Effects of Anticipatory Action on Human-Robot Teamwork

Efficiency, Fluency, and Perception of Team

Guy Hoffman
MIT Media Laboratory
20 Ames Street E15-468
Cambridge, MA 02139
guy@media.mit.edu

Cynthia Breazeal
MIT Media Laboratory
20 Ames Street E15-468
Cambridge, MA 02139
cynthiab@media.mit.edu

ABSTRACT

A crucial skill for fluent action meshing in human team activity is a learned and calculated selection of anticipatory actions. We believe that the same holds for robotic teammates, if they are to perform in a similarly fluent manner with their human counterparts.

In this work, we propose an adaptive action selection mechanism for a robotic teammate, making anticipatory decisions based on the confidence of their validity and their relative risk. We predict an improvement in task efficiency and fluency compared to a purely reactive process.

We then present results from a study involving untrained human subjects working with a simulated version of a robot using our system. We show a significant improvement in best-case task efficiency when compared to a group of users working with a reactive agent, as well as a significant difference in the perceived commitment of the robot to the team and its contribution to the team's fluency and success. By way of explanation, we propose a number of fluency metrics that differ significantly between the two study groups.

Categories and Subject Descriptors

I.2 [**Computing Methodologies**]: Artificial Intelligence

General Terms

Algorithms, Human Factors, Performance, Design, Experimentation

Keywords

Human-Robot Interaction, Anticipatory Action Selection, Fluency, Teamwork

1. INTRODUCTION

Two people repeatedly performing an activity together naturally reach a high level of coordination, resulting in a fluent meshing of their actions. In contrast, human-robot interaction is often structured in a stop-and-go fashion, inducing delays and following a rigid turn-taking pattern. Aiming to design robots that are capable peers in human environments, we try to attain a more fluent meshing of human and machine activity.

In recent years, the cognitive mechanisms of joint action have received increasing attention [20]. Among other factors, successful coordinated action has been linked to the formation of expectations of each partner's actions by the other and the subsequent acting on these expectations [14, 22]. We argue that the same holds for collaborative robots: if they are to go beyond stop-and-go interaction, agents must take into account not only past events and current perceived state, but also expectations of their human collaborators.

In this paper we present an adaptive anticipatory action selection mechanism for a robotic teammate. We discuss a cost-based framework for examining coordinated action in shared-location human-robot teamwork, and investigate our model in this scenario. We compare our framework to a purely reactive agent acting within a traditional perception-action loop, and — based on theoretical analysis — predict an improvement in efficiency.

We then present results from a study involving untrained human subjects working with a simulated version of a robot using our anticipatory system. We show a significant improvement in best-case task efficiency when compared to users working with a purely reactive agent. However, we were not able to show this difference being significant when measuring the mean score over repetitions. We attribute this to the small number of repetitions used in our study.

That said, we are not interested solely in efficiency, but also in the qualitative notion of fluency in coordinated action meshing, ultimately leading to more appropriate collaborative behavior. In a post-study survey we found a significant difference in the perceived contribution of the robot to the team's fluency and success, as well as its commitment to the team. Given there are no generally accepted measures of teamwork fluency, we compare three candidate metrics between the two conditions, finding the groups to differ significantly in two (time between human and robot action, and time spent in concurrent motion), but not in a third (human idle time).

The remainder of the paper is structured as follows: In Section 2 we briefly describe the cost-based Markov process in which our agent is set, and in Section 3 outline a reactive action-selection mechanism for an agent in this world. In Section 4 we introduce our adaptive cost-optimizing antic-

HRI'07, March 8-11, 2007, Arlington, Virginia, USA.

ipatory agent and analyze its behavior. Section 5 presents and discusses results from the human subject study; Section 6 discusses related work, and we conclude in Section 7 with future research directions.

2. WORLD DESCRIPTION

We model the team fluency problem as a discrete time-based deterministic decision process including two agents, a *robot* and a *human*, working together on a shared task.[1]

Both robot and human share a common workspace, which at any time point is in one of a finite number of states. The agents also have a number of states, which are different for the robot and the human. The robot can only perceive the state of the workspace if it is in a subset of states, called *perceptive* states. Human and robot have distinct abilities, described as two sets of actions. There is a transition function that maps certain state-action pairs to new states.

A central motivation of our model is to investigate aspects of *time* associated with actions of two collaborating agents. Therefore, state transitions are not atomic, and the decision to take a particular action does not result in an immediate state transition. Instead, moving between states takes time, and is associated with a known discrete cost, which is a function of the states before and after the action. This cost can be thought of as the 'distance' between states, or more generally — the duration it takes to transition between states. We denote the cost of transitioning between states s_k and s_l with $d(s_k, s_l)$.

