot question asking [27]. In addition, the large body of work on mixed-initiative control in the context of robot teleoperation [28] has some relevance to our work.

III. SYSTEM

To study different help trigger mechanisms, we develop an end-to-end system for joint task execution that allows a robot to perform object manipulation actions as well as monitor the execution of the same actions by a human. In this section we present the details of our system.

A. Platform

Our system is built around the PR2 robot platform (Fig. 1). PR2 has two 7 degrees-of-freedom arms giving it a large workable space for tabletop manipulation tasks. Each arm has 1 degree-of-freedom parallel-finger gripper that can grasp objects up to a width of 8cm. PR2’s arms are passively balanced and actuated with low-power motors, making it safe to work around humans. For perception, it has a Kinect sensor attached to the head that has a high-speed pan and tilt motion. Note that most of the system was designed independently of the platform while the action execution part was designed for and with the PR2.

B. Domain and Task Representation

We focus on joint preparation tasks. This category of tasks shares many properties of tasks previously studied in the context of human-robot collaboration (*e.g.*, circuit building [1], lego model assembly [23], food preparation [5], industrial assembly [4]), including partially ordered action sequencing and shared physical space. More specifically, we consider food tray preparation with n objects, m tray locations and three non-overlapping table regions. Objects can be uniquely recognized and their location is represented as a 2D coordinate on the table. We also represent relations among objects and targets with the three predicates `is-on(object, object)`, `is-at(object, location)`, and `is-in(object, region)`. Note that `is-on(object, object)` is detected directly through the perception module, while the two other predicates are inferred based on the task knowledge. The table is split into three regions based on who is allowed to manipulate: robot-only (near robot), human-only (near human), and both-allowed (middle). Task goals are represented as a conjunction of instantiated predicates; *i.e.*, the set of relations that need to be true.




Fig. 2. (a) Goal states for the two task categories used in our evaluation. (b) Pictorial description of a sample task instance (category Task B), used for explaining the task to participants in the user study.

Our experiments involve six specific tasks from two task categories (Tasks A and B) in slightly different domains. All tasks in the same task category have the same set of predicates in their initial state and goal descriptions; however specific tasks differ in the particular objects and locations with which the task is instantiated. Task A involves four objects to be placed in four target locations on the tray. Task B involves six objects to be arranged on two locations on the tray. The two task categories are described in Fig. 2a and individual task instances are shown in Fig. 4.

Both the human and the robot are assumed to have one task-relevant action: `pick-and-place(object, x, y)`. The x and y coordinates can be anywhere on the table, including particular tray locations or on other objects. The action is applicable for an agent (human or robot) only on objects whose current location is within the regions allowed to the agent. In our task scenarios, one object is initially placed in the robot-only region for both tasks; two objects are placed in the human-only region for Task B.

C. Robot Perception and Actions

The robot can segment and recognize tabletop objects using the point cloud obtained from the robot’s RGBD sensor. It uses the Point Cloud Library implementation of tabletop segmentation, which detects the table plane with the RANSAC algorithm. It then extracts a point cloud segment corresponding to each object on the table. If an object is inside or in contact with another object, they are segmented as one object with possibly multiple colors. The robot represents and recognizes objects based on their color and size extracted from the segmented point cloud. Color is discretized into six values and size in three values.

The robot’s pick-and-place action is parametrized with an *object* to be picked and a *location* at which the objects is to be placed. The action is defined as a sequence of poses relative to the object (pre-grasp, grasp, and lift poses) followed by poses relative to the target location (transfer, lower, and drop poses). While the overall action template remains the same, some of the poses in the action are tuned to the particular object being manipulated.

D. Joint Task Execution Model

The overall system for joint task execution is illustrated in Fig. 3. At the core of this system are two modules for (i) tracking the state of the task and anticipating future actions,

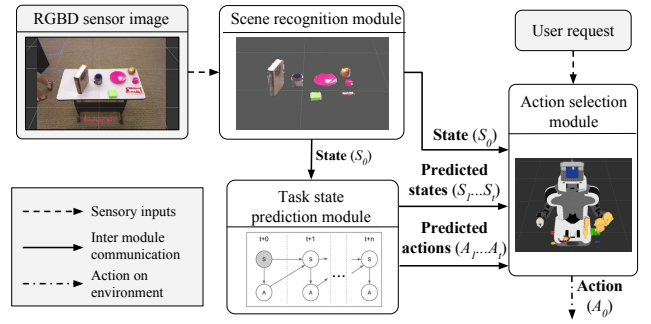


Fig. 3. Model for helping robots: recognizes the current environmental state, predicts the possible future states using a dynamic Bayesian network and generates actions to achieve the desired end-states.

and (ii) selecting a robot action based on the observed and anticipated states accreting to different help strategies. A more detailed description of these modules is given in the following.

