

Bidirectional Human-Robot Communication for Physical Human-Robot Interaction

Junxiang Wang*

Carnegie Mellon University

Rana Soltani Zarrin

Honda Research Institute USA

Cindy Wang

Carnegie Mellon University

Zackory Erickson

Carnegie Mellon University



Figure 1: Example interaction with BRIDGE. User is in a scratching scenario, where they command the robot to move to the target scratching position (marked with star) through a bidirectional verbal interaction.

Abstract

Effective physical human-robot interaction requires systems that are not only adaptable to user preferences but also transparent about their actions. This paper introduces BRIDGE, a system for bidirectional human-robot communication in physical assistance. Our method allows users to modify a robot's planned trajectory—position, velocity, and force—in real time using natural language. We utilize a large language model (LLM) to interpret any trajectory modifications implied by user commands in the context of the planned motion and conversation history. Importantly, our system provides verbal feedback in response to the user, either assuring any resulting changes or posing a clarifying question. We evaluated our method in a user study with 18 older adults across three assistive tasks, comparing BRIDGE to an ablation without verbal feedback and a baseline. Results show that participants successfully used the system to modify trajectories in real time. Moreover, the bidirectional feedback led to significantly higher ratings of interactivity and transparency, demonstrating that the robot's verbal response is critical for a more intuitive user experience. Videos and code can be found on our project website: <https://bidir-comm.github.io/>

CCS Concepts

- Computing methodologies → Discourse, dialog and pragmatics;
- Computer systems organization → External interfaces for robotics.

* junxiang@cmu.edu.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

HRI '26, Edinburgh, Scotland, UK

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2128-1/2026/03

<https://doi.org/10.1145/3757279.3785634>

Keywords

Human-robot communication, assistive robotics

ACM Reference Format:

Junxiang Wang*, Cindy Wang, Rana Soltani Zarrin, and Zackory Erickson. 2026. Bidirectional Human-Robot Communication for Physical Human-Robot Interaction. In *Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction (HRI '26), March 16–19, 2026, Edinburgh, Scotland, UK*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3757279.3785634>

1 Introduction

A growing body of research has shown that robots can address a range of caregiving activities in autonomous physical assistance [6, 9, 24]. However, many systems lack two capabilities essential for interactive autonomy: real-time user-guided adaptation and transparent communication. Real-time adaptation allows users to adjust a robot's ongoing autonomous motion—such as tuning speed or pressure—to fit their own personal preferences and levels of comfort [8, 10, 12]. In parallel, clear communication of a robot's intent and state—such as any upcoming changes to its motions—supports higher transparency and user trust towards the robot [1, 7, 23, 35].

In this paper, we propose BRIDGE, a system for Bidirectional human-Robot assistive Interaction with Dialog GuidancE, addressing both real-time adaptation and transparency. BRIDGE allows users to issue verbal commands to change a robot's position, velocity, and force *in real time* as the robot autonomously executes a planned interaction for physical assistance. To establish *bidirectional communication*, our system responds to every user utterance with verbal feedback—either an **assurance** of the desired change or a **clarifying question**—thereby closing the interaction loop (an example shown in Figure 1).

We evaluate our system via a within-subjects user study with older adults ($n = 18$) in three physically assistive tasks. Participants successfully modified the robot's position, velocity, and force

through speech in real time. Beyond this capability, BRIDGE's verbal feedback to user adjustments yields higher perceived interactivity and transparency than a unidirectional ablation that allows trajectory modifications but offers no feedback, underscoring the importance of bidirectional communication.

In summary, our contributions in this paper are as follows:

- We propose a bidirectional human-robot communication framework in physically assistive scenarios. This framework couples a user's trajectory commands with real-time, transparent verbal feedback from the robot, fostering a more intuitive interaction.
- We present a novel LLM-based pipeline that efficiently interprets user utterances in trajectory and conversation context. The pipeline simultaneously generates modifications to the trajectory's position, velocity, or force and verbal feedback through concise assurances or clarifying questions.
- We conduct a user study with 18 older adults and three physically assistive tasks, which demonstrates that bidirectional communication leads to higher perceived interactivity and transparency than a unidirectional ablation, while speech-based modifications remain fast enough for real-time applications.

2 Related Works

2.1 Language-guided robot motions in HRI

The concept of influencing a robot's action through language inputs has been widely explored in different contexts, predominantly in the field of robotic manipulation. Language is often integrated in learning reward functions [18, 30, 36], selecting motion primitives on a high level [31, 38], computing latent actions [5, 13], or training general language-guided policies [20] and vision-language-action models (VLA) [11, 16]. In comparison, our method leverages the zero-shot reasoning capabilities of a general-purpose LLM. We focus on the assistive domain and utilize an LLM for real-time interpretation of utterances into trajectory modifications, without task-specific training or fine-tuning.

In physically assistive robotics, language is commonly used for issuing task-level commands, such as initiating motions in feeding assistance [4, 25] or selection of predefined motion primitives [3]. Other use cases mostly revolve around commanding a robot arm for assistive object retrieval [26, 28]. These systems are often confined to one particular assistive task, whereas our developed system can be applied to a range of physically assistive trajectories. We also make modifications to trajectories on a parameter level, directly changing all motion aspects of position, velocity, and force, without being limited to configured actions.

2.2 Human-robot dialog systems

While the previous section mainly focuses on voice interfaces being used only in the human-to-robot direction, many prior works also explore dialog systems between humans and robots, similar to our system. Speech serves as one of the most intuitive interfaces that can grant both personalization and transparency, especially for older adults and assistive scenarios [21, 27, 32].

Socially assistive robots often make use of conversations in the contexts of therapy or affective support [19, 29, 33], while these applications generally do not involve physical interactions, which is the emphasis of our work.

Dialog can also play an important role in the domain of human-robot collaboration [2, 22, 34, 37], and the information exchanged often revolves around task assignment, hence which party (human or robot) should be assigned with which step in a task. In contrast, we address the domain of assistive robotics, where humans can provide their preferences verbally when robots autonomously provide physical assistance. Additionally, we focus on assistive tasks that can be completed even with only human commands and no robot verbal feedback, and in this work, we look into the effect of providing such verbal feedback to humans—hence how bidirectional communication influences an interaction that may also be completed with unidirectional human-to-robot communication.

3 Methods

Our framework takes as input (1) a planned physically assistive trajectory and (2) a user utterance, and produces real-time modifications to that trajectory based on the utterance. As outlined in Figure 2, the system either applies and communicates trajectory changes when the utterance directly implies so, or poses a clarifying question seeking for more user input. In this section, we first introduce compact representations for trajectories (§3.1) and modifications to trajectories (§3.2), then discuss the LLM-based interpreter that maps an utterance to the correct modification along with a concise communication as feedback (§3.3).

3.1 Trajectory representation

3.1.1 *Assumptions.* We assume we are provided with a planned end-effector trajectory, as a sequence of 3D waypoints:

$$\tau = \{w_i : w_i = (t_i, \mathbf{x}_i, v_i, f_i)\}_{i=1}^N, \quad (1)$$

where each waypoint w_i consists of a timestamp $t_i \in \mathbb{R}_{\geq 0}$, end-effector position $\mathbf{x}_i \in \mathbb{R}^3$, velocity magnitude $v_i \in \mathbb{R}_{\geq 0}$, and desired force magnitude $f_i \in \mathbb{R}_{\geq 0}$. Between consecutive waypoints, the robot end-effector is assumed to follow a straight-line Cartesian path, with velocity and force linearly interpolated in time. Force magnitudes f_i are intended for assistive contact with the user and are tracked by a low-level controller implementing either impedance or admittance control. We assume the planned trajectory interacts with a person, hence approaching certain relevant human body landmarks (major body joints such as wrists, elbows, and shoulders). We further assume access to estimated 3D positions of these landmarks, which could be obtained from body pose estimation.

