

Abstract

Robot Learning During Collaborations with Non-Expert Robot Users

Kathleen Coyle Candon

2026

As robots integrate into everyday life, they must not only be able to execute a variety of tasks, but must also align their behavior with the varying preferences of non-expert users with whom they collaborate. Because these preferences are often unknown prior to an interaction, robots must be able to learn from non-expert users during collaborations. However, traditional robot learning frameworks often overlook the fact that these collaborations occur within a dynamic social context – a complex interplay of factors related to the agents, their environments, and associations between them.

This dissertation argues that the social context is not merely a backdrop for interaction, but an active mechanism that can be leveraged and shaped to improve robot learning. We contribute a formal taxonomy for the “social context of a human-robot interaction” to provide a unified language for the field. We then interpret and influence the social context in different ways to improve robot learning.

We demonstrate that robots can actively influence the social context to influence how human collaborators provide explicit feedback through the framing and timing of reminders. Then, we show that we can extract additional informative signals that humans “leak” as part of the social context during interactions. We develop models that incorporate this socially contextual information to more accurately predict human preferences over agent behaviors. Additionally, we facilitate combining these types of feedback by contributing the REACT database, a multimodal collection of human reactions and evaluative feedback

captured over time during two different interaction scenarios. Finally, we propose a novel mathematical formulation of human preferences over a collaboration and introduce a model to learn from both explicit and implicit human feedback, combining different signals from the social context of the collaboration.

Ultimately, this work demonstrates that by actively engaging with and exploiting the social context of a human-robot interaction, robots can leverage more feedback from non-expert users without increasing human's teaching burden during human-robot collaborations.

Robot Learning During Collaborations with Non-Expert Robot Users

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Kathleen Coyle Candon

Dissertation Director: Marynel Vázquez

September 2026

Copyright © 2026 by Kathleen Coyle Candon

All rights reserved.

Contents

Dedication	x
Acknowledgements	xi
List of Figures	xvi
List of Tables	xxiii
1 Introduction	1
1.1 Outline	3
2 Related Work	8
2.1 Robot Learning	8
2.1.1 Explicit Human Feedback	9
2.1.2 Implicit Human Feedback	9
2.2 Human-Robot Collaboration (HRC)	11

2.3	Non-Expert Robot Users	12
2.4	Synthesis	13
3	The Social Context of Human-Robot Interactions	14
3.1	Overview	15
3.2	Review Methodology	18
3.3	The Term “Social Context” in HRI	20
3.3.1	Explicit definitions of the term “social context”.	20
3.3.2	Broader usage of the term “social context”	23
3.4	A Conceptual Model of the Social Context of a Human-Robot Interaction	24
3.4.1	Terminology	25
3.4.2	Selection criteria for analysis of proposed definition of social context	30
3.4.3	Example use-case scenarios	30
3.4.4	Types of socially contextual information	34
3.5	Discussion	40
3.5.1	Planning for human-robot interactions	41
3.5.2	Robot behavior generation during interactions	42
3.5.3	Post-interaction analyses	45
4	Soliciting Explicit Signals	47

4.1	Related Work	49
4.2	Interaction Task: Space Invaders	51
4.3	Method	53
4.3.1	Study Design	54
4.3.2	Hypotheses	55
4.3.3	Setup	57
4.3.4	Procedure	58
4.3.5	Dependent Measures	59
4.3.6	Participants	60
4.3.7	Manipulation Checks	61
4.4	Results	63
4.4.1	Rate of Feedback	63
4.4.2	Elapsed Time to the Next Feedback	66
4.4.3	Reasons for Providing Feedback	67
4.4.4	Perceptions of the Feedback Process	68
4.4.5	Perceptions of the Robot	68
4.5	Discussion	70
4.6	Limitations and Future Directions	72
4.7	Conclusion	72

5	Interpreting Facial Expressions as Implicit Signals	74
5.1	Related Work	76
5.2	Method	78
5.2.1	Data Collection	79
5.2.2	Preference Prediction Task	81
5.3	Results	87
5.3.1	Naive Baselines for Predicting Preferences	88
5.3.2	Nonverbal Reactions Can Help Model Preferences	88
5.3.3	Incorporating Additional Context	89
5.4	Discussion	92
5.5	Limitations and Future Work	94
5.6	Conclusion	95
6	Two Datasets for Analyzing Implicit and Explicit Signals Together	97
6.1	Related Work	99
6.1.1	Existing Datasets	99
6.1.2	Interaction History	100
6.2	The REACT-Nao Dataset	101
6.2.1	Data Collection	101

6.2.2	Data Processing	102
6.2.3	Results	103
6.3	The REACT-Shutter Dataset	104
6.3.1	Data Collection	104
6.3.2	Data Processing	105
6.3.3	Results	106
6.4	Discussion	107
6.5	Conclusion	108
7	Combining Implicit Signals Via Physical Actions with Explicit Feedback	109
7.1	Related Work	111
7.2	Problem Setup: Learning Preferences Over a Human-Robot Collaboration	111
7.3	Preference Learning from Implicit and Explicit Feedback (PIE)	114
7.3.1	Implication of Explicit Human Feedback	116
7.3.2	Implication of Implicit Human Feedback	117
7.3.3	Estimating Belief Over Preference Weights	118
7.4	Evaluation in Simulation	120
7.4.1	The Pizza-Making Task	121
7.4.2	Simulating the Human in the Collaboration	122

7.4.3	Evaluation Setup	124
7.4.4	Results	124
7.5	Real-World Evaluation	129
7.5.1	Experimental Protocol	129
7.5.2	Results	131
7.6	Discussion	132
8	Conclusion	134
8.1	Contributions	134
8.2	Open Questions and Opportunities for Future Research	136
8.2.1	Deeper Understanding of When Humans Choose to Give Feedback to Robots	136
8.2.2	Evaluation of Feedback Quality	137
8.2.3	Broader Signals of Implicit Human Feedback	137
8.2.4	Verbal Feedback	138
8.2.5	Differentiated Interpretation of Different Types of Feedback	138
8.2.6	Other Ways to Solicit Feedback	138
8.2.7	Robot Learning During Group Human-Robot Collaborations	139
8.2.8	Summary	139

In loving memory of Grammy and Gigi — thanks for the kindling.

“For the mind does not require filling like a bottle, but rather, like wood, it only requires kindling to create in it an impulse to think independently and an ardent desire for the truth.”

— PLUTARCH

Acknowledgments

I am deeply grateful for the support and encouragement I receive not only over the course of this PhD journey, but throughout my life. While this was a six year process, the foundation was laid much earlier. I am so very lucky.

I feel incredibly fortunate to have been advised by Marynel Vázquez. I barely knew what human-robot interaction was when I started this endeavor, but you have been a constant wealth of knowledge, inspiration, and optimism. Thank you for your kindness, your care, and your belief in me. I aspire to be the mentor you are.

I was also fortunate to be unofficially co-advised by Brian Scassellati (Scaz), and I am a better thinker for it. Thank you for being honest with me, introducing me to new board games, and pushing me to defend my research convictions.

I am indebted to Tesca Fitzgerald for serving on my dissertation committee, participating in my area exam, and graciously chatting with me the many times I popped into your office. I want to especially thank you for the whiteboard sessions regarding your INQUIRE formalism and for all of your advice during the job search process. I am grateful to Julie Shah for serving as my external reader. Your kindness and encouragement have been inspiring, and I hope to pay that forward throughout my career. Thank you to Julian Jara-Ettinger for serving on my area exam committee, for the insightful chats about cognitive science, and for organizing the Metazoa reading group. I thank Janet Kayfatz for your encouragement and help in honing my storytelling.

Thank you to the staff who keep the Yale Computer Science Department running smoothly, especially Nancy Pellegrino, Judi Paige, Annette Durso, and Alicia Vignola. Your kindness and prompt responses to my many emails never went unnoticed. I also want to thank the faculty who made AKW such a welcoming place to work, especially Ted Kim,

Holly Rushmeier, Zhong Shao, Alex Wong, Lin Zhong, and Steve Zucker, who were always friendly faces in the halls.

Thank you to my incredible labmates in the Interactive Machines Group: Houston Claire, Nathan Tsoi, Sydney Thompson, Qiping Zhang, Austin Narcomey, Sasha Lew, and Etiosa Omeike. This journey would not have been possible without your feedback, insights, ideas, hard work, and camaraderie. I am especially indebted to Sasha for your thoughtfulness, intellect, and incomparable robot engineering skills. Thank you for all of the delicious baked goods, constant supply of lab snacks, and helpful metaphors. I also had the pleasure of collaborating with and learning from Chayan Sarkar and Sarah Gillet during their visits to Yale. Sarah – soon we’ll be old friends catching up at HRI conferences; I can’t wait. Sydney – thank you for the many lunches and your genuine friendship throughout this endeavor – I am grateful to have been in the trenches with you.

I am lucky to have had a second lab community in the Social Robotics Lab: Nicole Salomons, Meiyang Qin, Jake Brawer, Tim Adamson, Nicholas Georgiou, Debasmita Ghose, Emmanuel Adenieran, Rebecca Ramnauth, Kayla Matheus, Ellie Mamantov, Fern Limprayoon, Shuangge Wang, Liam Merz Hoffmeister, Moon Kim, Hend Rasheed. It was a pleasure to “go do good stuff” each week with you all. Ellie – thank you for asking me to grab lunch many summers ago and for your regular check-ins after leaving New Haven. I am grateful for your friendship and I treasure our TV nights, shopping excursions, and nights eating takeout.

I am grateful to my other robotics colleagues in the Inquisitive Robotics Lab and APOLLO Lab, especially Anjiabei Wang, Ulas Berk Karli, Anushka Potdar, Ziyao Shang-guan, and TJ Vitchutripop. I particularly look back fondly on the ski trip to Winter Park after HRI 2024 in Boulder.