1) *Task state prediction module*: Our system uses Dynamic Bayesian Networks (DBN) to predict future states and robot actions that lead to those states. DBNs are multi-time-slice Bayesian networks where variables are connected to one another over adjacent time steps as well as within the same time step. They are computationally efficient generalization of hidden Markov models and have been used to model multi-modal robot behavior in uncertain environments (e.g., [29]).

For this study, each time slice of the DBN contains a state and an action node, corresponding to two multinomial discrete random variables S and A . S can be one of all possible states $\{s_0, s_1, \dots, s_N\}$ that are distinct according to the defined predicates for a finite set of objects and named locations (Sec. III-B). In other words, two states in which an object’s position is different but both positions are not at a named location (e.g., `is-at(object1, location1)=F`) are considered the same discrete state. The variable A is one of all possible action instances $\{a_0, a_1, \dots, a_M\}$ that involve the combination of all objects and named locations in the environment, regardless of whether they are available to the human or the robot. S_0 represents the current observed state of the environment and S_t the predicted state at time t . Within a single time slice, the state influences the action. Between consecutive time slices, the state and action from the previous state, S_t and A_t , influence the next state S_{t+1} .

The DBN encodes the task knowledge in the conditional probabilities $P(A_t|S_t)$ which represent an action policy that the robot could use if it were to execute the task on its own. Since the tasks are known a priori in our scenario, these conditional probabilities were computed based on the known task structure (Sec. III-B), assuming each path for completing the task is equally likely. The conditional probabilities $P(S_{t+1}|S_t, A_t)$ encode the environment and action dynamics and were determined empirically. Future states and actions are predicted by computing the marginal probabilities $P(A_t)$ and $P(S_t)$ using Bayesian inference. The predictions are then sent to the action selection module.

2) *Action selection module*: The action selection module implements a policy that specifies what the robot should do at each time step. If the robot were to execute the task completely on its own, this module would directly return one of the possible actions predicted by the DBN immediately after every action. During joint task execution, on the other hand, the robot’s policy needs to account for the human’s direct input or their actions that result in changes in the world state. We implement three policies that differ in terms of *when* a robot action is triggered.

- **Human-initiated Help (H)**: The first policy gives complete control of robot actions to the user. The robot performs an action only when the user explicitly says “Robot, can you help me?”
- **Robot-initiated Reactive Help (R)**: In the second policy, robot actions are initiated by the robot when it detects that help is needed. The robot monitors the human’s task execution and tries to detect when one of the next states predicted by the DBN is not reached within an expected time window, indicating a delay or difficulty in the task progress.
- **Robot-initiated Proactive Help (P)**: The third policy involves performing actions whenever they are possible. However, different from a robot-only task execution, the robot takes into account human actions that might be *in progress* before a stable environmental state is reached. If at least one executable action exists that does not conflict with the human actions, the trigger is initiated.

When a robot action has been triggered, the robot selects the specific action to perform based on the state of the environment according to its task execution policy encoded by the DBN. Given alternative actions with the same utility, the robot prefers actions that involve objects and targets that are closest to one of the robot’s grippers (right or left). The robot always uses the nearest gripper to manipulate objects.

IV. USER STUDY

The help trigger mechanisms described in Sec. III-D2 are expected to yield different joint task execution dynamics. Furthermore, each mechanism on its own can result in a wide variety of behaviors depending on the particular user. For example, when interacting with the *human-initiated* policy, users may request help at every step or only when they need it. When interacting with the *robot-initiated proactive* policy, they might select their own actions such that the robot has many opportunities to help or they might (unintentionally or intentionally) block the robot’s actions. The differences across and within each policy can reflect on objective task execution measures, as well as the user’s subjective attitude towards the robot. To investigate these differences, we performed a user study that allows us to (i) characterize people’s behaviors while interacting with each policy, and (ii) compare the alternative policies for triggering robot help.

