Observer-Aware Action Sequence Planning for Human-Robot Collaboration

Chang Liu

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2016-115
http://www.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-115.html

May 20, 2016

Acknowledgement

The content in this report is mainly based on the joint work with Jaime F. Fisac and Jessica B. Hamrick, under the supervision of Prof. Anca D. Dragan, J. Karl Hedrick, Thomas L. Griffiths and S. Shankar Sastry. I would like to express my sincerest gratitude to all of them for their kind help and guidance. I also would like to thank ONR MURI project (ONR-N000141310341) who supports this work. Any opinions, findings, conclusions, or recommendations in this paper are solely those of the authors and do not necessarily reflect the position or the policy of the sponsors.

Acknowledgement

The content in this report is mainly based on the joint work with Jaime F. Fisac and Jessica B. Hamrick, under the supervision of Prof. Anca D. Dragan, J. Karl Hedrick, Thomas L. Griffiths and S. Shankar Sastry. I would like to express my sincerest gratitude to all of them for their kind help and guidance.

I also would like to thank ONR MURI project (ONR-N000141310341) who supports this work. Any opinions, findings, conclusions, or recommendations in this paper are solely those of the authors and do not necessarily reflect the position or the policy of the sponsors.

Abstract

Collaboration between humans and autonomous agents requires the ability to infer and adapt to other agents' plans while effectively conveying one's own intent. In some cases, our teammates' actions early on can give us a clear idea of what the remainder of their plan is, that is, what action sequence we should expect; in others, they might leave us less confident, confused or even lead us to the wrong conclusion.

In this work, we use a Bayesian model of how people make such predictions in order to facilitate the interpretation of robot plans by human collaborators. We subsequently propose the concept of t-predictability to quantitatively describe an action sequence in terms of its easiness for expressing the entire plan. A t-predictable planner is then developed to generate action sequences that purposefully maximize the expected accuracy and confidence with which human observers can predict the overall plan from only the initial few actions. Through an online experiment and an in-person user study with physical robots, we find that t-predictable planner outperforms a traditional optimal planner in objective and subjective collaboration metrics. We believe that t-predictability will play a significant role for improving human-robot collaboration.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

With robots becoming more capable and versatile, they are now stepping into mixed workspaces shared with human beings. For example, in manufacturing, human workers and industrial robots work side-by-side to collaboratively assemble products. The popularity of robots has led to increasing interests in human-robot collaboration [1, 2, 3].

Achieving seamless team collaboration requires team members to understand the intent and anticipate needs of other teammates [4]. For a team composed of only robots, this is often achieved by constantly exchanging information among team members through, for example, wireless communication networks. For human-robot collaboration, however, frequent and effective communication is more difficult. Reflecting upon the history of human-automation systems, there have been some serious records of failures due to ineffective interaction between machines and their operators [5, 6, 7]. The most common causes for this are mode confusion and "automation surprises", i.e. *misalignments* between what the automated agent is planning to do and what the human believes it is planning to do.

Our aim in this work is to eliminate such misalignments and facilitate human-robot collaboration: we want humans to be able to infer what a robot is planning to do during a collaborative task by observing robot's actions. We choose motion as the communication channel for conveying a robot's plan. This is an important topic since motion is a natural way of expressing intent and plan in reality. However, it is less investigated in robotics; most of works on motion planning in robotics focus on the functional aspects of motion, such as moving along the minimum-energy trajectory or a collision-free path.Researchers have also studied speech for effective communication. However, language has its own ambiguities and can lead to miscommunication. In addition, people often make inferences based on actions besides the language, whether intended or not. The discrepancy between the intent expressed by language and by actions can cause confusion and impair collaboration. Moreover, incessant explanations can be aggravating. Therefore, investigating the intent conveyance via actions is a natural and important topic for human-robot collaboration.

Conveying intent via motion enables the human to understand her partner's intent through natural interaction, adapt her own actions to be compatible with the robot's, and more effectively achieve common goals[8, 9, 10]. Traditionally, works on human-robot collaboration have focused on inferring either human's next action or human's overall plans, and adapting the robot's plan in response [11, 12]. Effectively communicating a robot's next goal based on its ongoing motion has also been explored in robotics recently [13], which proposed algorithms for computing *legible* and *predictable* robot actions. However there has been little investigation into effectively conveying robots' high-level plans (action sequences) to humans in order to facilitate collaboration in mixed human-robot teams.

In this work, we introduce a framework that explicitly predicts the inferences that human observers are likely to make for the robot's plan based on its initial actions, and incorporates them into the planning process in order to generate plans that are more easily and unambiguously interpretable by human collaborators. Such framework relies on the connection between models of human planning and that of human inference that was proposed in cognitive science [14]. It models the human as a noisy observer who will infer an agent's plan with exponentially decaying probability as its cost increases, which we expect that a similar mechanism will also apply to many peer-to-peer collaboration scenarios.

We focus on communicating the *sequences of future actions* from the *initial actions* and a *known goal*. Such situations usually appear in *task planning*, where the overall goal of the task is clear (completing all subgoals), but the sequence of actions the robot will take to achieve the goal is not. The key difference between our work and previous works is the richness of information that needs to be conveyed. Prior works have developed algorithms for generating trajectories that communicate the robot's *overall goal* from the ongoing *motion* [15, 13, 16, 17] while our work conveys the sequences of remaining actions.

1.2 Literature Review

1.2.1 Intent Recognition

In human-robot collaboration, the ability to infer an agent's intent composes a significant element of understanding its needs and generating collaborative actions accordingly. The Bayesian approach for intent recognition, which relates the observation of an agent's behavior with its unobservable intent, has been adopted in plan recognition [18], cognitive science [14] and the perception of human action [19]. Some typical formalisms based on Bayesian approach include Hidden Markov

Models [20, 21], Dynamic Bayesian Networks (DBNs) [22] and Markov Decision Processes (MDPs) [23].

The key component in Bayesian approach is the relation describing an agent's action generation given its intent. The Boltzmann policy [14], or the Luce-Shepard choice rule [24], is one of the most basic models for such choice behavior. It asserts that an agent is noisily optimal in behavior selection: the closer an action is to the optimal one in some metric (e.g. reward), the higher probability such action will be taken by the agent.

1.2.2 Intent Expression

Researchers have long studied how people use nonverbal communications in their interactions with one another [25, 26]. By making the robots' intended actions more readily apparent to their partners, we can improve people's abilities to coordinate their actions with that of robots. This line of research has led to works on anticipatory motion [35] and readable behavior [34].

Recent works have investigated different techniques for effectively expressing robots' intents, including motion trajectory planning [27], designing gesture [28], gaze [29, 30] and orientation [31]. Especially, people have incorporated animation principles [32, 33] to produce intentional motion, which focused on pre- and post- action expressions of forethought and reaction to show robots' "thinking of action" [34], and attracting observer attention via exaggerated motion synthesis [35].

1.2.3 Observer-Aware Motion Planning

The motion planning problem focuses on generating a sequence of actions or trajectories that achieves some objectives while satisfying certain constraints. Traditional works on robotic motion planning focused on functional motion such as achieving the minimum travel distance from an initial configuration to the final one while obeying the kinodynamic constraints of a robot.

Recent progress in human-robot interaction community has given rise to producing motion that is mindful of observer inferences. Such motion planners reason about the inferences that humans will make when observing the robot's behavior. The understanding of a robot's intent is important for seamless human-robot collaboration.

There exist two types of complementary inference that humans can make to relate an agent's actions and goals: action-to-goal and goal-to-action [36]. The "action-to-goal" inference refers to the observer's ability to infer an agent's goal state based on the understanding of the function of an action; the "goal-to-action" inference, on the contrary, refers to an observer's ability to predict the actions that an agent will take given the knowledge of its goal. These two types of inference leads to the ideas of *predictability* and *legibility* of motion [37].

Predictability usually refers to the property of motion that matches an observer's expectation when the goal of the motion is known to the observer. Some related works include [38, 39]. Legibility implies that an observer can understand or recognize the intent of the robot by the knowledge of its action, when the observer does not know the robot's intent a priori. Some recent works on legible motion planning include [28, 40, 27]. Previous works usually focus on the motion-level planning for intent conveyance about a single goal. Here, we study the predictability of a plan on task level, which consists of multiple subgoals.

1.2.4 Task Planning in Human-Robot Collaboration

Task planning in human-robot collaboration depends on the relation between the human and the robot during the interaction. Assistant robots are focused on the current task the human is conducting. When a need for assistance is detected, the robot plans for the best action to assist the human, given the current situation and human's intended task [41, 23]. Multi-modal communication is usually used for establishing and improving the understanding of human's needs [42].

