Verbally Soliciting Human Feedback in Continuous Human-Robot Collaboration: Effects of the Framing and Timing of Reminders

Kate Candon
Yale University
New Haven, CT, USA
kate.candon@yale.edu

Helen Zhou
Yale University
New Haven, CT, USA
helen.zhou@yale.edu

Sarah Gillet
KTH Royal Inst. of Tech.
Stockholm, Sweden
sgillet@kth.se

Marynel Vázquez
Yale University
New Haven, CT, USA
marynel.vazquez@yale.edu

ABSTRACT

Humans expect robots to learn from their feedback and adapt to their preferences. However, there are limitations with how humans provide feedback to robots, e.g., humans may give less feedback as interactions progress. Therefore, it would be advantageous if robots could influence humans to provide more feedback during interactions. We conducted a 2x2 between-subjects user study ($N = 71$) to investigate whether the framing and timing of a robot's reminder to provide feedback could influence human interactants. Human-robot interactions took place in the context of Space Invaders, a fast-paced and continuous collaborative environment. Our results suggest that reminders can influence the amount of feedback humans provide to robots, how participants feel about the robot, and how they feel about providing feedback during the interaction.

Figure 1: Experimental setup for our study.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in collaborative and social computing**; **Empirical studies in HCI**.

KEYWORDS

human-robot interaction, explicit feedback, robot reminders

ACM Reference Format:
Kate Candon, Helen Zhou, Sarah Gillet, and Marynel Vázquez. 2023. Verbally Soliciting Human Feedback in Continuous Human-Robot Collaboration: Effects of the Framing and Timing of Reminders. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23), March 13–16, 2023, Stockholm, Sweden.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3568162.3576980

1 INTRODUCTION

Human-Robot Interaction (HRI) has long acknowledged the importance of creating robots that can adapt to individual preferences [2, 70], which are often learned through human feedback [21]. Research has investigated robot adaptation in a variety of settings, such as while collaboratively building a toolbox [54], during tutoring sessions [18], or in prehabilitation exercises [80]. Common types of feedback in robot learning include evaluative feedback [46], demonstrations [66], corrections [5], and comparisons [61].

Unfortunately, there are shortcomings in how humans naturally provide feedback [16, 53, 72]. Notably, humans tend to give less feedback as an interaction progresses [52]. Also, research suggests that users tend to stop providing feedback once they are satisfied with an agent's performance [39]. As robots enter more collaborative interactions, humans will likely provide even less feedback if they are preoccupied with their own actions.

In situations where humans are not providing enough feedback during an interaction, a robot could remind them to provide feedback. This strategy could work well because robots can influence human behavior, as demonstrated in a wide range of work (e.g., [19, 24, 27, 34, 48, 65]). However, it then becomes essential that the robot does not annoy the human with reminders, thus making it important to understand how to make the reminders impactful.

We conducted a study to investigate how robots should remind humans to give evaluative feedback in fast-paced, cooperative interactions. Participants played a collaborative game with a robot, as shown in Fig. 1. The novelty of this interaction setup was twofold. First, the game was a continuous task, more similar to autonomous driving [62] than more typical turn-based interactions in collaborative robotics [22, 55, 68]. Second, the human interactant had additional objectives other than solely providing feedback to the robot (as is common in robot learning [4, 12, 13, 36, 47]). Overall, the interaction was naturalistic from the perspective that both the human and robot were busy with their own agenda.

Our study focused on investigating two factors that could influence humans receiving feedback reminders from robots: 1) the framing of the robot's utterances (highlighting the robot individually or its human-robot team); and 2) the timing of reminders (relative to a situation in which the robot changed its behavior in the game). Our results suggest that highlighting the individual robot versus the human-robot team in reminders can influence how participants feel about the robot and about providing feedback during the interaction. Also, the timing of reminders can impact when participants provide feedback about the robot's performance. Our

HRI '23, March 13–16, 2023, Stockholm, Sweden

ACM ISBN 978-1-4503-9964-7/23/03.
https://doi.org/10.1145/3568162.3576980

work provides interesting insights on how to design robots that learn from humans in realistic, continuous collaboration scenarios.

2 BACKGROUND

Research has investigated robots that request feedback from humans, including what kind of queries to ask [9, 12, 26], how often to query a user [10], or how to account for a human's ability to provide useful information [8]. Further, Ho et al. [33] studied how to build mental models of humans to determine how to ask for feedback, and Jeon et al. [40] provided a framework to enable agents to combine multiple types of feedback. This line of work often investigates turn-based tasks and/or users that are only focused on providing feedback (e.g., [12, 68]). To complement this research, we study general reminders for feedback in continuous collaborations. This is important because the time when a robot asks for feedback is not necessarily the best time for the human to provide feedback.

To the best of our knowledge, general feedback reminders have not been explicitly studied before in HRI; nevertheless, prior work provides insights on trade-offs when requesting feedback. For example, robots must be able to ask for feedback without annoying the human [31] or asking too many questions [79]. Thus, robots have the difficult task of ensuring requests or reminders for input are frequent enough to be useful, but not too incessant [7]. One approach is to identify opportune times for interruptions [1], such as by modeling user attention [43]. It is also important to try to maximize the benefit from an interruption when a disruption is necessary. Thus, another approach is to study the way in which robots should remind humans to provide feedback most efficiently.

While current work typically studies how to leverage the ways in which humans naturally communicate when teaching a robot [15, 25, 42, 71], we are interested in understanding if robots can influence how much feedback humans provide. Close to our work, Rogers and Howard [59] found that an agent's embodiment influenced how much reward or punishment humans provided in a financial advisement scenario. Additionally, there is evidence that people provide more frequent feedback when an agent chooses bad actions [45]. However, our goal is to elicit more feedback without harming performance. The next sections describe related work on two specific aspects of feedback reminders relevant to our study.

Robot Framing in Communication: Robot communicative signals are able to influence human actions and perceptions of a robot (e.g., in one-on-one settings [24, 34, 58, 65, 81] and in groups [19, 23, 27, 48, 67]). One interesting aspect of robot communication is how the robot frames itself relative to others. For example, whether a robot framed itself as competitive or relationship-oriented impacted how much participants looked at and supported the robot in a card game [56]. Additionally, how a robot attributed blame amongst a group influenced how much humans trusted the robot [29, 41, 76]. Close to our work, Salomons et al. [63] found that whether the robot referred to itself as a peer or as a teacher affected how much humans learned over the course of an interaction. This corpus of work inspired us to investigate:

Research Question 1: *Will framing feedback in a reminder as helping the team versus helping the individual robot influence how humans provide feedback or feel about the interaction?*

Timing of Robot Actions: Another important factor of reminders is timing. Timing can be critical in human-robot communication [17]. For instance, the time when a robot helps a human can impact the human's perception of the robot [6]. The timing of robot actions can also affect the fluency of human-robot interactions [11, 38]. Consequently, we asked:

Research Question 2: *How does the timing of a feedback request influence when the human provides feedback?*

Because prior work has shown that humans not only provide feedback in response to past actions, but also to guide future behavior [44, 71], we investigated the above question in relation to an important change in robot behavior during interactions.

3 INTERACTION TASK: SPACE INVADERS

Typically, when a robot learns from a human, the human's only objective is to teach the robot (e.g., [46, 64, 73]). Also, tasks are usually turn-based, where the robot takes an action and then waits for the human to provide feedback [77]. However, everyday interactions are more fast-paced and involve competing priorities. Thus, we chose to study feedback reminders in a two-player Space Invaders game, requiring continuous and fast-paced decision-making and action. The game was inspired by prior work on ad-hoc cooperation [49] and unexpected help from a virtual agent [14].

In our version of Space Invaders, a human controlled a purple spaceship that spawned on the left side of the game screen and the robot controlled the spaceship that spawned on the right side (Fig. 1). Rows of enemies appeared at the top of the screen and moved downwards until they were destroyed or reached the bottom of the screen. The participant and robot had one team score and received points for destroying enemies. Both players started the game with four lives and lost a life when they collided with an enemy or a bullet. The game ended when all enemies were destroyed, when both players lost all their lives, or when an enemy reached the bottom of the screen.