Thus when, at time t, an agent x decides to take an action on a certain state s_k, the world will be in the resulting state s_l only at time $t + d(s_k, s_l)$. While the other agent may take more actions during this time, the next time step at which agent x will be able to take another action is $t + d(s_k, s_l)$. For sake of simplicity, we will sometimes denote $d(s_k, s_l)$ as d_{kl}.

2.1 The Factory World

In our experiments we use a simulated factory setting (Figure 1). The goal of the team is to assemble a cart made of a *Body*, a *Floor*, two *Axle*s, and four *Wheel*s. The various parts have particular ways to be attached to each other — the *Body* is welded to the *Floor*, *Axle*s are riveted to the *Floor* and *Wheel*s are attached to *Axle*s using a wrench of matching color. A *component* is a partially assembled cart segment that includes one or more individual parts attached to each other, for example $Axle + Body + Floor$.

The labor is divided between the human and the robot: the human has access to the individual parts, and is capable of carrying them and positioning them on the workbench. The robot is responsible for fetching the correct tool and applying it to the currently pertinent component configuration in the workbench. Each part has a stock location (with an infinite supply of parts), and each tool has a storage location, to which it has to be returned for the robot to be able to find it again. The workbench can, at any one time, contain at most two components.

The above-described framework encompasses a state-space of 2,160,900 distinct states.

[1]In this paper we include an abridged description of the model, only as needed for the understanding of the subsequent analysis and study.

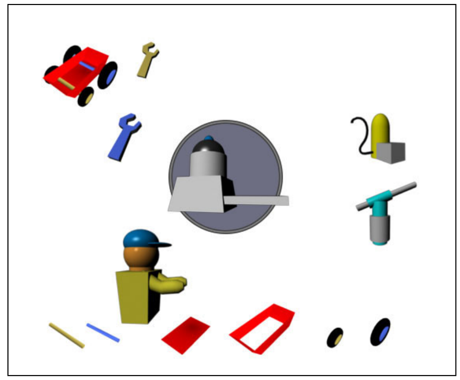

Figure 1: Simulated factory setting with a human and a robot building carts, while sharing a workbench (gray circle), but dividing their tasks. The robot has access to the tools (right and top-left of workbench), whereas the human is responsible to bring the parts (below the workbench). Top left shows a completed cart.

The robot's actions include going to each of the tool's locations, $PickUp$, $PutDown$, and Use. The latter only changes the workbench if the robot is holding a tool and is located at the workbench. The human's action space is similar, but with two more location actions, and no Use action. The duration cost of a state transition that involves navigation is the distance between the previous and the new location. The duration cost for state transitions involving the inventory of an agent, or changes to the workbench, is 1 in this implementation, but could theoretically be different for each tool.

The robot can perceive the state of the workbench only when it is located in it. Workbench state changes that happen while the robot is in any other state are not applied to its internal representation. Moreover, we assume that the robot has a function Φ that maps the workbench state to the appropriate tool required to bond the two components on the workbench. For example: $\phi(< Floor + Axle1, Wheel1 >) = Wrench1$. This can be a lookup table, or a simple decision process.

3. REACTIVE AGENTS

A baseline agent that is purely responsive to its environment and internal state, can be defined by an action policy that waits in the workbench when $\Phi(WorkBench) = \emptyset$, and fetches tool x, uses it, returns it, and returns to the workbench when $\Phi(WorkBench) = x$.

The obvious fallacy of this policy occurs when the same tool is needed twice in a row (which can happen with the wheels and axles, in the factory domain), resulting in a superfluous sequence of returning and then fetching the same tool. The naïve policy can therefore be improved by delaying the decision to return a tool until the state of the workbench changes. This prevents the agent from returning a tool before it is certain that it is not needed again in the next step. We call this policy *conservative tool return*.

However, it is straightforward to demonstrate that there is a negative impact of the "conservative tool return" strategy in the case where the next tool needed is different than the current tool. Note that the cost effect of conservative tool return is dependent not only on the known world configuration, but also on the turnaround time of the human action h, a quantity that can not be known but only estimated by the robotic agent. Additionally, the overall expected cost effect is dependent on the probability distribution on the workbench configuration over time. It therefore makes sense to discuss an action selection policy based on these factors, which is the topic of the following section. We will then frame the two reactive policies discussed here as a subset of the proposed anticipatory policy.

4. ANTICIPATORY ACTION SELECTION

As discussed in the introduction, humans are remarkably adaptive and increasingly effective when performing repetitive trials of an identical task collaborating with a consistent teammate. The use of educated anticipatory action based on expectations of each other's behavior may be a key ingredient in the achievement of this action fluency. In this section we will attempt to adopt this insight in the human-robot interaction domain within the discussed framework.

A necessary assumption for anticipatory action selection in our agent is that the human collaborator will follow a roughly consistent action pattern, i.e. will make similar decisions under similar circumstances.