For a human-robot team sharing a task, the planning usually involves each agent to decide and adapt their own collaborative actions based on other agents' behavior [43]. Tasks involving multiple subgoals, such as fetching objects from different positions and delivering them to specified location, usually requires effective assignment and scheduling of subgoals. Approaches have been developed for efficient online computation and adjustment of sequence of subgoals [44].

Robot's reasoning about human's intent also plays a significant role in task planning. Besides the intent recognition approaches described in Section 1.2.1, researchers have also investigated reasoning approached based on logic [45] and learned human models [46]. Such approaches have shown to be effective in enhancing task planning for human-robot collaboration.

1.3 Goals and Contributions of the Report

This report makes an initial effort in defining a framework that explicitly accounts for human's inference of robot's plan, and generating robot plans that are unambiguously intelligible by human partners. The report makes the following contributions:

- **Defining t-predictability.** In this work, we define a property of a robot plan that we refer to as *t-predictability*: a plan is t-predictable if a human can infer the robot's future actions in a task from having observed only the first t actions.

- **An algorithm for generating t-predictable plans.** Building on Bayesian inference models that have been used for action interpretation and plan recognition [47, 14], we propose a model for human's inference of robot's future plans. We then propose a planning algorithm that generates t-predictable plans, that is, the plans that a human can easily infer the rest of actions from the initial ones.

- **Extensive user studies.** We test the effects of t-predictable plans on human's inference via an extensive online user study and investigate the implications on human-robot collaboration via an in-person study.

  In the online user study, participants first observe a robot's partial plan, and then predict the order of its remaining actions. It was found that participants performed significantly better at anticipating the correct order when the robot is planning for t-predictability. A very high correlation (0.88) was spotted between our model's prediction of the probability of success and the participants' actual success rate.

  The in-person experiment involved human participants in a collaborative task with the robot. We analyzed the advantages of t-predictabilityon both objective and subjective collaboration metrics. Experiment results indicate that participants were more effective at completing tasks and preferred to work with a t-predictablerobot than with an optimal robot.

1.4 Outline

Chapter 2 introduces the concept of t-predictability and formulates the t-predictable planner. Chapter 3 presents the experimental design and results for the online study and Chapter 4 is for the in-person user study. Chapter 5 concludes the report with some ideas of our future work.

Chapter 2

Conveying Intent via Planning of Action Sequences

In this chapter, we consider the conveyance of robot intent to a human observer via planning of action sequences. Different from traditional works that unambiguously express an agent's *goal* via the design of motion trajectory, our work focuses on that, when the *goal* is given, the whole action sequence is clearly conveyed when only the partial sequence is shown. Image watching a robot starting to execute a task. Our goal is to make the robot do so in a way that makes it clear how the robot will finish the task, i.e., what is the remaining sequence of actions. We propose the concept of t-predictability and formulate the planning of t-predictable sequences as an optimization problem over the space of all possible plans.

2.1 Definition of t-predictability

We consider a multi-step task with an overall goal G (such as visiting all targets in an area) that can be achieved through a finite-horizon sequence of T actions. To quantitatively describe the property of a sequence in terms of the easiness of inferring an agent's plan, we define the concept of t-predictability . Let $\mathcal{A}$ denote the space of all possible action sequences of length T that achieve the goal.

Definition (t-predictability) A sequence of actions $\mathbf{a} = [a_1, a_2, ..., a_T] \in \mathcal{A}$ that achieves an overall goal G is t-predictable if an observer can accurately infer $[a_{t+1}, ..., a_T]$ after observing $[a_1, ..., a_t]$, and knowing the overall goal G.

Definition (t-predictable planner) A t-predictable planner generates the plan that will maximize the probability that a human observer with knowledge of the goal G will correctly predict all subsequent actions after observing the first t actions.

Note that for $t = 0$, the t-predictability of a plan simply becomes its *predictability*, that is, the ease with which the entire sequence of actions can be inferred with knowledge of the overall goal G [48].

A t-predictable planner thus returns the optimal action sequence $\mathbf{a}^* \in \mathcal{A}$ that maximizes the conditional probability that a human observer makes correct prediction:

$$\mathbf{a}^* = \arg\max_{\mathbf{a}\in\mathcal{A}} P(a_{t+1}, .., a_T|S, G, a_1, .., a_t) \tag{2.1}$$

with S being the starting state. This is equivalent to:

$$\mathbf{a}^* = \arg\max_{\mathbf{a}\in\mathcal{A}} \frac{P(a_1, ..., a_T|S, G)}{\sum_{[\tilde{a}_{t+1},...,\tilde{a}_T]} P(a_1, ..., a_t, \tilde{a}_{t+1}, ..., \tilde{a}_T|S, G)}. \tag{2.2}$$

To compute this, we need a model of $P(\mathbf{a}|S, G)$, which we discuss in the next section.

Remark t-predictability is a measure of the degree of confidence with which an observer will accurately infer a plan, which means we can quantitatively computes the t-predictability score of any sequence. The t-predictable planner returns the action sequence with highest t-predictability score among all possible sequences. It is possible that highly t-predictable sequence may not exist in some cases, which means there exist at least two sequences with close t-predictability score. However, in many other cases, highly t-predictable sequence exists and optimizing for t-predictability is highly useful, as evidenced by our experiment results.

2.2 Boltzmann Noisy Rationality

In psychology and cognitive science, Boltzmann probabilistic models of human noisy optimality have been used in the context of *goal* inference through inverse action planning [14]. It assumes that the human observer expects the other agent to be noisily optimal in taking actions to achieving a goal. We adopt an analogous model for modeling the inference of *action sequences*: the human is now modeled as expecting the robot to be noisily optimal, taking approximately the optimal sequence of actions to achieve G. This is actually consistent with recent cognitive research on human understanding of complex plans [49], although our approach takes a higher-level perspective and considers each step in the overall plan as the basic action.

We consider the *open* Traveling Salesman Problem, in which there is no restriction on the final location after visiting all targets. Each action consists of reaching a target $i \in \{1, .., T\}$ with location x_i in a metric space (χ, d). For the purpose of notational simplicity, we let $\mathcal{A} = \{1, \ldots, T\}$ and applying a_i at location $x_{a_{i-1}}$ triggers the transition to location x_{a_i}. The cost of an action sequence $\mathbf{a} \in \mathcal{A}$ is the cumulative distance traveled to visit all target locations in the sequence from a fixed initial state S with location x_S to complete the goal G from the space of goals $\mathcal{G}$.

We define optimality via some cost function $c : \mathcal{A} \times \mathcal{S} \times \mathcal{G} \to \mathbb{R}^+$, mapping each action sequence for a particular goal and from a starting state to a scalar cost. In this work, we use path length (travel distance) for c:

$$c(\mathbf{a}, S, G) = d(x_{a_1}, x_S) + \sum_{j=2}^{T} d(x_{a_j}, x_{a_{j-1}}), \tag{2.3}$$

where x_l, $l \in \{S, a_1, \ldots, a_T\}$ denotes the state of the robot; $d : \chi \times \chi \to \mathbb{R}^+$ is the cost function for state transition. The overall goal for the robot is to sequentially reach all targets T in the scene.

Applying the Bolzmann policy [14] based on c, we get:

$$P(\mathbf{a}|S, G) = \frac{e^{-\beta c(\mathbf{a}, S, G)}}{\sum\limits_{\tilde{\mathbf{a}} \in \mathcal{A}} e^{-\beta c(\tilde{\mathbf{a}}, S, G)}}. \tag{2.4}$$

Here $\beta > 0$ is termed the *rationality coefficient*. As $\beta \to \infty$ the probability distribution converges to one for the optimal sequence and zero elsewhere, corresponding to the case of a rational agent. As $\beta \to 0$, the probability distribution becomes uniform over all possible sequences $\mathbf{a}$ and the agent is indifferent; in such case, no intent information can be obtained by merely observing the agent's actions.

Remark The method proposed in this chapter is a general framework for generating predictable paths. We choose the TSP scenario since it is a classical problem that people are familiar with. In fact, humans are known to perform remarkably well at the TSP for as many as 20 targets [50]. Therefore we expect that humans can nicely infer the robot's plan by conducting mental simulation. Using TSP also allows us to isolate and measure the effects of t-predictability without confounding factors arising from structural complexity.