The participant used the right and left arrow keys to move within the bounds of the screen and pressed the spacebar to shoot. They provided explicit, evaluative feedback to the robot by pressing the up arrow (positive feedback) or down arrow (negative feedback) on their keyboard. When participants pressed the up or down arrows, "good job" or "bad job" text appeared on the screen to ensure participants were aware that their feedback was received.

Robot Gameplay Strategies. Space Invaders allowed us to create three visually different gameplay strategies for the robot's spaceship based on when the spaceship travelled to the left side of the screen (the participant's side). The strategies helped familiarize participants with the game dynamics and study the effects of timing on feedback requests.

1) Uncooperative strategy: The robot only destroyed enemies on the right side of the screen. Because the robot could shoot slightly faster than the human, the robot always destroyed all of the enemies on its own side before the participant destroyed the enemies on their side. Once all of the enemies on the right side were destroyed, the robot waited for the participant to finish destroying the enemies on the left side. For games in which the robot utilized the uncooperative strategy, the robot's spaceship was dark grey, as shown in Fig. 2(a).

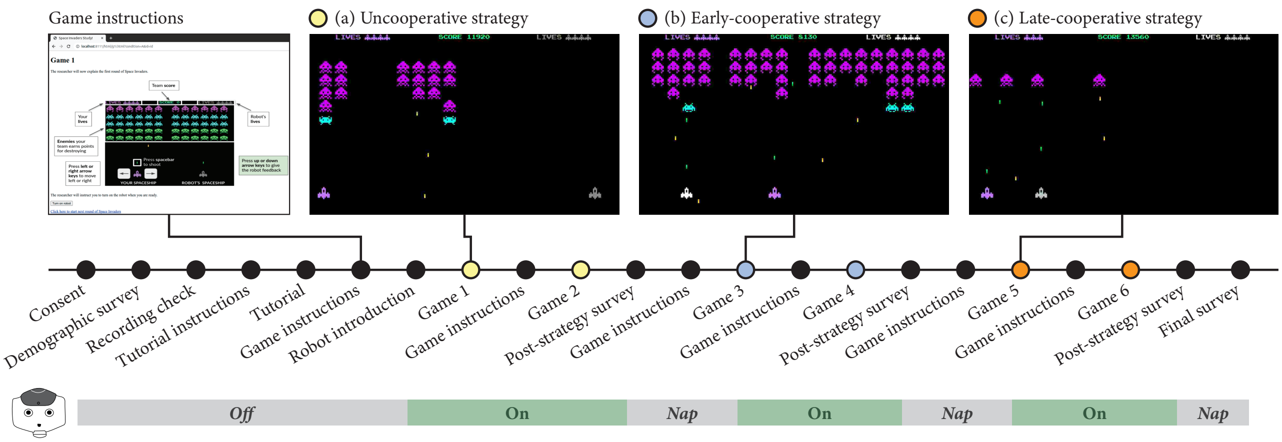

Figure 2: Experiment timeline. The robot reminded participants to provide feedback about the robot's performance in games 3 and 4. The first image shows the game instructions that were shown to the participant before each round of Space Invaders (best viewed in digital form). The other images show the three gameplay strategies for the robot's spaceship, as described in Sec. 3. The robot's spaceship was dark grey (a), white (b), or light grey (c). The bottom set of blocks shows the state of the robot.

2) Early-cooperative strategy: The robot went over to the left side of the game screen to help the participant destroy enemies on three visits during the game. For games in which the robot utilized this strategy, the robot's spaceship was white, as shown in Fig. 2(b).

The first visit to the participant's side of the screen was central to our study manipulations. When using the early-cooperative strategy, the robot emphasized the first visit by announcing "Look we/I are/am destroying enemies on the left side of the screen!" We wanted to ensure that participants noticed that the robot moved to the left side of the screen, exhibiting a new gameplay behavior.

3) Late-cooperative strategy: The robot only went over to the left side of the game screen to help destroy enemies after all of the enemies on the right side were destroyed. For games in which the robot utilized this strategy, the robot's spaceship was light grey. Fig. 2(c) depicts the late-cooperative strategy.

Implementation. We implemented the game with browser-based client technologies and a Python server. We used the Robot Operating System (ROS) [57] to provide game information to the robot. The supplementary material provides more implementation details.

4 METHOD

We conducted a user study to investigate the effects of *how* and *when* a robot reminded a participant to give feedback about the robot's behavior. Participants played six games of Space Invaders with a Nao robot, as in the timeline of Fig. 2. The participants were asked to help train the robot to be a good teammate by providing positive and negative feedback. The robot exhibited three different gameplay strategies, each for two games: uncooperative, early-cooperative, and late-cooperative (as described in Sec. 3). The uncooperative robot strategy served to familiarize participants with the game while playing with the robot. The early-cooperative strategy was the main focus of our study. The robot only reminded participants to provide feedback to the robot in these two games. The late-cooperative strategy was included in our study to evaluate if effects of our experimental manipulations, which are explained next, persisted in later interactions. These three gameplay strategies were intended to highlight changes in the robot's behavior, rather than being independent variables themselves.

4.1 Study Design

To investigate the research questions outlined in Sec. 2, we designed a 2x2 between-subject study with Framing (Individual vs. Team) and Timing (Before vs. After) as independent variables. The robot reminded the participant to provide feedback once in the third and fourth games of Space Invaders experienced in the study. The feedback reminders varied by:

Framing of utterances: We varied how the robot verbally referred to itself during gameplay using "I" vs. "we" pronouns. With the ***Individual*** framing, the robot referred to itself using the first-person, singular pronoun "I", e.g., "I'm ready to play" and "Remember to give feedback so I am a better player!" These utterances referred to the individual robot and focused the reminder on improving its gameplay. With the ***Team*** framing, the robot referred to itself using the first-person, plural pronoun "we", e.g., "We're ready to play" and "Remember to give feedback so we are a better team!" In the Team framing, the reminder was focused on improving the human-robot team, rather than the individual robot.

Timing of the reminder: We also varied when the robot reminded the participant to give feedback relative to changing its gameplay behavior. In particular, the robot's spaceship began playing Space Invaders on the right side of the screen. At three different points during the early-cooperative games, the robot's spaceship crossed over to the left side of the screen in order to help the participant. Our manipulation focused on the first of the robot's visits to the left side of the screen in both games 3 and 4. As explained in Sec. 3, the first crossover was announced with "Look we/I are/am destroying enemies on the left side of the screen!" With the ***Before*** reminder, the robot reminded the participant to give feedback before its spaceship crossed over to their side of the screen and announced the new behavior. With the ***After*** reminder, the robot's spaceship crossed over to participant's side of the screen, announced the new behavior, and then reminded the participant to give feedback once it was back on the right side of the screen.

The text bubbles and timelines in Fig. 3 illustrate the difference between the Before and After reminders for the Team framing. See our supplementary video for examples of experimental conditions.

4.2 Hypotheses

We hypothesized that our independent variables would have an effect on when participants provided feedback during the collaboration, and on how they reported feeling about the robot and the interaction. Specifically, in response to *RQ1*, we hypothesized:

H1a. Humans will give more feedback during the interaction with the Team framing than with the Individual framing.

H1b. Humans will feel more positive about giving feedback and about the robot with the Team framing than the Individual framing.

H1a and H1b were motivated by the psychology literature. By using the "we" pronoun, the robot stressed that the participant and the robot belonged to the same group. These feelings of group membership have been found to increase helping behaviors [51] and perceived responsibility for helping [50]. In our study, the participant helped the robot by providing feedback so that the robot could learn to be a better teammate in the future. Further, prior HRI work found that participants perceived a robot that expressed group-based emotions as more likeable and trustworthy than a robot that expressed individual-based emotions [20].

With respect to *RQ2*, we hypothesized:

H2a. Humans will give more feedback with the Before reminder than with the After reminder.

H2b. Humans will give feedback more quickly with the After reminder than with the Before reminder.

H2a and H2b were motivated by prior work on robots guiding human attention [35, 69, 82]. Also, humans provide feedback both in response to past actions and to guide future behavior [44, 71].

4.3 Setup

The experiment was conducted in a small office on a university campus in the United States. The room contained a table with a computer screen and a tablet. The participant sat in an office chair facing the computer screen, and the robot was on the table next to the participant. The physical setup is illustrated in Fig. 1.