A. Study design

We performed a within participants study with one independent variable (robot helping behavior) with three conditions: H, R, P (Sec. III-D2). In each condition, participants performed two tasks with the robot, one from each category (Task A and B). The order of the three conditions were counterbalanced.

B. Setup

The robot was placed in front of a 68cm high table. Participants sat across the table. The table top was separated into three zones as shown in Fig. 4. Participants were asked not to touch objects that are in the red zone (near the robot). Similarly the robot could not enter the blue zone (near the human). Both were allowed to manipulate objects in the middle zone. In the middle of the table there was a tray with four target positions.

Tasks were explained to participants with a one page pictorial description involving (i) the set of objects and targets involved in the task and (ii) the final state of the tray when the task is complete. An example task description is shown in Fig. 2b. An additional small table was placed to the right of the participant. Printed task descriptions were placed on this table, together with a tablet for logging task steps (see Sec. IV-C) and a laptop for responding to our questionnaire. The complete setup can be seen in Fig. 1.

C. Procedure

Participants were recruited from a campus and nearby neighborhoods through mailing lists. Interested individuals signed up for a 45 minutes time slot in advance. When participants arrived at their scheduled study time, we first explained the purpose of the study and asked them to sign a consent form. Then they were taken to the participants seat, introduced to the robot and the workspace, and given an overview of the procedure.

Next, the robot was activated and participants performed a practice task (Fig. 4a). They were explained what the task is using the corresponding pictorial description. The robot made a specific sound to indicate that it was ready. Participants were told that they can start the task when they hear this sound. They were told to perform one step of the task and then log the step on the tablet. The logging was done throughout the study as a mechanism to space human actions apart and give the robot an opportunity to detect intermediate states of the task. Each log required indicating who performed the step (human or robot), the two letter identifier for the object involved (as indicated in the task description), and the one letter identifier for the target position where the object was placed. The second step of the task was performed by the robot to familiarize participants with the robot’s motion. The robot made another sounds when it detected the task completion. Participants were told that they will perform similar tasks together with the robot in three conditions where the robot’s behavior will be different.

Next we moved on to the actual study. For each condition, the experimenter first gave condition specific instructions. In the human-initiated help (H) condition, participants were told




Fig. 4. Particular instances of the tasks used in the user study: (a) practice task, (b-d) three instances of Task A, and (e-g) three instances of Task B performed by participants in the three different conditions.

that they can request the robot’s help by saying “Robot, can you help me?” In the other conditions (R and P), they were told that the robot will decide when and how to help out with the task. Then the experimenter set up the initial state of the first task, told participants to start when they hear the robot sound, and left them alone with the robot. The experimenter came back to set up the next task after the robot detected that the task was complete. After completing both tasks in the same condition, participants were asked to respond to the condition-specific questionnaire. After all three conditions were complete, participants responded to additional questions drawing comparisons between the three conditions. At the end participants were thanked for participating and given the promised compensation of 10 USD equivalent gift card.

D. Measurements

The study was recorded from two cameras; one mounted on the robot’s head and another overseeing the workspace together with the robot and the participant. In addition, we logged the progression of tasks and robot actions with timestamps throughout the study. From the study logs we extracted the task completion time and the number of actions performed by each agent. From the videos we extracted quantitative measure that characterized each participant and the robot behaviors. The coding was performed by two coders ($IRR^1 = 0.72$), including one without prior knowledge of the study.

To compare the three conditions subjectively from the user’s perspective, we administered several questions after each condition as well as at the end. First we asked an open ended question to elicit the participants own description of the robot’s assistance behavior. Another question asked them to describe their strategy. Then we asked a set of Likert scale questions, similar to those commonly used in human-robot collaboration research [30]. These questions addressed the user’s perception of: the robot’s helpfulness, its awareness of the human and task progress, its contribution to the task, team fluency and efficiency, and naturalness of the interaction (see questions in Fig. 7). Additional questions at the end asked a forced ranking of the three conditions and open ended questions about perceived distinction between the two robot-initiated conditions and how different behaviors would be combined in an ideal interaction.