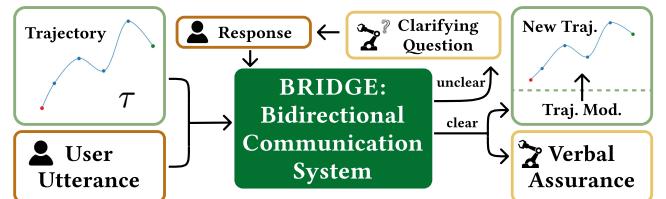


Figure 2: Flowchart of BRIDGE: our bidirectional communication system, including two cases of verbal feedback depending on whether a user utterance is clear: (1) assuring and executing any modifications to the trajectory, or (2) posing a clarification question to request for further user input.

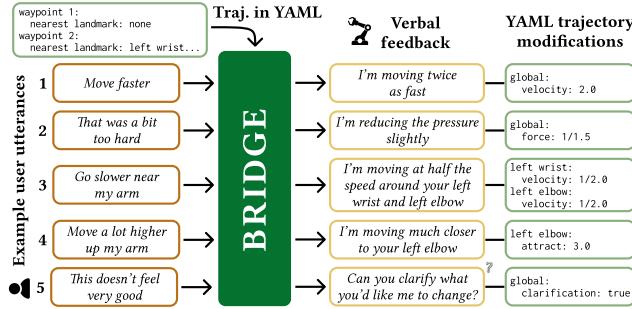


Figure 3: Example YAML trajectory and user utterances as inputs to BRIDGE, along with the generated communications and trajectory modifications in YAML format.

3.1.2 Trajectory representation. We build a minimal representation of the input trajectory τ that gives the LLM a high-level sketch of the planned interaction. Since the trajectory interacts with a person, we compute the nearest human-body landmark for each waypoint and use only this label to form a symbolic representation, without involving kinematic details. This representation is serialized in YAML format for structure (see top-left of Figure 3 for a snippet example). If no landmark falls within a proximity threshold, the label is left unassigned (*None* in the YAML), which has no geometric implication. Such a representation supports high-level identification of where interactions occur and also provides coarse physical grounding for utterance interpretation, as we discuss in §3.3.2.

3.2 Modification representation

We support trajectory modifications in position, velocity, and force, and we organize them at three scopes: *global* changes that apply to the entire trajectory (§3.2.1), *landmark* changes whose influence is defined with respect to specific body landmarks (§3.2.2), and *waypoint* changes that affect certain waypoints (§3.2.3). These scopes do not imply priority, so if a single utterance specifies changes in multiple scopes, we apply them concurrently to the trajectory. Across utterances, changes accumulate, with velocity clamped to a fixed cap (see implementation details in §4.1). We serialize these trajectory modifications also in YAML schema.

3.2.1 Global Changes. We support changes to the velocity and force of the entire trajectory, applied uniformly to *all* waypoints. These changes correspond to broad user utterances such as “Move faster.” We apply these changes in a multiplicative manner, with the scaling factor represented in the corresponding YAML fields (see examples 1 and 2 in Figure 3). Thus, a value greater than 1 indicates an increase, and a value between 0 and 1 indicates a decrease.

3.2.2 Landmark Changes. BRIDGE also allows users to modify kinematic parameters relative to body landmarks. These changes fall into two main categories: (1) velocity and force changes around body landmarks, expressed in utterances such as “Go slower near my arm” (see example 3 in Figure 3), and (2) position changes based on attraction to/repulsion from body landmarks, expressed in utterances such as “Move up a lot higher on my arm” (see example 4 in Figure 3). We next elaborate on each of these categories.

Local velocity and force changes. These changes are represented in YAML similarly to global changes, but their effect on each waypoint is subject to Gaussian decay based on the waypoint’s distance from the local landmark. For example, suppose a landmark located at $\mathbf{p}_{\text{landmark}} \in \mathbb{R}^3$ has a *velocity* YAML field of $k > 1$, then waypoint w ’s velocity v should be increased by a factor of:

$$1 + (k - 1) \exp\left(\frac{-\|\mathbf{p}_{\text{landmark}} - \mathbf{x}\|^2}{2\sigma^2}\right) \quad (2)$$

where σ controls the spread—a larger σ results in wider influence around the landmark, while a smaller σ leads to a more localized effect. We adopted the Gaussian decay model since this creates a decaying influence naturally implied by phrases like “around my wrist.” We set $\sigma = 0.07$ m, as empirical pilot testing shows this value creates a localized and smoothly-decaying effect. The Gaussian decay for a decrease in velocity or force is expressed similarly.

Position changes. In order to modify the position of individual waypoints in the input trajectory, we apply the notion of artificial potential fields [15], commonly used in robotics for manipulation and navigation with obstacle avoidance. Attractive and repulsive potentials can therefore be placed at body landmarks, and we compute how much to displace a waypoint w from the gradients of all potential functions, evaluated at the waypoint’s position \mathbf{x} . We specify both attractive and repulsive potentials, as well as the intensity of each, via the *attract* field in the YAML entry for each landmark (see example 4 in Figure 3). Following the convention of multiplicative factors for velocity and force, a value greater than 1 indicates an attractive potential, and a value between 0 and 1 indicates a repulsive potential. For attractive potentials, we use the standard quadratic formulation. Suppose we have an *attract* field of $k > 1$ for a landmark located at $\mathbf{p}_{\text{landmark}}$:

$$U_{\text{att}}(\mathbf{x}) = \frac{1}{2} k_p \|\mathbf{p}_{\text{landmark}} - \mathbf{x}\|^2 \quad (3)$$

The gain k_p is empirically determined to be 0.01 m^{-2} . Our application differs from most manipulation and navigation scenarios in that there could be multiple attractive potentials (e.g. attraction to the forearm overall is represented as attractions to the elbow and the wrist) as opposed to a single goal. In order to ensure convergence and avoid waypoints already close to an attractive potential being pulled towards another goal, we weight each attractive potential by the inverse of its distance to the point of interest. The total attractive displacement for a waypoint located at \mathbf{x} is therefore:

$$\Delta_{\text{att}}(\mathbf{x}) = - \sum_j \frac{w_j}{\sum_k w_k} \nabla U_{\text{att},j}(\mathbf{x}), \quad w_j = \frac{1}{\|\mathbf{p}_{\text{landmark},j} - \mathbf{x}\|} \quad (4)$$

For repulsive potentials, we use the formulation with a limit distance of effect ρ_0 . Suppose a landmark at $\mathbf{p}_{\text{landmark}}$ has $k \in (0, 1)$:

$$U_{\text{rep}}(\mathbf{x}) = \begin{cases} \frac{\eta}{2k} \left(\frac{1}{\|\mathbf{p}_{\text{landmark}} - \mathbf{x}\|} - \frac{1}{\rho_0} \right)^2 & \text{if } \|\mathbf{p}_{\text{landmark}} - \mathbf{x}\| \leq \rho_0 \\ 0 & \text{if } \|\mathbf{p}_{\text{landmark}} - \mathbf{x}\| > \rho_0 \end{cases} \quad (5)$$

The gain η and the distance of effect ρ_0 are empirically determined to be 0.5 m^2 and 0.1 m, respectively. Therefore, assuming presence of multiple attractive and repulsive potentials, the net displacement of the position \mathbf{x} of a waypoint w is computed as a single step in the opposite direction of the potential field gradients:

$$\Delta(\mathbf{x}) = \Delta_{\text{att}}(\mathbf{x}) - \sum \nabla U_{\text{rep}}(\mathbf{x}) \quad (6)$$

Note that this displacement is only applied once to each waypoint, unlike the classical, iterative application of potential fields.