We used the Nao robot by Softbank Robotics for our study. Nao is a humanoid robot. It is 22.6 inches tall, though it sat for the entirety of our study. The Nao was fully autonomous and controlled by the Python SDK for Naoqi on a computer running ROS. The robot spoke to the participant on set occasions throughout the interaction. We implemented a basic idling behavior where the Nao moved its head slightly every eight to fifteen seconds during the Space Invaders games so that it would seem attentive and engaged.

4.4 Procedure

Fig. 2 summarizes the sequence of events in a study session. After giving informed consent, participants filled out the pre-interaction demographics survey, which also included personality data via the Ten Item Personality Measure (TIPI) [28] and the Berkeley Expressivity Questionnaire (BEQ) [30].

The experimenter then instructed the participant to enter the office, sit at the computer, and complete a webcam check to ensure that the recording was working. Next, the experimenter explained the setup and controls for the Space Invaders game, including how to give positive or negative feedback to the robot. The participant was told the robot was still off, so the robot's spaceship would not move or shoot during the tutorial that followed. The experimenter stayed in the room while the participant completed the tutorial.

After the tutorial, the experimenter asked the participant to help train the robot and reminded the participant that the robot was their teammate. The experimenter stated: *"The robot already knows how to play the game, but not how to be a good teammate to you. You should give the robot feedback so that it learns to play in the way you like."* Participants were informed that the robot would not be adjusting its behavior based on feedback provided during the game, but that feedback would be used to improve robot behaviors in the future. The experimenter instructed the participant to turn on the robot, and the robot introduced itself. The participant then began the first game of Space Invaders with the robot.

The participant played six games of Space Invaders in total. One game of Space Invaders took on average 96.15 seconds ($SE = 0.69$). The first two games were with the uncooperative strategy, the middle two games were with the early-cooperative strategy, and the last two games were with the late-cooperative strategy. After each pair of games with a specific strategy, the participant answered a brief set of post-strategy survey questions. Finally, the participant answered a set of survey questions about the entire interaction. In order to reduce the likelihood that the participant interacted with the robot while answering survey questions, the robot stated *"I'm going to take a nap now while you answer some questions."*

At the end of the study, participants were compensated US$10. The interaction lasted approximately 35 minutes. The protocol was reviewed by our Institutional Review Board and refined via pilots.

4.5 Dependent Measures

We considered both objective and subjective measures in our study. For analysis of participant-provided feedback, we analyzed game logs for up and down button presses. Important game events included when the Before and After reminders would have been and when the robot announced its new behavior. Unless otherwise noted, survey questions were scored on a 7-point agreement scale with 1 being "strongly disagree" and 7 being "strongly agree."

Rate of feedback: We calculated how many times participants provided feedback with button presses via game logs. To account for varying game length, we computed feedback signals per minute (fspm). We analyzed the rate of feedback across entire games, as well as in ten second windows after important events in games 3 and 4, as depicted by the W measures in Fig. 3.

Elapsed time to feedback: We analyzed the number of seconds from game events to when the participant next provided feedback with the up and down keys. Fig. 3 shows the elapsed time measures (E) we analyzed and how they differ between Timing conditions.

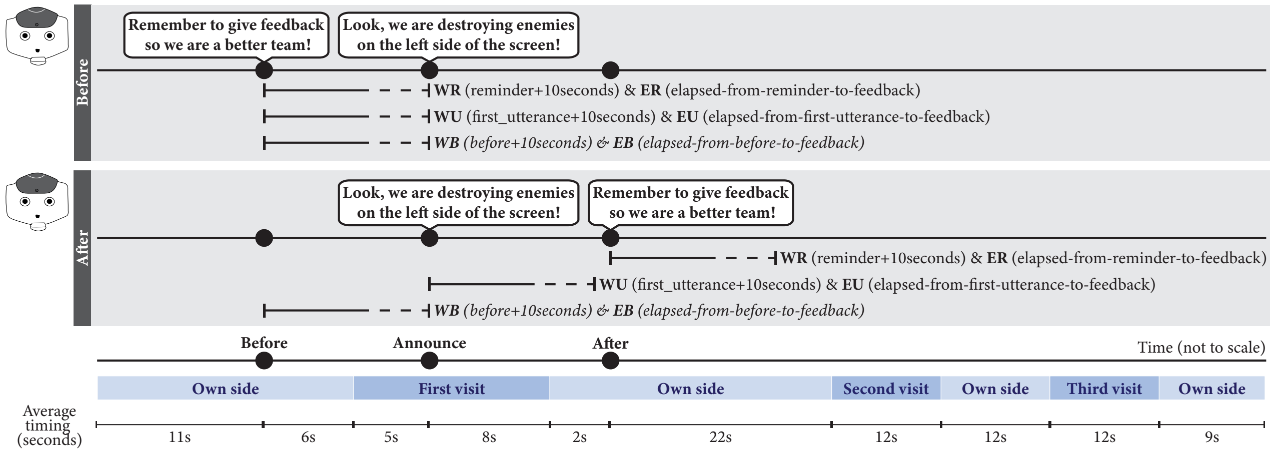

Figure 3: Robot behavior and related measures in the early-cooperative strategy (games 3 & 4) for the two timings of feedback (Before & After). Measures include ten-second windows (W) and elapsed time to next feedback press by participant (E) measures from game events. Italicized labels signify static measures across Timing conditions; non-italicized labels signify measures that differ across Timing conditions. Example shows Team framing, but measures were the same for Individual framing.

Feedback process: After completing all six games, participants were asked a series of questions about the process of providing feedback. They provided free text responses and rated how strongly they agreed it was difficult or distracting to give feedback, and if they thought they were able to give the robot helpful feedback.

Perceptions of Robot: After playing two games with each robot gameplay strategy, participants rated statements about the robot. The statements included if the robot was helpful, proficient at the game, or annoying, and if the participant liked the robot's behavior.

Table 1: Participant demographics by condition.

Framing	Timing	N	#Males	#Females	Age ($\mu \mp \sigma$)
Individual	Before	18	8	10	23.78 ∓ 6.11
Individual	After	18	9	9	26.50 ∓ 9.70
Team	Before	17	8	9	23.82 ∓ 4.57
Team	After	18	8	10	23.78 ∓ 3.57
All		71	33	38	24.48 ∓ 6.42

4.6 Participants

Our study had a total of 72 participants, with 18 participants in each of the four conditions. One participant in the Team-Before condition was excluded because they continuously provided feedback in all rounds of Space Invaders and their survey responses were inconsistent with the provided instructions. Thus, our final participant pool had 71 total participants. Participants were recruited via flyers, online postings, and word of mouth. They were required to be at least 18 years of age, be fluent in English, and have normal or corrected-to-normal hearing and vision.

Table 1 summarizes participant demographics. On average, participants indicated using a computer daily ($M = 1.08, SD = .50$) and playing video games between once a week and once a month ($M = 4.27, SD = 1.62$). Specific to Space Invaders, 21% reported playing the game before, 49% reported never having played the game, and 30% were not sure. The majority of participants (65%) reported that they interacted with robots less than once a month.

4.7 Manipulation Checks

4.7.1 Framing of utterances. In the final set of survey questions, we asked participants to rate the frequency that the robot referenced itself and the team (with 1 being "never" and 7 being "always"). We used a standard least squares model considering Framing, Timing, and their interaction as main effects. Participants in the Individual conditions stated that the robot referenced itself significantly more frequently ($M = 4.61, SD = .26$) than participants in the Team conditions ($M = 2.54, SD = .26$), $F(1, 67) = 31.56, p < .0001$. On the other hand, participants in the Individual conditions stated that the robot referenced the team significantly less frequently ($M = 2.78, SD = 0.29$) than participants in the Team conditions ($M = 4.80, SD = 0.29$), $F(1, 67) = 24.63, p < .0001$. These results suggest that our Framing manipulation was effective.

