

Show Me What You Can Do: Capability Calibration on Reachable Workspace for Human-Robot Collaboration

Xiaofeng Gao¹, Luyao Yuan, Tianmin Shu, Hongjing Lu, and Song-Chun Zhu¹

Abstract—Aligning humans’ assessment of what a robot can do with its true capability is crucial for establishing a common ground between human and robot partners when they collaborate on a joint task. In this work, we propose an approach to calibrate humans’ estimate of a robot’s reachable workspace through a small number of demonstrations *before* collaboration. We develop a novel motion planning method, reachability-expressive motion planning (REMP), which jointly optimizes the physical cost and the expressiveness of robot motion to reveal the robot’s reachability to a human observer. Our experiments with human participants demonstrate that a short calibration using REMP can effectively align a non-expert user’s estimation of the robot’s reachability with its true capacity. We show that this calibration procedure not only results in better user perception, but also promotes more efficient human-robot collaborations in a subsequent joint task.¹

Index Terms—Human-aware motion planning, human factors and human-in-the-loop, human-robot collaboration.

I. INTRODUCTION

ONE of the main challenges in Human-Robot Interaction is that the capacity of the robot perceived by the human partner may not be consistent with its actual capacity [1]–[3]. Such discrepancy may lead to overuse or misuse of the robot. Particularly, in an ad-hoc teaming setting where humans do not have prior experience with their robot partners, the consequence caused by such discrepancy could be detrimental to the team collaboration [4].

In this work, our key insights to address this challenge are two-fold: i) humans’ perception of the capability of a robot can

Manuscript received September 9, 2021; accepted January 12, 2022. Date of publication January 25, 2022; date of current version January 31, 2022. This letter was recommended for publication by Associate Editor K. S. Haring and Editor G. Venture upon evaluation of the reviewers’ comments. This work was supported by DARPA under Grant XAI N66001-17-2-4029. (*Corresponding author: Xiaofeng Gao.*)

Xiaofeng Gao and Luyao Yuan are with the Center for Vision, Cognition, Learning, and Autonomy (VCLA), University of California, Los Angeles, Los Angeles, CA 90095 USA (e-mail: xfgao@ucla.edu; yuanluyao@ucla.edu).

Tianmin Shu is with the Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: tshu@mit.edu).

Hongjing Lu is with the Department of Psychology, University of California, Los Angeles, Los Angeles, CA 90095 USA (e-mail: hongjing@ucla.edu).

Song-Chun Zhu is with the Beijing Institute for General Artificial Intelligence (BIGAI), Beijing 100080, China, with the School of Artificial Intelligence and Institute for Artificial Intelligence, Peking University, Beijing 100871, China, and with the Department of Automation, Tsinghua University, Beijing 100871, China (e-mail: sczhu@bigai.ai).

Digital Object Identifier 10.1109/LRA.2022.3144779

¹Code and demos are available at <https://xfgao.github.io/calib2022ral/>.

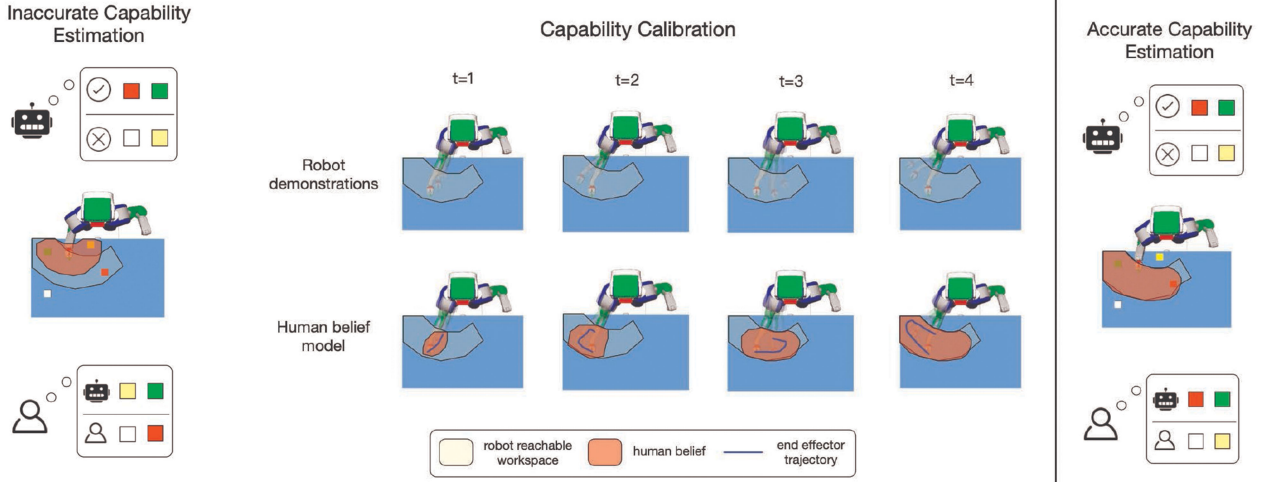
be calibrated by observing its behavior, e.g., robot demonstrating its motion trajectories in pursuit of certain goals, and ii) calibrating the perceived robot capability improves the quality of subsequent human-robot collaboration.