The agent thus models the workbench as a first-order Markov Process.[2] The probability of the workbench state at time t, σ_t^w, is thus conditional on σ_{t-l}^w and denoted as

$$p_{i|j}^w \equiv Pr(\sigma_t^w = s_i | \sigma_{t-l}^w = s_j)$$

The agent can learn the parameters of this Markov process using a naïve Bayesian estimate. To do this, the agent keeps a one-step history of the state transitions of the workbench. A change from state s_j to state s_i increases the counter $n_{i|j}$. Consequently, $p_{i|j}^w$ is computed as

$$p_{i|j}^w = \frac{n_{i|j}}{\sum_{k=1}^{|\Sigma_W|} n_{k|j}}$$

However, in order to estimate the cost of preemptive action as described in the following section (which is ill-defined for non-constructive workbench states), and also to reduce the decision state space, the robot in our factory domain can alternatively model the probability of the tool needed based on the previous state: if $Q(x) = \{s_i : \phi(s_i) = x\}$ is the set of workbench states that warrant tool x, the new probability model learned is now

$$p_{x|j} \equiv Pr(\sigma_t^w \in Q(x) | \sigma_{t-l}^w = s_j)$$

We estimate this model as follows: a change from state s_j to state $s_i \in Q(x)$ increases the counter $n_{x|j}$. Using a Laplace correction [15] of 1, $p_{x|j}$ is then estimated by

$$p_{x|j} = \frac{n_{x|j} + 1}{\sum_{k=1}^{|Tools|} n_{k|j} + 1}$$

[2] A presumably more realistic model would be to view the collaboration as a Hidden Markov Model, with the human state transitions being hidden, and the workbench transitions being the evidence layer of the model. However, since many of the human's state transitions do not affect the workbench state, and the probability of workbench transitions conditional on the human state transitions $Pr(\sigma_t^w = s_i | \sigma_t^h = s_j)$ are not independent of σ_{t-l}^w, it is unclear whether such a model would indeed be of value in our domain, and is therefore left to future investigation.

4.1 Action Selection

As the agent only perceives the workbench state (and therefore information about the transition distribution) when it is in the workbench state, it makes sense to make decisions in terms of *action sequences*. The acquisition of these sequences is beyond the scope of this paper, but suffice to say that in our scenario the agent needs only to consider action sequences that begin and terminate while it is in the workbench state.

In the discussed factory domain we can identify four proto-sequences: (1) Pick up a tool and use it; (2) Return a tool and return to workbench; (3) Return a tool, bring a new tool, and use it; (4) Do nothing and wait.

The action selection process operates as follows: at any time the robot is in the workbench state, it evaluates the cost of each of the proto-sequences. Proto-sequence 1 needs to be grounded for each tool and proto-sequence 3 needs to be grounded for each of the currently not held tools. Given the probability distribution, the robot can compute the expected cost for choosing each of the strategies, and selects a grounded sequence optimizing for cost. Note that in the cost for proto-sequences 1–3, we assume that h is smaller than twice the distance between the workbench and any tool. Also note that the cost in our calculations includes performing the correct action afterwards.

Denoting the current state of the workbench s_j, and the workbench position 0, the expected duration cost of proto-sequence 1–3 are as follows:

$$Cost_1(x) = p_{x|j}(2d_{0x} + 2) + \sum_{y \neq x} [p_{y|j}(3d_{0x} + d_{xy} + d_{0y} + 4)]$$

$$Cost_2(x) = \sum_{y=1}^{|Tools|} [p_{k|j}(2d_{0x} + 2d_{0y} + 3)]$$

$$Cost_3(x, y) = p_{y|j}(d_{0x} + d_{xy} + d_{y0} + 3) + \sum_{z \neq y} [p_{z|j}(d_{0x} + d_{xy} + 2d_{y0} + d_{yz} + d_{z0} + 5)]$$

The above costs are derived as follows: for proto-sequence 1, the expected cost is made up of (a) the expected cost if the anticipated tool is correct, i.e. twice the distance from the workspace to the tool position d_{0x}, plus the cost of pickup, plus the cost of using the tool $= 2d_{0x} + 2$; and (b) the expected cost in case the anticipated tool is incorrect — three times d_{0x}: getting the anticipated tool and returning it, plus the distance to the correct tool d_{xy}, plus the distance back to the workspace d_{0y}, plus two pickups, one put-down, and one tool operation $= 3d_{0x} + d_{xy} + d_{0y} + 4$. The costs for proto-sequences 2 and 3 can be similarly derived.