2.3 t-Predictability Optimization

The solution to Equation (2.2) can be computed by enumeration in problems with small T. For larger problems, approximation methods may be used. Incorporating

Equation (2.4) into Equation (2.2), it follows that

$$\mathbf{a}^* = \arg\max_{\mathbf{a}\in\mathcal{A}} \frac{\exp\big(-\beta c([a_{t+1},...,a_T],a_t,G)\big)}{\sum\limits_{[\tilde{a}_{t+1},...,\tilde{a}_T]\in\mathcal{A}_{\mathbf{a}}^t}\exp\big(-\beta c([\tilde{a}_{t+1},...,\tilde{a}_T],a_t,G)\big)}$$
$$= \arg\max_{\mathbf{a}\in\mathcal{A}} \frac{\prod\limits_{j=t+1}^{T}\exp\big(-\beta d(x_{a_j},x_{a_{j-1}})\big)}{\sum\limits_{[\tilde{a}_{t+1},...,\tilde{a}_T]\in\mathcal{A}_{\mathbf{a}}^t}\prod\limits_{j=t+1}^{T}\exp\big(-\beta d(x_{\tilde{a}_j},x_{\tilde{a}_{j-1}})\big)}, \tag{2.5}$$

with $\mathcal{A}_{\mathbf{a}}^t$ denoting the set of all permutations of the $T-t$ targets remaining after eliminating the first t targets in $\mathbf{a}$ from the original set $\{1,...,T\}$.

Our insight is that initial actions can be used to clarify what future actions will be. We find that in many situations, the robot can select initial actions that might seem somewhat surprising at first, but that make the remaining sequence of actions trivial to anticipate (or "auto-complete").

Chapter 3

Evaluation via Online Experiment

We have set up an online experiment to test the effectiveness of our t-predictable planner in expressing the robot's plan. To be specific, we designed a web-based virtual human-robot collaboration experiment where the human had to predict the behavior of three different robot avatars that uses different planners. Participants first watched the robots move to a number of targets (either zero, one, or two) and then had to predict the sequence of remaining targets the robot would complete. The experiment was conducted via the Amazon Mechanical Turk using the psiTurk experimental framework [51]. We describe the setup of the experiment and analyze the results in the following sections.

3.1 Experiment Setup

3.1.1 Independent Variables

We manipulated two variables: the t-predictable planner with t={0,1,2}, and the number of observed targets k, with $k = \{0, 1, 2\}$.

Planner. There were three different planners that differed in their optimization criteria, which was the number t of targets assumed known to the observer. A participant interacted with three robot avatars, each using one of the following three planners:

Optimal (0-predictable): This robot chooses the minimum-cost action sequence, starting from the initial location and visiting all target locations once; that is, the "traditional" solution to the open TSP. This robot can be equivalently thought of as solving Equation (2.5) for $t = 0$.

1-predictable: This robot chooses the action sequence that solves Equation (2.5) for $t = 1$; the sequence might make an inefficient choice for the first target in order to make the sequence of remaining targets very clear.

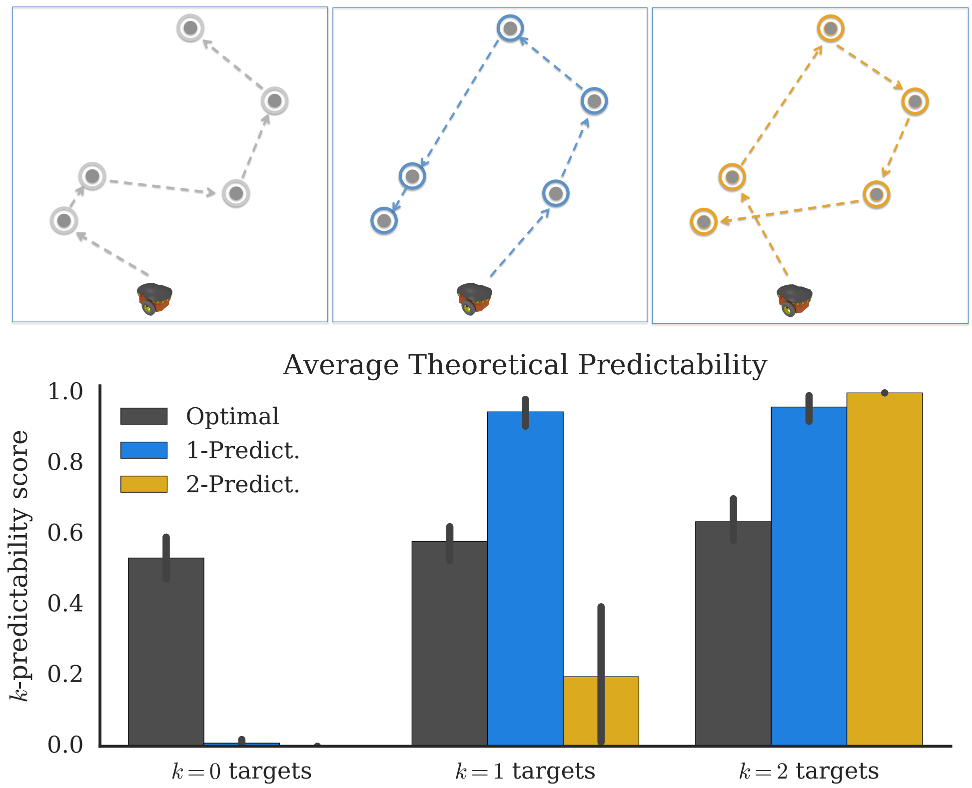

Figure 3.1: **Average theoretical predictability.** Top: a typical task layout where the t-predictable planners generate three different sequences for different t's. Bottom: the average theoretical predictability of task layouts used in the online experiment under different numbers of observed targets (k) of sequences generated by different t-predictable planners. In this plot, the lowest error rates occur when $k = t$.

2-predictable: This robot chooses the strategy solving Equation (2.5) for $t = 2$; the sequence might make an inefficient choice for the first *two* targets in order to make the sequence of remaining targets very clear.

Number of observed targets. Each subject was shown the first $k \in \{0, 1, 2\}$ targets of the robot's chosen sequence in each trial and were asked to predict the remainder of the sequence. This variable was manipulated between participants (a between-subject variable); thus, a given participant always saw the same number of targets (k) on all trials.

Example. Figure 3.1 plots the average k-predictability score, which corresponds to the conditional probability in Equation (2.2), for each planner and the number of shown targets across different stimuli. The upper plot shows an example of $0, 1$ and 2-predictable action sequences for the given target layout. The 0-predictable planner is best in the $k = 0$ case, the 1-predictable is best for $k = 1$, and 2-predictableis (marginally) best (almost perfect) for $k = 2$ in terms of the

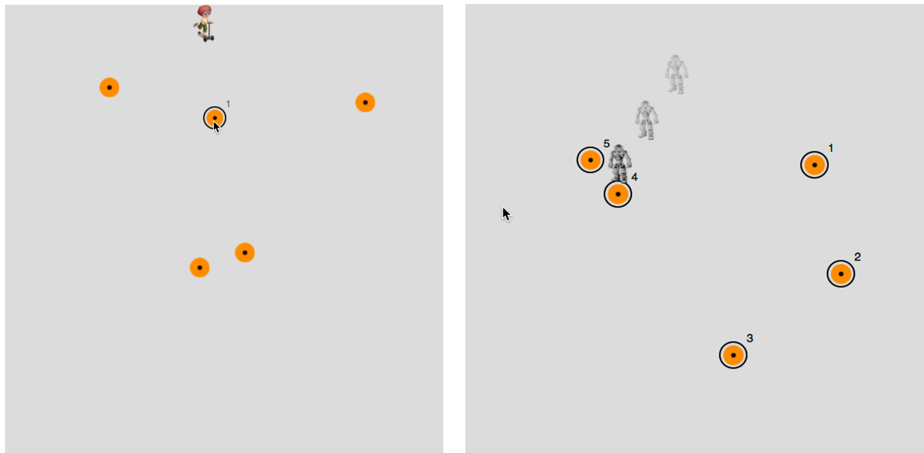

Figure 3.2: **Examples of training and experimental phase.** The left subplot shows the participant can click on a target and the human avatar will move to it. The right subplot shows that the participant first predicts action sequence (numbers indicating the order), then the robot avatar moves to these targets showing its actual sequence.

predictability score. We hypothesize that our results with real users will show the similar trends. Note that the 2-predictable robot might choose a really inefficient sequence for the first 2 targets in order to make the rest maximally clear.

3.1.2 Procedure

The experiment was divided into two phases: a training phase to familiarize participants with TSPs and how to solve them, and an experimental phase. We additionally asked participants to fill out a survey at the end of the experiment.