We focus on a case study as shown in Fig. 1, where a human user and a robot share the same workspace, and they must take turns picking up all objects as fast as possible. As the robot can only reach part of the workspace due to its mechanical limits, the human partner needs to pick up the objects that the robot can not reach to achieve maximum efficiency in completing this joint task. We introduce capability calibration as shown in Fig. 1 b, where we allow the robot to show a small number of demonstrations. After watching each demonstration, the human can estimate the robot’s capability accordingly. The goal is to come up with motion plans to pragmatically demonstrate the robot’s capability.

To achieve a sample-efficient calibration, we propose reachability-expressive motion planning (REMP), a novel planning algorithm that models perceived robot capability as a human’s belief over a robot’s reachable workspace, and integrates the belief update into motion planning by introducing an additional cost in trajectory optimization. As a result, REMP can generate a series of expressive trajectories for different robots to showcase their reachability to users. We conducted a user study in which participants i) first observed several robot demonstrations, then ii) estimated where the robot could reach, and iii) proceeded to work with the same robot in a joint task: picking up all objects in the shared workspace as fast as possible. We find that i) REMP significantly increases the accuracy of humans’ reachability estimation, ii) the subsequent human-robot collaboration benefits from a successful calibration, iii) users perceive the robot as more predictable and reliable.

II. RELATED WORK

Perceived robot capability: Various works have studied how humans perceive a robot’s capability, differentiating between social and physical capabilities [5]. In prior work, social capabilities were defined as a robot’s ability to communicate and interact with humans [6], and physical capabilities were defined on a set of tasks a robot can successfully perform [7], such as lifting different objects on the table [8], [9], searching and firefighting in various weather and fire conditions [10]. In these works, robots’ capabilities of different tasks are estimated separately based on experience counts of action outcomes – a higher success rate indicates a stronger capability in a task [11]. Since the capability modeled in these works are highly task dependent, the user’s knowledge of a robot’s capability in one task can not be easily



(a) Inaccurate capability estimation can lead to failure in collaboration. (b) In this example, the human is supposed to pick up the white and the yellow cubes and let the robot collect the red and the green ones.

Fig. 1. (a) Consider a collaborative table clearing task, where the robot has a limited capability and cannot reach the yellow and white objects. Inaccurate estimation of the robot's reachable workspace would harm collaboration: users who incorrectly estimate that the robot can reach the yellow object would assign it to the robot, resulting in a worse teaming performance. (b) We propose capability calibration, where the robot uses its motion to demonstrate its capability before collaboration.

generalized to the knowledge of its capability in other tasks. In contrast, our work focuses on physical capabilities that serve as a basis for a robot to achieve success in a wide range of tasks. In particular, we focus on *reachability*, which is one of the most fundamental physical capabilities for robots. By understanding a robot's reachability, users can better assess its overall capability in various tasks where reaching is involved. Given an arbitrary task, the user can decide whether the robot can successfully perform it based on perceived reachability.

Robot expressive motions: When deploying robots in real-world settings that are beyond factory environments, functional motions only designed to accomplish tasks are inadequate for human users to correctly understand the robots and establish effective collaborations [12]. It is equally important to convey the rationality and intent of a robot through its motion [13], [14]. To generate such motions, prior work formulated and optimized the legibility of trajectories via functional gradient descent [15], [16]. Similar ideas were also adopted to study robots' expression of emotion [17] and style [18]. To express robot (in)capability, prior work used repetitive motions, either generated by simple heuristics [19] or hand-crafted for each task [20]. In contrast, [21] proposed a trajectory optimization-based method that maximizes the similarity between motions and would-be successful executions. Our work takes one further step in this direction: we i) model how humans update their beliefs of the capability of a robot given the observed robot motions and ii) integrate the belief update process into trajectory optimization to generate new motions that can optimally improve humans' beliefs.

III. CAPABILITY CALIBRATION

We propose a capability calibration framework (as shown in Fig. 1 b) to align a human user's understanding of a robot's capability with the ground truth, where the user can watch a small number of demonstrations of her robot partner before

they work together. In this section, we introduce our approach to generate such demonstrations that can efficiently reveal the robot's reachability. We show how this calibration could be applied to collaboration in Section IV.

A. Calibrating Reachable Workspace

In this sub-section, we define the reachability calibration task. In Section III-B, we describe how human belief would be updated before a new trajectory can be generated. In Section III-C, we propose REMP, which enables the robot to generate one trajectory showing its reachable workspace based on a simulated human belief that models what the human has already known about the robot. In Section III-E, we reify our capability calibration framework by combining REMP and task planning. We begin with some notations.

The robot's ground truth reachability is defined as $f : \mathcal{X}_{ws} \rightarrow \{0, 1\}$, i.e. whether a target position x in the workspace \mathcal{X}_{ws} is reachable by the end-effector according to the robot's kinematic constraints. Meanwhile, we assume the human is maintaining a belief $b_h^t : \mathcal{X}_{ws} \rightarrow [0, 1]$, modeling how likely a target is reachable after observing a robot trajectory $\xi_{1:N}^t \in \Xi$ with length N at time $t \in \{1 \dots T\}$. After observing all the robot demonstrations, human's final belief would become $b_h^T = \tau(b_h^0, Z)$, with τ denoting the belief transition from the initial guess b_h^0 to the final b_h^T with all seen demonstrations $Z = \xi^{1:T}$. In addition, we define $\mathcal{X}_{rs} \subseteq \mathcal{X}_{ws}$ as the robot's reachable workspace. $\phi_{ee} : \mathcal{Q} \rightarrow \mathcal{X}_{rs}$ is the forward kinematic function of the end-effector, generating its position given a configuration.