Action sequence 4 is unique insofar as it is dependent not only on the state transitions in the workbench, but also on the behavior of the human teammate. If the human's next workbench-changing action is at time $t + h$, the cost of waiting is the cost of performing the correct action with complete confidence, plus h. For the case that the robot is holding a tool z:

$$Cost_4 = p_{z|j} + \sum_{y \neq z} [p_{y|j}(d_{0z} + d_{zy} + d_{y0} + 3)] + h$$

For the case that the robot is not holding a tool:

$$Cost_4 = \sum_{y=1}^{|Tools|} [p_{y|j}(2d_{0y} + 2)] + h$$

However, Since h is not directly accessible to the robotic agent, its estimate can be used as a confidence parameter, adjusting between an aggressively anticipatory behavior and a more cautious approach.

Using the above notation, we can now rephrase the previously discussed reactive agent behaviors. The naïve agent's policy can be viewed as selecting proto-sequence 2 whenever it is holding a tool in the workbench, and selecting proto-sequence 1 whenever a tool is warranted. The agent employing conservative tool return can be rephrased as selecting proto-sequence 4 whenever no tool is warranted, and selecting proto-sequence 1 or 3 if a workbench state warrants a tool.

4.2 Performance with a Human Teammate

In a theoretical analysis we find that using the anticipatory action selection mechanism initially results in a performance that is either identical or slightly less efficient than that of a reactive agent. However, in repeating trials with a consistent human the anticipatory behavior becomes increasingly rewarding, quickly outperforming the reactive agent. For example, using $h = 250$ in the factory scenario, we usually see the agent outperforming the reactive agents within 2 trials, and converging into full anticipatory behavior within 10 trials. A more detailed comparative analysis of the described strategies is currently underway in a separate publication.

In actual trial runs with an experienced and consistent human teammate, we can see evidence to that effect. Whereas the reactive agent with conservative tool return remains constant at a construction cost[3] of circa 800, the anticipatory adaptive agent shows a significant improvement after the first trial and again at the sixth trial, finally settling at a lower per-cart construction cost of circa 650 (see: Figure 2).

Finally, note that inconsistency on the human teammate's part delays the anticipatory behavior of the agent, resulting in slower convergence into a fluent and efficient activity pattern.

5. HUMAN SUBJECT STUDY

To investigate the effect of adaptive anticipatory action selection, we conducted a human subject study. We expected to see an increase in efficiency as predicted by the theoretical analysis, as well as an increase in the perceived contribution of the robot to the team's fluency and success.

5.1 Experimental Design

We recruited 32 participants (15 female) from the MIT community through email solicitation and posters. Participants arrived at our laboratory and were arbitrarily assigned to one of two experiment conditions. Subjects in *Group A* interacted with a reactive agent using the "conservative tool return" policy; those in *Group B* interacted with an anticipatory agent.

[3]The cost units, when measured with a human teammate, are in simulation frames, running at 30 frames per second.

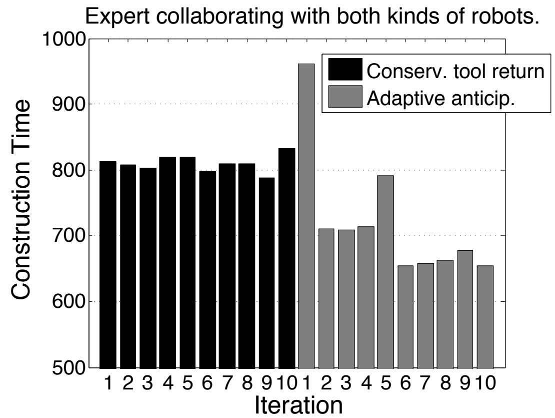

Figure 2: Change in per-cart construction time with an expert consistent human vis-a-vis the reactive agent (left) and the adaptive anticipatory agent (right).

All the participants received the following identical instructions (edited for brevity, omitting user interface instructions):

> In this study you play a video game. This game has two characters, Symon, a forklift-like robot in a cart factory and an avatar representing you, the human. Symon is surrounded by four tools: the welder, the rivet gun and two wrenches. The human is surrounded by six kinds of cart components: a floor, a body, two kinds of axles, and two kinds of wheels. In the center of the screen is a round workspace.
>
> In this game your goal is for the human-robot team to build 10 carts. Each of the team members has their own role in this joint effort. The human's role is to bring components to the workspace, the robot's role is to attach the car parts using the tools. The following tools attach the following components:
>
> 1. The wrench attaches a wheel to the matching color axle
> 2. The welder attaches the floor to the body
> 3. The rivet gun attaches both the axles to the floor
>
> A complete car has one floor, one body, two axles (one of each kind), and four wheels (two of each kind).
>
> The robot can only use a tool if there are exactly two cart parts in the workspace. Each of these parts can be made up of more than one simple components. For example - the workspace could contain one part made up of an axle with two

wheels, and one part made up of a floor with a body attached to it. In this case the robot could use the rivet gun. If there are more or less than two parts in the workspace, the robot can't do anything.