¹Cohen’s kappa.

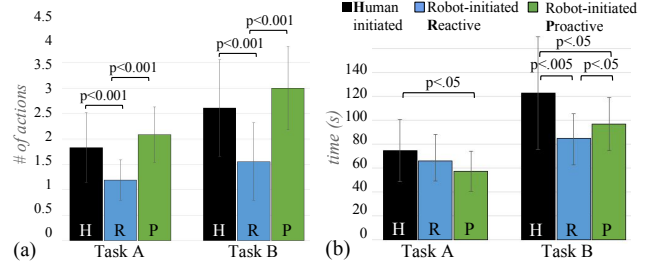


Fig. 5. (a) Number of actions performed by the robot for each task category in each condition. (b) Task completion time for each task category in each condition. Error bars represent standard deviation.

V. FINDINGS

Our study was completed by 18 participants (9 females, 9 males, ages 18 to 35). This section presents our findings based on data collected from these participants. A one-way ANOVA was conducted to compare the effect of conditions H, R and P on the different objective metrics. We performed post-hoc tests (two-tailed paired-t-test) to explore differences between pairs of conditions.

A. Objective metrics

We first examine common task and collaboration metrics. Fig. 5a shows the average number of task actions performed by the robot in each condition (Task A: $F(2, 51) = 15.35, p < .001$; Task B: $F(2, 51) = 13.24, p < .001$) and Fig. 5b shows the overall task completion times by the human-robot team (Task A: $F(2, 51) = 2.20, p = .12$; Task B: $F(2, 51) = 6.15, p < .005$). Fig. 6(a-d) show the breakdown of task completion times into robot-only, human-only, concurrent, and no motion segments and Fig. 6(e-f) separately show the human idle time and robot idle time. The results of the ANOVA for results in Fig. 6 are as follows: (a): (Task A: $F(2, 51) = 7.64, p < .005$; Task B: $F(2, 51) = 12.62, p < .001$); (b): (Task A: $F(2, 51) = 10.51, p < .001$; Task B: $F(2, 51) = 0.61, p = .55$); (c): (Task A: $F(2, 51) = 4.41, p < .05$; Task B: $F(2, 51) = 2.79, p = .07$); (d): (Task A: $F(2, 51) = 4.78, p < .05$; Task B: $F(2, 51) = 6.97, p < .005$); (e): (Task A: $F(2, 51) = 2.73, p = .070$; Task B: $F(2, 51) = 9.22, p < .001$); (f): (Task A: $F(2, 51) = 5.32, p < .01$; Task B: $F(2, 51) = 5.61, p < .01$).

1) *Proactive versus Reactive*: First we focus on the comparison of robot-initiated help strategies. Proactive help results

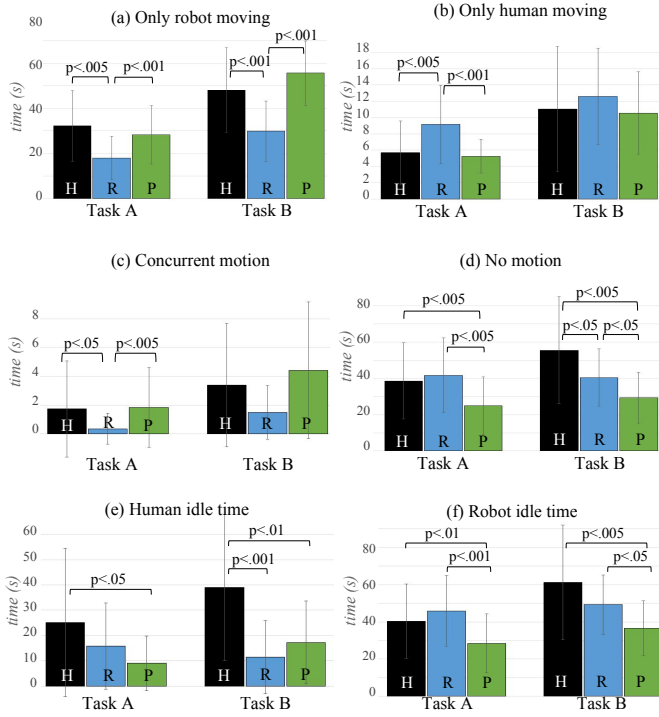


Fig. 6. Breakdown of task completion times into (a) robot-only, (b) human-only, (b) concurrent, and (d) no motion time segments. These include only motion related to the joint task. (e) Human idle time. This excludes the time during which the human is performing their secondary task of *logging* task actions. (f) Robot idle time.