Using notations above, we can formalize capability calibration as an optimizing problem over a set of trajectories, Z , whose cardinality may not necessarily be known in advance. The goal is to reduce the mismatch between the robot's ground truth capability and the user's final belief:

$$\begin{aligned} & \arg \min_{Z \in \Xi^*} Cost(Z) \\ \text{s.t. } & \sum_{x \in \mathcal{X}_{ws}} \left| \tau(b_h^0(x), Z) - f(x) \right| < \epsilon, \end{aligned} \quad (1)$$

where Ξ^* uses the Kleene star to represent all possible sequences of robot motions, $Cost(\cdot) : \Xi^* \rightarrow \mathbb{R}$ is a function to evaluate the overall cost for trajectories. The condition means the user has a reachability estimation close enough to the robot's true capability.

One intuitive cost is the total length of all the trajectories in Z , optimizing which is equivalent to minimizing the cardinality of Z when the trajectories all have similar length. In this letter, we maintain the homogeneity of trajectories by further regulating the start configurations of all trajectories to be in a set of configurations $S \subset \mathcal{Q}_0$ and target positions in a set of positions $G \subset \mathcal{X}_{rs}$.

Due to the size of motion space, an exact solution to (1) is intractable. Thus, we adopt an incremental update: we keep generating new trajectories ξ^t until the user's belief is sufficiently aligned with the robot's capability. Every time we want to expand Z , we first select a pair of starting configuration and target position $(q^t, x_r) \in S \times G$ and then generate a motion using it. We term the former as the task planning problem and the latter as motion planning.

B. Human Belief Model

Our objective is to make people without any knowledge about robotics easily understand the true capacities of a robot. Thus, our human belief model attempts to capture what a novice user may think about a robot's reachability after watching its trajectories. We assume human updates its belief on an interested point x in the workspace after observing a new robot trajectory ξ . Intuitively, if a point is close to the visited positions in an observed trajectory, the human observer would consider it more likely to be reachable. We model the belief update process as an iterative Bayesian inference beginning from a uniform prior:

$$\tau(b^t, \xi^t)(x) = b_h^{t+1}(x) \propto b_h^t(x) p(\xi^t | x) \quad (2)$$

and $p(\xi^t | x) = e^{-\gamma d(\phi(\xi^t), x)}$, where $d(\phi(\xi), x)$ captures the distance between the trajectory ξ and the interested position x . The hyperparameter γ defines how much the human extrapolates the observed trajectory to the points nearby: a large γ means that such extrapolation mainly happens to the point which is very close to the trajectory. In particular, we use the end-effector position ϕ_{ee} as the feature, and compute the squared euclidean distance between the interested position and the closest end-effector position in the trajectory:

$$d(\phi(\xi), x) = \min_i \|\phi_{ee}(\xi_i) - x\|^2. \quad (3)$$

The design of our distance function is motivated by the fact that, given a trajectory, it is straightforward for users to focus on the robot's end-effector which is central to the task, while trying to estimate its reachable workspace.

C. REMP: Reachability-Expressive Motion Planning

Expressing robot reachability is more than randomly moving the end-effector to somewhere in its reachable workspace. Our insight is that it is essential to understand what the human

Algorithm 1: REMP.

- 1 Given a target position x_r and a starting configuration q^t , human belief b_t ;
 - 2 Generate trajectory ξ^t based on b_h^t, q^t , (4);
 - 3 Update human belief b_h^{t+1} using ξ^t , (2);
 - 4 **return** b_h^{t+1}, ξ^t
-

already knows or does not know about the robot, so that every demonstration can communicate as much information to the human as possible. We believe this can be formulated as an optimization problem: finding a new trajectory that would minimize the misalignment between the ground truth reachability and human's updated estimation. We capture the misalignment using a cost function $c(\xi, b_h^t, f)$ and formulate the optimization problem as the following:

$$\begin{aligned} \xi^t = \arg \min_{\xi} & c(\xi, b_h^t, f) + \frac{1}{\lambda} \sum_{i=1}^N \|\xi_{i+1} - \xi_i\|^2, \\ \text{subject to } & \phi_{ee}(\xi_n) = x_r, \text{ collision-free}(\xi). \end{aligned} \quad (4)$$

The first term is an expressiveness cost and the second term is a smoothness cost commonly seen in trajectory optimization. The trajectory at the t -th step is generated by minimizing the sum of the two costs, subjecting to a constraint that requires the end effector to reach a target position x_r at the end of the trajectory.