Your goal is to build cars in the least amount of time. A cart's construction time is measured from the moment the first part is dropped in the workspace and until the cart is completed. You can always see your best score and your last score, as well as the all-time best score, in the corner of the screen.

The instructions were phrased so as to imply the importance of the team as a joint performing entity. To control for instruction bias, neither group was told whether the robot will adapt to their behavior. Before beginning the experiment, participants were allowed to practice for an unlimited amount with the software, set to their assigned experimental condition.

5.2 Results

Of the participants, five had to be eliminated from the study. Two violated the experimental protocol, one experienced a software crash, one was significantly inattentive, resulting in scattered behavior, and for one subject the logging functionality was not working, resulting in a loss of data. This left us with 27 subjects, 14 in Group A and 13 in Group B. All 32 completed a post-study survey regarding their experience.

Table 1: Total cart completion metrics for untrained human subjects in the reactive (Group A) and adaptive anticipatory condition (Group B). We compare each subject's best score in ten trials, mean score over ten trials, and tenth trial.

Score	Group A		Group B		
metric	mean	std.dev.	mean	std.dev	T(25)
Best	1091.6	200.5	930.1	105.6	2.59; $p < 0.02$
Mean	1423.5	328.6	1233.3	227.5	1.73; not signif.
Final	1182.4	274.3	1030.7	154.8	1.75; not signif.

Table 1 shows total cart construction measures for the population. Cost units are in simulation frames at 30 frames per second.

Each subject's best performance is significantly better at a confidence level of 98% in the adaptive anticipatory case compared to the reactive case. Measuring the mean construction time over ten trials, as well as the time for construction of the tenth cart, we find the subjects in the anticipatory case to be better (at $p < 0.1$), but not significantly at a 95% confidence level. We believe that this is in part due to the fact that several subjects in Group B took a number of inconsistent trials to identify that the robot was adaptive, leading to a convergence to a stable construction pattern only in the last few carts (see also: 5.3.2). According to this hypothesis, both the mean and the final cart construction cost would be significantly lower in the anticipatory case if there were more trial runs per subject.

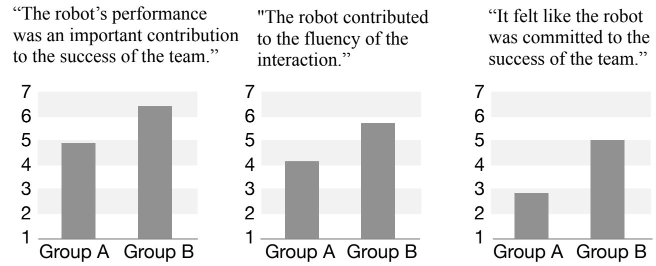

Figure 3: Self-report differences between participants in the *reactive* (Group A) and the *anticipatory action* (Group B) conditions. All differences are significant at $p < 0.01$.

5.2.1 Survey

In the post-experimental survey, we found significant differences between participants in the two groups (Figure 3). On a seven-point Likert scale, subjects in the anticipatory action agent "Group B" selected a significantly higher mark than those in the reactive agent "Group A" when asked whether:

- **"The robot's performance was an important contribution to the success of the team."**: Group A: **4.88** [SD=1.71] ; Group B: **6.38** [SD=1.2]; T(30)=2.87; $p < 0.01$
- **"The robot contributed to the fluency of the interaction."**: Group A: **4.13** [SD=1.54]; Group B: **5.69** [SD=1.4]; T(30)= 2.99; $p < 0.01$
- **"It felt like the robot was committed to the success of the team."**: Group A: **2.8** [SD=2.0]; Group B: **5.0** [SD=1.73]; T(30)= 3.21; $p < 0.005$

The two groups did not differ significantly when subjects were asked whether they themselves were "committed to the success of the team", or whether they "trusted the robot to do the right thing at the right time." Both groups averaged between 6 and 7 on these two questions.

5.2.2 Measures of Fluency

In an attempt to ground the subject's perceptions of fluency as well as those of the robot's appropriate teamwork in behavioral terms, we try to measure the fluency of the teams. However, while there is a body of work measuring verbal fluency, there are no generally accepted measures of fluency in shared-location joint action. In this work, we propose three fluency metrics, and compare the mean performance of the two groups along these measures.