In the training phase, subjects controlled a human avatar (Figure 3.2). They were instructed to click on targets in the order that they believed would result in the quickest path for the human avatar to visit all of them. The human avatar would start moving to a target in the straight line after each click and "capture" the selected target upon its arrival. The target would be considered as "completed" and be removed from the display.

For the second phase of the experiment, participants first saw a robot avatar move to either $k = 0$, $k = 1$, or $k = 2$ targets using one of the t-predictable plan ($t = 0, 1$ or 2). After moving to these targets, the robot paused so that participants could predict the remaining sequence of targets by clicking on the targets in the order in which they believed the robot would complete them. Afterwards, participants were presented with an animation showing the correct sequence that the robot followed to move to the rest of the targets, determined by the corresponding planner.

Layout Generation. We chose $\beta = 1$ and randomly generated 270 layouts, each layout displaying a square domain with five or six targets in the form of orange circles with a black dot at the center (Figure 3.2). We selected 176 layouts for which the optimal sequence was different between all three planners so that the stimuli would be distinguishable. Further, valid layouts were sorted by the theoretical predictability score in 1-predictability to 2-predictability, to avoid scenarios where the information gain was marginal. Some scenarios were also discarded when the trajectory of the robot closely approached a target without capturing it, to avoid confounds. We finally selected 15 trials based on the aforementioned criteria.

Stimuli. There were a total of 60 trials, consisting of four repetitions of 15 unique target layouts. The first 15 trials were used in the training phase for the participant to become acquainted with the game setup. The following 45 trials were repetitions of the original 15, presented in a random order, with each of the three robots. The trials were grouped so that each participant observed the same robot for three trials in a row before switching to a different robot. In the training trials, the avatar was a gender-neutral cartoon of a person on a scooter, and the robot avatars were images of the same robot in different poses and colors (either red, blue, or yellow) which were counterbalanced across participants.

Controlling for Confounds. We take the following ways to control for confounds:

- counterbalance the colors of the robots for each planner;
- use a human avatar in the practice trials;
- randomize the trial order;
- include the attention checks (described below).

Attention Checks. After reading the instructions, participants were given an attention check in the form of two questions asking them the color of the targets and the color of the robot that they *would not* be evaluating. At the end of the experiment, we also asked them whether they had been paying attention to the difference in helpfulness between the three robots. Data associated with participants who failed the attention checks were not used for analysis.

3.1.3 Dependent Measures

Objective measures. We recorded the proportion of correct predictions of the robot's sequence of targets out of all 15 trials for each planner, resulting in a measure of *error rate*. We additionally computed the *Levenshtein distance* between predicted and actual sequences of targets. This is a more precise measure of how similar participants' predictions were to the actual sequences produced by the planner.

Subjective measures. After every ninth trial of the experiment, we asked participants to indicate which robot they preferred working with. At the end of the experiment, each participant was also asked to complete a questionnaire (adapted

Table 3.1: Subjective Measures

t-predictability
1.0 The (color) robot's sequence of tasks is easy to predict.
1.1 The (color) robot's sequence of tasks is easy to predict after seeing the first task.
1.2 The (color) robot's sequence of tasks is easy to predict after seeing the first two tasks.
2.0 I find it hard to anticipate the order in which the (color) robot will complete the tasks.
2.1 I find it hard to guess the rest of the (color) robot's sequence after seeing it do the first task.
2.2 I find it hard to guess the rest of the (color) robot's sequence after seeing it do the first two tasks.
Consistency
1.1 The (color) robot behaves in a coherent way for the first task.
1.2 The robot behaves in a coherent way for the first two tasks.
2. The start of the (color) robot's sequence (before my prediction) does not make any sense.
Capability
1. The (color) robot seems to know what it's doing.
2. The (color) robot chose its actions poorly.
Helpfulness
1. Overall, I feel that the robot's choice of sequence made my prediction job easier.
2. The robot's choices were confusing and made my prediction job less straightforward.

from [40]) to evaluate their perceived performance of three robots. The questionnaire, as shown in Table 3.1, consists of questions evaluating human's perception of the robots in terms of t-predictability, consistency, capability and helpfulness. A 7-level Likert scale was used to represent degrees from "Strongly agree" to "Strongly disagree". There was an open question (optional) at the end for participants to leave comments about their feelings of the robot.

Since each participant only experience a fixed k, the number of observed targets, in all trials, different questions were designed for k and presented to corresponding people. These is shown in questions numbered 1.x and 2.x in"t-predictability" and "Consistency" categories in Table 3.1. In addition, "Consistency" questions were only asked to participants who saw either one or two targets ($k = 1$ or 2). The "(color)" was replaced by the corresponding robot's color in the actual questionnaire given to participants.

3.1.4 Hypotheses

H1 - Comparison with Optimal. *When showing 1 target, the* 1*-predictable robot will result in lower error than the optimal baseline. When showing 2 targets, the* 2*-predictable robot will result in lower error than the optimal baseline.*

H2 - Generalization. *The error rate will be lowest when* $t = k$*: the number of targets shown,* k*, equals the number of targets assumed by the* t*-predictable planner,* t*.*

H3 - Preference. *The perceived performance of the robots is highest when* $t = k$*.*

3.1.5 Participants

We recruited a total of 242 participants from Amazon's Mechanical Turk, among which 42 participants were excluded from analysis due to failing the attention checks, leaving a total of $N = 200$ participants whose data we used. All participants were treated in accordance with local IRB standards and were paid \$1.80 for an average of 22 minutes of work, plus an average bonus of \$0.47. Bonuses could range from \$0.00 to \$1.35 depending on performance. In the training trials, participants could get a \$0.03 bonus on each trial if they chose the shortest possible sequence; a \$0.02 bonus if they chose a sequence within the top 5% shortest sequences; or a \$0.01 bonus if they chose a sequence within the top 10% shortest sequences (relative to all possible sequences). They received no bonus if they chose a sequence that was longer than the top 10% shortest sequences. In the experimental trials, participants could get a \$0.02 bonus on each trial if they correctly predicted the robot's sequence of targets on that trial.

3.2 Results

Model validity

We first looked at the validity of our model of t-predictability with respect to people's performance in the experiment. We computed the k-predictability scores for each task layout under each planner and number of targets the users observed. We also computed people's actual prediction accuracy on each of these layouts under each condition, averaged across participants.

We computed the Pearson correlation between k-predictability scores and participant accuracy, finding a correlation of $\rho = 0.88$ 95% CI $[0.81, 0.91]$; the confidence interval around the median was computed using 10,000 bootstrap samples (with replacement). This surprisingly high correlation strongly suggests that our *model* of how people predict action sequences of other agents correlates with their *actual* predictions.

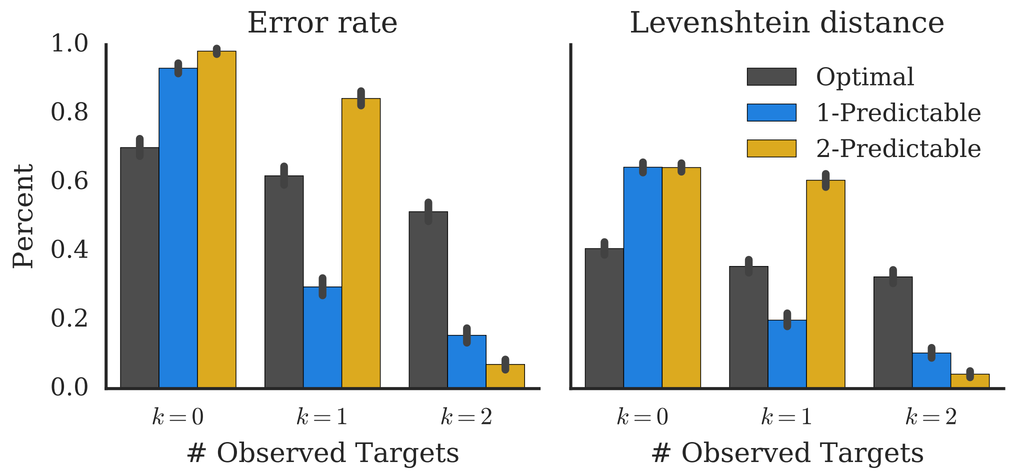

Figure 3.3: **Error rate and edit distance.** The left subplot shows the average proportion of incorrect predictions for different numbers of observed targets (k) of sequences generated by different t-predictable planners. The right subplot shows the average Levenshtein distance between predicted sequences and actual sequences. In both plots, the lowest error rates occur when $k = t$.