Assuming each point in the trajectory contributes to the cost independently, we design the cost function based on a value $v_i(\xi_i, b_h^t, f)$, which represents the degree of alignment between human's estimation and the robot's ground truth reachable workspace:

$$\begin{aligned} c_b(\xi, b_h^t, f) &= \alpha \sum_{i=1}^N v_i(\xi_i, b_h^t, f) \\ &= \alpha \sum_{i=1}^N e^{\beta(b_h^t(\phi_{ee}(\xi_i)) - f(\phi_{ee}(\xi_i)))} \end{aligned} \quad (5)$$

A small value v_i suggests that the human observer is underestimating the robot's capability at ξ_i . In that case, we want to facilitate calibration by encouraging the robot to move to ξ_i . On the other hand, we would see a large v_i if the human is over-estimating the capability. In that case, it is beneficial for the robot to avoid reaching points near ξ_i . The hyperparameter α and β control how aggressive the trajectory would be in expressing the capability. We call this cost function c_b , which captures human updated belief. Note that the intuition is if the observer previously underestimates the reachability of a point x , $b_h^t(x) - f(x)$ will be negative and give low cost for trajectories covering x . Hence, trajectories passing through underestimated points are more likely to be chosen. Trajectories including overestimated points, on the contrary, have larger costs and are less likely to be selected.

Static human model: Our key intuition is the human would update its belief of the robot's reachability after observing each trajectory. To test it, we also design a fixed cost function as baseline, assuming an underlying uniform belief model $\forall x, b_{static}(x) = b_0$. The corresponding cost function under the assumption of a static human model is c_s . Note that this




Fig. 2. Simulated human estimation of robot A's reachability map, after observing each demonstration generated by Algorithm 2, measured by Intersection of Union (IoU) between the human estimation and the ground truth. Robot A is a 2-link arm with link lengths 0.1.

Algorithm 2: REMP-T

- 1 Given a list of target position and starting configuration pairs $(x, q)^{1:T}$ and human belief b_h ;
 - 2 **for** $t = 1, \dots, T$ **do**
 - 3 $b_h, \xi^t = \text{REMP}(x^t, q^t, b_h)$
 - 4 **return** $b_h, \xi^{1:T}$
-

Algorithm 3: Calibration on Reachable Workspace.

- 1 Given a set of target positions G , starting configurations S , initial human belief b_0 ;
 - 2 **for** $t = 1, 2, \dots$, **do**
 - 3 $\forall x, b_h^0(x) \leftarrow b_0$;
 - 4 Let $\sigma(\zeta) = \sum_{x \in \mathcal{X}_{ws}} |f(x) - \text{REMP-T}(\zeta, b_h^0)[b_h]|$;
 - 5 $\zeta_t^* \leftarrow \arg \min_{\zeta \in (G \times S)^t} \sigma(\zeta)$;
 - 6 $\sigma_t \leftarrow \sigma(\zeta_t^*)$;
 - 7 **if** $\sigma_t - \sigma_{t-1} < \varepsilon$ **then**
 - 8 **return** $\text{REMP-T}(\zeta_t^*, b_h^0)$
-

baseline generates functional motions that solely aim to finish the physical task of reaching the target. We envision that in reality, users may also learn from these physical motions the robot's capability by interacting with the robot on some tasks, but such learning is not as efficient as the learning in a dedicated calibration phase.

D. Generating Reachability-Expressive Trajectories

Implementation: We implemented our framework using TrajOpt [22] on two kinds of simulated robots in OpenRAVE [23], including a manipulator with 2 links and a PR2 robot. For the 2-link arm, we manipulated its joint limits and link lengths to allow it to have a variety of two-dimensional reachable workspaces. These serve as testing cases for our framework, as we want to study how well the framework copes with reachable workspaces of different sizes and shapes. For the PR2 robot, we didn't do such manipulations since the goal here is to see how practical it is to apply the framework to real robot manipulators. Without loss of generality, we focus on the right arm of the PR2 robot. In practice, we use grid search to find hyperparameters that generate trajectories to maximize the accuracy of reachability estimation in simulation, as described in Section IV-C.


Algorithm 4: Calibration with Fixed T .

- 1 Given a set of target positions $G = \{x_1, \dots, x_N\}$, number of trajectories T , starting configuration set S , initial human belief b_0 ;
 - 2 $\forall x, b_h^0(x) \leftarrow b_0, \delta \leftarrow \infty$;
 - 3 **for** $\kappa \in T - \text{combination}(G \times S)$ **do**
 - 4 $b_h \leftarrow b_h^0$;
 - 5 **for** $t \leftarrow 1$ **to** T **do**
 - 6 $b_h, \xi^t \leftarrow \text{REMP}(\kappa^t(G), \kappa^t(S), b_h)$;
 - 7 $\sigma = \sum_{x \in \mathcal{X}_{ws}} |b_h(x) - f(x)|$;
 - 8 **if** $\sigma < \delta$ **then**
 - 9 $\delta \leftarrow \sigma$;
 - 10 $\xi^{1:T} \leftarrow \xi^{1:T}$;
 - 11 **return** $\xi^{1:T}$
-

Qualitative behaviors: Fig. 2, Fig. 3(a) and Fig. 3(b) show the trajectories generated by the cost functions c_b and c_s for the robots by running REMP iteratively using the updated belief, following Algorithm 2. As both c_b and c_s assume a uniform belief on the robot's reachable workspace at the beginning, the first trajectories generated by these cost functions are almost identical. Starting from the second trajectories, we find that the ones generated using c_b can cover a large part of the robot's reachable workspace. On the contrary, trajectories generated by c_s are more sensitive to the physical cost. Overall, It is clear that REMP accommodates human belief at each time step and tries to traverse uncovered regions to better express the robot's reachability.