Concurrent motion In post-experiment interviews, some of our participants noted a sense that the team was well synchronized when "both team members were constantly in motion". We measured the percentage of

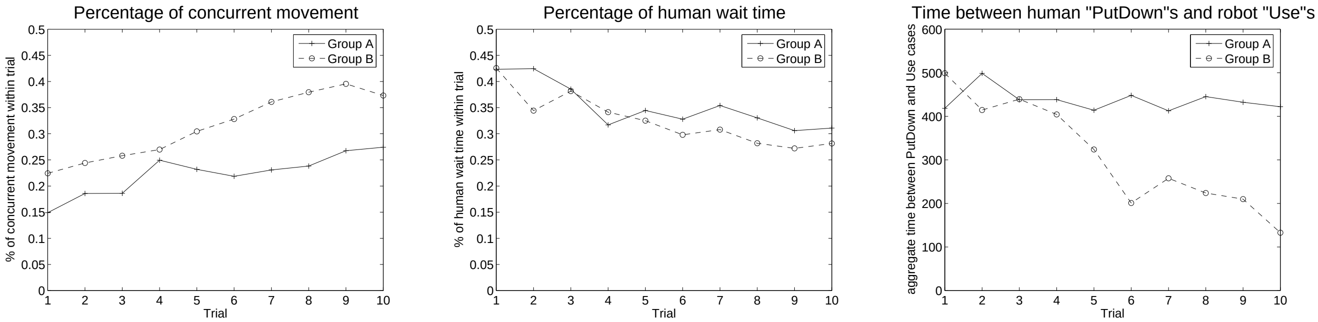

(a) (b) (c)

Figure 4: Three measures of fluency per cart averaged over study groups A (reactive) and B (anticipatory). (a) percentage of concurrent motion within trial; (b) percentage of human idle time; (c) aggregate time between human *PutDown* to robot *Use* delay.

frames within each trial in which both human and robot were in motion (i.e. in transition between two location-based internal states), and found those to be significantly different between the two groups (A: **0.23** [SD=0.08]; B: **0.32** [SD=0.08]; T(25)=3.11; $p < 0.005$). Figure 4(a) shows the mean percentage of concurrent motion for each of the 10 trials, averaged over subjects in each group. The graph shows that while the percentage of concurrent motion is improving for both groups, it does so at a higher rate in the anticipatory action condition.

Human idle time Another candidate for a measure of fluency is the amount of time the human spent waiting for the robot. We postulated that if the human was to spend much time waiting, it would feel like the team was not working fluently. However, we found no significant difference between the two groups in terms of the percentage of human waiting time (Figure 4(b)). Both groups seemed to decrease the human waiting time at an approximately equal rate, and with similar results. This is probably a result of the human adaptation to the robot's behavior.

Time between human and robot action A final measure of fluency is the time between the human's *PutDown* action and the robot's subsequent *Use* action. We found this measure to be significantly lower for Group B (A: **436.78** [SD=48.8]; B: **310.64** [SD=78.84]; T(25)=5.04; $p < 0.001$), and more decidedly so for the second half of each subject's trial sequence, after the robot has adapted to the human's construction pattern (A: **432.07** [SD=49.23]; B: **205.08** [SD=125.38]; T(25)=6.28; $p < 0.001$) (Figure 4(c)). In the reactive case, there is virtually no improvement across trials.

5.3 Discussion

The open-ended segment of the post-experiment questionnaire reveals a qualitative difference between the two conditions. Several subjects in Group B noticed the anticipatory behavior and remarked on it positively, e.g.: "it was nice when [the robot] anticipated my next move", or "[the] robot's anticipation of my actions was impressive and exciting". Negative remarks in Group B usually referred to a desire for even more anticipatory behavior, such as "[the robot] could do better by getting the first tool before/while I take the first part, because it was a consistent process and could be predicted", or "the robot should watch what I'm grabbing in advance."

Somewhat surprisingly, many subjects in Group A — without having been informed that the study was related to anticipatory action or that the robot was meant to be adaptive — noted with frustration that the robot did not predict their actions. We view this tendency as indicative of the fact that adaptiveness and anticipatory action are natural expectations of a robotic teammate in a repetitive task. Quotes from Group A included: "I was hoping that the robot would learn to anticipate more", "I expected more predictive behavior from the robot", "[the] robot was not able to anticipate [the] human's actions", and "it might have been more efficient if after a few carts the robot could pick up on the order in which i was bringing in the parts and be prepared with the equipment to join it."

Group A's positive comments regarding the robot's performance were limited to remarks shaped by a low level of expectation from the agent: "The robot seemed to do what was expected", "the robot did not mess up", and "the robot was highly responsive and never let the human down with its predictability," were representative responses in this condition.

5.3.1 Notions of Teamwork

It is interesting to note that several subjects in Group A noted that the team felt "lopsided", that "the human was the one who strategized, the robot just sat there", that the human "was more important than the robot", and that "the team's performance was highly dependent on human innovation". Subjects in this group concluded that "the robot seemed more like an assembly tool than a team member", that they "didn't see the robot as a team player", that the robot was used "as a tool", and one subject said that they "didn't get a sense that the robot really cared about the success of the team." In contrast, in Group B only one subject noted that they "felt that the success or failure of the task was [their] responsibility." Conversely, one other stated that they “trusted [the robot] more over time, as it seemed to anticipate what [they were] going to do.” The rest of the subjects in Group B did not address the balance of the team, the issue of trust, or that of commitment, in any way.