in the robot having a greater contribution to the task, as indicated by the significantly higher number of actions performed by the robot (Task A: $p < .001$, Task B: $p < .001$) (Fig. 5a). This is also reflected in the significantly lower robot idle times for the proactive robot (P) as compared to the reactive robot (R) (Task A: $p < .001$, Task B: $p < .05$) (Fig. 6f). The average number of actions performed by the reactive robot was around 1 (Task A: $M = 1.17$, $SD = .38$, Task B: $M = 1.56$, $SD = .76$), which is the minimum number of actions required by the robot. Whereas, the proactive robot performed around 2 (Task A) and 3 (Task B) actions (Task A: $M = 2.17$, $SD = .48$, Task B: $M = 3.00$, $SD = .82$), which are about half of the actions needed to complete the task. This finding is expected and confirms that our model produced the intended behavior.

Despite the difference in the number of robot actions, there was no significant difference in the total task durations in Task A (Task A: $p = .12$) and little difference in task B. A potential reason for this could be lack of parallelization between human and robot actions. However, the significant increase in the *concurrent* human-robot motion (Fig. 6c) in the proactive condition indicates that parallelization did indeed happen at least in Task A (Task A: $p < .005$). In addition, the total task durations appeared to be greatly influenced by the difference in human and robot action speeds as humans are several orders of magnitude faster at pick-and-place actions. Hence they were not slower in completing the overall task

in the reactive condition. Despite this difference, human idle times were not significantly higher in the proactive robot condition (P) (Fig. 6e).

2) *Human-initiated versus Robot-initiated*: Next, we look at comparisons between the human-initiated help (H) condition and robot-initiated help conditions to characterize how people chose to get help from the robot when they had control. From Fig. 5a, we see that the number of actions performed by the robot in the H condition was about half of all task actions, as in the P condition. The number of actions performed by the robot was significantly higher than in the R condition (H-R - Task A: $p < .001$, Task B: $p < .001$). It resulted in significantly higher concurrent motions in Task A for the H condition compared to the R conditions (H-R - Task A: $p < .05$) (Fig. 6c). We believe that it is because participants asked for help and then started doing their own actions as soon as they understood the robot's intention. This is similar to the P condition, where participants briefly waited until they recognized what the robot was doing and then acted. The added waiting time in the H condition was reflected in overall task completion times (Fig. 5b), which was significantly higher than in the R condition for Task B (H-R - Task B: $p < .005$) and in the P condition for both tasks (H-P - Task A: $p < .05$, Task B: $p < .05$). This was also reflected in the human idle times (Fig. 6e) which was highest for the H condition in both tasks (H-R - Task A: $p = .27$, Task B: $p < .001$; H-P - Task A: $p < .05$, Task B: $p < .01$). We noticed that one participant made the robot do all actions for Task 1; two participants made the robot do all possible actions for Task 2 in the H condition. This contributed to the high human idle time and task completion time, while making the variance in this condition high.

B. Subjective metrics

Participant responses to the Likert-scale questions are summarized in Fig. 7. The inter-condition differences were analyzed using the Wilcoxon signed rank test², which is a standardly used non-parametric test. As suggested in [31] and to avoid family-wise errors, we grouped the seven scales into two sub-scales representing the quality of interaction (Fig. 7a) and the system performance (Fig. 7b). There were no statistically significant differences between the human-initiated help (H) and proactive robot (P) conditions in any of the sub-scales, despite the differences observed in objective metrics (e.g., the task completion time shown in Fig. 5b) between these two conditions. Subjective ratings of the quality of interaction appeared to be correlated with the number of actions performed by the robot (Fig. 5a), rather than the overall task efficiency (Fig. 5b). The reactive robot (R) condition was rated significantly lower than the other two (H and P) conditions, indicating that participants agreed significantly more that the quality was better in the H and P conditions (see Fig. 7a). Whereas the significant differences were observed in the quality of interaction, participants did not rate differently the system performance. It seems they did not attribute the

²We also conducted parametric tests and obtained similar results.