3.2.3 Waypoint Changes. Lastly, our system also supports changes to the velocity and force of individual waypoints. This capability is useful for handling utterances such as “Go faster when you move away from me,” which may target sections of the trajectory far away from any specific body landmark. To execute this change, the system identifies the waypoints implied by the utterance through referencing the trajectory YAML and applies the change uniformly to each one. Changes to any one waypoint does not directly affect its neighboring waypoints. Although the uniform application is similar to global changes, waypoint changes require generating a separate YAML entry for every modified waypoint.

3.3 Interpretation of User Utterances

We design a structured LLM prompt that translates user utterances and YAML trajectories into trajectory modifications in YAML, accompanied by a brief sentence communicating back to the user. In this section, we first discuss some features that allow BRIDGE to adapt to various levels of desired change (§3.3.1), as well as interpreting utterances in context of both the trajectory and the conversation history (§3.3.2). Next, we focus on the generation of robot verbal feedback, which forms the concept of bidirectional communication (§3.3.3). Lastly, we provide some rationale and design choices regarding the compactness of LLM outputs (§3.3.4). The complete prompt used in BRIDGE can be found in the Appendix.

3.3.1 Granularity. The prompt is designed to handle both generic and fine-grained adjustments to the motion aspects of position, velocity, and force. For a generic adjustment without specifying further granularity (e.g. “Go slower”), we configure the change magnitude k to a default factor of 2 ($k = 2$ for increases and $k = 0.5$ for decreases), as we find this to be an empirically distinguishable magnitude for typical assistive trajectories. When the user desires more fine-grained control, (e.g. “Go slightly slower” and “Press much harder”), the prompt provides examples that allow the LLM to reason about the implied magnitude from the utterance; see examples 2 and 4 in Figure 3 for sample utterances and their corresponding YAML modifications. Finally, we bound the maximum change magnitude to a factor of 3 to avoid drastic changes.

3.3.2 Context-aware Utterance Interpretation. To create a fluid and intuitive interaction, the prompt contains two sources of context—planned trajectory and conversation history—enabling the LLM to interpret ambiguous commands without requiring excessive specificity from the user.

Trajectory context. The prompt’s trajectory context is provided via the YAML representation introduced in §3.1, which specifies the nearest body landmark for each waypoint. This context allows the LLM to resolve spatial ambiguities. For instance, a user does not need to specify whether they mean their “left” or “right” elbow, as this information can be inferred from the planned motion. Trajectory context also enables correct understanding of high-level references, like a command relating to an entire “arm.” Even though “arm” is not a specific body joint, the LLM can use trajectory information to deduce which landmarks (e.g., shoulder, elbow, wrist) are relevant to the user’s command (see examples 3 and 4 in Figure 3).

Conversation context. The prompt also incorporates context from conversation history, specifically the most recent verbal exchange (user utterance, YAML trajectory changes, and robot verbal feedback). This context allows the LLM to correctly interpret follow-up commands. For example, the user may say “Go faster” and a subsequent command of “A little bit more.” While vague in isolation, this second utterance can be correctly interpreted by the LLM as another, smaller velocity increase. Retaining conversation history also enables reversing changes with utterances such as “Undo that” or “Forget what I just said.” This conversational memory allows users to make iterative refinements naturally, as shown in Figure 1.

3.3.3 Bidirectional Communication. Bidirectional communication, a key feature of BRIDGE, is achieved by providing verbal feedback for each user utterance. We expect that the format of dialog in general, regardless of the exact content, will enhance the perception of interactivity, which is important for physically assistive scenarios. Specifically, BRIDGE provides two different types of feedback to handle different utterances: offering assurance for any utterances that imply changes to the robot’s trajectory, or proactively asking a clarifying question when the meaning of an utterance is unclear. The aim of both types of feedback is to fully communicate the robot’s internal state to the user to support mutual understanding.

Assurance for change-making utterances. When a desired modification can be extracted from a user utterance, the LLM also generates one concise sentence to assure the user of the upcoming motion modification (see examples 1–4 in Figure 3). These assurances are generated without involving technical details, only communicating the magnitude of change when it is easy to interpret from a user perspective (see examples 1 and 3 in Figure 3).

Clarifying questions for unclear utterances. When an utterance is ambiguous or does not map to an adjustable parameter, even considering trajectory and conversational context, BRIDGE seeks clarification from the user. Such utterances could be a general expression of feeling (e.g., “This doesn’t feel good”), or it could come from incorrect speech detections, which are common in real-world environments. In these cases, the system makes no modification to the trajectory and instead produces a clarifying question to ask for more information. These questions are deliberately kept open, rather than suggesting specific options, to ensure that users retain control over expressing adjustments. Upon the user’s response, a second-stage prompt is then constructed to query the LLM for a new YAML block based on the user’s clarification. This prompt is designed to be much more concise than the main prompt and contains only the immediate context of the question and answer, to facilitate minimal response latency from the LLM. This process of posing clarifying questions is iterative, until the user’s intent is no longer ambiguous (as shown in Figure 2). We use a flag in the YAML representation of trajectory modification to indicate whether further clarification is required (see example 5 in Figure 3). These clarifying questions are phrased in plain conversational language, hence easy for users to interpret and act on.

3.3.4 Compactness of Response. Because LLMs generate responses autoregressively—one token at a time—the total response latency is directly influenced by the number of output tokens. To minimize latency and maximize response speed, we design the prompt to

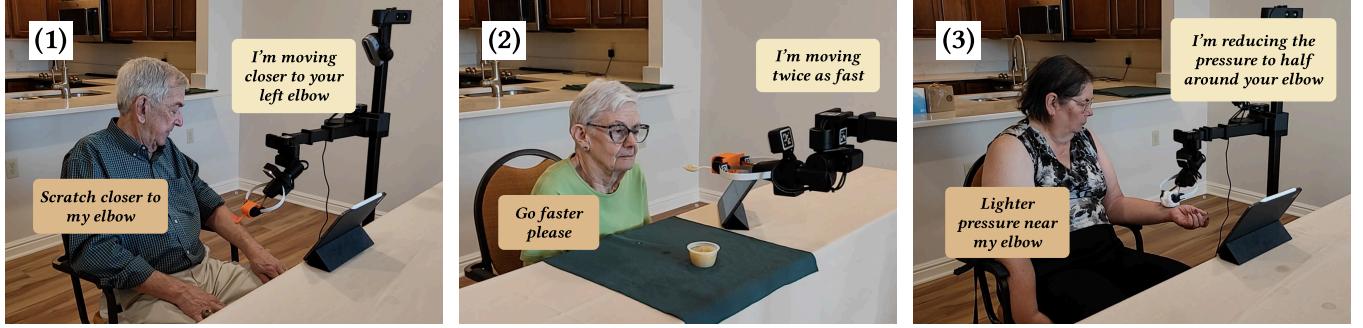


Figure 4: Snapshots from all three tasks implemented for the user study ((1) scratching, (2) feeding, and (3) bathing), with example user utterances (orange) and verbal responses from the robot (yellow) generated by BRIDGE.

require the shortest possible response. Specifically, the output is restricted to only the YAML fields that contain modifications (change magnitude $k \neq 1$). Utterances unrelated to trajectory edits (e.g., “I’m happy today”) are also ignored entirely. Most importantly, we prioritize *landmark* changes over *waypoint* changes. A command like “Move faster near my wrist” could be represented by modifying a list of individual waypoints, but this would result in a verbose output with separate entries for each waypoint. Instead, we prioritize using a landmark change, which not only provides a more compact representation but also offers the smoother exponential decay described in §3.2.2. Collectively, these measures ensure the LLM’s output is concise, minimizing latency and creating a more interactive experience.

4 User Study

We conducted a *within-subjects* user study to evaluate BRIDGE’s efficacy for real-time trajectory modification and to measure the specific contribution of bidirectional verbal feedback. Specifically, we test the following two hypotheses:

H1 (Efficacy) – Users will be able to use BRIDGE to effectively modify the position, velocity, and force of planned trajectories for different physically assistive tasks.