E. Planning for Start and Target Pairs

We have shown how REMP can generate an expressive trajectory given a starting configuration and a target position. For a better capability calibration, we also want to optimize the number of trajectories as well as the sequence of starting configurations and target positions. As outlined in Algorithm 3, this could be achieved by Task and Motion Planning (TAMP) [24], where the plans of start and target pairs come from task planning and the trajectories for a given pair comes from REMP. In Algorithm 3, we solve (1) in an incremental manner. Namely, we keep increasing the cardinality of Z until user's belief is aligned with the robot's actual capability. For each size of Z , we find the best sequence of starting configurations from S and target positions from G (line 5 of Algorithm 3). To avoid trajectories that are too short or uninformative, we set S as the set of configurations

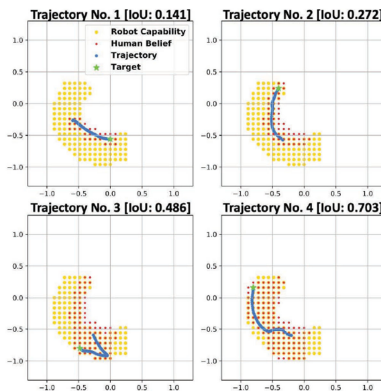


(a) Robot B is a 2-link arm with link lengths 0.13 and 0.07.

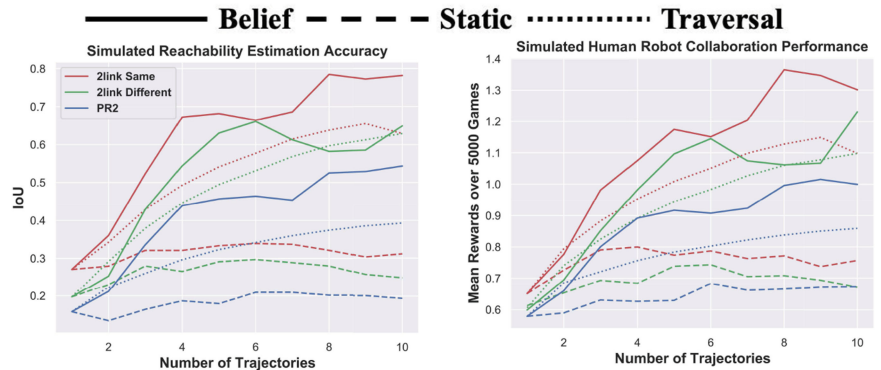


(b) Robot C is a PR2 robot. In this work, we consider the reachable workspace of its right arm.

Fig. 3. Visualization of the robot reachable workspace and the trajectories generated by cost function c_b (*belief*) and c_s (*static*). (a) and (b) show the results for Robot B and Robot C respectively. It can be seen that the *belief* trajectories cover broader regions of the reachable workspace and new trajectories tend to visit areas that haven't been covered by their predecessors. The red dots, corresponding to Fig. 5, represent the points we use to query the users in our experiments.



(a) Trajectories generated by task and motion planning and the simulated reachability estimation given observed trajectories.



(b1) Simulated reachability estimation accuracy, (b2) Simulated human-robot collaboration performance measured by Intersection of Union between the human reachability estimation and the ground truth. Higher value indicates better estimation. (b2) performance measured by averaged rewards acquired by the group. Every point on the curves for *traversal* is the mean of 100 trajectories.

Fig. 4. (a) Combining REMP with task planning, we can optimize the starting and target positions for better calibration. (b) Simulation results of reachability estimation and collaboration performance. Starting and target positions are chosen greedily.

near the neutral configuration of the robot and G as the set of positions far away from the neutral end effector positions:

$$S = \{q, \forall q | q - q_{neutral} | < a_1\} \quad (6)$$

$$G = \{x, \forall x \min_{q \in S} |x - \phi_{ee}(q)| > a_2\} \quad (7)$$

Finding Z incrementally, despite giving the exact optimum, can be time consuming. In practice, rather than demonstrating to humans constantly until converge, we can pre-define T to a reasonable number and find the optimal set of trajectories. In Algorithm 4, we assume fixed number of trajectories. We start from a uniform prior for the human belief, and update the belief w.r.t. (2). Fig. 4(a) depicts an example to optimize 4 trajectories that start from different configurations and reach

different targets. Note that Algorithm 4 plans by enumerating all possible combinations, but any stochastic planning approaches can be used to further accommodate resource constraints and task scalability.

IV. APPLYING REMP TO HUMAN-ROBOT COLLABORATION

In this section, we discuss how to apply REMP to human-robot collaboration *after* the calibration.

A. Collaborative Table Clearing

We design a human-robot collaboration task in a table clearing scenario, where some objects are scattered on a table and a robot

can assist the human with the object collection. The human and the robot take turns picking up the objects. In each step, the human collects first and the robot collects one of the remaining objects. The human can reach all of the objects, while the robot can only reach a subset of them. To finish the task as quickly as possible, the human and the robot need to split the work wisely, so that, in each round, the robot has some objects to pick up. The reward is calculated by the number of objects picked up and the time penalty.