To the undergraduate students I've had the privilege of mentoring, thank you for your fresh perspectives and the energy you brought to my research. In particular, thank you to my undergrad co-authors – Jeese Chen, Zoe Hsu, Yoony Kim, Lena Qian, Alyssa Quarles, Sidney Richardson, and Helen Zhou – this dissertation is stronger because of your contributions.

My time in CT was made infinitely better by my New Haven friends – Lily, Violet, Conor, Ethen, Gillie, Davis, Mason, Michelle, Susanna, Lauren, Sara Beth, Matt, Morgan, Ansong, Marco, John, and so many more – thank you for the laughter and joy. I am grateful for all of the pickleball, ski trips, trivia nights with Little Disco, ice cream stops at Elena's, softball games with the Rusty Gears, soccer games with Kick out the Hams, basketball games Boolean Ballers, and all of the other adventures we undertook.

One perk of taking some time to work between college and graduate school is that I had even more time to amass great friendships before undertaking this adventure. Thank you all for the fun and support over the years. Andrew – from riding bikes in preschool to showing me around the West Village, your friendship has been a mainstay in my life for as long as I can remember. Thank you for hopping on Metro North to provide company and DJ-ing services on many drives between CT and VT. Jenna – I am forever grateful that we both wore lacy socks to kindergarten. Kristen and Audrey – my musketeers from college; I'm always looking forward to our next trip. Jen and Devin – you both inspire me everytime we catch up. Skylar – thank you for understanding that I may not always send a text, but I will take the train to NYC to meet you for dinner. Friendship from your first job is a special bond; Shelbi and Katherine – I would redo those first two years of “adulting” all over again just to make sure our paths crossed. Jane – thank you for answering random phone calls and catching up like no time has passed.

I am indebted to Lisa, Lorie, and Ed, who opened their homes to me when I arrived in

New Haven in the middle of the pandemic with no friends. Thank you for all of the meals, games of cards and dominoes, and showing me the Carina Park pickleball courts.

To my extended family (Table 0), I am incredibly fortunate to belong to such a wonderful crew. To the Candons, I am grateful for the weeks spent on Lake Champlain, the many nights spent playing various card games, and all of the Thanksgiving memories. Nancy and Reets, my math teacher aunts – thank you for the math gene. Sue – thank you for your advice and the mai tais. Reilly – thank you for hosting me in Castlebar; I am so glad we’ve grown our friendship in adulthood. To the Ojalas, I am grateful for the party-to-plan-a-party parties, the lake days on Lake Bomoseen, uncountable games of Spud and Phase 10. Uncle John and Pam – thank you for raising two of my best friends and always making me feel at home in Proctor. Abbey and Greta – thank you for being the older sisters I didn’t have; I have looked up to you both for as long as I can remember. And Jackie – thank you for being my person. Whether we’re driving a campervan around New Zealand or watching Grey’s, I am so lucky to have you by my side.

Table 0. Evidence supporting my reputation as the “one with a big family”.

Group	Candon Family	Ojala Family	Thank you...
Aunts and Uncles	Dynan & Diane, Nancy & George, Jim & Betty, Reets & Ace, Mark & Debbie	Carla & Don, David, Peggy, Gigi & Johnny T, Gus & Jeanne, Fran & Mary, Uncle John & Pam	for being some of my biggest cheerleaders.
Cousins	Michael, Liz, Heather, Joe, Sue, Jimmy, Sean, Luke, Justin, Molly, Reilly	Susannah, Aili, Abbey, Greta, Jennifer, Carl, Christie, Lucas, Hannah, Jackie, Nick	for making me laugh like no one else can.
Cousin-in-Laws	Steve, Jen, Mark, Suzanne, Nicole, Eoghan	Adam, Brian, Tim, Mike, Anthony, Rachel, Madison, Pam, Richie, Dan, Dehlia	for bringing extra fun to the fam.
The Kiddos	Skylar, Kiley, Bode, Addy, Ace, Kenna, Cooper, Will, Mae, Eloise, Marky, Libby, JD, Sadie Jo	Sam, Hank, Vinny, Harper, Tilly, Russell, LT, Gianna, Lukey, Baby Bird, Baby Campbell	for the joy, the smiles, and your infectious curiosity.

Joe – thank you for being the fun one and reminding me to get out of my head when I need it. Some of my favorite New Haven memories coincide with your visits, especially the Harvard-Yale games, the rainy rodeo party, and the mechanical bull.

My deepest thanks go to Mom and Dad; I could not have done this without you. Even if I didn't always believe it, I knew you two thought I could do anything. Thank you both for your unwavering support and unconditional love.

Funding: This research was made possible by the generous support of the National Science Foundation (NSF), Tata Sons Private Limited, Tata Consultancy Services Limited, and Titan for the funding that enabled the research presented in this dissertation. I gratefully acknowledge the pivotal role that these institutions play in funding scientific advancement; their support was instrumental in enabling the research presented in this dissertation.

List of Figures

1.1	Illustrative human-robot collaboration. This dissertation explores how robots can use a novel definition of the “social context of a human-robot interaction” to extract additional feedback from collaborations for learning. We explore this through three different channels. (a) First, we explore solicited explicit signals via robot reminders for more feedback (Chapter 4). (b) Then, we analyze leaked implicit signals via human nonverbal communicative signals (Chapter 5). Third, we interpret the physical actions a human takes during collaborations as implicit signal and combine those signals with explicitly provided feedback via evaluative button presses (Chapter 7).	4
3.1	(a) Paper selection process for review. (b) Histogram of 320 papers using the term “social context” (or close variants like “socially contextual”) in the title or body of the paper. Best viewed in color.	20

3.2	<p>Examples of relevant entities for two different human-robot interactions: one between a firefighter and a robot, and another between a robot guide and a person in an office. Figures (a) and (d) illustrate the interactions at a given time; the images were created with Gemini’s Generative AI. The diagram in (b) shows the relevant entities for (a), including the robot (R_1) and the firefighter (H_1). The firefighter teleoperates the robot that is in E_1, while being physically located in E_2. Dashed edges have relevant agent-environment associations, while solid edges indicate relevant agent-agent associations. The diagram in (c) shows the interaction some time after (b), when the firefighter and robot team find another person (H_2). This illustrates how social context can be dynamic – the relevant elements of the interaction can vary over time, so the attributes that influence human-robot interactions are also variable. In (d), a robot (R_2) brings a person (H_3) to an office, where there are two other people (H_4 and H_5) who are not considered part of the interaction of interest. Hence, they are visualized in (e) with a lighter color-scheme. However, attributes of H_4 and H_5, like whether they are attending to the R_2 or H_3, influence the interaction of interest. As a result, the attributes of H_4 and H_5 are part of the social context of the interaction of interest.</p>	28
3.3	Taxonomy for socially contextual information of an environment.	36
3.4	Taxonomy for socially contextual information of an agent	38
3.5	Taxonomy for socially contextual information of an association	39
4.1	Experimental setup for our study.	48

4.2	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. 4.2. 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.	52
4.3	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.	57
4.4	Rate of feedback in three windows described in Fig. 4.3 by Timing. Unit is feedback signals per minute (fspm).	64
4.5	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.	69
5.1	We collected data from an online human-agent interaction to investigate the usefulness of including human nonverbal signals into models for predicting which agent behavior participants preferred. In our study, the participants were recorded while playing the game with an autonomous agent via their webcams.	75

5.2	Participants were introduced to the game as shown above. The explanation included explicit pointers to the scores and number of remaining lives for each player, where the enemies would appear, and how to move the spaceship and shoot bullets.	78
5.3	Best F_1 -Scores for each input combination. F_1 -Score was calculated over a confusion matrix derived from individual predictions of 194 folds of LOOCV. Dots on left indicate which information was considered in the model. Machine learning algorithms (SVM, RF, MLP, KNN) are ordered left-to-right in decreasing order of highest F_1 -Score. The darkest bar highlights the highest F_1 -Score for each algorithm. Dotted lines separate input combinations into: Survey only (original baseline), Survey + Nonverbal (new baseline), Survey + Nonverbal + one type of additional interaction context, and Survey + Nonverbal + both types of additional interaction context.	90
5.4	Feature importance by type of input feature.	92

5.5	Visualization of data for participant P0354Z. Each plot shows game data (top) and webcam data (bottom). Game data includes six game events: co-player destroys enemy on the left side of the game screen (CL) or right side (CR), co-player dies (CD), participant destroys enemy on the left side (PL) or right side (PL), and participant dies (PD). The orange highlight identifies when the co-player is on the participants' side (left side). Webcam data includes predicted action units (AU) by OpenFace: AU02 (outer brow raiser), AU05 (upper eyelid raiser), AU12 (lip corner puller), and AU15 (lip corner depressor). The x-axis corresponds to frame number as the game progresses from start to finish. All participants consented to having their images shared publicly.	94
6.1	Overview of REACT. In REACT-Nao, people played a collaborative video game with a Nao robot (a). In REACT-Shutter, participants interacted with a Shutter robot during a photography task (d). For both datasets, we captured images of participants throughout the interaction (b,e) and provide facial analyses of the images (c,f).	98
6.2	Mean of sum of AU values during robot visits in REACT-Nao. Error bars are standard error. Letters (A,B,C) denote statistical significance. If visits do not share a letter, there is a statistically significant difference between values.	103
6.3	Median of sum of AU values over the photography interaction. Error bars are standard error. Trend line is a linear regression model with a 95% confidence interval.	106