4.7.2 Timing of reminders. After the third and fourth games of Space Invaders, the survey asked participants to identify when the robot reminded them to give feedback. In the Before conditions, 68% of participants correctly answered "before the robot said that it was destroying enemies on the left side of the game screen", 26% participants answered incorrectly, and 6% participants answered that they did not remember the ordering. In the After conditions, 75% participants correctly answered "after the robot said that it was destroying enemies on the left side of the game screen", 11% participants answered incorrectly, and 14% participants answered that they did not remember the ordering. This suggests that our Timing manipulation was perceived effectively by most participants.

Importantly, the difference in the Timing independent variable was not evident until games 3 and 4. However, an REML analysis showed that Timing had a significant effect on the rate at which participants provided feedback in games 1 and 2 ($p = .003$), even though this manipulation was not yet evident. This led us to investigate and identify four covariates through correlation analyses: amount of feedback provided in the tutorial ($r(142) = .46, p < .0001$), time to first button press in first game ($r(142) = .33, p < .0001$), participant agreeableness ($r(142) = -.31, p = .0002$), and positive expressivity ($r(142) = .19, p = .02$). These covariates accounted for the difference by Timing in games 1 and 2. Therefore, all statistical analyses in Sec. 5 include these covariates. We also confirmed that significant differences in the manipulation checks described for Timing and Framing persisted after the addition of the covariates.

5 RESULTS

This section presents our results based on the measures described in Sec. 4.5. Unless otherwise noted, we used linear mixed models estimated with Restricted Maximum Likelihood (REML) analyses [74] via JMP Pro [37] to statistically examine survey data and participant feedback. In these analyses, Framing (Individual or Team) and Timing (Before or After) were considered as main effects, and participant ID was a random effect. When the measures were repeated by game number or gameplay strategy, we included Game Number or Gameplay Strategy as a main effect. We also set our selected covariates from Sec. 4.7.2 as fixed effects. We conducted post-hoc Tukey Honestly Significant Difference (HSD) tests or post-hoc Student's t-tests as appropriate.

5.1 Rate of Feedback

First, we present results of analyzing the rate of feedback across all games and in specific windows of time within games (as in Fig. 3).

5.1.1 All games. Across all six games of Space Invaders, participants provided an average of 8 feedback signals per minute (fspm) ($M = 8.02, SE = 0.71$). This ranged from 0 to 160 fspm, with a median value of 4.04 fspm. A REML analysis, including Game Number as a main effect, showed no significant effects by Framing, Timing, or their interaction. The REML analysis did show a significant difference by Game Number, $F(5, 340) = 2.33, p = 0.0423$, but a post-hoc Tukey HSD test showed no significant differences. When considering average feedback across the full length of the games, it is likely that differences from our manipulations in games 3 and 4 were diluted through the whole interaction. Thus, we also looked at the rate of feedback in specific windows of time within games.

5.1.2 Specific windows. As discussed in Sec. 4.1, the robot reminded participants to provide feedback during games 3 and 4. In the ten seconds after the reminder (WR), the rate of feedback varied significantly based on the Timing of the reminder, $F(1, 63) = 5.71, p = .02$, and on the Game Number, $F(1, 68) = 7.43, p = .008$. Participants provided more frequent feedback in WR when the reminder was Before the robot changed its behavior ($M = 22.03, SE = 2.92$) than After ($M = 11.99, SE = 2.87$), as shown in Fig. 4. Participants also provided more feedback in Game 3 ($M = 19.78, SE = 2.24$) than Game 4 ($M = 14.24, SE = 2.24$) in this window (WR).

Because the window after the reminder (WR) had the robot's spaceship in different parts of the game screen based on the Timing of the reminder, we evaluated the rate of feedback in other windows to further investigate the influence of our manipulation. First, we compared the rate of feedback in the ten seconds after when the

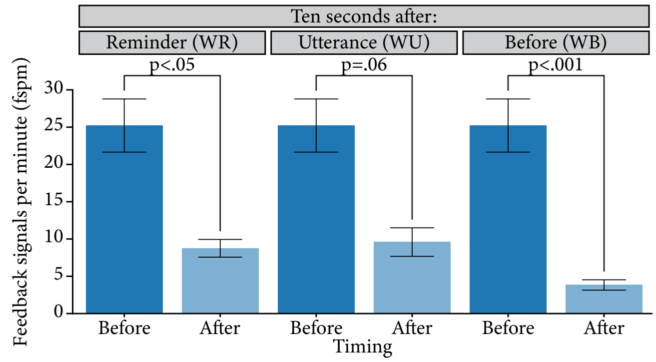

Figure 4: Rate of feedback in three windows described in Fig. 3 by Timing. Unit is feedback signals per minute (fspm).

Before reminder would have been between the Timing conditions (WB), and found a significant difference, $F(1, 63) = 13.38, p = .0005$. Participants in the Before conditions provided more frequent feedback ($M = 22.27, SE = 2.93$) during this window than participants in the After conditions ($M = 6.85, SE = 2.88$) who did not receive the reminder at the start of this window (WB). For WB, the actions of the robot's spaceship were consistent between Timing conditions, so we can assume that the difference is due to the presence of the reminder in the Before conditions. Therefore, it is unlikely that the difference in the rate of feedback that we saw before for WR was due solely to the actions of the robot's spaceship in the game, which differed between the ten seconds following the Before reminder and the After reminder.

Second, we compared the rate of feedback in the ten seconds after the first utterance (WU) and found a trend for Timing having an effect on the results (p=.06). The participants in the Before conditions provided feedback at a rate of 21.74 fspm ($SE = 3.11$) while participants in the After conditions had a rate of 13.10 fspm ($SE = 3.07$). This suggests that for WR, the increased amount of feedback with the Before conditions was not just due to the novelty of the robot speaking for the first time in the Before conditions.

Had the rate of feedback not been higher in WB with the Before reminder than without the Before reminder (due to the participant being in the After conditions), it could be argued that the Before reminder happened to occur at a point in the game when participants were inherently more likely to provide feedback. However, because participants with the Before reminders provided more feedback than participants with the After reminders in both WB and WR, we conclude that the Before reminder increased participants' feedback.

Timing did not have a significant effect on the rate of feedback during the second and third visits in the early-cooperative games, nor during the end of games 5 and 6, when it became evident that the robot had a new gameplay strategy. The Framing of robot utterances (Team vs. Individual) had no significant effect on the rate of feedback provided in any window-based measure.

5.2 Elapsed Time to the Next Feedback

We investigated if there were differences in the elapsed time between when the robot reminded participants to give feedback and when participants next provided feedback via the up or down arrow keys (ER). An REML analysis revealed a significant difference by Timing, $F(1, 63.91) = 4.38, p = .04$. Participants with the Before reminder ($M = 3.24, SE = 1.03$) provided feedback more quickly than with the After reminder ($M = 6.35, SE = 1.02$). There were no other significant effects on the ER measure.

Similar to the secondary analyses for the rate of feedback in specific windows, we again analyzed other elapsed-time measures to evaluate the influence of our manipulation. First, we compared the elapsed time from when the Before reminder would have been across both conditions (EB). We found that the elapsed time varied significantly by Timing for the EB measure, $F(1, 63.74) = 44.29, p < .0001$. Participants with the Before reminders ($M = 3.20, SE = .94$) provided feedback significantly more quickly when the Before reminder was uttered than when the Before reminder was not uttered (because participants instead received the After reminder) ($M = 12.23, SE = .92$). This suggests that the reminder did influence how quickly the participant provided feedback, and the difference observed for ER was not just due to the position of the robot's spaceship, which was the same for both Timings in EB. Additionally, the interaction between Timing and Game Number had a significant effect on elapsed time to feedback in EB, $F(1, 67.91) = 8.54, p = .005$. The post-hoc test showed that Game 3 ($M = 2.56, SE = 1.16$) and Game 4 ($M = 3.84, SE = 1.18$) with the Before reminder led to faster feedback than Game 4 with the After reminder ($M = 10.07, SE = 1.14$). Also, these three combinations (Before-3, Before-4, and After-4) had significantly faster feedback than Game 3 ($M = 14.39, SE = 1.14$) with the After reminder.