B. Human and Robot Policy

After observing robot expressive demonstrations and updating the belief with (2), the human is assumed to act in an approximately rational way with respect to the current estimation of the robot capability, b_h^t . We use a Boltzmann noisily-rational human decision model [25], assuming the human is more likely to help the robot with its unreachable objects based on the human's current reachability estimation. Since we want to emphasize the effect of the calibration, we use a simple uniform robot policy in the simulation, i.e., it would randomly pick up objects it can reach, and do nothing if no objects are reachable.


C. Simulation Results

Using the behavior model described in the previous sections, we simulated with 3 robots A, B and C with different configurations and reachability: (i) A is a 2-link arm where each link is of equal length, (ii) B is a 2-link arm where the length of its first link (0.13) is larger than the length of the second (0.07), (iii) C is a PR2 robot. The *belief* and *static* methods in the legend correspond to the definition in Section III-B. In addition, we implemented a *traversal* baseline, where the robot moves its end-effector to traverse the workspace to demonstrate its reachability. From the starting position, the end effector moves to unreached waypoints in its reachable workspace one by one. The number of waypoints we sampled corresponds to the number of trajectories in *belief* and *static*.


Fig. 4(b1) and Fig. 4(b2) shows the quantitative results of capability calibration and human-robot collaboration. The result suggests that as the robot shows more demonstrations, the human has a better understanding of its capability and collaborates with it more effectively for all baselines. Looking at the sample efficiency, we notice that without modeling human belief changes, the improvement is quite slow and limited: a large number of demonstrations need to be observed before calibration is achieved. On the contrary, trajectories generated by our proposed REMP algorithm keep providing new information to the user. As a result, the user's estimation accuracy increases much faster for *belief* compared to the baselines. There is fluctuation when many trajectories are shown, due to the limited memory of our human model.

V. USER STUDY

As we have shown the effectiveness of REMP in simulation, we now turn to investigate how much it helps users work with robots in a user study. This study was certified as exempt from IRB review per 45 CFR 46.104 category 3 by the UCLA Institutional Review Board on 9/4/2020.



(a) 36 points for 2-link arms.



(b) 25 points for PR2.

Fig. 5. To evaluate users' estimation of the robot's reachable workspace, we sample query points in the workspace and ask users to select points that they think the robot's end effector can reach. These points correspond to the red dots in Fig. 3(a) and Fig. 3(b).

TABLE I
SURVEY STATEMENTS TO EVALUATE REACHABILITY, PREDICTABILITY,
RELIABILITY AND TRUST TOWARD ROBOTS

1. It is easy to tell where the robot's hand can reach.
2. The robot behaves in a predictable manner.
3. I can rely on the robot to function properly despite its limited capability.
4. I trust the robot.

A. Experiment Design

Participants: We recruited 202 participants (37% Female, median age 34) from Amazon Mechanical Turk.

Materials: During the study, participants interact with the three robots in the simulation as described in Section IV-C. We measure how well participants can understand the robot's capability and how such understanding can help them in the collaboration, as well as their self-reported perception of the robot. To measure **capability understanding**, we ask users to choose positions that they think the robot can reach from a number of object queries, as shown in Fig. 5. We record their selections and compare them with the ground truth. For **collaboration task performance**, we use the accumulated reward of the team as a measure. To measure the **perception of the robot**, we ask participants to rate statements listed in Table I on a 7-point Likert scale labeled from "strongly agree" to "strongly disagree," after they have finished interaction with a robot. Inspired by [26], the statements shown in Table I are designed to evaluate their subjective understanding of the robot in different aspects, including reachability, predictability, reliance and trust.

Calibration task: In the calibration task, the participant would be randomly assigned to an experiment group, and observe T demonstrations. Based on the simulation results in Fig. 4(b), we believe showing more demonstrations would generally lead to a better calibration. Considering the limited time of participants in our online study, however, we cannot use an arbitrarily large T . From simulation, we witness the most significant improvement during the first 4 trajectories, thus we control the number of demonstrations $T = 4$ in practice. After seeing all demonstrations, participants are asked to estimate the robot's reachable workspace by choosing positions that they think the robot can reach from a number of object queries, as shown in Fig. 5.

Collaboration task: In the collaboration task, participants are asked to perform an online table clearing task together with the same robot they have just been calibrated. As discussed in Section IV-A, the task required the team to clear all four objects




Fig. 6. User study results. Here we report means and standard errors. * indicates statistical significant pairs ($p < .05$).

on the table. During each time step, the participant would pick up an object first before the robot makes its decision. The team would get rewarded based on how fast they take all the objects. Since two of the objects cannot be reached by the robot, to get the maximum accumulated reward (+2), the participant needs to pick up objects that cannot be reached by the robot. Failure to do so would result in the team getting a lower reward (0).

Experiment conditions: Like simulations, we varied types of motion users observed in the user study, i.e., the *belief*, *static* and *traversal* methods defined in Section IV-C.

Design: The robot types are within-subject: participants interacted with all three robots. Demonstrations are between-subject: participants only saw demonstrations from one of the three experiment conditions when interacting with a robot.