7.1	We study the problem of learning human preferences over human-robot collaborations. We propose to leverage human actions as implicit feedback that can help the robot learn preferences over the collaboration, alongside traditional explicit feedback. We evaluate this idea in a laboratory, where participants assemble pizzas with a robot and can give explicit feedback by pressing buttons.	110
7.2	Results for RQ1: L_2 error with respect to the ground truth preference weights (top) and conflict percentage (bottom) as the interaction progresses. Lines represent average L_2 error and conflict percentage, and shading represents std. error across 100 interactions. Columns show results considering different β values, where $\beta_H = \hat{\beta}_{imp} = \hat{\beta}_{exp} = \beta$ for $\beta \in \{0.1, 1, 10\}$. . .	123
7.3	Results for RQ1: L_2 error at the end of the interactions based on Feedback Modality and Rationality ($\beta = \beta_H = \hat{\beta}_{imp} = \hat{\beta}_{exp}$) . Error bars are std. error. Bars labeled with different letters (A-F) have significantly different error based on a Tukey HSD post-hoc test. See the text for more details. .	126
7.4	Results for RQ2: Average L_2 error as the robot learns with PIE. The two lines in the plots indicate whether the robot’s assumption about the rationality of the human’s action ($\hat{\beta}_{imp}$) match the oracle’s behavior (β_H): solid lines correspond to $\hat{\beta}_{imp} = \beta_H$; dashed lines correspond to an erroneous assumption $\hat{\beta}_{imp} \neq \beta_H$. In the top row, when $\hat{\beta}_{imp} \neq \beta_H$, $\hat{\beta}_{imp} = 10$. In the bottom row, when $\hat{\beta}_{imp} \neq \beta_H$, $\hat{\beta}_{imp} = 1$. Shaded areas (visible when zoomed in) are the std. error in 100 interactions.	127

7.5 L_2 error with respect to the ground truth preference weights (left) and conflict percentage (right) as the interaction progresses. Lines represent average L_2 error and conflict percentage, and shading represents the standard error across the 21 participants in the real-world evaluation. 131

List of Tables

3.1	Categorization of the usage of the term “social context” in the HRI literature.	22
4.1	Participant demographics by condition.	61
5.1	Description of Open Face 2.0 [21] attributes used and which summary statistics (standard deviation (stdev), mean, and/or maximum value) from the capture frames were included.	84
5.2	Best F_1 -Scores for each classifier for Survey and Survey+Nonverbal input combinations. F_1 -Score was calculated over a confusion matrix derived from individual predictions of 194 folds of LOOCV.	90

6.1 Comparison of related available datasets. “Interactive task” indicates whether the human is actively interacting with the robot. “Additional task(s)” indicates if the participant had additional tasks other than just providing feedback to the robot (e.g., playing game in REACT-Nao). “Evaluative feedback” refers to if the dataset includes explicit, evaluative feedback about the robot from the participant throughout the interaction (either live or through annotations). The “Context” columns describe what additional context is provided in the dataset: Environment (e.g., location of enemies in REACT-Nao); Human (e.g., whether human spaceship moved left or right in REACT-Nao); Robot / agent (e.g., actual text of robot utterances in REACT-Shutter). 99

Chapter 1

Introduction

Historically, the field of robotics envisioned robots doing tasks that are defined by the “3Ds”: dirty, dull, and dangerous [237]. In these situations, the robot operates as an autonomous tool, often in isolation, and success can be measured by objective functional metrics, such as speed or task completion. For example, a manufacturing robot operates in a safety cage and can learn to place a screw through millions of isolated repetitions.

However, as robotics transitions into everyday human environments, the “3Ds” are being replaced by tasks that are inherently social and collaborative. In these settings, success is subjective. A robot assisting a human in a kitchen must align its behavior with the personal preferences of its partner. Because these preferences are individual and often unknown prior to an interaction, robot behavior cannot be pre-programmed. Thus, robots must be able to *learn* to adapt their behavior *during collaborations* with partners who are *not necessarily experts* in robot learning.

In order to learn, robots need feedback. The American Psychological Association defines **feedback** as “information about a process or interaction provided to the governing system or agent and used to make adjustments that eliminate problems or otherwise optimize

functioning” [343]. Robots can get this information from both **explicit** and **implicit** human feedback. We define explicit feedback as signals that are directly communicated with the specific intent of providing information to guide the robot’s performance. Specifically, this dissertation will focus on feedback conveyed through direct channels, such as button presses for evaluative feedback. On the other hand, implicit human feedback encompasses a wide range of communicative signals inadvertently conveyed during interactions, without the explicit intent to convey information about the robot’s performance. One key difference between explicit and implicit human feedback lies in the fact that the meaning of explicit feedback is apparent, whereas the meaning of implicit feedback must be inferred.

Robots typically learn from humans via explicit feedback, such as evaluative feedback [190], preferences [38], demonstrations [112], or corrections [137]. However, relying solely on these types of explicitly provided feedback poses challenges in human-robot interactions. Most notably, explicit feedback from a human teacher is sparse. Expanding from humans solely as teachers to humans as active collaborators exacerbates this sparsity problem. Additionally, explicit feedback can take time and attention away from the person’s own tasks during a collaboration with a robot and interrupt the flow of an interaction.

These challenges with human feedback make it difficult to apply standard learning approaches to human-robot collaborations. Traditional robot learning has demonstrated that success can be achieved through **volume**: high-fidelity simulations and “robot farms” provide the millions of iterations required for reinforcement learning. However, these approaches fail when a robot is trying to learn *from a human* for two reasons. First, while we can simulate the physics of a pick-and-place task, we cannot yet simulate the nuance and unpredictability of human behavior. Recent work has explored utilizing large-language models (LLMs) to help model humans, to varying success [368]. However, it is unreasonable to expect that a simulation would be able to simulate any one particular

person’s preferences and behavior. Second, unlike a robot as part of a “robot farm”, a human collaborator cannot be expected to endure endless repetitions of a task for the learning process. A proposed solution is to deploy robots out across thousands of homes, but that is not going to help adapt to **individual proclivities**.

In other words, we cannot solve the data scarcity problem in human-robot interactions with volume alone. We need a way to extract more information from fewer interactions.

We argue that the key to this data problem lies in the **dynamic social context** of human-robot collaborations – a complex interplay of factors related to the agents, their environments, and the associations between them. This social context is viewed as something that makes learning in human-robot collaborations more difficult, but we argue that it is actually an opportunity. This dissertation argues that **the social context is not merely a backdrop for interaction, but an active mechanism that can be leveraged and shaped to improve robot learning**. By viewing the interaction through a formal taxonomy of social context, we can enable robots to actively influence, interpret, and integrate human feedback – turning real collaborations into data-rich learning opportunities. Figure 1.1 illustrates an example human-robot collaboration and outlines how we can use social context to extract feedback to facilitate better robot learning.

1.1 Outline

This dissertation begins with an overview of relevant background, highlighting the challenges with robot learning during collaborations with non-expert robot users. We then formally define the “social context of human-robot interactions” and show different ways to interpret and influence the social context to improve robot learning.




Figure 1.1: Illustrative human-robot collaboration. This dissertation explores how robots can use a novel definition of the “social context of a human-robot interaction” to extract additional feedback from collaborations for learning. We explore this through three different channels. (a) First, we explore solicited explicit signals via robot reminders for more feedback (Chapter 4). (b) Then, we analyze leaked implicit signals via human nonverbal communicative signals (Chapter 5). Third, we interpret the physical actions a human takes during collaborations as implicit signal and combine those signals with explicitly provided feedback via evaluative button presses (Chapter 7).

(CHAPTER 2) RELATED WORK. We begin with an overview of relevant background in human-robot interaction and the three key building blocks scoping this dissertation: robot learning, human-robot collaborations, and non-expert robot users. This chapter includes a discussion of challenges when learning during collaborations with non-expert users. We describe related approaches in this area and highlight the key limitations that we aim to address.

(CHAPTER 3) THE SOCIAL CONTEXT OF HUMAN-ROBOT INTERACTIONS. In this chapter, we begin by surveying the HRI literature for existing definitions and uses of the term “social context”, highlighting the opportunity for consistent language. Drawing from this survey, we propose a conceptual model for describing the social context of a human-robot interaction. We demonstrate the utility of this model by applying it to existing work,

and provide a taxonomy for the different attributes of the social context. We discuss a range of attributes of social contexts that can help researchers before, during, and after interactions. We conclude with a discussion of open research questions related to understanding and modeling the social contexts of human-robot interactions. This work is forthcoming in Volume 9 of the Annual Review of Control, Robotics, and Autonomous Systems [327]:

“The Social Context of Human–Robot Interactions,” S. Thompson, K. Candon, and M. Vázquez. In Annual Review of Control, Robotics, and Autonomous Systems, 9. 2025.

(CHAPTER 4) SOLICITING EXPLICIT SIGNALS. To demonstrate how robots can actively shape the social context to solicit more explicit feedback for learning, we present a user study exploring how a robot can remind a person to provide feedback during a collaboration. The social context includes attributes such as the robot’s utterances, the human’s actions, and the human’s perceptions about the robot. By manipulating the framing and timing of reminders, the robot is able to influence how the human provides explicit feedback, which could be used for learning. This contribution provides empirical evidence for our core thesis: by strategically acknowledging and manipulating the attributes that define an interaction’s social context, a robot can shape human behavior to influence the amount of feedback it receives during ongoing collaborations. This experiment was published in the 2023 ACM/IEEE International Conference on Human-Robot Interaction [57]:

“Verbally soliciting human feedback in continuous human-robot collaboration: Effects of the framing and timing of reminders,” K. Candon, H. Zhou, S. Gillet, and M. Vázquez. In Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2023.

(CHAPTER 5) INTERPRETING FACIAL EXPRESSIONS AS IMPLICIT SIGNALS. Another

way to leverage the social context to facilitate better robot learning during collaborations is to interpret additional information that participants “leak” as feedback about robot behavior. Using data collected while participants played an online game with a robot, we evaluated whether the inclusion of nonverbal human signals, as well as additional context (e.g., via game or personality information), led to improved prediction of user preferences between agent behaviors compared to explicitly provided survey responses. Our results suggest that nonverbal communication, a common type of human implicit feedback, can aid in understanding how people want computational agents to interact with them. These findings were published in the 2023 International Conference on Autonomous Agents and Multiagent Systems [56]:

“Nonverbal Human Signals Can Help Autonomous Agents Infer Human Preferences for Their Behavior,” K. Candon, J. Chen, Y. Kim, Z. Hsu, N. Tsoi, and M. Vázquez. In Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2023.