Accuracy

To determine how similar people's predictions of the robots' sequences were to the actual sequences, we used two objective measures of accuracy: first, overall error rate (whether they predicted the correct sequence or not), as well as the Levenshtein distance between the predicted and correct sequences. These measures are shown in Figure 3.3.

As the two measures have qualitatively similar patterns of result, and the Levenshtein distance is a more fine-grained measure of accuracy, we performed quantitative analysis only on the Levenshtein distance. We constructed a linear mixed-effects model with the number of observed targets k (k from 0 to 2) and the planner for t-predictability (t from 0 to 2) as fixed effects, and trial layout as random effects.

This model revealed significant main effects of the number of observed targets ($F(3, 10299) = 1894.75$, $p < 0.001$) and planner ($F(2, 42) = 6.59$, $p < 0.01$) as well as an interaction between the two ($F(4, 10299) = 554.00$, $p < 0.001$). We ran post-hoc comparisons using the multivariate t adjustment.

Comparing the planners across the same number of targets, we found that in the 0-targets condition the optimal (or 0-predictable) robot was better than the other two robots; in the 1-target conditon, the 1-predictable robot was better than the other two; in the 2-target prediction, the 2-predictable robot was better than the optimal robot, but only slightly better than the 1-predictable robot. All differences expect the last are with $p < .001$.

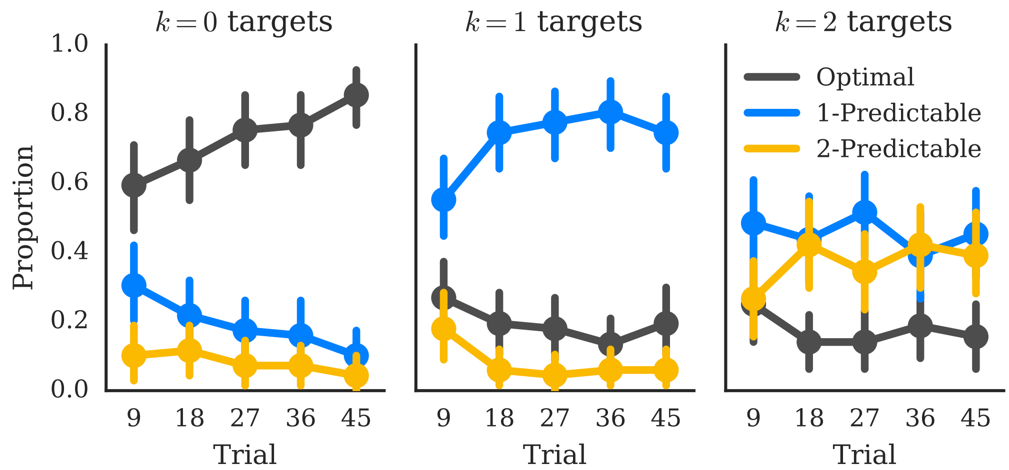

Figure 3.4: **Preferences over time.** Participants prefer the 0-predictable (optimal) robot for $k = 0$ and the 1-predictable robot for $k = 1$, as well as $k = 2$: despite performing slightly better with it, they are frustrated by how inefficient its first 2 actions are.

Comparing the performance of a planner across the number of targets, we found significant differences in all contrasts, with one exception: the accuracy when using the optimal planner was not significantly different when seeing 1 target vs 2 targets ($t(10299) = 2.647$, .1149).

Overall, these results support our hypotheses **H1** and **H2**, that accuracy is highest when t used in the planner equals k, the number of observed targets.

Preferences over time

Figure 3.4 shows the proportion of participants choosing each robot planner as a function of trial. We constructed a logistic mixed-effects model for binary preferences (where 1 meant the robot was chosen, and 0 meant it was not) with planner, number of observed targets, and trial as fixed effects and participants as random effects.

Using Wald's tests, we found a significant main effect of planner ($\chi^2(13.66) = 2$, $p < 0.01$) and trial ($\chi^2(24.68) = 2$, $p < 0.001$). We detected only a marginal effect of number of targets ($\chi^2(4.67) = 2$, $p = 0.097$). However, there was a significant interaction between planner and number of targets ($\chi^2(20.26) = 4$, $p < 0.001$). We also found interactions between planner and trial ($\chi^2(24.68) = 2$, $p < 0.001$) and between number of targets and trial ($\chi^2(16.07) = 2$, $p < 0.001$), as well as a three-way interaction ($\chi^2(39.43) = 4$, $p < 0.001$).

Post-hoc comparisons with the multivariate t adjustment for p-values indicated that for the 0-targets condition, the optimal robot was preferred over the 1-predictable robot ($z = 13.22$, $p < 0.001$) and the 2-predictable robot ($z = 14.56$, $p < 0.001$). For the 1-target condition, the 1-predictable robot was preferred over the optimal robot

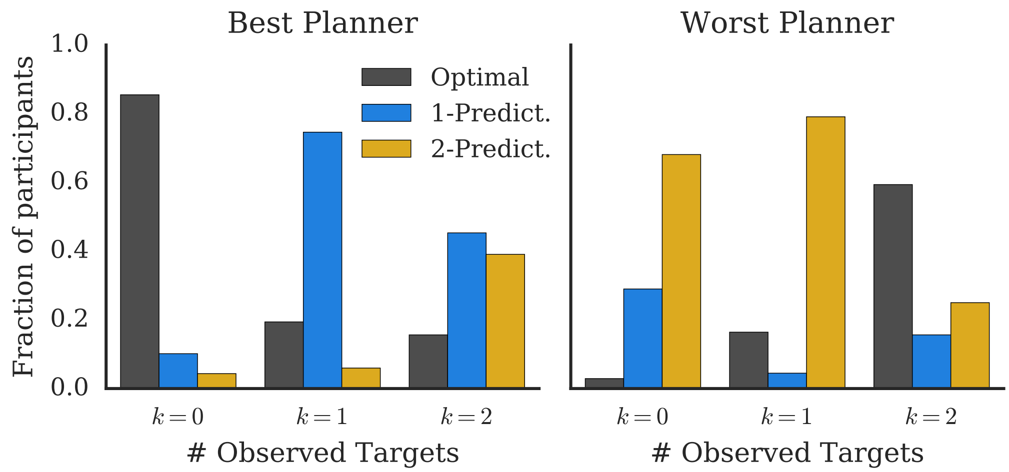

Figure 3.5: **Final rankings.** Participants ranked the planners differently depending on how many targets were observed. For $k = 0$, people preferred the optimal planner; for $k = 1$ and $k = 2$, they preferred the 1-predictable planner.

($z = 12.97$, $p < 0.001$) and the 2-predictable robot ($z = 14.00$, $0 < 0.001$). In the two-task condition, we did not detect a difference between the two 1-predictable and 2-predictablerobots ($z = 2.26$, $p = 0.29$), though both were preferred over the optimal robot ($z = 7.44$, $p < 0.001$ for the 1-predictable robot and $z = 5.40$, $p < 0.001$ for the 2-predictable robot).

Overall, these results are in line with our hypothesis **H3** that the perceived performance is highest when t used in the planner equals k, the number of observed targets. This is the case for $k = 0$ and $k = 1$, but not $k = 2$: even though users tended to perform better with the 2-predictable robot, its suboptimal actions in the beginning frustrated the users, e.g. *"This robot mostly starts out in the worst way and then goes in weird directions but eventually starts to make sense.*[1]*"*.

Final rankings

The final rankings of "best robot" and "worst robot" are shown in Figure 3.5. We used the following procedure to analyze these rankings. For each participant, we assigned each robot a score based on their final rankings. The best robot received a score of 1; the worst robot received a score of 2; and the remaining robot received a score of 1.5. We constructed a logistic mixed-effects model for these scores, with planner and number of observed targets as fixed effects, and participants as random effects; we then used Walds tests to check for effects.

We found significant main effects of planner ($\chi^2(41.38) = 2$, $p < 0.001$) and number of targets ($\chi^2(12.97) = 2$, $p < 0.01$), as well as an interaction between them ($\chi^2(88.52) = 4$, $p < 0.001$). We again performed post-hoc comparisons using the multivariate t adjustment. These comparisons indicated that in the 0-target con-

[1]This excerpt comes from a comment in the questionnaire.

dition, people preferred the optimal robot over the 1-predictable robot ($z = 3.46$, $p < 0.01$) and the 2-predictable robot ($z = 5.60$, $p < 0.001$). In the 1-target condition, there was a preference for the 1-predictable robot over the optimal robot, however this difference was not significant ($z = 2.18$, $p = 0.27$). The 1-predictable robot was preferred to the 2-predictable robot ($z = 6.54$, $p < 0.001$). In the 2-target condition, both the 1-predictable and 2-predictable robots were preferred over the optimal robot ($z = 4.85$, $p < 0.001$ for the 1-predictable robot, and $z = 3.85$, $p < 0.01$ for the 2-predictable robot), though we did not detect a difference between the the 1-predictable and 2-predictable robots robots themselves ($z = -1.33$, $p = 0.84$). Overall, these rankings are in line with the preferences over time.