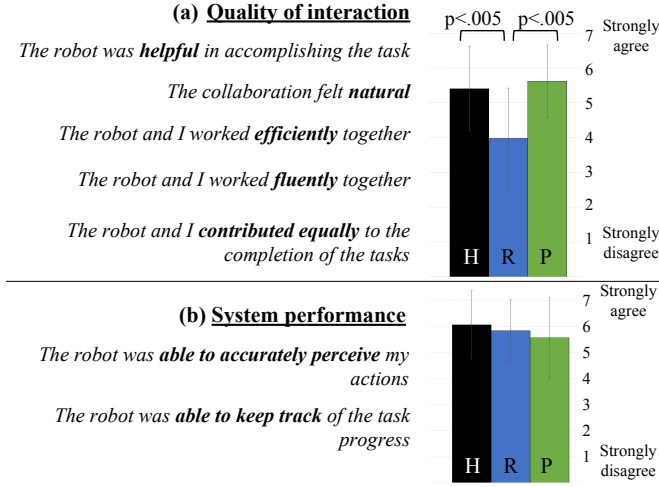


Fig. 7. Mean Likert-scale ratings in questionnaire responses. Significant differences according to Wilcoxon signed rank tests are indicated with p -value ranges.

robot's behavior in the R condition to its inability to perceive the human or keep track of task progress. In the forced ranking question administered at the very end of the study, 72% of participants (13/18) indicated P as their *most* preferred behavior, while 22% (4/18) indicated H and only 6% (1/18) indicated R. 78% of participants (14/18) indicated R as their *least* preferred behavior, with 17% (3/18) for H and 6% (1/18) for P. The question yielded a clear ranking of the three conditions as $P > H > R$ from most preferred to least preferred. Furthermore, in a separate two-choice questions, 67% of participants (12/18) indicated they prefer letting the robot take initiative, while the remaining 33% said they preferred having control over the robot's actions. These results demonstrate that although there were no significant differences between the H and P conditions in the Likert-scale ratings, people are more likely to prefer P over H in favor of the improved objective metrics (Sec. V-A).

C. Perceived differences of robot strategies

An open-ended question asked participants to describe the differences between the two conditions R and P in which the robot decided when to act, if they noticed any difference at all. All participants reported that they noticed a difference. The reactive robot was perceived as "slow" and characterized as "lazy" and "hesitant" by some of the participants. The proactive robot, on the other hand, was perceived "fast" and "pro-active". Descriptions of the perceived robot behaviors were accurate; for example:

- M, 35: "... [P] felt more natural to have unprompted collaboration while I was performing the task, rather than the robot waiting for me to finish as it did during [R]."
- M, 20: "[P] was more **proactive** in its help ... [R], by contrast, would only complete actions that I was unable to complete."
- M, 22: "[In P] the robot took the **initiative** a lot more than [R]"

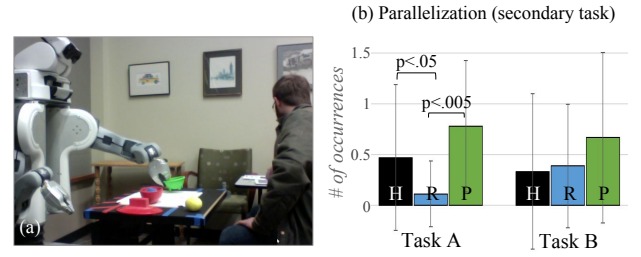


Fig. 8. Occurrences of parallelization of the human's secondary task (logging) with the primary task (tray preparation task) being performed by the robot.