Second, because the After reminder was the second utterance of the manipulation, we also conducted an REML analysis on the elapsed time between the robot's first utterance of the manipulation in games 3 and 4 and when the participant next provided feedback (EU). Again, there was a significant difference in the elapsed time by Timing, $F(1, 63.76) = 13.93, p = .0004$. Participants with the Before reminders ($M = 3.33, SE = .84$) provided feedback more quickly after the first utterance than with the After reminders ($M = 7.81, SE = .82$). This result suggests that it was not only that participants responded to the robot saying something in the middle of the game, but that the reminder itself was important. There were no other significant differences.

5.3 Reasons for Providing Feedback

Participants predominantly provided positive feedback to the robot: 83.5% of all participant feedback across all six rounds of Space Invaders was positive. Participants were asked to select all reasons that they gave feedback. Reasons from most to least commonly selected were: "when the robot was on the left side of the screen" (87%), "when the robot was trying to help" (83%), "when the robot was on the right side of the game screen" (70%), "when the robot was not helping" (56%), "when the robot was not being efficient" (35%), "randomly" (30%), and "when the robot lost a life" (11%). Additionally, ten participants (14%) selected "Other". When asked to elaborate on why they chose "Other", nine of the ten participants provided another rationale for giving positive feedback. Of the nine, four participants said they provided feedback when the robot was performing better than they were. For example, P124 wrote *"I would look over and see the robot had done a better job of destroying enemies than I did, and I gave it positive feedback based off of that."* The other five positive reasons were not relative to the participant, but just that the robot was doing well in general, e.g., *"when I saw it was shooting with high frequency"* (P105) or *"whenever it finished clearing its side"* (P186). The one negative reason that was provided was *"when the robot was on the left side of the screen but there were still enemies on the right side of the screen"* (P102).

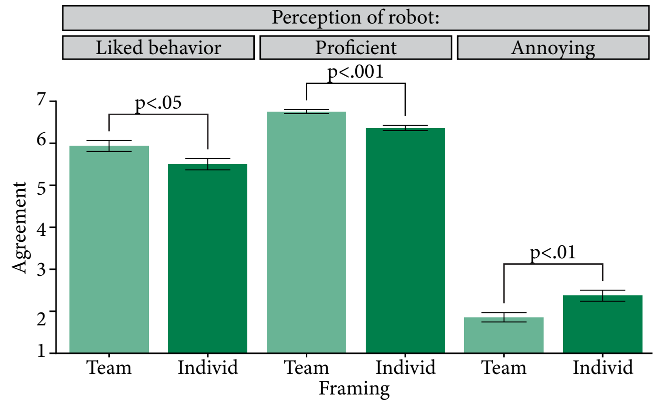

Figure 5: Participant agreement with "I liked the behavior of the robot in the game" (Liked), "The robot was proficient at the game" (Proficient) and "The robot was annoying" (Annoying) on a 7-point responding format.

5.4 Perceptions of the Feedback Process

We next analyzed post-interaction survey questions about the feedback process. The REML analysis showed that how strongly participants agreed that the feedback they provided was helpful varied significantly by Framing, $F(1, 63) = 6.42, p = .01$. Participants in the Team conditions more strongly agreed ($M = 5.72, SE = .24$) that they were able to give helpful feedback to the robot than participants in the Individual conditions ($M = 4.87, SE = .23$). Neither Timing nor the interaction between Framing and Timing had a significant effect on this measure. There were no significant differences by Framing, Timing, or their interaction on how strongly participants agreed that giving feedback was distracting or difficult.

5.5 Perceptions of the Robot

We conducted REML analyses for the post-strategy survey measures about participant perceptions of the robot. Because survey questions were after two games with a specific strategy, robot Gameplay Strategy was included as a main effect. However, given that it is not the focus of this paper, we do not include results for differences by robot Gameplay Strategy due to space constraints.

An REML analysis showed a significant difference by Framing in how much participants liked the behavior of the robot in the game, $F(1, 63) = 4.74, p = .03$. Participants that experienced the Team framing ($M = 5.96, SE = .15$) liked the robot more than those that experienced the Individual framing ($M = 5.47, SE = .15$). The analysis also showed significant differences by Framing in how proficient ($F(1, 63) = 16.11, p = .0002$) and annoying ($F(1, 63) = 7.57, p = .008$) the participants found the robot. The Team framing led to the robot being perceived as more proficient ($M = 6.77, SE = .07$) and less annoying ($M = 1.82, SE = .15$) than the Individual framing (proficiency: $M = 6.35, SE = .07$; annoyance: $M = 2.40, SE = .15$).

We found no other significant effects of Framing, Timing, or their interaction on perceptions of the robot.

6 DISCUSSION

Our first hypothesis (H1a) was not supported. The framing of utterances did not significantly impact the rate of feedback within games nor in the relevant windows of time that we analyzed.

H1b was supported as the participants felt more positively about giving feedback and about the robot when the reminder was framed as helping the team compared to when it was framed as helping the individual robot. Participants that experienced the Team framing more strongly agreed that they were able to give helpful feedback (Sec. 5.4). This could be advantageous for future human-robot interactions because individuals may continue to provide feedback throughout longer interactions if they feel that the feedback they are providing is worthwhile. The Team framing also made the robot seem more proficient and less annoying, and participants reported that they liked the robot's behavior more compared to the Individual framing (Sec. 5.5). While the Framing manipulation did not appear to influence participant actions, it did influence how participants felt about the interaction. Our results reinforce prior work that shows that even a difference of just a few words in how a robot communicates with users matters [20, 63].

We found partial support for H2a, which stated that participants with Before reminders would provide more feedback. While the difference was not significant when we considered full games, participants did provide more feedback in the ten seconds after the Before reminder than in the ten seconds following the After reminder. We suspect this difference was because the robot guided the human's attention to its new behavior with the Before reminder, whereas there was not a novel behavior following the reminder in the After conditions. Based on the results in Sec. 5.1, we are led to believe that a reminder before the robot changes its behavior is more fruitful than a reminder after the change in behavior.

We did not find support for H2b, but instead found evidence that suggests a reverse effect. We hypothesized that participants would give feedback more quickly when the reminder was after the change in behavior. Instead, we found that participants more quickly provided feedback when the reminder was before it was apparent the robot was trying a new gameplay behavior. Whether the goal is to increase the amount of feedback provided or to decrease the elapsed time until the robot receives feedback, the Before reminder appears advantageous based on our study results.

Importantly, participants provided less feedback in the ten seconds after the reminder to give feedback in Game 4 than in Game 3 (Sec. 5.1). This difference highlights the importance of novelty and underscores the importance of understanding how feedback reminders in HRI can be most effective, because it appears that reminders become less meaningful as they are repeated (as in [52]).

Our findings are limited to evaluative feedback. We chose to focus on this type feedback because it required minimal interruption to the participant's own task. However, we posit that our results will transfer to other types of feedback, but would need to study this in future work. In this regard, we suspect that with other kinds of feedback (like corrections), our results may even be stronger than in this study because humans would likely have to focus more on the process of providing feedback for these other types.

7 LIMITATIONS AND FUTURE DIRECTIONS

Our work was limited in several ways, which highlight opportunities for further research. First, our study was conducted in the context of a Space Invaders game. Future research should investigate if the proposed methods for eliciting human feedback are generalizable to other interactions, especially tasks involving more physical manipulation by the robot, e.g., robots learning how to cook with users [60, 78], build physical objects [3, 32], or deliver parts in assembly lines [75]. Second, it is possible that participants were less sensitive to the robot's behavior because its actions changed a virtual environment, not the physical state of the world, even though the robot was situated next to them. Third, in our study, the robot already knew how to play Space Invaders, so participant feedback was for the purpose of fine-tuning collaborative behaviors. It would be interesting to investigate how feedback reminders influence participants when the robot has no prior knowledge of how to perform a task. Another limitation is that the algorithm that determined when a robot reminded participants to provide feedback was based on heuristics and fixed. Future work should investigate how to adapt the framing and timing of reminders to the behavior of users. Finally, our work studied the quantity of feedback provided, but it will be important for future work to study the quality of the feedback provided by humans.