(CHAPTER 6) TWO DATASETS FOR ANALYZING IMPLICIT AND EXPLICIT SIGNALS TOGETHER. We have shown that by examining interactions through the lens of the social context, both implicit and explicit human signals can be extracted as feedback for learning during human-robot collaborations. To facilitate future research in this direction, we contribute the REACT database, a collection of two datasets of human-robot interactions that display users’ natural reactions and explicitly provided feedback to robots during a collaborative game and a photography scenario. We analyze the datasets to show that the interaction history is an important attribute of the social context that can influence human reactions to robots. The dataset and analyses were presented as a short contribution at the 2024 ACM/IEEE International Conference on Human-Robot Interaction [58]:

“REACT: Two Datasets for Analyzing Both Human Reactions and Evaluative Feedback to Robots Over Time,” K. Candon, N. Georgiou, H. Zhou, S. Richardson, Q. Zhang, B. Scassellati and M. Vázquez. In Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2024.

(CHAPTER 7) COMBINING IMPLICIT SIGNALS VIA PHYSICAL ACTIONS WITH EXPLICIT FEEDBACK. In this chapter, we present an approach for learning human preferences during a human-robot collaboration from implicit and explicit feedback together. We evaluate our approach in simulations and with real users in a cooking scenario. Our simulation results indicate that combining multiple modalities of human feedback improves a robot’s ability to estimate human preferences over the collaboration, with a similar trend observed in real-world evaluations. These findings highlight a promising direction for enabling robots to adapt to a user’s preference model more quickly, thereby reducing the amount of time a person must spend teaching a robot. This approach was published in the 2026 ACM/IEEE International Conference on Human-Robot Interaction [59]:

“Learning Human Preferences Over a Human-Robot Collaboration Based on Explicit and Implicit Human Feedback,” K. Candon, Q. Zhang, A. Lew, H. Claire, L. Qian, A. Quarles, C. Sarkar, and M. Vázquez. In Proceedings of the 2026 ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2026.

(CHAPTER 8) CONCLUSION. We conclude the dissertation by summarizing the main contributions and discussing the implications of these findings. We reflect on the limitations and identify directions for future research.

Chapter 2

Related Work

This chapter reviews related work on robot learning, especially as it relates to human-robot collaborations and non-expert robot users. Our work builds on these efforts.

2.1 Robot Learning

An overarching goal of Artificial Intelligence (AI) is to create autonomous, social agents capable of helping people in their daily lives [61, 194, 314, 110, 42]. While robots promise a future where they will help us with many physical and social tasks in human environments, current interactions are often scripted and rigid, lacking the adaptability typical of human social encounters. As robots transition from controlled industrial settings into social environments, such as homes, the tasks that they perform will become increasingly subjective and driven by personal preferences [248, 38]. Thus, it becomes infeasible to pre-program every variation of a household task. Rather, it is essential to make robots better at learning directly from human partners [5].

The Human-Robot Interaction (HRI) community has long acknowledged the importance

of creating robots that can adapt to individual preferences [320, 4]. Enabling robots to understand and align their behavior with human preferences can result in enhanced efficiency and safety [235] as well as higher user satisfaction [10]. This adaptation is often directly learned through human feedback [88]. Research has investigated robot adaptation in a variety of settings, such as while collaboratively building a toolbox [236], during tutoring sessions [80], or in prehabilitation exercises [360].

2.1.1 Explicit Human Feedback

One common approach to robot learning is to collect explicit signals about user preferences (e.g., [98, 287]). Typically, humans teach robots through a variety of explicit feedback modalities. Common explicit modalities include evaluative feedback [190], demonstrations [12, 294], corrections [221, 19], rankings [242], and comparisons [284, 75]. For these explicit feedback modalities, the implications of the feedback for robot learning are clear. Recent, unified frameworks for learning can extract information from multiple feedback modalities [174, 113], potentially resulting in better and faster preference learning.

2.1.2 Implicit Human Feedback

Implicit human feedback encompasses a wide range of spontaneous behavior and reactions humans provide during interactions [291]. These signals are implicit because they require interpretation as they are not necessarily intended for teaching the robot. Nonetheless, these signals convey information about the human's state, intentions, or preferences [111], and about how they perceive social encounters [357, 182]. It is generally agreed upon that effective social agents must be able to analyze, comprehend, and respond to nonverbal cues [77].

This line of research is motivated by the knowledge that people are often able to adapt their behavior to other people by interpreting social cues provided in human-human interactions. For example, expert teachers are able to recognize students' affective states so they can then adjust the pace and content of the learning material [208]. Elder adults infer which caregivers are available to provide assistance based on body orientation, head position, and gaze [363]. Individuals select acceptable partners for different social goals based on observable nonverbal behavior [292]. But in human-agent interaction scenarios, it is unclear whether humans would provide useful nonverbal cues in response to an agent's actions. People may behave differently when interacting with an autonomous agent than when interacting with another person [11, 224, 92, 118, 275, 93, 90]. For instance, people have been found to provide fewer nonverbal signals when interacting with a robot if there are no other humans present [126] and show fewer physiological signs of emotion in response to agents compared to humans [286, 344].

To address utilizing implicit human feedback, affective computing [115, 266] and social signal processing [347] have studied how to create computational models to interpret human nonverbal reactions. These fields have utilized a diverse set of implicit cues, including gaze [241], facial expressions [159, 89, 370, 64], and physiological signals (e.g., EEG, GSR) [167, 340].

While these methods are compelling, this work is typically focused on specific affective states or event detection. For instance, researchers have leveraged nonverbal human signals to adapt the behavior of in-home devices [364, 85], improve generative deep learning models [170], and better support virtual negotiations [348]. Gaze has also been studied as a replacement for a "wake word" for a smart speaker [232] or to provide clues about a person's intent [3], and facial and bodily expressions have been used to measure engagement in game tasks [130]. Guerdan et al. [136] use human responses to explanations provided by an agent,

focusing on specific affective states such as perceived challenge or competence. Leite et al. [207] proposed an empathetic social robot that behaved differently based on perceived user boredom, interest, or frustration. Also, nonverbal human signals have been used to detect agent errors [350, 305, 195]. Others have inferred user states or assessed robot performance based on interaction dynamics, timing, or hesitations [333, 371]. Relatedly, Learning from Observation (LfO) focuses on robots learning tasks by watching human actions [159], but typically aims for skill acquisition or goal inference rather than understanding preferences about the interaction itself.

Despite some success, interpreting these cues is still challenging. Different cultures or situations can result in similar nonverbal cues, so these cues may have different meanings depending on the context in which they are generated [168, 24, 56]. It can be difficult to attribute feedback to particular triggers; for example, when a person is engaged in solving a task with an agent, are human social cues due to the agent's behavior or due to the person's own actions in the task? Also, social cues may have different meanings in different situations. Smiles, for instance, are often considered an expression of enjoyment [117]. Nonetheless, people have been shown to smile in response to robot mistakes [126], for which we may naturally expect expressions of surprise, disappointment or disapproval instead.

2.2 Human-Robot Collaboration (HRC)

In interactive robot learning, the human commonly serves as a dedicated teacher, whose only role is to provide explicit feedback (e.g., [113, 293, 191]). However, this paradigm is impractical for real-world human-robot collaborations, where the human is typically engaged in the task itself and does not focus solely on teaching. Expecting continuous explicit

feedback is not only impractical, but can also lead to frustration and disengagement [323]. Relying solely on explicit feedback can be burdensome for the human, particularly within a collaboration where the person is also focused on task execution [48].

Collaboration involves working within a shared space towards common goals. Typically, HRC leverages the complementary strengths of both parties in a collaboration; for example, pairing human dexterity and judgment with robotic precision and repeatability [298, 163]. This synergy often results in specialized roles for the interactants, particularly in common applications such as manufacturing [229] and assembly [373]. In these settings, robots collaborate with humans to perform repetitive tasks such as placing and sealing screws [200] or assembling objects from a collection of parts [139, 307]. More recently, collaborations are expanding into more varied tasks, such as cooking together [129]. As robots enter less structured settings, such as hospitals [331], homes [193, 144], and hospitality environments [302, 216], the success of collaboration depends less on mechanical throughput and more on how well the robot aligns with their human partner’s personal preferences. For these collaborations to be effective, robots must be able to learn how humans want them to collaborate, in a way that is natural and intuitive for the humans.

2.3 Non-Expert Robot Users

In addition to difficulties from shifting to collaborative settings, there are inherent shortcomings in how humans naturally provide feedback [322, 220, 72]. Notably, humans tend to give less feedback as an interaction progresses [214]. Also, research suggests that users tend to stop providing feedback once they are satisfied with an agent’s performance [166]. As robots enter more collaborative interactions, humans will likely provide even less feedback if they are preoccupied with their own actions.

Research has investigated robots that request feedback from humans, including what kind of queries to ask [49, 36, 114], how often to query a user [50], or how to account for a human’s ability to provide useful information [37]. Further, Ho et al. [145] studied how to build mental models of humans to determine how to ask for feedback, and Jeon et al. [173] provided a framework to enable agents to combine multiple types of feedback.

2.4 Synthesis

Robot learning has established methods for adapting robot behaviors based on human input. However, this line of work largely assumes large volumes of clean, expert-provided feedback, where the expert is solely teaching the robot. Recent work has started to challenge these assumptions, but typically treats the human teacher as a passive oracle that can provide feedback, but not a dynamic social agent that can be influenced and can leak extra informative signals. Our work addresses the challenges in extending robot learning to collaborations with non-expert users by contributing a fresh lens with which to view the dynamic social context.