An informal analysis of the questionnaire suggests similar results as those obtained from the aforementioned measures. Thus we have omitted the analysis of the survey results in this report.

Participant comments

Some participants left interesting comments that reflected their feelings about each robot. These comments are consistent with the subjective metrics and gives us more details about how perceived the robot's actions.

For example, for $k = 0$ (optimal) and $t = 0$, some comments include "I felt like it had a clear way of doing things." and "Generally logical with a few blips here and there"; when $t = 1$, people said that "I saw the pattern was there, but I couldn't find it. Makes me want to keep trying"; however, for $t = 2$, some people obviously got irritated, saying that "Confusing and frustrating trying to predict this guy" and "I hate that stupid robot".

For $k = 1$ and $t = 0$, some participant complained that "Thinks like a dysfunctional computer"; when $t = 1$, people had good impression for the robot as expected, saying that "Logical, similar to me, easy to understand"; when $t = 2$, robot's performance disappointed participants, who said "I want to beat this robot against a wall".

For $k = 2$ and $t = 0$, some participants felt bad about the robot, saying that "I feel like maybe I'm a dumb human and the red robot might be the most efficient, because I have no idea. It frustrated me"; when $t = 1$, people seemed to like the robot, saying that "I like the little guy, he thinks like me"; for $t = 2$, participants liked this robot who said "It was real easy to predict".

Summary

Experimental results show that our t-predictable planner worked as expected, with the t-predictable robots leading to the highest user prediction accuracy given the first t targets. However, focusing on just 2-predictable sometimes frustrated our users due to the the unexpected choice of the first two targets. Overall, we believe t-predictability will be important in a task for all ts, and hypothesize that optimizing for a weighted combination between optimality and t-predictability would perform best in practice.

Chapter 4

Evaluation via In-Person Study

Online experiments have shown strong support for the proposed method's ability to produce t-predictable sequences that clearly convey robot's future plan. We next ran an in-person study to further test the implications of t-predictability. Participants used a smartphone to operate a remote-controlled Sphero BB-8 robot, and had to predict and adapt to the actions of an autonomous Pioneer P3-DX robot in a collaboration scenario (Figure 4.1). A projector was used to display scenarios on the ground. Both robots were tracked using 12 infrared VICON cameras.

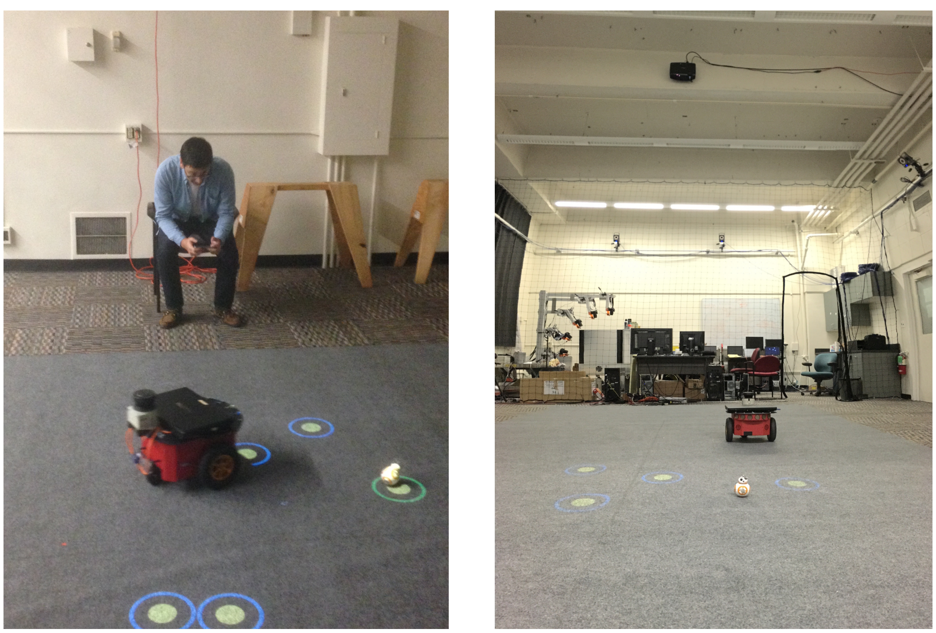

Figure 4.1: **The experiment setup.** Left: a participant is remotely controlling the Sphero BB-8 to clear the target that he believes will be third in the P3-DX's sequence. Right: the VICON system (infrared cameras) and the projector.

4.1 Experiment Setup

4.1.1 Independent Variables

We manipulated one single variable, *predictability*, as a within-subjects factor. Having confirmed the expected effects of 1-predictable and 2-predictable planners, and given the good overall performance of the 1-predictable planner across different conditions, we decided to omit the 2-predictable agent and focus on testing the implications of 1-predictable with respect to optimal in a more immersive collaboration context.

4.1.2 Procedure

At the beginning of the experiment, participants were told that they and their two robot friends were on a secret mission to deactivate an artifact. In each trial, the autonomous P3-DX would navigate sequentially to each of 5 power sources and deactivate it; however, security sensors would activate at each power source after 3 or more had been powered down. The subject's mission was to use BB-8 to jam the sensors at the third, fourth and fifth power sources before the P3-DX arrived at them, by steering BB-8 into the corresponding sensor for a short period of time. If the P3-DX arrived at a jammed sensor in a different order than expected, all other previously jammed sensors would be reactivated, which made anticipating the correct sequence important to ensure success.

In the initial practice phase, participants had a chance to familiarize themselves with the objective and rules of the task, as well as the BB-8 teleoperation interface. Afterwards, there were two blocks of 4 trials. In each block, the subject would collaborate with the P3-DX under a different task planner; participants were told that two different robot "personalities" would be loaded onto the P3-DX. We counterbalanced the order of the conditions.

Stimuli. Each of the 5 power sources (targets) in each trial was projected onto the floor as a yellow circle, using an overhead projector. Each circle was initially surrounded by a projected blue ring representing a dormant sensor. Both the circle and the ring were eliminated when reached by the P3-DX. When the P3-DX reached the third target, the blue circles turned red symbolizing their switch into active state. Whenever BB-8 entered a ring, the ring would turn green for 2 seconds and then disappear indicating successful jamming. If the P3-DX was ever on a red ring, a large red rectangle would be projected symbolizing capture and the trial would end in failure. Conversely, if the P3-DX completed all 5 targets without entering a red ring, a green rectangle would indicate successful completion of the trial. Projections and P3-DX motion were automatically controlled using ROS [52].

Layout Generation. The 4 layouts used were taken from the larger pool of 15 layouts in the online experiment. There was a balance between layouts where online participants had been more accurate with the optimal planner, with the 1-predictable planner, or similarly.

Controlling for Confounds. We take the following methods to control for confounds:

- counterbalance the order of the two planners;
- use a practice layout where both robots showed equivalent action sequence;
- randomize the trial order.

4.1.3 Dependent Measures

Objective measures. We recorded the number of successful trials for each subject and robot planner, as well as the number of trials where participants jammed targets in the correct sequence.

Subjective measures. After every block of the experiment, each participant was also asked to complete a questionnaire, as shown in Table 4.1, to evaluate their perceived performance of the P3-DX robot. At the end of the experiment, we asked participants to indicate which robot (planner) they preferred working with.

4.1.4 Hypotheses

H4 - Comparison with Optimal. *The* 1*-predictable robot will result in more successful trials than the optimal baseline.*

H5 - Preference *Users will prefer working with the* 1*-predictable robot.*

4.1.5 Participants

We recruited a total of 14 participants from the University of California, Berkeley, with various backgrounds. Participants were treated in accordance with local IRB standards and were paid $10 for participation. The study took about 30 minutes on average.