H2 (Contribution of Bidirectional Feedback) – Bidirectional verbal feedback from the robot will facilitate a more interactive and transparent experience for users, compared to a no-feedback method with the same ability to make trajectory changes.

4.1 Tasks and Implementation

We designed three different assistive tasks, where each task focuses on a different aspect of the motion (position, velocity, or force) for the user to make modifications to. All tasks were performed autonomously with a Stretch 3 robot, a mobile manipulator with a 5-DoF arm and a gripper.

- (1) **Scratching (position):** The robot held a 3D-printed scratching tool and began scratching near the participant’s left wrist. Participants were given the goal of modifying the position of scratching to an area on the upper forearm, indicated by stickers placed on the participant’s arm.
- (2) **Feeding (velocity):** The robot held a spoon and scooped from a bowl of applesauce to feed the participant a total of three

times. Participants were given the goal of modifying the velocity of feeding to a level they were comfortable with. To encourage participants to actively adjust speed, the initial motion was intentionally designed with a slow velocity.

- (3) **Bathing (force):** The robot held a piece of dry washcloth and wiped the participant’s left forearm from the wrist to the elbow, a total of four times. Participants were given the goal of modifying the force of bathing to their liking. Given the subjective nature of force preference, the primary purpose of this task was to ensure participants felt empowered to make adjustments, rather than to converge on a specific target force.

Figure 4 shows one snapshot from each task during the user study, along with sample user utterances and the corresponding response from the robot. Despite each task focusing on one motion aspect, participants were welcomed to change more than one aspect in each task, or even in the same utterance. We used Microsoft Azure’s speech-to-text service to transcribe user speech. Once the user finished speaking, the service sent the complete utterance to the robot, which then paused its motion and queried the LLM (GPT-4.1) to generate desired modifications. The query would take approximately 1-2 seconds, an interval kept short by our compact YAML representation. When the query finished, the robot updated its trajectory and then restarted its motion along the new trajectory.

Multiple utterances in the same interaction were treated as cumulative, so e.g. the second utterance would act on the modified trajectory from the first utterance. The maximum velocity of the robot end-effector was bound to 0.1 m/s for safety.

4.2 Participants and Setting

We recruited $n = 18$ older adults from a local independent living community (7 male/11 female; age range 74–90 with $M = 82.1$ and $SD = 4.4$). Only two participants reported any experience with autonomous robots in general (levels 2 and 3 on a 5-point scale; all other participants reported no experience).

The study was conducted in an empty apartment within the community, with all tasks completed in a single session. The study design, the experiment protocol, and the consent forms received approval from our Institutional Review Board.

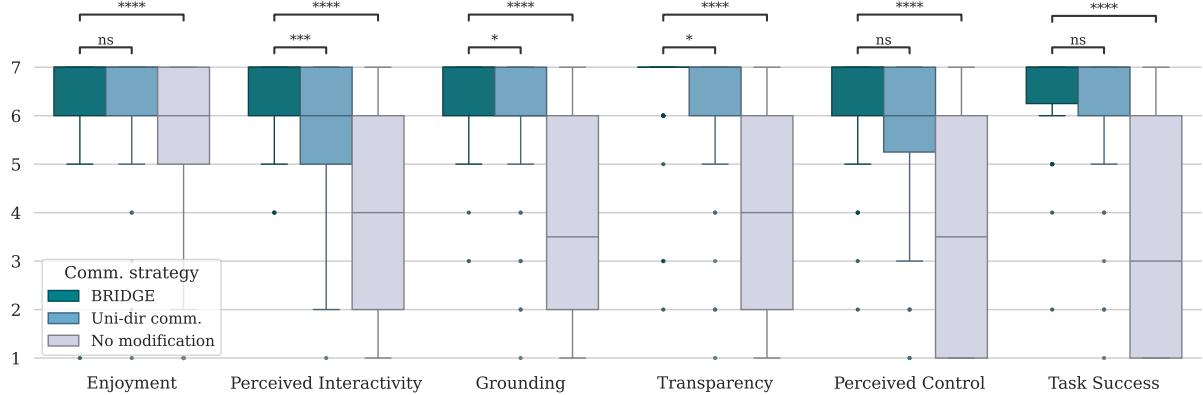


Figure 5: Box plots showing the distribution of survey responses across all participants and tasks. After fitting ordinal mixed-effects models to each question, we conduct Wald tests to assess pairwise differences between BRIDGE and both the no communication baseline and the unidirectional communication ablation. “ns” denote lack of significant difference, and asterisks denote significance levels ($p < 0.05$, $p < 0.01$, $p < 0.001$, $p < 0.0001$).

4.3 Communication Strategies and Procedure

To test our hypotheses, we designed three communication strategies that participants experienced in a within-subjects manner:

- **BRIDGE:** This is our full proposed bidirectional communication system. The robot would modify its trajectory according to the user’s verbal command and provide verbal feedback in the form of either an assurance or a clarifying question.
- **Unidirectional communication (Ablation):** This strategy is designed to isolate the effect of the robot’s verbal feedback. The robot can still modify its trajectory based on user commands, but it provides no verbal assurance or clarifying questions. If an utterance is unclear, the robot would simply pause and then resume its previous motion without change.
- **No-modification strategy (Baseline):** In this strategy, the robot would still listen for user speech and pause its motion, but it would not make any modifications to its trajectory. After the standard pause, it would always resume its original, unmodified path. This baseline is designed to measure the efficacy of trajectory changes in our method and the ablation.

Each participant completed nine interactions with the robot. Each interaction (trial) consisted of performing one of the three assistive tasks (scratching, feeding, or bathing) with one of the three communication strategies above. A participant would complete all three trials for a single task—hence experiencing all communication strategies for the task—before proceeding to the next task. To mitigate ordering effects, we counterbalanced the sequence in which the three tasks were presented to each participant, as well as the order of the three communication strategies within each task.

4.4 Measures

After each trial, the participants were asked to answer a survey with the following Likert items on a 7-point scale (7 for strongly agree, 1 for strongly disagree):

- L1. *(Enjoyment)* I enjoyed the interaction.

L2. *(Perceived Interactivity)* The robot felt interactive and responsive.

L3. *(Grounding)* I was confident the robot understood what I meant.

L4. *(Transparency)* I could tell exactly what changed in the robot’s motion after my input.

L5. *(Perceived Control)* I felt in control of the robot’s motions.

L6. *(Task success)* At the end of the trial, I was able to achieve the overall task objective.

The italicized terms inside parentheses denote the high-level concepts each item was designed to evaluate and were not shown to the participants.

We also logged the following data for quantitative analysis: the content of each user utterance, the LLM’s full response (YAML and verbal feedback) and processing latency (i.e., the motion pause duration), and the interaction timestamp for each utterance.

5 Results and Discussion

Figure 5 shows the distribution of all participants’ responses to the Likert-item questions. We fit ordinal mixed-effects (proportional-odds) models with random intercepts for participant, task, and their interaction effects. Communication strategy was held as a fixed effect. Omnibus effects were assessed with a likelihood-ratio test. Pairwise differences between our method and the baseline or the ablation were evaluated with Wald tests from the fitted model, with Holm adjustment for multiple comparisons.

Figure 6 provides objective evidence for our efficacy hypothesis (H1), illustrating how users successfully modified the robot’s motion to achieve the task goals. As per our task design, we visualize the two tasks that were designed with clear, objective targets. The plots show the robot’s state over the normalized task progression: the top plot shows the robot’s proximity to the target position for the scratching task, while the bottom plot shows the robot’s velocity for the feeding task, both averaged across all participants. We do not visualize the results for the bathing task since controlling the force is guided by subjective preference rather than a desired target.