Chapter 3

The Social Context of Human-Robot Interactions

This chapter builds off of work that is forthcoming in Volume 9 of the Annual Review of Control, Robotics, and Autonomous Systems [327]:

“The Social Context of Human–Robot Interactions,” S. Thompson, K. Cannon, and M. Vázquez. In Annual Review of Control, Robotics, and Autonomous Systems, 9. 2025.

The Human-Robot Interaction (HRI) community often highlights the social context of an interaction as a key consideration when designing, implementing, and evaluating robot behavior. Unfortunately, researchers use the term “social context” in varied ways. This can lead to miscommunication, making it challenging to draw connections between related work on understanding and modeling the social contexts of human-robot interactions. To address this gap, in this chapter, we survey the HRI literature for existing definitions and

uses of the term “social context”. Then, we propose a conceptual model for describing the social context of a human-robot interaction. We apply this model to existing work, and we discuss a range of attributes of social contexts that can help researchers plan for interactions, develop behavior models for robots, and gain insights after interactions have taken place. We conclude the chapter with a discussion of open research questions in relation to understanding and modeling the social contexts of human-robot interactions.

The contents of this chapter are forthcoming in Volume 9 of the Annual Review of Control, Robotics, and Autonomous Systems [327].

3.1 Overview

The social context of human-robot interactions is key for the design, evaluation, and automatic generation of appropriate robot behavior [26, 320]. It is generally accepted that the social context shapes how humans interpret signals [347], expect others (including robots) to act [105], and behave themselves [17]. In the field of Human-Robot Interaction (HRI), prior work typically investigates some element of robot’s understanding of the social context (e.g., [22]), how robots should interact or generate suitable behaviors within a particular social context (e.g., [240, 187]), and/or human experience with a robot acting within a social context (e.g., [189]). Further, academic venues have long invited work that examines interactions in their social context [184], calling for work that is ecologically valid [149] and addresses complex situated encounters [40].

Unfortunately, the use of the term “social context” in HRI is often overloaded or underspecified. For example, the term “social context” may refer to specific circumstances in which an interaction takes place or a particular application domain (such as healthcare

or home robotics). To further complicate the matter, the term “social context” is often shortened to “context”, which is even more overloaded in practice. “Context” can refer to information passed to an algorithm to interpret sensor observations (e.g., map of environment [346]) or to generate robot behavior (e.g., interaction states provide context to generate locomotion [219]). Papers in the HRI literature commonly use the term “social context” without describing what it means. Readers experienced in HRI can perhaps make an informed guess; but for others, like new HRI researchers, the term could get lost in translation, becoming jargon rather than clarifying the key ideas. The lack of precision makes it difficult to identify, relate, and synthesize gaps in existing work about social contexts and their effects on human-robot interactions. What should we investigate next about the social context of a human-robot interaction? What information about the social context should a robot use for reasoning when deployed in a new scenario? Without having a clear operational definition of what constitutes the social context of a human-robot interaction, it is unlikely that we will be able to answer these types of questions effectively.

Luckily, other related fields have grappled with the problem of conceptualizing the context of human-human and human-computer interactions. For instance, Social Psychology has long studied the concept of social context when investigating individual human behavior and human-human interactions. Stangor [304] laid the foundation for the scientific study of how social factors influence individual and group behavior. Lewin [212] posited that behavior is a function of individual characteristics and the surrounding social environment, $Behavior = f(Person, Environment)$, emphasizing the dynamic interplay between people and their social environments. More recent work by Argyle et al. [14] has conceptualized social contexts for interactions in terms of social situations. Social situations are characterized by the presence of two or more people who interact together and are shaped by a variety of factors, including the roles of the participants, the norms

that govern behavior, people's goals or motives, and their forms of communication. As an example of the value of such conceptualization, social situations have recently been adapted to HRI, helping create more varied simulation environments for social robot navigation research [335].

In the field of Human-Computer Interaction, Dey [94] provided an operational definition of context for ubiquitous computing: "*Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves.*" This definition served to scope context: For an indoor tour guide application on a mobile device, is the local weather context? According to the definition in [94], information about the weather is not context because the application is being used indoors and the weather does not affect it. However, information about the presence of other people who take part in the tour with the user is context, because they could affect which sites the user visits while using the application. Based on this notion of context, Dey went on to characterize context-aware computing, helping application builders more easily determine what features their applications should support and what context is critical to support the features. There is similar, practical value in better conceptualizing social context in the Human-Robot Interaction field.

Riek and Robinson [274] introduced an initial conceptual model of social context for researchers "concerned with the automatic analysis of (and response to) human behavior". They defined social context as: "*the environment, E, where a person, P, is situated, with four factors that may influence P's behaviors. These factors include the situational context, P's current social role in E, the cultural conventions of both E and P, and the social norms of E.*" O'Connor and Riek [254] expand upon this definition to provide formalisms for this social context and tease apart the conceptual ideas in practice, and

Nigam and Riek apply this model to robotics [251]. We propose a model that builds upon these initial ideas, while focusing specifically on the social context of human-robot interactions.

In this survey, we first review existing definitions and the use of the term “social context” in HRI. The goal is to contrast different perspectives on how the community conceives of social context. Then, we propose an explicit definition of and a conceptual model for the *social context of a human-robot interaction*. Our conceptual model is inspired by the perspectives about context and social context described previously, but is designed specifically for HRI, considering the social nature of human-robot interactions, the importance of relationships within these interactions, and the unique characteristics of robots that differentiate them from other computing technology. To explain our conceptual model, we provide examples of how it can be applied to prior work in HRI and provide a taxonomy for socially contextual information that highlights the diversity of factors it can include.

Overall, with this review and our proposed conceptual model, we provide a pathway to think in an explicit and structured manner about social contexts in HRI. We end the chapter by discussing practical uses of our conceptual model and open questions.

3.2 Review Methodology

We performed a systematic literature review with two main goals in mind: 1) better understanding the uses of the terms “social context” in the HRI literature (Sec. 3.3); and 2) defining and validating our proposed conceptual model for the social context of a human-robot interaction (Sec. 3.4).

We reviewed publications in well-established academic proceedings for HRI work from the years 2012-2023: ACM/IEEE Int’l Conf. on Human-Robot Interaction (HRI), IEEE Int’l Conf. on Robotics and Automation (ICRA), IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems (IROS), IEEE Int’l Conf. on Robot and Human Interactive Communication (RO-MAN), and ACM Trans. on Human-Robot Interaction (THRI).

The selected proceedings comprised of 27,843 articles, as shown in Figure 3.1(a). We filtered the set of papers for work that discussed aspects of social context in human-robot interactions using string matching. We searched for text matching the regular expression: `social(?:ly)?[\ ']?context(?:ual)?`. The result was a corpus of 320 papers, a portion of which were short contribution papers because some conference proceedings combine short and full papers into the same volume. Figure 3.1(b) shows the number of papers in this corpus, organized by year. There is a clear increase in the use of the term “social context” and its close variants (per the regular expression), suggesting that the concept of social context is becoming increasingly important in HRI.

The 320 papers included many papers using the term “social context” without it being a particular focus of the paper; thus, we identified papers that had at least one match to the regular expression for “social context” in the title or had three or more matches in the body of the paper. This resulted in 58 papers, in which 257 sentences included the term “social context” in the papers’ main content, excluding abstracts and titles. In the next section, we analyze these 58 papers to understand how the robotics community uses the term “social context”. Later, we narrow down this collection to papers that focus on real-world interactions to explain our proposed conceptual model for the social context.

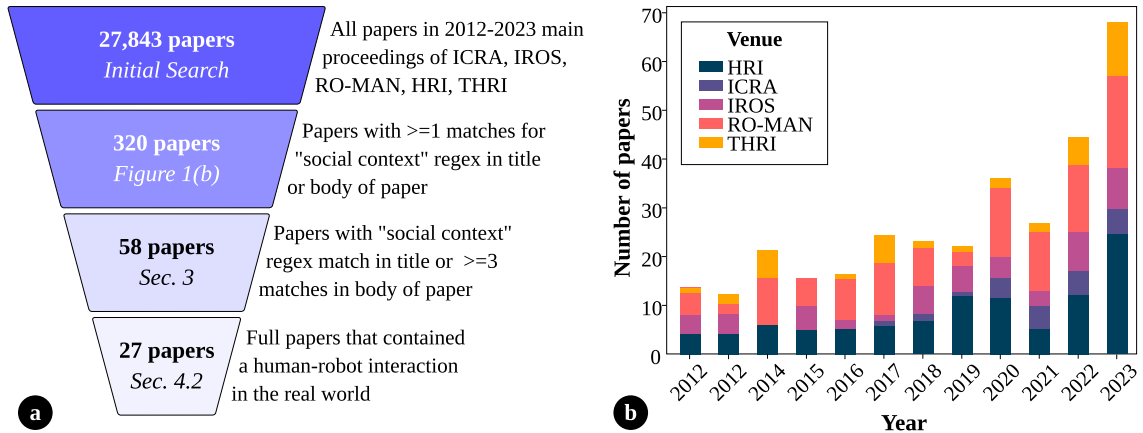


Figure 3.1: (a) Paper selection process for review. (b) Histogram of 320 papers using the term “social context” (or close variants like “socially contextual”) in the title or body of the paper. Best viewed in color.

3.3 The Term “Social Context” in HRI

This section first discusses different explicit definitions that have been provided for the term “social context” among the 58 papers considered in our review that had at least one match to the regular expression for “social context” in the title, or had three or more matches in the body of the paper. Then, it discusses broader usage of the term in these 58 papers.