5.3.2 *Effect of Repetition Size*

As noted in Section 5.2, we believe that the relatively minor improvement in mean task efficiency through anticipatory action is related to the small number of repeating trials in the experiment. Appraisal of server logs, as well as user testimony, reveals that in many cases subjects experimented with various construction strategies in the first few runs, which caused the Bayesian model to converge more slowly. This seemed to be particularly true when subjects noticed that the robot changed its behavior, causing them to experiment with different construction sequences in an attempt to reveal the robot’s modus operandi. One reason for this behavior was the experiment’s insistence on identical instructions for both groups, not revealing that the robot would adapt to the human’s consistent behavior. Several subjects explicitly noted that the team would have performed better had they known in advance that the robot learned to anticipate their actions. Another possible way to counter this effect would be to discount the learning over time (see also: Section 7).

5.3.3 *Effect of “Best Score” Indicator*

We also believe that the display of the game’s all-time “Best Score” in the user interface was detrimental to the experiment as it might have caused subjects to experiment with different strategies instead of forming a consistent behavior pattern. Originally intended to motivate subjects to faster performance, the exceedingly good record time (only possible with a well-adapted agent) provoked subjects to question their strategy attaining a significantly worse score, and subsequently to change it several times over the course of the experiment.

6. RELATED WORK

Most work related to joint action — whether in philosophy, psychology, or artificial intelligence — has been concerned with a goal-oriented view of the problem, paying little attention to the *quality* of action meshing and fluency of teamwork, both as it is perceived by the team members, and as it effects quantitative measures of the collaboration.

In this body of work, joint action is usually described as solving a problem where the participants share the same goal and a common plan of execution. Grosz pointed out, in this context, that collaborative plans do not reduce to the sum of the individual plans, but consist of an interplay of actions that can only be understood as part of the joint activity [9].

In Bratman’s detailed analysis of Shared Cooperative Activity he defines certain prerequisites for an activity to be considered shared and cooperative [3]. He stresses the importance of mutual responsiveness, commitment to the joint activity, and commitment to mutual support. Supporting Bratman’s guidelines, Cohen and Levesque propose a formal approach to building artificial collaborative agents [18]. Their notion of joint intention is viewed not only as a persistent commitment of the team to a shared goal, but also implies a commitment on part of all its members to a mutual belief about the state of the goal. These principles have been used in a number of human-robot teamwork architectures [11, 1].

Much work has been done in the field of Discourse Theory, investigating discourse as a collaborative activity. Grosz and Sidner have analyzed the structure of discourse and subsequently modeled shared plans as a separate *extension*, rather than a *composition* of simple, single-agent plans [10]. Later work has further elaborated the workings of collaborative discourse, in terms of plans, beliefs, goals, and actions (e.g. [17, 2]). Collaborative discourse systems have been developed and implemented on screen-based and robotic dialog systems, taking into account both the verbal and the nonverbal aspects of discourse (e.g. [19, 21]). Still, the question of fluency in action meshing has not been part of this corpus. Moreover, as these works focused mainly on linguistic dialog, they have not addressed the case of nonverbal shared-location teamwork, or the improvement thereof through repetitive joint execution of a task.

Human-robot teamwork has also remained mostly in the turn-taking domain. Some have studied a robotic arm assisting a human in an assembly task [13]. Their work addressed issues of vision and task representation, but does not investigate joint adaptation, and does not address the timing issue. Other work studies human-robot collaboration with an emphasis on dialog and control, aimed primarily at teleoperation [7, 12] . Some frame human-robot collaboration in the context of mixed-initiative control and shared autonomy, arbitrating between the robot’s autonomy and direct human control, but also fail to address the question of shared-location fluency [5, 8].

Some work in shared-location human-robot collaboration has been concerned with the mechanical coordination of robots in shared tasks with humans (e.g. [23]). This work is predominantly concerned with single-action control and safety issues.

We have previously presented work in shared-location human-robot teamwork, investigating the role of nonverbal behavior on teamwork [11, 4]. While this task-level work included turn-taking and joint plans, anticipatory action and fluency have not been addressed.

Timing and synchronization have been reviewed on the motor level in the context of a human-robot synchronized tapping problem [16]. Anticipatory action, without relation to a human collaborator, has been investigated in robot navigation work, e.g. [6].

7. CONCLUSION

We have presented work investigating the effect of adaptive anticipatory action on the efficiency and fluency of action in human-robot teamwork. We hope to initiate an interest in the question of shared-location action timing, and have presented initial results on both the theoretical analysis of this method and its effect on untrained humans.