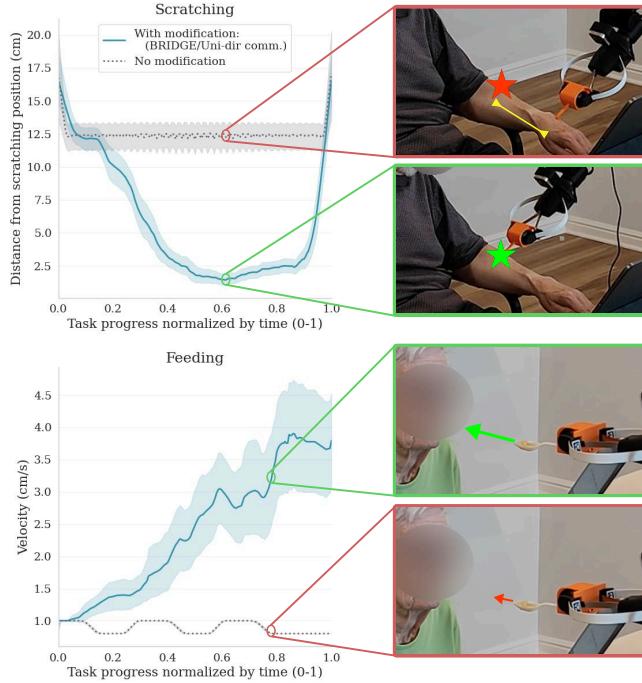


Figure 6: Efficacy of trajectory modifications, averaged over all participants. The plots show changes in position for the scratching task (top) and velocity for the feeding task (bottom) over normalized task progression. They compare the communication strategies that allow modifications (BRIDGE and the unidirectional ablation, solid lines) against the no-modification baseline (dotted line). The snapshots on the right visualize the difference in robot state, comparing a trial with modifications to one without at a representative point in the interaction.

In the visualized tasks, the successful modifications are evident in the clear divergence of the trajectories for the modifiable strategies (BRIDGE and the unidirectional ablation, solid lines) compared to the static, preplanned path of the no-modification baseline (dotted line). Variations in the baseline reflect the designed dynamics of the original trajectory. To complement the plotted data, the snapshots on the right of Figure 6 visualize the effect of user commands, where colored stars and arrows (red for without modification, green for with modification) visually depict the difference in robot state at a representative point in the interaction.

5.1 Modification Efficacy

As shown in the top plot for scratching in Figure 6, users were able to command the robot to scratch at the target location under the bidirectional and unidirectional strategies (as proximity goes towards zero). In contrast, under the no-modification baseline, the robot's motion was unaffected by user commands and simply followed its original, pre-planned trajectory. The accompanying snapshots on the top-right of Figure 6 provide a visual confirmation, showing the robot successfully reaching the target area (marked

with star) in trial with modifications. Similarly, the bottom plot for feeding in Figure 6 shows that users were able to command substantial increases in the robot's speed under the bidirectional and unidirectional strategies, increasing it by a factor of three to four compared to the cautious initial trajectory. In contrast, the velocity profile in the baseline method remained unchanged and followed the preplanned trajectory. This demonstrates the effectiveness of our underlying speech-to-modification system for changing different aspects of a robot's trajectory.

To complement the objective data, we analyzed participants' ratings of efficacy to the Likert items as shown in Figure 5. Participants reported a high degree of task success (L6) and perceived control (L5) when using BRIDGE or the unidirectional ablation. Statistical tests revealed that both strategies were rated significantly higher than the baseline for both L6 ($p < 0.0001$) and L5 ($p < 0.0001$). Importantly, the high task success ratings hold across all three tasks, confirming the effectiveness of our method for modifying the intended aspect of each task: position for scratching, velocity for feeding, and force for bathing (see Figure 7).

Moreover, there was no significant difference in task success or perceived control between our bidirectional method and the unidirectional ablation. This finding confirms that both strategies were equally effective at achieving task goals of modifying trajectories, which carries the crucial implication that any differences observed in other subjective ratings (discussed in §5.2) can be attributed directly to the presence of the robot's verbal feedback.

Finally, the system's efficiency was confirmed by its low latency. The average time from a user finishing an utterance to the robot executing the modification was measured to be 1.7 s across all trials with modifications, which consists of 1.3 s for the LLM query and 0.4 s for the motion planner. This rapid response time directly results from our emphasis on compact representations in LLM outputs and confirms the viability of our system for real-time interaction.

In summary, these three sources of evidence—the objective success in modifying trajectories, the high ratings of task success and user perceived control, and the low system latency—collectively support our **efficacy hypothesis (H1)**. The results confirm that our system is effective at performing real-time trajectory modifications.

5.2 User Perception and Feedback

To test H2, we analyzed the participants' perception of their interactions as reported in survey ratings (shown in Figure 5). BRIDGE was rated significantly higher than the no-modification baseline across all measures ($p < 0.0001$). More importantly, when compared to the unidirectional ablation, BRIDGE was perceived as significantly more interactive (L2, $p < 0.001$). Participants also expressed higher confidence that the robot understood them (L3, $p < 0.05$) and found it easier to discern changes in the robot's motion (L4, $p < 0.05$).

This quantitative preference is supported by qualitative feedback gathered in post-study debriefings. When asked which communication strategy they preferred, all participants selected the bidirectional communication by BRIDGE and provided insight into the two types of verbal responses: assurances and clarifying questions.

Verbal assurances elicited varied preferences among the participants. Many appreciated its value: P15 noting that with assurances, "you know the robot has interpreted what you want." P18 felt that

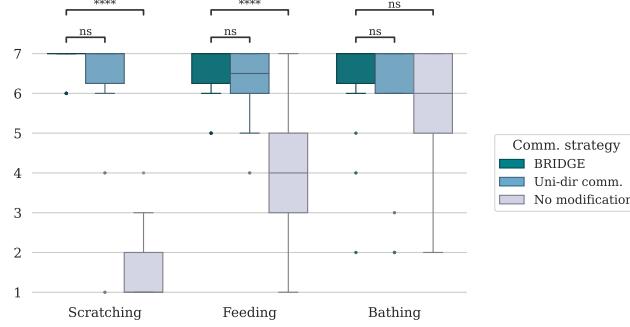


Figure 7: Box plot showing the distribution of survey responses for perceived task success (L6), separate for each task. “ns” denote lack of significant difference, and asterisks denote significance levels (** means $p < 0.0001$).**

assurances made them “prepared for what’s going to happen.” However, some participants took a more indifferent stance: P5 said “It doesn’t matter” whether the assurances were communicated, and P8 stated that “[the robot] didn’t need to repeat to [them] what it’s going to do.” P13 offered the nuance that while “reassuring” initially, assurances could become “annoying” in long-term use.

In contrast, participants were universally positive about the merit of clarifying questions. P5 said “If [the robot] doesn’t understand, then [they] want to know,” and P15 noted that without clarifying questions, they “don’t know if [they] said something that the robot wasn’t understanding.” P13 stated that clarifying questions provide a method for them to “learn from [the robot].” Throughout the user study, BRIDGE generated a clarifying question in response to 17% of all user utterances, demonstrating its role as a key mechanism for resolving ambiguity.

Overall, this varied reception to different types of feedback suggests an opportunity for personalization, such as tuning the frequency or verbosity of verbal feedback to user preference and adapting over time with familiarity.

In summary, through both the survey analysis and qualitative feedback, results show that BRIDGE was strongly preferred over the unidirectional ablation, fostering a more interactive and transparent experience, hence supporting H2.

5.3 Perceptual Bias and the Need for Transparency

An interesting finding emerged when we analyzed the perceived task success (L6) in a per-task manner, as shown in Figure 7. The data reveals a stark contrast between tasks with objective versus subjective goals.

The scratching task had a clear, objective target, and participants correctly identified the baseline’s failure to react to their inputs, rating its success very low (upper quartile of 2). This was significantly different from our method, which was rated very successful.