4.2 Results

Successful completions. We first looked at how often participants were able to complete the task with each robot. We constructed a logistic mixed-effects model for completion success with predictability as a fixed effect and participant and task layout as random effects. We found a significant effect of predictability ($\chi^2(11.17) =$

Table 4.1: Subjective Measures

Legibility
1. The robot was moving in a way that helped me figure out what order it was planning to do the targets in.
2. The robot's initial actions made it clear what the rest of the plan was.
3. I found it hard to predict the last three targets in the right order.
4. Overall, I felt that the robot's choice of sequence made my prediction job easier.
Capability
1. The robot seemed to know what it was doing.
2. The robot chose its actions poorly.
Fluency
1. The robot and I worked fluently together.
2. I found it difficult to collaborate with the robot in a fluent way.
Trust
1. The robot was trustworthy.
2. I trusted the robot to do the right thing at the right time.
Predictability
1. I often found the robots overall sequence of targets confusing.
2. I found the robot's first target choice confusing.
3. The robot chose a reasonable sequence of targets.

1, $p < 0.001$), with the 1-predictable robot yielding more successful completions than the optimal robot ($z = 3.34$, $p < 0.001$). This supports **H4**.

Prediction accuracy. We also looked at how accurate participants were at predicting the robots' sequence of tasks, based on the order in which participants jammed tasks. We constructed a logistic mixed-effects model for prediction accuracy with predictability as a fixed effect and participant and task layout as random effects. We found a significant effect of predictability ($\chi^2(9.49) = 1$, $p < 0.01$), with the 1-predictable robot being more predictable than the optimal robot ($z = 3.08$, $p < 0.01$).

Robot preferences. We asked participants to pick the robot they preferred to collaborate with. We found that 86% ($N = 12$) of participants preferred the predictable robot, while the rest ($N = 2$) preferred the optimal robot. This result is significantly different from chance ($\chi^2(1) = 7.14$, $p < 0.01$). This supports **H5**.

Perceptions of the collaboration. Participants' perceptions of the robots' behavior are shown in Figure 4.2. As in the survey results from the online experiment, we analyzed these perceptions by averaging each participant's responses to the individual questions for each robot and measure, resulting in a single score per participant, per measure, per robot. We constructed a linear mixed-effects model for the survey responses with predictability and measure type as fixed effects, and with participants as random effects. We found a main effect of predictability ($F(2, 117) = 16.417$,

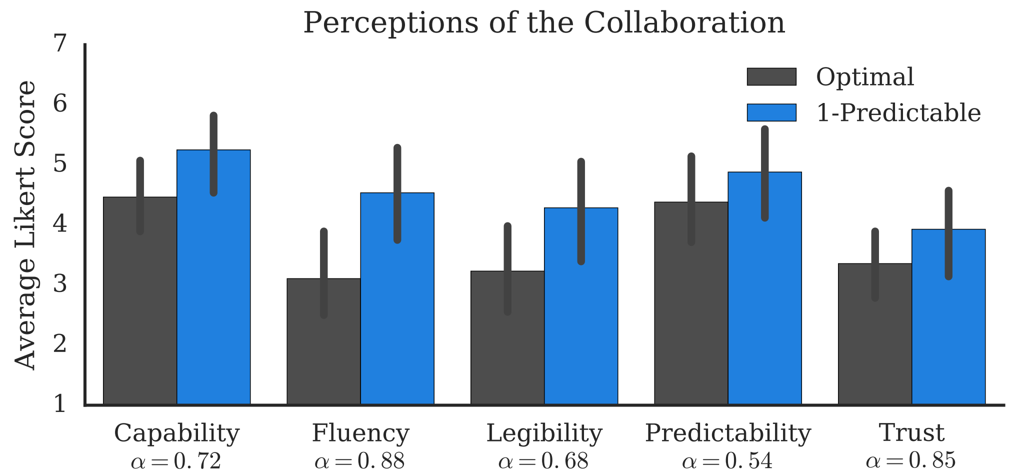

Figure 4.2: **Perceptions of the collaboration.** Over all measures, participants ranked the 1-predictable planner as being preferable to the optimal planner.

$p < 0.001$) and measure ($F(4, 117) = 5.4456$, $p < 0.001$). Post-hoc comparisons using the multivariate t method for p-value adjustment indicated that participants preferred the predictable robot over the optimal robot ($t(117) = 4.052$, $p < 0.001$) by an average of 0.87 ± 0.21 SE points on the Likert scale.

Participant comments. Based on participants' feedback, the BB-8 robot takes some efforts to control, which negatively affected their interaction with the P-3DX robot. However, they still noticed the difference between the two "personalities" of P-3DX. Some said when working with the optimal robot, "bb8 is a little hard to control within the time constraint" while in the legible robot, he/she commented that "I feel easier in the second set of trials". One participant also said that "the second robot seemed clever than the first one. Again, it was hard for me to direct the same robot. I was kind of more focused on the small robot than the big robot to get it go where I wanted". The comments support our hypothesis that the legible robot is easier and desirable to collaborate with. In our future in-person study, we will eliminate the negative effects introduced by the control of a robot to get a more "clean" experiment results.

Summary.

We conducted an in-person user study and the results show that our t-predictability planner worked as expected, with 1-predictable robots leading to the highest user prediction accuracy and the perceived collaboration. The comments from participants also shows people's preference over the 1-predictable robot. The user study further confirms our belief in the importance of t-predictability in human-robot collaboration.

Chapter 5

Conclusions and Future Work

The popularity of robots has led to increasing interest in research on human-robot collaboration. The misalignments between human's belief and the automation's actual plan have unfortunately caused serious issues. In this work, we proposed the concept of t-predictability that quantitatively describes a plan in terms of the easiness that a human observer can accurately predict it after seeing only part of the plan. We then developed the theoretical formulation for computing action sequences, called the t-predictable planner, based on a Boltzmann model of human inference. The t-predictable planner enables a robot to generate a t-predictable plan that a human can confidently infer the rest of the plan after observing the first t actions. We tested the ability to make plans t-predictable in a large-scale online experiment in the setting of an open Traveling Salesman Problem, in which subjects' predictions of the robot's action sequence significantly improved, as well as the perceived robot's performance. The experiment results also show high correlation between the theoretical and the empirical error rates in plan prediction. In an in-person study, we investigated the effects of the t-predictability on human-robot collaboration, indicating that t-predictable can lead to significant objective and perceived improvements compared to traditional optimal planning.

Works in this report have opened several directions for further investigation. First, this work focuses on the plan expression on the task planning level (i.e. action sequences). However, intent conveyance on motion level, such as the "legible" trajectory of a mobile robot, can contain valuable information related to the robot's plan. A combination of task-level and motion-level planning to express a robot's intent deserves further study. Second, in this work, we compare the effects of $t = 0, 1, 2$. However, the choice of t highly depends on the total number of actions to conduct. Finding the optimal t is also an interesting topic to investigate. Third, in our experiment, we choose $\beta = 1$ for generating t-predictable action sequences. In the future work, fitting β using actual human data will be valuable. Lastly, we should note that a human's understanding of robot's plan does not necessarily lead to her adaptation to it. The conditions under which a human decides to adapt is worth investigating. Moreover, the robot can incorporate such knowledge into its planning of actions to further improve human-robot collaboration.

Bibliography

[1] Terrence Fong, Illah Nourbakhsh, Clayton Kunz, Lorenzo Flückiger, John Schreiner, Robert Ambrose, Robert Burridge, Reid Simmons, Laura M Hiatt, Alan Schultz, and Others. The peer-to-peer human-robot interaction project. *AIAA Space*, 2005(September):1–11, 2005.

[2] M Bernardine Dias, Balajee Kannan, Brett Browning, E Gil Jones, Brenna Argall, M Freddie Dias, Marc Zinck, Manuela M Veloso, and Anthony J Stentz. Sliding autonomy for peer-to-peer human-robot teams. In *10th Intelligent Conference on Intelligent Autonomous Systems (IAS 2008), Germany*, pages 332–341, 2008.

[3] Julie Shah and Cynthia Breazeal. An empirical analysis of team coordination behaviors and action planning with application to human-robot teaming. *Human factors*, 52(2):234–245, 2010.

[4] John Yen, Xiaocong Fan, Shuang Sun, Timothy Hanratty, and John Dumer. Agents with shared mental models for enhancing team decision makings. *Decision Support Systems*, 41(3):634–653, 2006.

[5] Duane T. McRuer. Pilot-Induced Oscillations and Human Dynamic Behavior. (July), 1995.

[6] N. B. Sarter and D. D. Woods. Team play with a powerful and independent agent: operational experiences and automation surprises on the Airbus A-320. *Human factors*, 39(4):553–569, 1997.

[7] M. Saffarian, J. C. F. de Winter, and R. Happee. Automated Driving: Human-Factors Issues and Design Solutions. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 56(1):2296–2300, 2012.