3.3.1 Explicit definitions of the term “social context”.

In the group of 58 papers, only 4 papers explicitly defined the “social context”:

- Huang and Mutlu [155] stated that the “*social context or setting might characterize the physical environment (e.g., a domestic environment or a workplace), the organization of the interaction (e.g., dyadic interaction or group setting), the relative statuses of the participants (e.g., a supervisor or a subordinate), and the roles of the participants (e.g., a speaker or a bystander).*”

- Nigam and Riek [251] defined “*the social context for an agent (robot), P , in a given environment, E , as the disjoint union of several subsets: the situational context as a function of E , the social role of P in E , P ’s cultural norms (irrespective of E); E ’s cultural norms (irrespective of agents in P); and the social norms for P in E .*”
- Zaech et al. [367] stated that “*social context... comprises of the positions and velocities of the other agents.*”
- Lubber et al. [223] used “*the angle of approach $\alpha_{i,j}$ between the two [agents] π_i and π_j as the criterion to quantify and distinguish what [they] define here to be a social context.*”

Huang and Mutlu [155] and Nigam and Riek [251] suggested that the social dynamics between participants and the physical environment play a fundamental role in the social context of an interaction. For social robot navigation, Zaech et al. [367] and Lubber et al. [223] proposed that social context is comprised of the physical behavior of agents, including their relative physical state. Together, the definitions from this subset of the HRI literature suggest:

1. Aspects of an *environment* can be an important component of the social context of a human-robot interaction (e.g., home, work, cultural norms of a location, etc.).
2. Attributes of individual *agents* are an important part of the social context (e.g., their conversational role, their position, their velocity, etc.).
3. *Associations* between agents can also be part of the social context (e.g., the relative orientation between two agents, job hierarchy, etc.).

Table 3.1: Categorization of the usage of the term “social context” in the HRI literature.

Category	Stats [†]	Example	Citations
State of society	4p (6.9%) 7s (2.7%)	“We believe that one key step HRI researchers can take to center the <i>social context</i> is to include a societal implication consideration section in all papers” [259]	[259, 95, 359, 282]
Domain	3p (5.2%) 3s (1.2%)	“Robots were previously built to be used in <i>social contexts</i> with members of the public, including healthcare, education, and robots used at home” [7]	[7, 67, 263]
Task	6p (10.3%) 15s (5.8%)	“... goals are more important for a specific <i>social context</i> . For instance, if a robot were deployed in a service role that involved interacting with members of the public (e.g., museum tour guide, reception waiter, etc) ...” [46]	[337, 84, 101, 342, 41, 46]
Social Setting	35p (60.3%) 104s (40.5%)	“The parallels between being excluded by robots and being excluded by humans ... , suggests that [robot-robot-human interaction] experiences have the potential to form a powerful <i>social context</i> that impacts humans’ emotions and behavior” [103]	[244, 243, 315, 367, 337, 223, 47, 246, 342, 328, 67, 202, 22, 15, 103, 188, 95, 104, 218, 160, 175, 265, 76, 203, 329, 256, 74, 131, 7, 16, 234, 41, 362, 169, 345]
Physical Setting	4p (6.9%) 4s (1.6%)	“... our robot explored three types of social contexts on our college campus: study areas, dining areas, and lobby areas, across both the student center and library” [140]	[140, 259, 329, 362]
Social & Physical Setting (Explicit)	19p (32.8%) 30s (11.7%)	“Both the limited roles of participants and the confines of the experimental environment present quite a different <i>social context</i> from that in which robots are eventually meant to operate.” [202]	[244, 353, 251, 84, 328, 67, 202, 22, 15, 104, 218, 175, 265, 261, 181, 176, 282, 119, 68]
Social & Physical Setting (Implicit)	33p (56.9%) 82s (31.9%)	“The researchers are immersed within the <i>social context</i> they study, while being aware of the mutual influence researcher and participant have on each other and therefore keeping some distance to the people they study [22].” [250]	[244, 353, 315, 116, 53, 84, 258, 47, 101, 342, 67, 202, 22, 15, 188, 95, 259, 257, 143, 181, 263, 250, 359, 176, 106, 329, 256, 74, 119, 68, 362, 46, 169]
Other	7p (12.1%) 12s (4.7%)	“Discriminating and following others’ gaze direction is an essential component of establishing a common <i>social context</i> and it is pivotal to the ability to infer others’ mental states” [265] “Support for this hypothesis would indicate that the agency ascribed to the robot is of key importance in determining whether human interactants will heed its protests, whereas lack of support would indicate the importance of other factors, such as <i>social context</i> ” [47]	[47, 101, 175, 257, 265, 181, 250]

[†] The number and percentage of both papers (p) and sentences (s) using the term “social context”. There are 257 sentences from 58 papers.

3.3.2 Broader usage of the term “social context”

In the corpus of 58 papers, the term “social context” was used in a variety of ways without providing an explicit definition. Some usages of the term were broad, like using “social context” to describe a problem area or domain [67]. Other usages were more specific, such as using “social context” as a synonym for “social norm” [188]. Even within the same paper, authors sometimes used the term in different ways, referring to “social context” at different levels of specificity. Because usage varied so widely, it is difficult to suggest a single meaning or implicit definition underpinning the term in the current literature.

In order to better understand how people currently use the term “social context” in HRI, we systematically examined the uses of the term across the 58 papers. In an initial inspection, we found 6 common usages: *state of society*, *domain*, *task*, *social setting*, *physical setting*, *social and physical setting*. We then classified each of the 257 sentences that matched the regular expression for “social context” in the 58 papers into one of these 6 categories. Anything that did not clearly fit into one of these categories was classified as “Other”. Table 3.1 summarizes statistics and examples for each category, which we also describe below:

State of society. “Social context” refers to the broader state of society or of the world.

Domain. “Social context” is used to explicitly refer to a broad HRI domain, like “healthcare” or “entertainment”. The domain could include many different tasks or environments.

Task. “Social context” describes a specific application (e.g. robot tutoring), task (e.g. learning to read), or interaction scenario (e.g. robot asks a child to read something to it).

Social setting. The term “social context” describes the social setting, which might include beliefs, social norms, roles, expectations, or group membership.

Physical setting. The term “social context” describes the physical setting or environment of an interaction, without mentioning social aspects.

Social & physical setting. “Social context” references both the social and physical environment of an interaction. In total, 41 papers used the term “social context” in this manner across 112 sentences. If both the physical and social setting were referenced explicitly, the paper was further categorized as “social & physical setting (explicit)” (19 papers, 30 total sentences). If the references to the physical or social setting were implicit, then the paper was categorized as “social & physical setting (implicit)” (33 papers, 82 total sentences). Implicit cases include situations where the physical environment was mentioned in nearby sentences, but not the sentence that contained the term “social context”.

Other uses. The intended meaning of the term “social context” was unclear or did not fit one of the above categories.

The wide range of usages for the term “social context” motivated us to conceptualize a model, specific to human-robot interactions, that could serve to connect different perspectives in the literature.

3.4 A Conceptual Model of the Social Context of a Human-Robot Interaction

We propose a conceptual model for the *social context of a human-robot interaction*. Our goal in creating this model was twofold. First, we wanted to provide an explicit definition of “social context” specifically for human-robot interactions that connects the literature. We demonstrate how to apply our conceptual model to a variety of use-case scenarios and provide taxonomies for different types of attributes of social contexts drawing upon

the literature. Second, we wanted to provide a tool – the conceptual model – to facilitate planning for interactions, generating behavior, and analyzing interactions after they have occurred. We discuss these practical implications in Section 3.5 along with future research.

The proposed conceptual model is specifically designed and scoped to describe the social context of a human-robot interaction of interest, where the relevant human(s) and robot(s) act as embodied agents that perform actions in an environment, potentially influencing each other and changing the physical state of the world. For the purposes of the proposed model, a robot is embodied. Physical embodiment makes a robot inherently different from other computing technologies, as discussed in the book “Human-Robot Interaction: An Introduction” [26]. While Reeves and Nass’s Media Equation [271] suggests that people will act in a similar way with technology as with each other, this is not always true for robots [105].

3.4.1 Terminology

Because our definition of the social context of a human-robot interaction is dependent on an interaction of interest, we first define what we mean by a human-robot interaction:

Definition 1. *A **human-robot interaction** is an exchange between agents, which must include at least one human and at least one robot. At its core, the interaction corresponds to a sequence of actions taken by the agents in a given environment, which are related to the task of the interaction (or the goals that each agent aims to accomplish). The interaction has temporal bounds that define when it begins and ends.*

Human-robot interactions can be *explicit* (e.g., as in conversations, robot tutoring settings, etc.) or *implicit* [177], occurring without the explicit command or awareness of the human(s) involved in the interaction. For example, common implicit interactions in the

social robot navigation literature involve having a robot navigate alone nearby people (e.g., as in [231]). The people are not engaged in co-navigation or in an explicit information exchange with the robot, but they still adapt their actions to the robot as needed.

Human-robot interactions can be *one-on-one* interactions between a robot and a human only, or they can be *multi-party*, involving more agents [296, 124]. They can also change in size over time. The task or agent goals in an interaction can be the same for all agents or different. This can result in *collaborative* interactions [147], *mixed-motive* interactions [54], or *adversarial* interactions [99]. Also, interactions can have varied length, from brief interactions (like accidental encounters in a given physical space) to longer term (like a robot helping a person practice exercising). In general, we are not concerned with defining these aspects of interactions, but let HRI practitioners and researchers decide what human-robot interaction is of interest, including the relevant agents, task, and temporal bounds.

Given a human-robot interaction of interest, we define the social context of the interaction as follows:

Definition 2. *The social context of a human-robot interaction is the set of attributes of the relevant agents, of their environment(s), and of their associations that influence the interaction.*

Inspired by how Dey [94] defines “context” for context-aware computing and existing uses of the term “social context” in HRI (Section 3.3), we consider the *attributes* in the social context of a human-robot interaction to be information that characterizes the relevant agents, environment(s) and associations between them. The main requirement for these attributes to be part of the social context is that they influence the interaction either *directly*, by having an effect on the sequence of actions of the interactants, or *indirectly*, by having



an effect in the interacting agents which, subsequently, influences their actions.