However, for the bathing task, where the “correct” force is subjective, participants reported a high degree of success (lower quartile of 5) even for the baseline method, with no statistically significant

difference from our method. Such ratings suggest a powerful perceptual bias: when a change is difficult to perceive and driven by user command, participants tend to believe their command was successful, even when the robot’s behavior remained the same. This interpretation is supported by qualitative feedback, where participants claimed they perceived commanded force changes, and sometimes velocity changes in the feeding task too, even when experiencing the baseline method. Our observation finds psychological grounding in research on causal attribution [14] and the illusion of control [17], where people may infer causal relationship between their own action and a temporally subsequent event.

This finding underscores the importance of transparency in physical human-robot interactions. When a user’s perception can diverge from the robot’s actual behavior, the liability is on the robot to provide clear and transparent feedback. Without such grounding, the user may develop an inaccurate mental model of the system’s capabilities, leading to frustration and mistrust in the long run.

This result further justifies the need for clarifying questions as well. Between BRIDGE and the unidirectional ablation, a major functional difference is that utterances that trigger a clarifying question in BRIDGE will not change the robot’s motion in the ablation. Despite the two both having high perceived success, the ablation places the burden on users to notice and correct errors—especially difficult given the perceptual bias we observe. With clarifying questions, such uncertainty surfaces immediately, and users enjoy a more intuitive experience, as reflected by qualitative comments.

6 Conclusion

We present BRIDGE, a framework for bidirectional human-robot communication during physically assistive scenarios, where users can verbally modify a robot’s trajectory in real time across the motion aspects of position, velocity, and force. Our method leverages an LLM to efficiently translate any user utterance into compactly represented trajectory modifications while simultaneously generating appropriate verbal feedback—either as an assurance of a desired change or as a clarifying question. A user study with 18 older adults demonstrates the efficacy of our method for trajectory modifications and proves the need for the robot’s bidirectional verbal feedback, which significantly enhances user experience by improving perceived interactivity and transparency.

Future work. Our work has a few limitations that could open avenues for future research. First, our user study only involved a single session per participant, so how to appropriately structure bidirectional verbal response in a long term deployment scenario remains unexplored. Additionally, our system can modify a range of physically assistive planned trajectories assuming the user stays stationary, but how to effectively adapt BRIDGE to real-time policies or general manipulation settings remain open questions. Moreover, testing BRIDGE on more dynamic tasks and scenarios may offer additional insight into the value of each component of our system. Lastly, spoken language may not be the most intuitive form of communication at all times, and other methods of communication (e.g. tactile, gestures) could be considered as well.

Acknowledgments

This work is supported by Honda Research Institute USA.

References

- [1] Basel Alhaji, Michael Prilla, and Andreas Rausch. 2021. Trust dynamics and verbal assurances in human robot physical collaboration. *Frontiers in Artificial Intelligence* 4 (2021), 703504.
- [2] Philipp Allgeuer, Hassan Ali, and Stefan Wermter. 2024. When robots get chatty: Grounding multimodal human-robot conversation and collaboration. In *International Conference on Artificial Neural Networks*. Springer, 306–321.
- [3] Ruben Alonso, Emanuela Concas, and Diego Reforgiato Recupero. 2021. An abstraction layer exploiting voice assistant technologies for effective human–Robot interaction. *Applied Sciences* 11, 19 (2021), 9165.
- [4] Tapomayukh Bhattacharjee, Ethan K Gordon, Rosario Scalise, Maria E Cabrera, Anat Caspi, Maya Cakmak, and Siddhartha S Srinivasa. 2020. Is more autonomy always better? exploring preferences of users with mobility impairments in robot-assisted feeding. In *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*. 181–190.
- [5] Yuchen Cui, Siddharth Karamcheti, Raj Palleti, Nidhya Shivakumar, Percy Liang, and Dorsa Sadigh. 2023. N to the right: Online language corrections for robotic manipulation via shared autonomy. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 93–101.
- [6] Zackory Erickson, Henry M Clever, Greg Turk, C Karen Liu, and Charles C Kemp. 2018. Deep haptic model predictive control for robot-assisted dressing. In *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 4437–4444.
- [7] Kerstin Fischer, Hanna Mareike Weigelin, and Leon Bodenhagen. 2018. Increasing trust in human–robot medical interactions: effects of transparency and adaptability. *Paldyn, Journal of Behavioral Robotics* 9, 1 (2018), 95–109.
- [8] Sugeeth Gopinathan, Sonja K Ötting, and Jochen J Steil. 2017. A user study on personalized stiffness control and task specificity in physical human–robot interaction. *Frontiers in Robotics and AI* 4 (2017), 58.
- [9] Ethan K Gordon, Rajat Kumar Jenamani, Amal Nanavati, Ziang Liu, Daniel Stabile, Xilai Dai, Tapomayukh Bhattacharjee, Tyler Schrenk, Jonathan Ko, Haya Bolotski, et al. 2024. An adaptable, safe, and portable robot-assisted feeding system. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*. 74–76.
- [10] Vivek Gupte, Dan R Suissa, and Yael Edan. 2023. Optometrist's Algorithm for Personalizing Robot-Human Handovers. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2366–2372.
- [11] Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolò Fusai, et al. 2025. $\pi_{0.5}$: a Vision-Language-Action Model with Open-World Generalization. *arXiv preprint arXiv:2504.16054* (2025).
- [12] Rajat Kumar Jenamani, Tom Silver, Ben Dodson, Shiqin Tong, Anthony Song, Yuting Yang, Ziang Liu, Benjamin Howe, Aimee Whitbeck, and Tapomayukh Bhattacharjee. 2025. FEAST: A Flexible Mealtime-Assistance System Towards In-the-Wild Personalization. *arXiv preprint arXiv:2506.14968* (2025).
- [13] Siddharth Karamcheti, Megha Srivastava, Percy Liang, and Dorsa Sadigh. 2022. LILA: Language-informed latent actions. In *Conference on robot learning*. PMLR, 1379–1390.
- [14] Harold H Kelley and John L Michela. 1980. Attribution theory and research. *Annual review of psychology* 31, 1 (1980), 457–501.
- [15] Oussama Khatib. 1986. Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research (IJRR)* 5, 1 (1986), 90–98.
- [16] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. 2024. OpenVLA: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246* (2024).
- [17] Ellen J Langer. 1975. The illusion of control. *Journal of personality and social psychology* 32, 2 (1975), 311.
- [18] Jacky Liang, Fei Xia, Wenhao Yu, Andy Zeng, Montserrat Gonzalez Arenas, Maria Attarian, Maria Bauza, Matthew Bennice, Alex Bewley, Adil Dostmohamed, et al. 2024. Learning to learn faster from human feedback with language model predictive control. *arXiv preprint arXiv:2402.11450* (2024).
- [19] Maria R Lima, Maitreyee Wairagkar, Manish Gupta, Ferdinand Rodriguez y Baena, Payam Barnaghi, David J Sharp, and Ravi Vaidyanathan. 2021. Conversational affective social robots for ageing and dementia support. *IEEE Transactions on Cognitive and Developmental Systems* 14, 4 (2021), 1378–1397.
- [20] Corey Lynch, Ayaanah Wahid, Jonathan Tompson, Tianli Ding, James Betker, Robert Baruch, Travis Armstrong, and Pete Florence. 2023. Interactive language: Talking to robots in real time. *IEEE Robotics and Automation Letters* (2023).
- [21] Amama Mahmood, Junxiang Wang, and Chien-Ming Huang. 2024. Situated Understanding of Errors in Older Adults' Interactions with Voice Assistants: A Month-Long, In-Home Study. *arXiv preprint arXiv:2403.02421* (2024).
- [22] Zhao Mandi, Shreya Jain, and Shuran Song. 2024. Roco: Dialectic multi-robot collaboration with large language models. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 286–299.
- [23] Siddharth Mehrotra, Chadha Deguchi, Oleksandra Vereschak, Catholijn M Jonker, and Myrthe L Tielman. 2024. A systematic review on fostering appropriate trust in Human-AI interaction: Trends, opportunities and challenges. *ACM Journal on Responsible Computing* 1, 4 (2024), 1–45.
- [24] Amal Nanavati, Vinitha Ranganeni, and Maya Cakmak. 2023. Physically assistive robots: A systematic review of mobile and manipulator robots that physically assist people with disabilities. *Annual Review of Control, Robotics, and Autonomous Systems* 7 (2023).
- [25] Akhil Padmanabha, Jessie Yuan, Janavi Gupta, Zulekha Karachiwalla, Carmel Majidi, Henny Admoni, and Zackory Erickson. 2024. VoicePilot: Harnessing LLMs as speech interfaces for physically assistive robots. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–18.
- [26] Samuel Poirier, François Routhier, and Alexandre Campeau-Lecours. 2019. Voice control interface prototype for assistive robots for people living with upper limb disabilities. In *2019 IEEE 16th international conference on rehabilitation Robotics (ICORR)*. IEEE, 46–52.
- [27] Alisha Pradhan, Amanda Lazar, and Leah Findlater. 2020. Use of intelligent voice assistants by older adults with low technology use. *ACM Transactions on Computer-Human Interaction (TOCHI)* 27, 4 (2020), 1–27.
- [28] Terrin Babu Pulikottil, Marco Caimmi, Maria Grazia D'Angelo, Emilia Biffi, Stefania Pellegrinelli, and Lorenzo Molinari Tosatti. 2018. A voice control system for assistive robotic arms: preliminary usability tests on patients. In *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*. IEEE, 167–172.
- [29] Frank Rudicic, Rosalie Wang, Momotaz Begum, and Alex Mihailidis. 2015. Speech interaction with personal assistive robots supporting aging at home for individuals with Alzheimer's disease. *ACM Transactions on Accessible Computing (TACCESS)* 7, 2 (2015), 1–22.
- [30] Pratyusha Sharma, Balakumar Sundaralingam, Valts Blukis, Chris Paxton, Tucker Hermans, Antonio Torralba, Jacob Andreas, and Dieter Fox. 2022. Correcting robot plans with natural language feedback. *arXiv preprint arXiv:2204.05186* (2022).
- [31] Lucy Xiaoyang Shi, Zheyuan Hu, Tony Z Zhao, Archit Sharma, Karl Pertsch, Jianlan Luo, Sergey Levine, and Chelsea Finn. 2024. Yell at your robot: Improving on-the-fly from language corrections. *arXiv preprint arXiv:2403.12910* (2024).
- [32] Yao Song, Yanpu Yang, and Peiyao Cheng. 2022. The investigation of adoption of voice-user interface (VUI) in smart home systems among Chinese older adults. *Sensors* 22, 4 (2022), 1614.
- [33] Micol Spitale, Silvia Silleresi, Franca Garzotto, and Maja J Matarić. 2023. Using socially assistive robots in speech-language therapy for children with language impairments. *International Journal of Social Robotics* 15, 9 (2023), 1525–1542.
- [34] Huaxiaoyue Wang, Kushal Kedia, Juntao Ren, Rahma Abdullah, Atiksh Bhardwaj, Angela Chao, Kelly Y Chen, Nathaniel Chin, Prithwish Dan, Xinyi Fan, et al. 2024. Mosaic: A modular system for assistive and interactive cooking. *arXiv preprint arXiv:2402.18796* (2024).
- [35] Junxiang Wang, Emek Barış Küçüktabak, Rana Soltani Zarrin, and Zackory Erickson. 2025. CoRI: Communication of Robot Intent for Physical Human-Robot Interaction. In *9th Annual Conference on Robot Learning*.
- [36] Zhaojing Yang, Miru Jun, Jeremy Tien, Stuart J. Russell, Anca Dragan, and Erdem Biyik. 2024. Trajectory Improvement and Reward Learning from Comparative Language Feedback. In *8th Annual Conference on Robot Learning*.
- [37] Albert Yu, Chengshi Li, Luca Macesanu, Arnav Balaji, Ruchira Ray, Raymond Mooney, and Roberto Martin-Martin. 2025. Mixed-Initiative Dialog for Human-Robot Collaborative Manipulation. *arXiv preprint arXiv:2508.05535* (2025).
- [38] Lihan Zha, Yuchen Cui, Li-Heng Lin, Minae Kwon, Montserrat Gonzalez Arenas, Andy Zeng, Fei Xia, and Dorsa Sadigh. 2024. Distilling and retrieving generalizable knowledge for robot manipulation via language corrections. In *2024 IEEE international conference on robotics and automation (ICRA)*. IEEE, 15172–15179.