[8] Michael Tomasello, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll. Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and brain sciences*, 28(05):675–691, 2005.

[9] Cordula Vesper, Stephen Butterfill, Günther Knoblich, and Natalie Sebanz. A minimal architecture for joint action. *Neural Networks*, 23(8):998–1003, 2010.

[10] Giovanni Pezzulo, Francesco Donnarumma, and Haris Dindo. Human sensorimotor communication: A theory of signaling in online social interactions. *PLoS One*, 8(11):e79876, 2013.

[11] Stefanos Nikolaidis and Julie Shah. Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy. *International Conference on Human-Robot Interaction*, pages 33–40, 2013.

[12] Chang Liu, Shih-Yuan Liu, Elena L Carano, and J Karl Hedrick. A framework for autonomous vehicles with goal inference and task allocation capabilities to support peer collaboration with human agents. In *ASME 2014 Dynamic Systems and Control Conference*, pages V002T30A005–V002T30A005. American Society of Mechanical Engineers, 2014.

[13] Anca Dragan and Siddhartha Srinivasa. Integrating human observer inferences into robot motion planning. *Autonomous Robots*, 2014.

[14] Chris L Baker, Rebecca Saxe, and Joshua B Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–49, dec 2009.

[15] Leila Takayama, Doug Dooley, and Wendy Ju. Expressing thought: improving robot readability with animation principles. *Proceedings of the 6th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 69–76, 2011.

[16] Daniel Szafir, Bilge Mutlu, and Terrence Fong. Communication of Intent in Assistive Free Flyers. *Proceedings of the 2014 {ACM/IEEE} International Conference on Human-robot Interaction*, 2(1):358–365, 2014.

[17] Michael J. Gielniak and Andrea L. Thomaz. Generating anticipation in robot motion. *Proceedings - IEEE International Workshop on Robot and Human Interactive Communication*, pages 449–454, 2011.

[18] Eugene Charniak and Robert P Goldman. A bayesian model of plan recognition. *Artificial Intelligence*, 64(1):53–79, 1993.

[19] Brian D Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, J Andrew Bagnell, Martial Hebert, Anind K Dey, and Siddhartha Srinivasa. Planning-based prediction for pedestrians. In *IROS 2009.*, pages 3931–3936. IEEE, 2009.

[20] Mark Gales and Steve Young. The application of hidden markov models in speech recognition. *Foundations and trends in signal processing*, 1(3):195–304, 2008.

[21] Richard Kelley, Alireza Tavakkoli, Christopher King, Monica Nicolescu, Mircea Nicolescu, and George Bebis. Understanding human intentions via hidden markov models in autonomous mobile robots. In *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*, pages 367–374. ACM, 2008.

[22] Oliver C Schrempf, David Albrecht, and Uwe D Hanebeck. Tractable probabilistic models for intention recognition based on expert knowledge. In *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, pages 1429–1434. IEEE, 2007.

[23] Alan Fern, Sriraam Natarajan, Kshitij Judah, and Prasad Tadepalli. A decision-theoretic model of assistance. In *IJCAI*, pages 1879–1884, 2007.

[24] R Duncan Luce. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2005.

[25] Michael Argyle. *Bodily communication*. Routledge, 2013.

[26] Erving Goffman et al. *The presentation of self in everyday life*. Harmondsworth, 1978.

[27] Daniel Szafir, Bilge Mutlu, and Terrence Fong. Communication of intent in assistive free flyers. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, pages 358–365. ACM, 2014.

[28] Michael J Gielniak and Andrea L Thomaz. Generating anticipation in robot motion. In *RO-MAN, 2011 IEEE*, pages 449–454. IEEE, 2011.

[29] Sean Andrist, Xiang Zhi Tan, Michael Gleicher, and Bilge Mutlu. Conversational gaze aversion for humanlike robots. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, pages 25–32. ACM, 2014.

[30] Bilge Mutlu, Fumitaka Yamaoka, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. Nonverbal leakage in robots: communication of intentions through seemingly unintentional behavior. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 69–76. ACM, 2009.

[31] Allison Bruce, Illah Nourbakhsh, and Reid Simmons. The role of expressiveness and attention in human-robot interaction. In *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*, volume 4, pages 4138–4142. IEEE, 2002.

[32] John Lasseter. Principles of traditional animation applied to 3d computer animation. In *ACM Siggraph Computer Graphics*, volume 21, pages 35–44. ACM, 1987.

[33] Frank Thomas, Ollie Johnston, and Frank. Thomas. *The illusion of life: Disney animation*. Hyperion New York, 1995.

[34] Leila Takayama, Doug Dooley, and Wendy Ju. Expressing thought: improving robot readability with animation principles. In *Proceedings of the 6th international conference on Human-robot interaction*, pages 69–76. ACM, 2011.

[35] Michael J Gielniak and Andrea L Thomaz. Enhancing interaction through exaggerated motion synthesis. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, pages 375–382. ACM, 2012.

[36] Gergely Csibra and György Gergely. obsessed with goals: Functions and mechanisms of teleological interpretation of actions in humans. *Acta psychologica*, 124(1):60–78, 2007.

[37] Anca D Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S Srinivasa. Effects of robot motion on human-robot collaboration. In *HRI 2015*, pages 51–58. ACM, 2015.

[38] Michael Beetz, Freek Stulp, Piotr Esden-Tempski, Andreas Fedrizzi, Ulrich Klank, Ingo Kresse, Alexis Maldonado, and Federico Ruiz. Generality and legibility in mobile manipulation. *Autonomous Robots*, 28(1):21–44, 2010.

[39] Rachid Alami, Alin Albu-Schäffer, Antonio Bicchi, Rainer Bischoff, Raja Chatila, Alessandro De Luca, Alfredo De Santis, Georges Giralt, Jérémie Guiochet, Gerd Hirzinger, et al. Safe and dependable physical human-robot interaction in anthropic domains: State of the art and challenges. In *Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on*, pages 1–16. IEEE, 2006.

[40] Anca Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha Srinivasa. Effects of robot motion on human-robot collaboration. In *HRI*, March 2015.

[41] Vaibhav V Unhelkar, Ho Chit Siu, and Julie A Shah. 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*, pages 82–89. ACM, 2014.

[42] Rachid Alami, Aurélie Clodic, Vincent Montreuil, Emrah Akin Sisbot, and Raja Chatila. Toward human-aware robot task planning. In *AAAI Spring Symposium: To Boldly Go Where No Human-Robot Team Has Gone Before*, pages 39–46, 2006.

[43] Wesley P Chan, Iori Kumagai, Shunichi Nozawa, Yohei Kakiuchi, Kenichi Okada, and Masayuki Inaba. Implementation of a robot-human object handover controller on a compliant underactuated hand using joint position error measurements for grip force and load force estimations. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 1190–1195. IEEE, 2014.

[44] Julie A Shah, Patrick R Conrad, and Brian C Williams. Fast distributed multi-agent plan execution with dynamic task assignment and scheduling. 2009.

[45] Kenton Williams and Cynthia Breazeal. A reasoning architecture for human-robot joint tasks using physics-, social-, and capability-based logic. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 664–671. IEEE, 2012.

[46] Stefanos Nikolaidis, Ramya Ramakrishnan, Keren Gu, and Julie Shah. Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 189–196. ACM, 2015.

[47] Eugene Charniak and Robert P Goldman. A Bayesian Model of Plan Recognition. *Artif. Intell.*, 64(1):53–79, 1993.

[48] Anca D. Dragan, Kenton C T Lee, and Siddhartha S. Srinivasa. Legibility and predictability of robot motion. *ACM/IEEE International Conference on Human-Robot Interaction*, pages 301–308, 2013.

[49] Ryo Nakahashi, Chris L. Baker, and Joshua B. Tenenbaum. Modeling Human Understanding of Complex Intentional Action with a Bayesian Nonparametric Subgoal Model. *arXiv preprint*, 2015.

[50] James N MacGregor and Yun Chu. Human performance on the traveling salesman and related problems: A review. *The Journal of Problem Solving*, 3(2):2, 2011.

[51] T. M. Gureckis, J. Martin, J. McDonnell, R. S. Alexander, D. B. Markant, A. Coenen, J. B. Hamrick, and P. Chan. psiTurk: An open-source framework for conducting replicable behavioral experiments online. *Behavioral Research Methods*, 2015.

[52] Morgan Quigley, Ken Conley, Brian P. Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y. Ng. Ros: an open-source robot operating system. In *ICRA Workshop on Open Source Software*, 2009.