8 CONCLUSION

We investigated the effect of general reminders for humans to provide feedback about a robot's behavior during continuous, collaborative interactions with a robot. Our experimental setup was valuable for investigating human feedback in HRI because while providing feedback, participants were also engaged in the Space Invaders task, which required continuous attention and action on their part. We found that by reminding participants to provide feedback before the Nao tried a new gameplay behavior, the robot could influence participants to provide feedback more quickly and more frequently. Although framing the feedback as helping the team during the reminder did not influence the amount of feedback provided by participants in our study, it did result in more positive feelings about the robot and the process of providing feedback. We hope that our findings encourage the HRI community to incorporate verbal reminders for feedback into interactions where a robot is learning from humans how to improve its behavior.

9 ACKNOWLEDGEMENTS

H. Zhou was partially supported by a Yale Summer Experience Award, and S. Gillet was funded by the Swedish Foundation for Strategic Research (SSF FFL18-0199). Also, this work was partially supported by the National Science Foundation (NSF) under Grant No. IIS-2106690. Any opinions, findings, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the NSF.

REFERENCES

[1] Piotr D. Adamczyk and Brian P. Bailey. 2004. If Not Now, When? The Effects of Interruption at Different Moments within Task Execution. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. Association for Computing Machinery, New York, NY, USA, 271–278. https://doi.org/10.1145/985692.985727

[2] Muneeb Ahmad, Omar Mubin, and Joanne Orlando. 2017. A Systematic Review of Adaptivity in Human-Robot Interaction. *Multimodal Technologies and Interaction* 1, 3 (2017). https://doi.org/10.3390/mti1030014

[3] Gopika Ajaykumar, Maia Stiber, and Chien-Ming Huang. 2021. Designing user-centric programming aids for kinesthetic teaching of collaborative robots. *Robotics and Autonomous Systems* 145 (2021), 103845.

[4] Brenna D. Argall, Eric L. Sauser, and Aude G. Billard. 2011. Tactile Guidance for Policy Adaptation. *Foundations and Trends® in Robotics* 1, 2 (2011), 79–133. https://doi.org/10.1561/2300000012

[5] Andrea Bajcsy, Dylan P. Losey, Marcia K. O'Malley, and Anca D. Dragan. 2017. Learning Robot Objectives from Physical Human Interaction. In *Proceedings of the 1st Annual Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 78)*. PMLR, 217–226.

[6] Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh PN Rao, and Minoru Asada. 2017. Efficient human-robot collaboration: When should a robot take initiative? *The International Journal of Robotics Research* 36, 5-7 (2017), 563–579.

[7] Adam Bignold, Francisco Cruz, Matthew E. Taylor, Tim Brys, Richard Dazeley, Peter Vamplew, and Cameron Foale. 2021. A conceptual framework for externally-influenced agents: An assisted reinforcement learning review. *Journal of Ambient Intelligence and Humanized Computing* (2021), 1–24.

[8] Erdem Bıyık, Dylan P. Losey, Malayandi Palan, Nicholas C Landolfi, Gleb Shevchuk, and Dorsa Sadigh. 2022. Learning reward functions from diverse sources of human feedback: Optimally integrating demonstrations and preferences. *The International Journal of Robotics Research* 41, 1 (2022), 45–67.

[9] Erdem B, Malayandi Palan, Nicholas C. Landolfi, Dylan P. Losey, and Dorsa Sadigh. 2019. Asking Easy Questions: A User-Friendly Approach to Active Reward Learning. In *Proceedings of the 3rd Conference on Robot Learning (Proceedings of Machine Learning Research, Vol. 100)*. PMLR, 1177–1190.

[10] Maya Cakmak, Crystal Chao, and Andrea L. Thomaz. 2010. Designing Interactions for Robot Active Learners. *IEEE Transactions on Autonomous Mental Development* 2, 2 (2010), 108–118. https://doi.org/10.1109/TAMD.2010.2051030

[11] Maya Cakmak, Siddhartha S. Srinivasa, Min Kyung Lee, Sara Kiesler, and Jodi Forlizzi. 2011. Using Spatial and Temporal Contrast for Fluent Robot-Human Hand-Overs. In *Proceedings of the 6th International Conference on Human-Robot Interaction (HRI '11)*. Association for Computing Machinery, New York, NY, USA, 489–496. https://doi.org/10.1145/1957656.1957823

[12] Maya Cakmak and Andrea L. Thomaz. 2012. Designing Robot Learners That Ask Good Questions. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '12)*. Association for Computing Machinery, New York, NY, USA, 17–24. https://doi.org/10.1145/2157689.2157693

[13] Sylvain Calinon and Aude Billard. 2007. Incremental Learning of Gestures by Imitation in a Humanoid Robot. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI '07)*. Association for Computing Machinery, New York, NY, USA, 255–262. https://doi.org/10.1145/1228716.1228751

[14] Kate Candon, Zoe Hsu, Yoony Kim, Jesse Chen, Nathan Tsoi, and Marynel Vázquez. 2022. Perceptions of the Helpfulness of Unexpected Agent Assistance. In *Proceedings of the 10th International Conference on Human-Agent Interaction* (Christchurch, New Zealand) *(HAI '22)*. Association for Computing Machinery, New York, NY, USA, 41–50. https://doi.org/10.1145/3527188.3561915

[15] Joyce Y Chai, Maya Cakmak, Candace Sidner, and J Lupp. 2017. Teaching robots new tasks through natural interaction. In *Interactive Task Learning: Agents, Robots, and Humans Acquiring New Tasks through Natural Interactions, Strüngmann Forum Reports, J. Lupp, series editor*, Vol. 26.

[16] Vivienne Bihe Chi and Bertram F. Malle. 2022. Instruct or Evaluate: How People Choose to Teach Norms to Social Robots. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI '22)*. IEEE Press, 718–722.

[17] Francesca Cini, Tommaso Banfi, Gastone Ciuti, Laila Craighero, and Marco Controzzi. 2021. The relevance of signal timing in human-robot collaborative manipulation. *Science Robotics* 6, 58, eabg1308. https://doi.org/10.1126/scirobotics.abg1308

[18] Caitlyn Clabaugh and Maja J. Matarić. 2016. Exploring elicitation frequency of learning-sensitive information by a robotic tutor for interactive personalization. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE Press, 968–973. https://doi.org/10.1109/ROMAN.2016.7745226

[19] Joe Connolly, Viola Mocz, Nicole Salomons, Joseph Valdez, Nathan Tsoi, Brian Scassellati, and Marynel Vázquez. 2020. Prompting Prosocial Human Interventions in Response to Robot Mistreatment. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*. Association for Computing Machinery, New York, NY, USA, 211–220. https://doi.org/10.1145/3319502.3374781

[20] Filipa Correia, Samuel Mascarenhas, Rui Prada, Francisco S. Melo, and Ana Paiva. 2018. Group-Based Emotions in Teams of Humans and Robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI '18)*. Association for Computing Machinery, New York, NY, USA, 261–269. https://doi.org/10.1145/3171221.3171252

[21] Yuchen Cui, Pallavi Koppol, Henny Admoni, Scott Niekum, Reid Simmons, Aaron Steinfeld, and Tesca Fitzgerald. 2021. Understanding the Relationship between Interactions and Outcomes in Human-in-the-Loop Machine Learning. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI '21)*. International Joint Conferences on Artificial Intelligence Organization, 4382–4391. https://doi.org/10.24963/ijcai.2021/599

[22] Anca D. Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S. Srinivasa. 2015. Effects of Robot Motion on Human-Robot Collaboration. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '15)*. Association for Computing Machinery, New York, NY, USA, 51–58. https://doi.org/10.1145/2696454.2696473

[23] Hadas Erel, Elior Carsenti, and Oren Zuckerman. 2022. A Carryover Effect in HRI: Beyond Direct Social Effects in Human-Robot Interaction. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI '22)*. IEEE Press, 342–352.

[24] Juan Fasola and Maja J Matarić. 2013. A Socially Assistive Robot Exercise Coach for the Elderly. *Journal of Human-Robot Interaction* 2, 2 (June 2013), 3–32. https://doi.org/10.5898/JHRI.2.2.Fasola

[25] Taylor A. Kessler Faulkner, Elaine Schaertl Short, and Andrea L. Thomaz. 2020. Interactive Reinforcement Learning with Inaccurate Feedback. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 7498–7504.