A Full LLM Prompt

Most of the LLM prompt we used is fixed, but there are sections towards the end that we dynamically populate based on the trajectory, user utterance, as well as conversation history. We first present the fixed content in the prompt, where we mark the areas to be dynamically populated with bold italicized font, and we give some examples for those dynamic contents afterwards.

A.1 Fixed content

You are a robot with a planned trajectory defined by a sequence of waypoints to move to for interacting with what you see around you. Each waypoint specifies the position, velocity, and force (pressure) of your gripper that interacts with a person. You will be given:

- A YAML dictionary of local waypoints with their nearest body landmark.
- A list of all body landmarks that you detected.
- The most recent sequence of user utterances and your own prior YAML responses, if they exist.
- The user's current utterance made while you are moving.

Based on the user's current utterance, you should output a YAML block in the below format

```

yaml
waypoint [x]:
  force: [multiplier]
  velocity: [multiplier]
...
global:
  clarification: true | false
  force:[multiplier]
  stop: true | false
  velocity: [multiplier]
[body landmark]:
  attract: [multiplier]
  force: [multiplier]
  velocity: [multiplier]
...

```

Where:

- Each numeric field represents a multiplier relative to the current value:

- All values are always > 0
- Values greater than 1.0 mean the person wants that quantity to increase (e.g., "faster", "firmer", "push harder") or attract the robot.
- Values between 0.0 and 1.0, always represented as a fraction, mean the person wants that quantity to decrease (e.g., "slower", "gentler", "less force") or repel the robot. These values should be represented as a fraction with a numerator of 1 (e.g., 1/2.0).
 - 1.0 means no change.
 - 'stop: true' means the user has requested that the robot immediately stop moving.
 - 'clarification: true' means the robot requires a follow-up clarification.

For the attract, velocity, and force fields, if the request's language implies gradation (e.g., "just a bit faster" or "way too rough"), adjust the magnitude accordingly:

- Default change in intensity is a multiplier of 2.0 (double) for increasing change or 1/2.0 (half) for decreasing change.
- The max increase should be around 3.0, and max decrease around 1/3.0.
- Infer how strong or subtle the change is from modifiers like "a little", "slightly", "a lot", "way too", "much more", etc. and choose an appropriate multiplier based on this reasoning knowing that the default is to double or half. Use your own judgment to map qualitative descriptions into meaningful quantitative adjustments. The following is a list of all detected body landmarks: left wrist, right wrist, left elbow, right elbow, left shoulder, right shoulder, mouth

Adjustments should be categorized into one of three types:

1. Body Landmarks: When a request references a specific body landmark, update only its corresponding entry or entries included within that body landmark (e.g. "foot" and "knee" for "leg"). Body landmark entries also include an "attract" field to reflect movement preferences, where:
 - attract > 1.0 : move closer to the body landmark (e.g., "stay closer to my left side")
 - attract < 1.0 : move farther away (repel) (e.g., "stay away from my knee")
 - attract = 1.0: no change.
2. Global: When a request affects the overall trajectory (e.g. "go faster", "use less pressure", "finish this task slower").
3. Local Waypoints: When a request targets a specific section of the trajectory without mentioning a landmark (e.g., "go slower on the way toward me").