D. Collaboration enhancing human behaviors

The differences in the objective and subjective task metrics can be further dissected by examining the occurrence of certain events. Firstly, as mentioned earlier, concurrent motion was significantly higher in the P and H conditions for Task A (Fig. 6c), which shows better team work took place in these conditions. Secondly, we observed another type of collaboration, which is participants tended to perform the logging task when the robot was performing an action on the tray as pictured in Fig. 8a (Task A: $F(2, 51) = 5.88, p < .01$; Task B: $F(2, 51) = 1.03, p < .36$). We observed higher occurrences of this type of parallelizations in condition P and H compared to R in Task A (H-R - Task A: $p < .05$, Task B: $p = .83$, H-P - Task A: $p = .18$, Task B: $p = .25$, R-P - Task A: $p < .005$, Task B: $p = .26$) (see Fig. 8b). Thirdly, we observed that people intuitively encouraged collaboration by often starting tasks with objects that were in the human-only region of the table. Participant descriptions of their strategies, in a free form question in the questionnaire, reflected their intent to enhance the collaboration; for example:

- F, 22: "[In R] I chose objects closest to me or that were obscuring the place of the objects needed to be. I also moved slower than I would without the robot to give it time to help."
- M, 19: "[In P] I moved objects from the blue zone into the collaboration zone, and placed objects in-between logging and [the robot's] actions."

VI. DISCUSSION

Our study demonstrated that the behavior of the proactive robot was similar to the behavior people asserted when they had control over the robot's actions. In turn, the similar high subjective rating of the proactive and the human-controlled robots could be partially ascribed to this similarity. Furthermore, we believe that the behavior that was common in these two conditions is similar to how a human would collaborate in the same role. Indeed, participants thought that the collaboration was most natural.

On the question of *whether* a robot should take initiative, our results demonstrate that the answer depends on the robot's behavior. People were happy to give away control if the robot is proactive, but they would rather have control if the robot is reactive. Given its other benefits in terms of objective task and

team metrics, this suggests that collaborative robots should be designed to always be proactive. In practice, this might not always be possible. Challenges such as partial task knowledge and uncertain perception might reduce the robot's ability to help the user when it is actually possible for it to help. While the simplistic help request used in our experiments would not be sufficient, enabling users to ask for particular types of help by commanding actions could result in more effective collaboration in such circumstances. Thus, the overall implication of our study is that mixed-initiative help triggers might be ideal for collaborations in realistic settings. We also believe that increasing proactivity over time, after observing the user's collaboration preferences (e.g., [4]) might further improve the collaboration.

VII. CONCLUSION

We address the questions of whether and when a robot should take *initiative* during joint human-robot task execution by comparing three initiative models to trigger robot actions: *human-initiated help*, *robot-initiated reactive help*, and *robot-initiated proactive help*. Through a user study (N=18) we demonstrate that people collaborate best with a proactive robot, yielding better team fluency and high subjective ratings. While they are willing to give control of initiative to a proactive robot, they prefer having control rather than working with a reactive robot that only helps when it is needed.

ACKNOWLEDGMENT

The authors would like to thank Dan Butler for helping in the implementation of the perceptual system. This study is partially supported by MEXT/JSPS Grants-in-Aid for Scientific Research (24119003, 24000012, 25700027) and JSPS Core-to-Core Program and by the ONR Science of Autonomy grant no. N000141310817 and NSF grant no. 1318733.