[26] Tesca Fitzgerald, Pallavi Koppol, Patrick Callaghan, Russell Quinlan Jun Hei Wong, Reid Simmons, Oliver Kroemer, and Henny Admoni. 2022. INQUIRE: INteractive Querying for User-aware Informative REasoning. In *6th Annual Conference on Robot Learning (Proceedings of Machine Learning Research)*. PMLR.

[27] Sarah Gillet, Maria Teresa Parreira, Marynel Vázquez, and Iolanda Leite. 2022. Learning Gaze Behaviors for Balancing Participation in Group Human-Robot Interactions. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI '22)*. IEEE Press, 265–274.

[28] Samuel D. Gosling, Peter J. Rentfrow, and William B. Swann Jr. 2003. A very brief measure of the Big-Five personality domains. *Journal of Research in Personality* 37, 6 (2003), 504–528.

[29] Victoria Groom, Jimmy Chen, Theresa Johnson, F. Arda Kara, and Clifford Nass. 2010. Critic, Compatriot, or Chump? Responses to Robot Blame Attribution. In *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI '10)*. IEEE Press, 211–218.

[30] James J. Gross and O.P. John. 1997. Revealing feelings: Facets of emotional expressivity in self-reports, peer ratings, and behavior. *Journal of Personality and Social Psychology* 72 (1997), 435–448.

[31] Balint Gucsi, Danesh S. Tarapore, William Yeoh, Christopher Amato, and Long Tran-Thanh. 2020. To Ask or Not to Ask: A User Annoyance Aware Preference Elicitation Framework for Social Robots. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE Press, 7935–7940. https://doi.org/10.1109/IROS45743.2020.9341607

[32] Bradley Hayes and Brian Scassellati. 2016. Autonomously Constructing Hierarchical Task Networks for Planning and Human-Robot Collaboration. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE Press, 5469–5476. https://doi.org/10.1109/ICRA.2016.7487760

[33] Mark K. Ho, James MacGlashan, Michael L. Littman, and Fiery Cushman. 2017. Social is special: A normative framework for teaching with and learning from evaluative feedback. *Cognition* 167 (2017), 91–106. https://doi.org/10.1016/j.cognition.2017.03.006

[34] Guy Hoffman, Jodi Forlizzi, Shahar Ayal, Aaron Steinfeld, John Antanitis, Guy Hochman, Eric Hochendoner, and Justin Finkenaur. 2015. Robot Presence and Human Honesty: Experimental Evidence. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '15)*. Association for Computing Machinery, New York, NY, USA, 181–188. https://doi.org/10.1145/2696454.2696487

[35] Mohammed Moshiul Hoque, Tomomi Onuki, Dipankar Das, Yoshinori Kobayashi, and Yoshinori Kuno. 2012. Attracting and Controlling Human Attention through Robot's Behaviors Suited to the Situation. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI '12)*. Association for Computing Machinery, New York, NY, USA, 149–150. https://doi.org/10.1145/2157689.2157729

[36] Auke Jan Ijspeert, Jun Nakanishi, and Stefan Schaal. 2002. Movement imitation with nonlinear dynamical systems in humanoid robots. In *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, Vol. 2. IEEE, 1398–1403.

[37] SAS Institute. 2021. JMP Pro version 15.0.0.

[38] Tariq Iqbal and Laurel D. Riek. 2016. Human Coordination Dynamics with Heterogeneous Robots in a Team. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction (HRI '16)*. IEEE Press, 619–620.

[39] Charles Isbell and Christian Shelton. 2001. Cobot: A Social Reinforcement Learning Agent. In *Advances in Neural Information Processing Systems*, Vol. 14. MIT Press. https://proceedings.neurips.cc/paper/2001/file/92bbd31f8e0e43a7da8a6295b251725f-Paper.pdf

[40] Hong Jun Jeon, Smitha Milli, and Anca Dragan. 2020. Reward-rational (implicit) choice: A unifying formalism for reward learning. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 4415–4426. https://proceedings.neurips.cc/paper/2020/file/2f10c1578a0706e06b6d7db6f0b4a6af-Paper.pdf

[41] Poornima Kaniarasu and Aaron M. Steinfeld. 2014. Effects of blame on trust in human robot interaction. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE Press, 850–855. https://doi.org/10.1109/ROMAN.2014.6926359

[42] Tasneem Kaochar, Raquel Torres Peralta, Clayton T. Morrison, Ian R. Fasel, Thomas J. Walsh, and Paul R. Cohen. 2011. Towards Understanding How Humans Teach Robots. In *Proceedings of the 19th International Conference on User Modeling, Adaption, and Personalization (UMAP'11)*. Springer-Verlag, Berlin, Heidelberg, 347–352.

[43] Taylor Kessler Faulkner, Reymundo A. Gutierrez, Elaine Schaertl Short, Guy Hoffman, and Andrea L. Thomaz. 2019. Active Attention-Modified Policy Shaping. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*. IFAAMAS, 728–736.

[44] Elizabeth S. Kim, Dan Leyzberg, Katherine M. Tsui, and Brian Scassellati. 2009. How People Talk When Teaching a Robot. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction (HRI '09)*. Association for Computing Machinery, New York, NY, USA, 23–30. https://doi.org/10.1145/1514095.1514102

[45] W. Bradley Knox, Brian D. Glass, Bradley C. Love, W. Todd Maddox, and Peter Stone. 2012. How Humans Teach Agents. *International Journal of Social Robotics* 4, 4 (2012), 409–421. https://doi.org/10.1007/s12369-012-0163-x

[46] W. Bradley Knox and Peter Stone. 2008. TAMER: Training an Agent Manually via Evaluative Reinforcement. In *2008 7th IEEE International Conference on Development and Learning*. IEEE, 292–297.

[47] Martin Kraft and Markus Rickert. 2017. How to Teach Your Robot in 5 Minutes: Applying UX Paradigms to Human-Robot-Interaction. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE Press, 942–949. https://doi.org/10.1109/ROMAN.2017.8172416

[48] Minae Kwon, Mengxi Li, Alexandre Bucquet, and Dorsa Sadigh. 2019. Influencing Leading and Following in Human-Robot Teams. In *Proceedings of Robotics: Science and Systems (RSS)*.

[49] Jamie Large, Graham Stodolski, and Marynel Vázquez. 2020. Studying Human-Agent Interactions in Space Invaders. In *Proceedings of the 8th International Conference on Human-Agent Interaction (HAI '20)*. Association for Computing Machinery, New York, NY, USA, 245–247. https://doi.org/10.1145/3406499.3418747

[50] Mark Levine, Amy Prosser, David Evans, and Stephen Reicher. 2005. Identity and Emergency Intervention: How Social Group Membership and Inclusiveness of Group Boundaries Shape Helping Behavior. *Personality and Social Psychology Bulletin* 31, 4 (2005), 443–453. https://doi.org/10.1177/0146167204271651 PMID: 15743980.

[51] Mark Levine and Kirstien Thompson. 2004. Identity, Place, and Bystander Intervention: Social Categories and Helping After Natural Disasters. *The Journal of Social Psychology* 144, 3 (2004), 229–245. https://doi.org/10.3200/SOCP.144.3.229-245 PMID: 15168427.

[52] Guangliang Li, Hayley Hung, Shimon Whiteson, and W. Bradley Knox. 2013. Using Informative Behavior to Increase Engagement in the Tamer Framework. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '13)*. IFAAMAS, 909–916.

[53] Robert Loftin, James MacGlashan, Bei Peng, Matthew E. Taylor, Michael L. Littman, Jeff Huang, and David L. Roberts. 2014. A Strategy-Aware Technique for Learning Behaviors from Discrete Human Feedback. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI'14)*. AAAI Press, 937–943.