Procedure: After a brief introduction, each participant is asked to interact with three robots A, B and C in random order. The purpose is to see how robot's physical configurations affect capability perception. During the interaction with each robot, the participants would first go through a calibration task before collaborating with the same robot on the table clearing task.

Hypotheses: We hypothesized the capability calibration framework benefits the users in the following aspects:

H1: Participants going through capability calibration in the *belief* condition would have a better understanding of the robot's capability, compared to those in the other conditions.

H2: Teams in the *belief* condition would perform better in the collaboration tasks than those in the other conditions.

H3: Participants in the *belief* condition would have a more positive perception of the robot, compared to those in the other conditions.

B. Result and Analysis

Capability understanding: We first analyzed the accuracy of the user's estimation of the robot's reachable workspace, by computing the intersection of union (IoU) between their responses and the ground truth. We performed a Kruskal-Wallis H-test of the IoUs using the type of motion as the independent variable. As a result, we found a significant effect for the motion ($\chi^2(2, 603) = 113.52, p < .001$). A post-hoc analysis with Mann-Whitney U test revealed that all three conditions are different from each other, with *belief* significantly better than *static* ($p < .001$) and *traversal* ($p < .001$). This confirms

our hypothesis *H1*. Fig. 6(a) shows the accuracy of participants' reachability estimation w.r.t different robots. On average, *belief* performs 65% better than *static* and 32% better than *traversal*. Compared to the simulation results in Fig. 4, the user study result follows relatively the same order for different conditions.

Task performance: We also analyzed the collaboration task performance. A Kruskal-Wallis H-test indicates that there is a statistically significant effect of the accumulated rewards between conditions ($\chi^2(2, 603) = 22.62, p < .001$). The post-hoc Mann-Whitney U test showed a significant difference between *belief* and *static* ($p < .001$). This partially supports our hypothesis *H2*. We didn't observe a significant difference between *belief* and *traversal*. Fig. 6(b) shows the task performance for different robots in three conditions. The Pearson correlation coefficient between reachability estimation and collaboration performance is $r = .203$ with p-value smaller than 0.001, indicating a positive correlation. This validates that calibrating perceived robot capability benefits the collaboration performance. Surprisingly, users in the *traversal* condition have a slightly higher reward when collaborating with the PR2 robot compared with those in *static*, even if their reachability estimation is less accurate, although the difference is not significant. This is probably due to the stochastic nature of the *traversal* baseline and specific object locations in our collaboration task Fig. 6(c).

Perception of robots: Finally, we analyzed participants' perception toward robots Fig. 6(c). Running a Kruskal-Wallis H-test, we found significant effects for reachability ($\chi^2(2, 603) = 20.39, p < .001$) and predictability ($\chi^2(2, 603) = 11.30, p = .003$). The post-hoc Mann-Whitney U test revealed significant difference between *belief* and *traversal* for reachability ($p < .001$) and predictability ($p < .001$), confirming *H3*. Overall, users tended to prefer *belief* over *static*, and *static* over *traversal*. This is unexpected for reachability, considering the fact that users are actually better at predicting robots' reachability in *traversal* than in *static*. The Pearson correlation coefficient between reachability rating and prediction accuracy is $r = .109$, indicating a weak positive correlation. Similarly, we observe a weak correlation $r = .019$ between self-reported reliance and users' actual ability to rely on the robot during collaboration. This suggests that there may be a discrepancy between the users' self-reported capability understanding and what they actually know about the robot.

In summary, we found that users in the *belief* condition had the most accurate estimation of the robots' capability, and

reported the robots in this condition as the most reliable, the most predictable, and the easiest to understand among all three conditions. Moreover, users working with the *belief* robots achieved a higher reward than those working with the *static* robots did. These objective and subjective results together suggest that our approach has an overall advantage for improving humans' understanding of robots as well as the quality of collaboration over the baselines.

VI. CONCLUSION

We have proposed an expressive robot motion planning algorithm, REMP, which can generate informative trajectories by integrate human belief update into trajectory optimization. Our experiments show that our approach can efficiently calibrate a user's perception of a robot's reachability and consequently improve human-robot collaboration.

In this work, we focused on the robot's spatial reachability. As reaching is one of the most basic tasks in human-robot interaction, we believe understanding reachability would greatly help users understand robot capacities in more complex tasks. Thus we view our work as a successful first step towards a more general capability calibration setting. In the future, it is possible to extend REMP for other capabilities. Our current work treats predictability and reliability as separate measures from trust. Considering the multidimensional nature of trust, better instruments can be used for a more comprehensive trust measure [27]. Also, due to online experiment constraints, we only investigated the reachability calibration problem on a 2D plane. We intend to generalize our algorithm to 3D environment in future work.