Several improvements to our method present themselves: in the discussed framework, the robot has no knowledge of the structure of the task. Domain-specific knowledge can decrease the action space at each decision point and fortify the accuracy of the probabilities of subsequent states. We believe that our system can also be made more robust by introducing a discount factor in the learned state transition distribution, making more recent moves by the human teammate more salient to the robot. Furthermore, the estimate of the human's turnaround time h should be state-specific and could be learned by the robot during the collaboration.

8. ACKNOWLEDGMENTS

This work is funded in part by the Office for Naval Research YIP Grant N000140510623.

9. REFERENCES

[1] R. Alami, A. Clodic, V. Montreuil, E. A. Sisbot, and R. Chatila. Task planning for human-robot interaction. In *sOc-EUSAI '05: Proceedings of the 2005 joint conference on Smart objects and ambient intelligence*, pages 81–85, New York, NY, USA, 2005. ACM Press.

[2] J. F. Allen. The TRAINS project: A case study in building a conversational planning agent. *Journal of Experimental and Theoretical AI (JETAI)*, 7:7–48, 1995.

[3] M. Bratman. Shared cooperative activity. *The Philosophical Review*, 101(2):327–341, 1992.

[4] C. Breazeal, C. D. Kidd, A. L. Thomaz, G. Hoffman, and M. Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Edmonton, Alberta, Canada, August 2005.

[5] D. J. Bruemmer, D. D. Dudenhoeffer, and J. Marble. Dynamic autonomy for urban search and rescue. In *2002 AAAI Mobile Robot Workshop*, Edmonton, Canada, August 2002.

[6] Y. Endo. Anticipatory and improvisational robot via recollection and exploitation of episodic memories. In *Proceedings of the AAAI Fall Symposium*, 2005.

[7] T. W. Fong, C. Thorpe, and C. Baur. Multi-robot remote driving with collaborative control. *IEEE Transactions on Industrial Electronics*, 2003.

[8] M. Goodrich, D. Olsen, J. Crandall, and T. Palmer. Experiments in adjustable autonomy. In *Proceedings of the IJCAI Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents*, 2001.

[9] B. J. Grosz. Collaborative systems. *AI Magazine*, 17(2):67–85, 1996.

[10] B. J. Grosz and C. L. Sidner. Plans for discourse. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in communication*, chapter 20, pages 417–444. MIT Press, Cambridge, MA, 1990.

[11] G. Hoffman and C. Breazeal. Collaboration in human-robot teams. In *Proc. of the AIAA 1st Intelligent Systems Technical Conference*, Chicago, IL, USA, September 2004. AIAA.

[12] H. Jones and S. Rock. Dialogue-based human-robot interaction for space construction teams. In *IEEE Aerospace Conference Proceedings*, volume 7, pages 3645–3653, 2002.

[13] H. Kimura, T. Horiuchi, and K. Ikeuchi. Task-model based human robot cooperation using vision. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS'99)*, pages 701–706, 1999.

[14] G. Knoblich and J. S. Jordan. Action coordination in groups and individuals: learning anticipatory control. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(5):1006–1016, September 2003.

[15] R. Kohavi, B. Becker, and D. Sommerfield. Improving simple bayes. In *Proceedings of the European Conference on Machine Learning*, 1997.

[16] T. Komatsu and Y. Miyake. Temporal development of dual timing mechanism in synchronization tapping task. In *Proceedings of the 13th IEEE International Workshop on Robot and Human Communication (RO-MAN 2004)*, September 2004.

[17] L. Lambert and S. Carberry. A tripartite plan-based model of dialogue. In *Proceedings of the 29th conference on Association for Computational Linguistics*, pages 47–54. Association for Computational Linguistics, 1991.

[18] H. J. Levesque, P. R. Cohen, and J. H. T. Nunes. On acting together. In *Proceedings of AAAI-90*, pages 94–99, Boston, MA, 1990.

[19] C. Rich, C. L. Sidner, and N. Lesh. Collagen: Applying collaborative discourse theory to human-computer collaboration. *AI Magazine*, 22(4):15–25, 2001.

[20] N. Sebanz, H. Bekkering, and G. Knoblich. Joint action: bodies and minds moving together. *Trends in Cognitive Sciences*, 10(2):70–76, 2006.

[21] C. Sidner, C. Lee, and N. Lesh. Engagement rules for human-robot collaborative interaction. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, volume 4, pages 3957–3962, 2003.

[22] M. Wilson and G. Knoblich. The case for motor involvement in perceiving conspecifics. *Psychological Bulletin*, 131:460–473, 2005.

[23] H. Woern and T. Laengle. Cooperation between human beings and robot systems in an industrial environment. In *Proceedings of the Mechatronics and Robotics*, volume 1, pages 156–165, 2000.