[54] Thibaut Munzer, Marc Toussaint, and Manuel Lopes. 2017. Preference Learning on the Execution of Collaborative Human-Robot Tasks. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE Press, 879–885. https://doi.org/10.1109/ICRA.2017.7989108

[55] Stefanos Nikolaidis, Anton Kuznetsov, David Hsu, and Siddharta Srinivasa. 2016. Formalizing Human-Robot Mutual Adaptation: A Bounded Memory Model. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction (HRI '16)*. IEEE Press, 75–82. https://doi.org/10.1109/HRI.2016.7451736

[56] Raquel Oliveira, Patrícia Arriaga, Patrícia Alves-Oliveira, Filipa Correia, Sofia Petisca, and Ana Paiva. 2018. Friends or Foes? Socioemotional Support and Gaze Behaviors in Mixed Groups of Humans and Robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI '18)*. Association for Computing Machinery, New York, NY, USA, 279–288. https://doi.org/10.1145/3171221.3171272

[57] Morgan Quigley, Brian Gerkey, Ken Conley, Josh Faust, Tully Foote, Jeremy Leibs, Eric Berger, Rob Wheeler, and Andrew Ng. 2009. ROS: an open-source Robot Operating System. In *ICRA Workshop on Open Source Software*, Vol. 3. 5.

[58] P.L. Patrick Rau, Ye Li, and Dingjun Li. 2009. Effects of Communication Style and Culture on Ability to Accept Recommendations from Robots. *Computers in Human Behavior* 25, 2 (March 2009), 587–595. https://doi.org/10.1016/j.chb.2008.12.025

[59] Kantwon Rogers and Ayanna Howard. 2021. Looks Can Permit Deceiving: How Reward or Punishment Decisions Are Influenced by Robot Embodiment. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21 Companion)*. Association for Computing Machinery, New York, NY, USA, 57–61. https://doi.org/10.1145/3434074.3447128

[60] Radu Bogdan Rusu, Brian Gerkey, and Michael Beetz. 2008. Robots in the Kitchen: Exploiting Ubiquitous Sensing and Actuation. *Robotics and Autonomous Systems* 56, 10 (2008), 844–856.

[61] Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. 2017. Active Preference-Based Learning of Reward Functions. In *Proceedings of Robotics: Science and Systems (RSS)*.

[62] Dorsa Sadigh, Shankar Sastry, Sanjit A. Seshia, and Anca D Dragan. 2016. Planning for Autonomous Cars that Leverage Effects on Human Actions. In *Proceedings of Robotics: Science and Systems*, Vol. 2. 1–9.

[63] Nicole Salomons, Kaitlynn Taylor Pineda, Adérónké Adéjàre, and Brian Scassellati. 2022. "We Make a Great Team!": Adults with Low Prior Domain Knowledge Learn More from a Peer Robot than a Tutor Robot. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI '22)*. IEEE Press, 176–184.

[64] Joe Saunders, Chrystopher L. Nehaniv, and Kerstin Dautenhahn. 2006. Teaching Robots by Moulding Behavior and Scaffolding the Environment. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction (HRI '06)*. Association for Computing Machinery, New York, NY, USA, 118–125. https://doi.org/10.1145/1121241.1121263

[65] Shane Saunderson and Goldie Nejat. 2019. How Robots Influence Humans: A Survey of Nonverbal Communication in Social Human–Robot Interaction. *International Journal of Social Robotics* 11, 4 (2019), 575–608. https://doi.org/10.1007/s12369-019-00523-0

[66] Mariah L. Schrum, Erin Hedlund-Botti, Nina Moorman, and Matthew C. Gombolay. 2022. MIND MELD: Personalized Meta-Learning for Robot-Centric Imitation Learning. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction (HRI '22)*. IEEE Press, 157–165.

[67] Sarah Strohkorb Sebo, Priyanka Krishnamurthi, and Brian Scassellati. 2020. "I Don't Believe You": Investigating the Effects of Robot Trust Violation and Repair. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI '19)*. IEEE Press, 57–65.

[68] Sarah Strohkorb Sebo, Margaret Traeger, Malte Jung, and Brian Scassellati. 2018. The Ripple Effects of Vulnerability: The Effects of a Robot's Vulnerable Behavior on Trust in Human-Robot Teams. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction (HRI '18)*. Association for Computing Machinery, New York, NY, USA, 178–186. https://doi.org/10.1145/3171221.3171275

[69] Xiang Zhi Tan, Sean Andrist, Dan Bohus, and Eric Horvitz. 2020. Now, Over Here: Leveraging Extended Attentional Capabilities in Human-Robot Interaction. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*. Association for Computing Machinery, New York, NY, USA, 468–470. https://doi.org/10.1145/3371382.3378363

[70] Andrea Thomaz, Guy Hoffman, and Maya Cakmak. 2016. Computational Human-Robot Interaction. *Foundations and Trends® in Robotics* 4, 2-3 (2016), 105–223. https://doi.org/10.1561/2300000049

[71] Andrea L. Thomaz and Cynthia Breazeal. 2006. Reinforcement Learning with Human Teachers: Evidence of Feedback and Guidance with Implications for Learning Performance. In *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 1 (AAAI'06)*. AAAI Press, 1000–1005.

[72] Andrea L. Thomaz and Cynthia Breazeal. 2007. Asymmetric Interpretations of Positive and Negative Human Feedback for a Social Learning Agent. In *The 16th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 720–725. https://doi.org/10.1109/ROMAN.2007.4415180

[73] Andrea L. Thomaz and Maya Cakmak. 2009. Learning about Objects with Human Teachers. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction (HRI '09)*. Association for Computing Machinery, New York, NY, USA, 15–22. https://doi.org/10.1145/1514095.1514101

[74] Robin Thompson. 1980. Maximum likelihood estimation of variance components. *Statistics: A Journal of Theoretical and Applied Statistics* 11, 4 (1980), 545–561. https://doi.org/10.1080/02331888008801560

[75] Vaibhav V. Unhelkar, Przemyslaw A. Lasota, Quirin Tyroller, Rares-Darius Buhai, Laurie Marceau, Barbara Deml, and Julie A. Shah. 2018. Human-Aware Robotic Assistant for Collaborative Assembly: Integrating Human Motion Prediction With Planning in Time. *IEEE Robotics and Automation Letters* 3, 3 (2018), 2394–2401. https://doi.org/10.1109/LRA.2018.2812906

[76] Diede P.M. Van der Hoorn, Anouk Neerincx, and Maartje M.A. de Graaf. 2021. "I Think You Are Doing a Bad Job!": The Effect of Blame Attribution by a Robot in Human-Robot Collaboration. In *Proceedings of the 2021 ACM/IEEE International*

Conference on Human-Robot Interaction (HRI '21). Association for Computing Machinery, New York, NY, USA, 140–148. https://doi.org/10.1145/3434073.3444681

[77] Guan Wang, Carl Trimbach, Jun Ki Lee, Mark K. Ho, and Michael L. Littman. 2020. Teaching a Robot Tasks of Arbitrary Complexity via Human Feedback. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*. Association for Computing Machinery, New York, NY, USA, 649–657. https://doi.org/10.1145/3319502.3374824

[78] Xiaofei Wang, Kimin Lee, Kourosh Hakhamaneshi, Pieter Abbeel, and Michael Laskin. 2022. Skill Preferences: Learning to Extract and Execute Robotic Skills from Human Feedback. In *6th Annual Conference on Robot Learning (Proceedings of Machine Learning Research)*. PMLR, 1259–1268.

[79] David Whitney, Eric Rosen, James MacGlashan, Lawson L. S. Wong, and Stefanie Tellex. 2017. Reducing Errors in Object-Fetching Interactions through Social Feedback. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE Press, 1006–1013. https://doi.org/10.1109/ICRA.2017.7989121

[80] Bryce Woodworth, Francesco Ferrari, Teofilo E. Zosa, and Laurel D. Riek. 2018. Preference Learning in Assistive Robotics: Observational Repeated Inverse Reinforcement Learning. In *Proceedings of the 3rd Machine Learning for Healthcare Conference (Proceedings of Machine Learning Research, Vol. 85)*. PMLR, 420–439. https://proceedings.mlr.press/v85/woodworth18a.html

[81] Julia L. Wright, Shan G. Lakhmani, and Jessie Y.C. Chen. 2022. Bidirectional Communications in Human-Agent Teaming: The Effects of Communication Style and Feedback. *International Journal of Human–Computer Interaction* 38, 18-20 (2022), 1972–1985. https://doi.org/10.1080/10447318.2022.2068744

[82] Zhou Yu, Dan Bohus, and Eric Horvitz. 2015. Incremental Coordination: Attention-Centric Speech Production in a Physically Situated Conversational Agent. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*. 402–406.