The agents in Definition 2 can be of type *robot*, *person*, or *other*. The robot(s) and human(s) that take part in the human-robot interaction of interest have attributes that are part of the social context of their interaction. For instance, this may include the embodiment of a robot, or a person's age, attention, etc. Other agents could include, for example, pets, whose attributes may be particularly relevant in home or assistive robotics applications.

The inclusion of “associations” in Definition 2 is based on the importance that previous work places on relationships between agents when defining “social context” for an agent in robotics [251, 254], as well as emphasis on these associations on context-aware computing [374]. In general, associations between agents and environments in Definition 2 can be of three types: agent-agent, agent-environment, and environment-environment associations. Because there is important relational information that can affect human-robot interactions, we make these associations first-class entities in our conceptual model for the social context. That is, we give associations the same level of importance as agent and environment entities. Similar to the latter entities, associations can have more than one attribute that is part of the social context of an interaction. For example, an association between a human and a robot could have information about roles (e.g., whether a person serves as a teacher for a robot) and human impressions about the robot (e.g., whether the person thinks the robot is acting competently). Also, an association between a person and an environment could contain information about the medium through which the person experiences that environment (e.g., in person or via a computer interface with a given set of attributes).

Across human-robot interactions, we find it natural to think about the relevant entities of the interaction via graph visualizations. Nodes in the graph can be of two types (agents or environments) and edges can be used to encode associations between the nodes. If a

Emergency Example



Office Example




Figure 3.2: Examples of relevant entities for two different human-robot interactions: one between a firefighter and a robot, and another between a robot guide and a person in an office. Figures (a) and (d) illustrate the interactions at a given time; the images were created with Gemini’s Generative AI. The diagram in (b) shows the relevant entities for (a), including the robot (R_1) and the firefighter (H_1). The firefighter teleoperates the robot that is in E_1 , while being physically located in E_2 . Dashed edges have relevant agent-environment associations, while solid edges indicate relevant agent-agent associations. The diagram in (c) shows the interaction some time after (b), when the firefighter and robot team find another person (H_2). This illustrates how social context can be dynamic – the relevant elements of the interaction can vary over time, so the attributes that influence human-robot interactions are also variable. In (d), a robot (R_2) brings a person (H_3) to an office, where there are two other people (H_4 and H_5) who are not considered part of the interaction of interest. Hence, they are visualized in (e) with a lighter color-scheme. However, attributes of H_4 and H_5 , like whether they are attending to the R_2 or H_3 , influence the interaction of interest. As a result, the attributes of H_4 and H_5 are part of the social context of the interaction of interest.

node or an edge appears in the graph, it is because it has relevant attributes that are part of the social context. For example, imagine an emergency scenario in which a firefighter (H_1) teleoperates a robot (R_1) that goes into a disaster area with fire, as in Figure 3.2a and Figure 3.2b. The firefighter teleoperates the robot through a computer interface and from a remote location to stay safe. In this scenario, the social context of the interaction between H_1 and R_1 includes attributes from two environments. First, the environment of the robot (E_1), which the firefighter accesses via a teleoperation interface, influences the commands that H_1 sends to the robot. Second, the environment in which the human is located (E_2) affects the human-robot interaction as well by influencing the human, e.g., to what extent it pays attention to the robot due to noise or other environmental distractions.

The attributes that are part of the social context of an interaction can change over time, and the relevant entities may also change. For instance, in the prior example of a firefighter-robot interaction, it could happen that the robot comes close to a victim (H_2), who then becomes part of the interaction as in Figure 3.2c. The attributes that characterize this person and that influence the interaction are now part of the social context.

There can also be additional agents that may not be part of the human-robot interaction of interest, but that have attributes that are part of the interaction's social context. For example, imagine that the interaction of interest is between a robot R_2 that guides a person H_3 to a room in an office environment, where other people work (as in Figures 3.2d and 3.2e). Then, whether these other nearby people attend to the interaction could influence the actions of R_2 and H_3 , e.g., making them speak in a lower volume. Thus, the attention of the other people is an attribute of the social context of the interaction of interest.

To more explicitly distinguish the possible attributes of all environments and all agents from those that are part of the social context of a given human-robot interaction, we define:

Definition 3. *Socially contextual information are the attributes in the social context of a human-robot interaction.*

In the remainder of this dissertation, we will often describe the social context of a human-robot interaction as a collection, or set, of socially contextual information.

3.4.2 Selection criteria for analysis of proposed definition of social context

To analyze the current literature through the lens of our proposed definition of social context, we created a subset of the papers selected for the literature review (per Section 3.2) which focused on aspects of the social context of an interaction. In particular, for each paper in the corpus of 58 papers identified with our matching criteria for the regular expression for “social context”, we read the abstract and checked the paper to determine if it was a full paper about a two-way interaction between at least one robot and at least one human. We did not include videos or extended abstracts in this subset of the papers because these shorter papers are typically about preliminary results. Also, we excluded papers that only discussed online human-robot interaction studies, where participants did not experience an interaction with a robot in the real world. This filtering process led to 27 papers.

3.4.3 Example use-case scenarios

This section illustrates how our proposed conceptual model for the social context of a human-robot interaction can be instantiated in specific scenarios. We categorized the 27 papers described in Section 3.4.2 based on how the paper’s main contribution related to our

concept of social context. The categories we identified were:¹

- **Study:** 10 papers described a study in which socially contextual information constituted independent and/or dependent variables.
- **Computational Models and Systems:** 9 papers described a computational model or system that used attributes of the social context of an interaction as inputs or estimated socially contextual information.
- **Design:** 5 papers discussed the design of human-robot interactions. These discussions included both influences of the social-context of a human-robot interaction and how socially contextual information may be influenced.
- **Survey:** 4 papers in our corpus were survey papers related to ideas captured by our definition of the social context of a human-robot interaction.

For the first three types of contributions described above, we discuss how our model fits the social context discussed in an illustrative paper. While not exhaustive, these use cases demonstrate the flexibility of our conceptual model and its applicability to existing work.

3.4.3.1 Use case 1: Effects of robot appearance

In “Actions Speak Louder Than Looks: Does Robot Appearance Affect Human Reactions to Robot Protest and Distress?”, Briggs et al. [47] investigated whether a robot’s appearance influenced how people responded to the robot verbally protesting a command. In this study, one robot built towers of colored cans. This robot was then removed, and participants were asked to instruct a second robot to knock down the can towers. The demolition robot protested the participant’s request.

¹The numbers per category sum to 28 as one paper fell into both the Study and into the Computational Model and Systems groups.

In applying our conceptual model to this study, we identified a variety of socially contextual information relevant to the interaction. First, the demolition robot's appearance was the independent variable in the paper, which they found to affect the human's perceptions about the robot's obligation to follow their commands. Thus, under our conceptual model, robot appearance is socially contextual information. Second, we consider the participant's commands to the demolition robot socially contextual information. Because the robot executes the participant's commands, the commands directly influence the interaction. Third, the state of the can tower is socially contextual information attributed to the environment. This is because the cans being stacked into a tower impacts the participant's understanding of the task instructions to knock them down. Conversely, the color of the cans is likely not socially contextual information in this case (per Definition 2) because there is no reason to believe that the cans' color affects the interaction.

Under our conceptual model, all positive results in HRI study papers, i.e., all confirmed factors that influence the interaction of interest directly or indirectly, are socially contextual information. These factors are the independent variables of the study, e.g., characteristics of agents (such as action, role, appearance, physical state, or mental state), of environments (like location) or associations between agents and/or environments (like relationships). Negative results indicate that no significant evidence exists that an attribute is socially contextual information.

3.4.3.2 Use case 2: Selecting listening behaviors

In "A Bayesian Theory of Mind Approach to Nonverbal Communication", Lee et al. [203] introduce a computational model for robot listening behaviors to indicate attentiveness. The robot's behavior is based on the storyteller's actions and a prediction of the storyteller's belief about the robot's attention.

A common assumption in computational models is that the inputs have an underlying causal relationship with the outputs. In this case, there are two models whose outputs are the storyteller's beliefs and the robot's action. For the models' inputs where the assumption holds, i.e. the causal relationship exists, the inputs are socially contextual information.²

Besides the behavior selection model, the robot's gaze and when it demonstrates a listening behavior is determined by a rule-based model that uses the storyteller's gaze, goals, and attributes of their speech including pitch, energy, pauses, and length. Because this rule-based model is encoding a causal relationship between the storyteller's behavior and the robot's actions, we consider these input attributes socially contextual information.

In addition to the models determining the robot's behavior, Lee et al. discuss specific attributes of the robot that, based on prior research, they think could influence the interaction, and carefully control them to prevent unwanted effects. Specifically, these attributes were robot's color, gaze, facial expressions, and utterances. These attributes are likely socially contextual information given prior HRI results.

3.4.3.3 Use case 3: Designing interactions for an airport

In "Design Methodology for the UX of HRI: A Field Study of a Commercial Social Robot at an Airport", Tonkin et al. [329] provide a methodology for designing human-robot interactions in public environments that create a positive user experience. Tonkin et al. outline the steps for designing such an interaction and mention several factors about environments, humans, and robots that they believe could impact how an interaction would unfold. They note that for different deployment locations (e.g., airports, hospitals, or train stations) the volume level can influence whether users can hear any of the robot's

²In practice, validating causal effects for a particular interaction can be difficult, and computational models can be negatively affected by spurious correlations in the data.

speech. For humans, they mention that internal state and role (e.g., visitor or staff) can give important insight to their needs, which also may be dependent on the environment. For the robot they note that its morphology, personality, task, voice, identity, gestures, and screen display can impact how useful and positive the users' experience is. Under our conceptual model of social context, we would consider each of these attributes likely candidates for socially contextual information.