REFERENCES

- [1] B. Hayes and B. Scassellati, "Effective robot teammate behaviors for supporting sequential manipulation tasks," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2013.
- [2] G. Hoffman and C. Breazeal, "Effects of anticipatory perceptual simulation on practiced human-robot tasks," *Autonomous Robots*, vol. 28, no. 4, pp. 403–423, 2010.
- [3] C. Perez-D'Arpino and J. Shah, "Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification," in *IEEE Intl. Conf. on Robotics and Automation*, 2014.
- [4] S. Nikolaidis and J. Shah, "Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy," in *ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI)*, 2013.
- [5] A. D. Dragan, S. Bauman, J. Forlizzi, and S. S. Srinivasa, "Effects of robot motion on human-robot collaboration," in *ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI)*, 2015.
- [6] C. Chao and A. Thomaz, "Controlling social dynamics with a parametrized model of floor regulation," *J. of Human Robot Interaction*, 2013.
- [7] J. Shah, J. Wiken, B. Williams, and C. Breazeal, "Improved human-robot team performance using chaski, a human-inspired plan execution system," in *ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI)*, 2011.
- [8] J. Mainprice, E. A. Sisbot, L. Jaillet, J. Cortés, R. Alami, and T. Siméon, "Planning human-aware motions using a sampling-based costmap planner," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2011.
- [9] J. Mainprice and D. Berenson, "Human-robot collaborative manipulation planning using early prediction of human motion," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2013, pp. 299–306.
- [10] E. A. Sisbot, A. Clodic, R. Alami, and M. Ransan, "Supervision and motion planning for a mobile manipulator interacting with humans," in *ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI)*, 2008.
- [11] B. Mutlu, A. Terrell, and C.-M. Huang, "Coordination mechanisms in human-robot collaboration," in *HRI Workshop on Collaborative Manipulation*, 2013.
- [12] A. S. Clair and M. Mataric, "How robot verbal feedback can improve team performance in human-robot task collaborations," in *Proceedings of the Tenth Annual ACM/IEEE Intl. conf. on Human-Robot Interaction*, 2015, pp. 213–220.
- [13] M. Awais and D. Henrich, "Proactive premature intention estimation for intuitive human-robot collaboration," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2012, pp. 4098–4103.
- [14] A. Moon, D. M. Troniak, B. Gleeson, M. K. Pan, M. Zheng, B. A. Blumer, K. MacLean, and E. A. Croft, "Meet me where i'm gazing: how shared attention gaze affects human-robot handover timing," in *ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI)*, 2014.
- [15] C.-M. Huang, M. Cakmak, and B. Mutlu, "Adaptive coordination strategies for human-robot handovers," in *Robotics: Science and Systems (RSS)*, 2015.
- [16] G. Hoffman and C. Breazeal, "Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team," in *ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI)*, 2007.
- [17] S. Nikolaidis, R. Ramakrishnan, K. Gu, and J. Shah, "Efficient model learning from joint-action demonstrations for human-robot collaborative tasks," in *ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI)*, 2015.
- [18] N. Jarrassé, J. Paik, V. Pasqui, and G. Morel, "How can human motion prediction increase transparency?" in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2008, pp. 2134–2139.
- [19] K. P. Hawkins, S. Bansal, N. N. Vo, and A. F. Bobick, "Anticipating human actions for collaboration in the presence of task and sensor uncertainty," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014.
- [20] W. Y. Kwon and I. H. Suh, "Proactive planning using a hybrid temporal influence diagram for human assistive robots," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 1785–1791.
- [21] N. P. Cuntoor, R. Collins, and A. J. Hoogs, "Human-robot teamwork using activity recognition and human instruction," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2012.
- [22] N. Najmaei and M. R. Kermani, "Prediction-based reactive control strategy for human-robot interactions," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2010, pp. 3434–3439.
- [23] K. Sakita, K. Ogawara, S. Murakami, K. Kawamura, and K. Ikeuchi, "Flexible cooperation between human and robot by interpreting human intention from gaze information," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, vol. 1. IEEE, 2004, pp. 846–851.
- [24] J. Baraglia, Y. Nagai, and M. Asada, "Prediction error minimization for emergence of altruistic behavior," *4th Intl. conf. on Development and Learning and on Epigenetic Robotics*, pp. 281–286, Oct. 2014.
- [25] M. C. Gombolay, R. A. Gutierrez, G. F. Sturla, and J. A. Shah, "Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams," *Robots: Science and Systems (RSS)*, 2014.
- [26] R. Groten, D. Feth, A. Peer, and M. Buss, "Shared decision making in a collaborative task with reciprocal haptic feedback-an efficiency-analysis," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2010, pp. 1834–1839.
- [27] M. Cakmak, C. Chao, and A. Thomaz, "Designing interactions for robot active learners," *IEEE Trans. Autonomous Mental Development*, vol. 2, no. 2, 2010.
- [28] T. Fong, C. Thorpe, and C. Baur, "Multi-robot remote driving with collaborative control," *Industrial Electronics, IEEE Transactions on*, vol. 50, no. 4, 2003.
- [29] C. Huang and B. Mutlu, "Learning-based modeling of multimodal behaviors for humanlike robots," in *ACM/IEEE Intl. Conf. on Human-robot Interaction (HRI)*, 2014.
- [30] G. Hoffman, "Evaluating fluency in human-robot collaboration," in *HRI Workshop on Human Robot Collaboration*, 2013.
- [31] J. Carifio and R. J. Perla, "Ten common misunderstandings, misconceptions, persistent myths and urban legends about likert scales and likert response formats and their antidotes," *Journal of Social Sciences*, vol. 3, no. 3, pp. 106–116, 2007.