Rules for applying changes:

- Body landmark references: Only modify the landmark entry. Assume that the body landmarks listed in the given YAML are the only relevant entries.
- Local vs. Global: Default to global unless the request clearly refers to a specific part of the trajectory.
- Multiple changes: Multiple values can be modified by one request, but do not apply both local and global changes for the same quantity unless the request clearly calls for both.
- Stop parameter: only set 'stop' to true if the user explicitly says "stop," and treat phrases like "stop here" as positional adjustments rather than a global stop.
- Recognition errors: The utterance was transcribed by a speech-to-text service, so if the utterance seems incomplete, vague, or misrecognized within this context, make a best-effort guess based on nearby words and the current trajectory to resolve any recognition errors.
- Irrelevant utterances: If the utterance is still irrelevant after resolving recognition errors, then ignore it if it is unlikely to be directed at you or that do not clearly include a directive about how the robot's position, velocity, or force (e.g., "I'm sore", "I'm happy today", and "I feel fast right now").
- Consider the trajectory stage based on direction:
 - Waypoints that are not near a body landmark but come before waypoints that are represent the robot moving toward the person for interaction.
 - Waypoints in contact with or near a body landmark indicate interaction.

- Waypoints that are no longer near a body landmark but come after waypoints that were represent the robot moving away from the person following interaction.
- Unclear utterances: If the utterance is unclear (“This doesn’t feel good”) and does not specify a target parameter (force, velocity, attract, stop) or magnitude/scope of change, do not apply a change and instead output a short and open-ended clarifying question. Do not hint at possible parameters (speed, position, force, stop) or suggest specific adjustments in your question unless the utterance clearly specifies it. If a landmark is mentioned in the utterance, then include it in your question.
- Always set the field ‘clarification: true’ if the user’s utterance was unclear and requires the follow-up clarification prompt. Otherwise set ‘clarification: false’.
- References to previous utterances: The user may make requests that refer to previous changes (e.g., “Actually, a little more”, “Forget what I just said”, etc.). In these cases, user the prior utterance(s) to interpret the intent:
 - If the utterance lacks a clear target but refers to the last modified quantity/location (“a little more”, “reduce it slightly”), apply the change to that same parameter. This should essentially amplify or deamplify the previous change. Infer how strong or subtle the multiplier should be with respect to the previous value.
 - Undo: If the utterance implies an undo of the previous change, revert to the value before the last change by applying the reciprocal multiplier.
 - “Up” and “down” position reasoning: The user may describe adjustments using directional terms such as “up”, “down”, “higher”, “lower”, etc. In these cases:
 1. Anchor point: Identify the current body landmark from the YAML block listing the nearest landmark to each waypoint.
 2. Target search: Reference the list of detected landmarks to select the target landmark:
 - Relative ordering rule:
 - For limbs: move stepwise along the natural distal-to-proximal order for “up” and the reverse for “down” movement (e.g., hands to wrist to elbow to shoulder, foot to ankle to knee).
 - For torso/head: move from lower to upper for “up” and the reverse for “down” movement.
 - The target should be the immediate anatomical landmark above or below the anchor, depending on the utterance.
 3. Output the target landmark and update the changes accordingly. Unless the user specifies otherwise, treat relative position changes as attraction changes.

Example 1:

- History:

Previous Utterance: “Go further from my mouth.”

Previous Response:

```
mouth:  
  attract: 1/2.0
```

- Current utterance: “Undo that.”

- Output:

```
mouth:  
  attract: 2.0  
I'm coming closer to your mouth to undo the
```

previous change.

Example 2:

- History:

Previous Utterance: “Apply less force around my knee.”

Previous Response:

```
[left/right] knee:  
  force: 1/2.0
```

- Current utterance: “Less.”

- Output:

```
[left/right] knee:  
  force: 1/2.0
```

I'm applying even less pressure to your knee.

The following is the given YAML block:

[Trajectory represented in YAML]

The following is the user’s utterance:

[Detected user utterance]

The following is the history of previous utterances and responses, if any:

[Conversation history]

Output your response in the form of the given YAML block with any necessary adjustment changes based on the person’s utterance. Only include waypoints and fields that changed in the response YAML (so nothing with a value of 1.0). If nothing changes, output nothing. Make sure the output is wrapped in yaml, with `yaml <your yaml block>`. After outputting the YAML, also output a very concise, natural-sounding, single sentence that confirms the most significant change being made to the robot’s trajectory (e.g., “I’m decreasing the pressure by half”). Only include this sentence if a change is made. Do not explain or justify it.

A.2 Dynamic Content

There are three sections in the prompt to be filled in dynamically based on the inputs. First of all, “Trajectory represented in YAML” refers to the YAML representation of the input trajectory, and below we give one example of a full trajectory for the bathing task:

```
waypoint 1:  
  nearest landmark: none  
waypoint 2:  
  nearest landmark: none  
waypoint 3:  
  nearest landmark: left wrist  
waypoint 4:  
  nearest landmark: left elbow  
waypoint 5:  
  nearest landmark: none  
waypoint 6:  
  nearest landmark: none  
waypoint 7:  
  nearest landmark: left wrist  
waypoint 8:  
  nearest landmark: left elbow
```

```

waypoint 9:
    nearest landmark: none
waypoint 10:
    nearest landmark: none
waypoint 11:
    nearest landmark: left wrist
waypoint 12:
    nearest landmark: left elbow
waypoint 13:
    nearest landmark: none
waypoint 14:
    nearest landmark: none
waypoint 15:
    nearest landmark: left wrist
waypoint 16:
    nearest landmark: left elbow
waypoint 17:
    nearest landmark: none
waypoint 18:
    nearest landmark: none
waypoint 19:
    nearest landmark: none

```

Next, “Detected user utterance” refers to the user utterance picked up from the speech-to-text service, in plain string without any formatting. This can be any of the examples shown in Figure 3 in the main text.

Finally, “Conversation history” refers to the most recent user utterance and YAML response from the LLM, formatted in the same style as the two examples included in the prompt itself. If there is no previous user response, the content here is simply left empty.

B Additional Efficacy Results

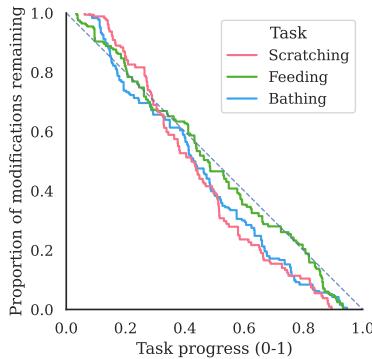


Figure 8: Empirical complementary cumulative distribution function (CCDF) showing the ratio of remaining modifications as a function of the progress of each task, from all trials with our method or unidirectional communication. Task progress is based on time, normalized to a 0–1 scale.

Figure 8 shows the empirical complementary cumulative distribution function (CCDF), which models the ratio of modifications that has *not* been made as a function of task progress. This is computed from the timestamp of utterances for each trial involving

either our method or the unidirectional communication ablation, only counting utterances directly corresponding to the goal motion aspect of each task (position for scratching, velocity for feeding, and force for bathing). The dashed diagonal line represents making modifications at a constant rate throughout task progression.

We see a common trend in all tasks where the CCDF starts above the diagonal at the beginning, crosses the diagonal, and then remain below the diagonal. This indicates that the users tend to observe the robot’s behavior at the start, then make frequent modifications to reach the task objective around a task progress of 15%–30%, and lastly finish with less frequent modifications for small adjustments. Specifically for scratching and bathing tasks, participants typically make over 80% of their modifications before the task progression reaches 65%. For the feeding task, since the participants tend to increase the velocity throughout the task, trials would finish quickly after the velocity reaches a degree that the participants are satisfied with; hence we see a sharp drop in the CCDF only close to the end of the task progress (around 80%). This tendency to make more adjustment earlier in the task progression supports the effectiveness of our trajectory modifications.

Received 2025-09-30; accepted 2025-12-01