REFERENCES

- [1] T. A. Stoffregen, K. M. Gorday, Y.-Y. Sheng, and S. B. Flynn, "Perceiving affordances for another person's actions," *J. Exp. Psychol.: Hum. Percept. Perform.*, vol. 25, no. 1, pp. 120–136, 1999.
- [2] A. Powers and S. Kiesler, "The advisor robot: Tracing people's mental model from a robot's physical attributes," in *Proc. 1st ACM SIGCHI/SIGART Conf. Hum.-Robot Interact.*, 2006, pp. 218–225.
- [3] S. R. Fussell, S. Kiesler, L. D. Setlock, and V. Yew, "How people anthropomorphize robots," in *Proc. 3rd ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2008, pp. 145–152.
- [4] S. V. Albrecht and P. Stone, "Reasoning about hypothetical agent behaviours and their parameters," in *Proc. 16th Conf. Auton. Agents MultiAgent Syst., Ser., Richland, SC: Int. Found. Auton. Agents Multiagent Syst.*, 2017, pp. 547–555.
- [5] E. Cha, A. D. Dragan, and S.S. Srinivasa, "Perceived robot capability," in *Proc. 24th IEEE Int. Symp. Robot Hum. Interactive Commun.*, 2015, pp. 541–548.
- [6] M. F. Jung, J. J. Lee, N. DePalma, S. O. Adalgeirsson, P. J. Hinds, and C. Breazeal, "Engaging robots: Easing complex human-robot teamwork using backchanneling," in *Proc. Conf. Comput. Supported Cooperative Work*, 2013, pp. 1555–1566.
- [7] F. A. Robinson, M. Velonaki, and O. Bown, "Smooth operator: Tuning robot perception through artificial movement sound," in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2021, pp. 53–62.
- [8] S. Nikolaidis, S. Nath, A. D. Procaccia, and S. Srinivasa, "Game-theoretic modeling of human adaptation in human-robot collaboration," in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2017, pp. 323–331.
- [9] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa, "Planning with trust for human-robot collaboration," in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2018, pp. 307–315.
- [10] Y. Xie, I. P. Bodala, D. C. Ong, D. Hsu, and H. Soh, "Robot capability and intention in trust-based decisions across tasks," in *Proc. 14th ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2019, pp. 39–47.
- [11] J. Lee, J. Fong, B. C. Kok, and H. Soh, "Getting to know one another: Calibrating intent, capabilities and trust for human-robot collaboration," in *IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, *arXiv:2008.00699*.
- [12] G. Venture and D. Kulić, "Robot expressive motions: A survey of generation and evaluation methods," *ACM Trans. Hum.-Robot Interact.*, vol. 8, no. 4, pp. 1–17, 2019.
- [13] D. Szafir, B. Mutlu, and T. Fong, "Communication of intent in assistive free flyers," in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2014, pp. 358–365.
- [14] G. Lemasurier *et al.*, "Methods for expressing robot intent for human-robot collaboration in shared workspaces," *ACM Trans. Hum.-Robot Interact.*, vol. 10, no. 4, pp. 1–27, 2021.
- [15] A. Dragan and S. Srinivasa, "Generating legible motion," in *Robotics: Science and Systems*, Jun., 2013.
- [16] F. Stulp, J. Grizou, B. Busch, and M. Lopes, "Facilitating intention prediction for humans by optimizing robot motions," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2015, pp. 1249–1255.
- [17] M. L. Felis, K. Mombaur, and A. Berthoz, "An optimal control approach to reconstruct human gait dynamics from kinematic data," in *Proc. IEEE-RAS 15th Int. Conf. Humanoid Robots*, 2015, pp. 1044–1051.
- [18] C. K. Liu, A. Hertzmann, and Z. Popović, "Learning physics-based motion style with nonlinear inverse optimization," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 1071–1081, 2005.
- [19] M. N. Nicolescu and M. J. Mataric, "Learning and interacting in human-robot domains," *IEEE Trans. Syst., Man, Cybern.-Part A: Syst. Hum.*, vol. 31, no. 5, pp. 419–430, Sep. 2001.
- [20] L. Takayama, D. Dooley, and W. Ju, "Expressing thought: Improving robot readability with animation principles," in *Proc. 6th Int. Conf. Hum.-Robot Interact.*, 2011, pp. 69–76.
- [21] M. Kwon, S. H. Huang, and A. D. Dragan, "Expressing robot incapability," in *Proc. ACM/IEEE Int. Conf. Hum.-Robot Interact.*, 2018, pp. 87–95.
- [22] J. Schulman, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, and P. Abbeel, "Finding locally optimal, collision-free trajectories with sequential convex optimization," in *Robotics: Science and Systems*, Berlin, Germany, vol. 9, no. 1. 2013, pp. 1–10.
- [23] R. Diankov, "Automated construction of robotic manipulation programs," Ph.D. dissertation, USA, Carnegie Mellon University, Robotics Institute, 2010.
- [24] L. P. Kaelbling and T. Lozano-Pérez, "Hierarchical task and motion planning in the now," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2011, pp. 1470–1477.
- [25] O. Morgenstern and J. Von Neumann, *Theory of Games and Economic Behavior*. Princeton, NJ, USA: Princeton Univ. Press, 1953.
- [26] M. Madsen and S. Gregor, "Measuring human-computer trust," in *Proc. 11th Australas. Conf. Inf. Syst.*, 2000, vol. 53, pp. 6–8.
- [27] B. F. Malle and D. Ullman, "A multidimensional conception and measure of human-robot trust in Trust," in *Human-Robot Interaction*. Amsterdam, The Netherlands: Elsevier, 2021, pp. 3–25.