3.4.4 Types of socially contextual information

This section presents taxonomies for socially contextual information. We built the taxonomies using papers described in Section 3.4.2. For each of the 27 papers, we identified the key attributes of the human(s), robot(s), environment(s), and their associations that were explicitly discussed in the paper. Then, we used an iterative and collaborative process to develop the taxonomies using affinity diagrams. We expected the 27 papers to provide good coverage for the range of socially contextual information typically considered in HRI; but, when appropriate, we expand with other examples to better convey the richness of the social context of human-robot interactions.

3.4.4.1 Environment attributes

As shown in Figure 3.3, we identified four main categories of socially contextual information describing the environment of a human-robot interaction:

Location. The location was the most common type of socially contextual information for environments. Papers described locations at varying levels of granularity. For example, some described the area in which interactions occur, such as public spaces [329]. Other

work focused on broader locations by noting the building in which an interaction takes place, such as a nursing home [67] or grocery store [250]. Even more specific were references to particular rooms, such as the lobby or activity area within an eldercare facility [68].

Objects. Humans and robots often engage in physical interactions that involve the manipulation of objects, making attributes of such task objects an important piece of information that can influence the interactions. An example is the location of cans discussed in Section 3.4.3.1. Similarly, humans and robots consider furniture [68] or other obstructions [175, 67, 362, 46] when deciding how to navigate within an interaction.

Behavior constraints. Environments of human-robot interactions can have implications for the behavior of agents, which are socially contextual information that constrains their actions. Two noted examples from our literature review include safety constraints [15] – which we view as hard constraints on the behavior of robots – and location-specific social norms [250] – which we consider soft constraints.

Physical properties. We consider the attributes of an environment that can be measured as their physical properties. When these attributes influence an interaction, they are socially contextual information. For example, papers referenced the layout of an environment as an attribute that can influence behavior. This layout was encoded via environmental maps [362] or referenced through elements such as hallway locations [175]. Other physical properties include room size and condition (e.g., if it needs cleaning) [202], the time of day [251], the brightness of the environment [319], or volume of an environment [329]. Though not directly referenced in the papers in our literature review, this category of environment attributes could include other properties such as temperature or humidity.

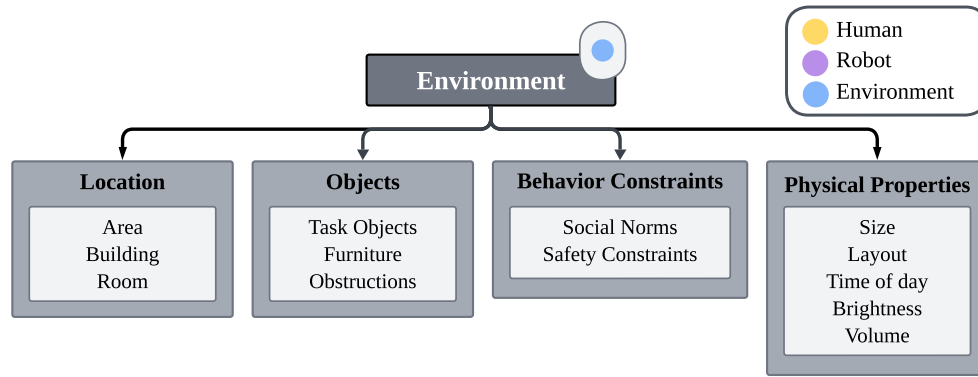


Figure 3.3: Taxonomy for socially contextual information of an environment. The blue symbol ● is used to convey that these attributes could be relevant for the blue environment nodes in Figure 3.2.

3.4.4.2 Agent attributes

The papers that we reviewed discussed an agent’s actions, attributes that are common population characteristics, physical state, appearance, and mental states. While some specific attributes were discussed more commonly for robots or for people in the selected papers, these five categories of socially contextual information are generally applicable to both. Figure 3.4 shows the derived taxonomy for agent attributes.

Actions. Action-related attributes can be categorized as either *communicative* or focused on *strategy*. Communicative attributes can be associated to *verbal* and *nonverbal* actions. The content of utterances [342, 265, 203, 256] and tone of voice [328, 329, 319] comprised the most common verbal attributes. Lee et al. also note how a human’s length of utterances or pauses in speech can influence the interaction [203]. Attributes of nonverbal actions commonly include information about gaze [67, 345, 265, 203], facial expressions [203], and gestures [103, 265, 160, 15, 329, 282]. Attributes related to an agent’s strategy include an agent’s goal [22, 362, 46], their task [15, 329, 328, 46, 250, 53], their planned trajectories [22], the consistency of their behavior [328], or their recovery strategy [328].

Common Population Characteristics. A variety of characteristics often used to describe populations – also known as demographics – were noted often for humans in the reviewed papers, although they could be applied to robots as well. For example, the papers that we reviewed noted the age [67, 202, 265, 250, 263, 203, 106], race [68], gender [202, 104, 265, 67, 263, 203, 106, 282, 68], diagnoses (e.g., autism [160], dementia [68], mental health conditions [202]), income [202], education level [202], and employment [202] of humans as factors that could influence interactions. While the selected papers did not discuss demographic characteristics of robots explicitly, it is common for people to assign gender and other similar attributes to robots as they anthropomorphize them. These attributes could influence robots' behavior (e.g., affecting how they communicate with people) or humans' behaviors (e.g., affecting mental models of the robots).

Physical state. While details of the general location of a human-robot interaction can be captured in socially contextual information for the environment, the physical state of an agent is often a piece of socially contextual information. This state often includes an agent's location, which can be represented in (x, y) -coordinates [175], but can have other abstractions. For instance, the physical state of a person could include their articulated body pose [175, 53] or orientation [345, 203]. Physical state could also include motion information [175].

Appearance. An agent's appearance can affect how other agents perceive them. The appearance of robots was often noted in the selected papers. For example, the general design [67, 263, 329], color [203], and morphology [328, 202, 103, 104, 250, 329, 282] of robots can affect interactions. Their screen display [329, 7, 250] or interface [160] can also influence how people interact with them. While the papers that we reviewed did not explicitly discuss the appearance of humans, this attribute can influence human-robot interactions, e.g., in terms of whether a robot can (re-)identify a person visually.

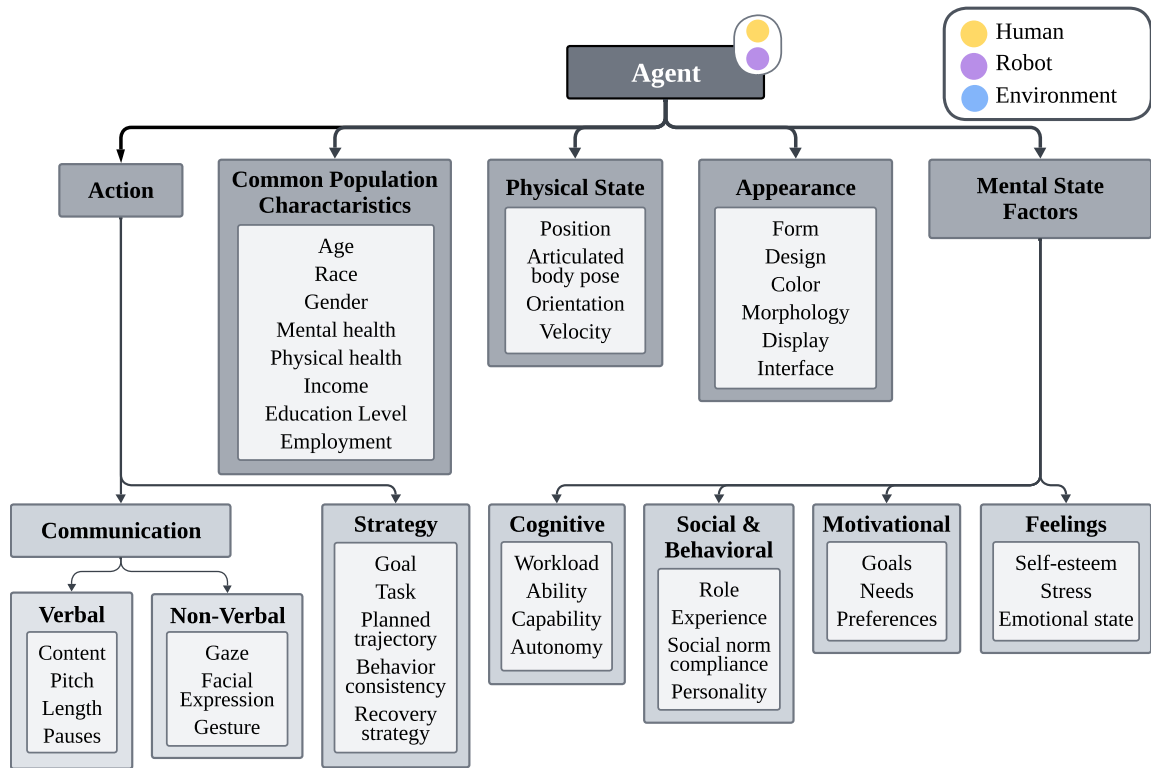


Figure 3.4: Taxonomy for socially contextual information of an agent. The yellow ● and purple ● symbols convey that these attributes could be relevant to the human and robot nodes in Figure 3.2.

Mental states. Information about the internal, intellectual activity of agents is important socially contextual information because their internal states drive behavior.³ We identified four subgroups of attributes in the papers that we reviewed. First, mental states can be *cognitive factors* such as cognitive workload [328, 282], abilities [250, 15], capabilities [15], or level of autonomy [202]. Second, mental states can be *social and behavioral* attributes. A common factor is the role of an agent in an interaction. For example, whether a human is a resident, staff member, or visitor at an eldercare facility may influence how a robot chooses to interact with them [329]. Similarly, whether a robot is an active or passive participant in a task (e.g., builder versus observer) can influence how a human might interact with said robot [47]. Additionally, an agent’s experiences (such as with technology

³We considered mental states to be attributed to an agent when they were about the agent itself, not about other entities. Later, we discuss mental states about others (e.g., beliefs) as associations.

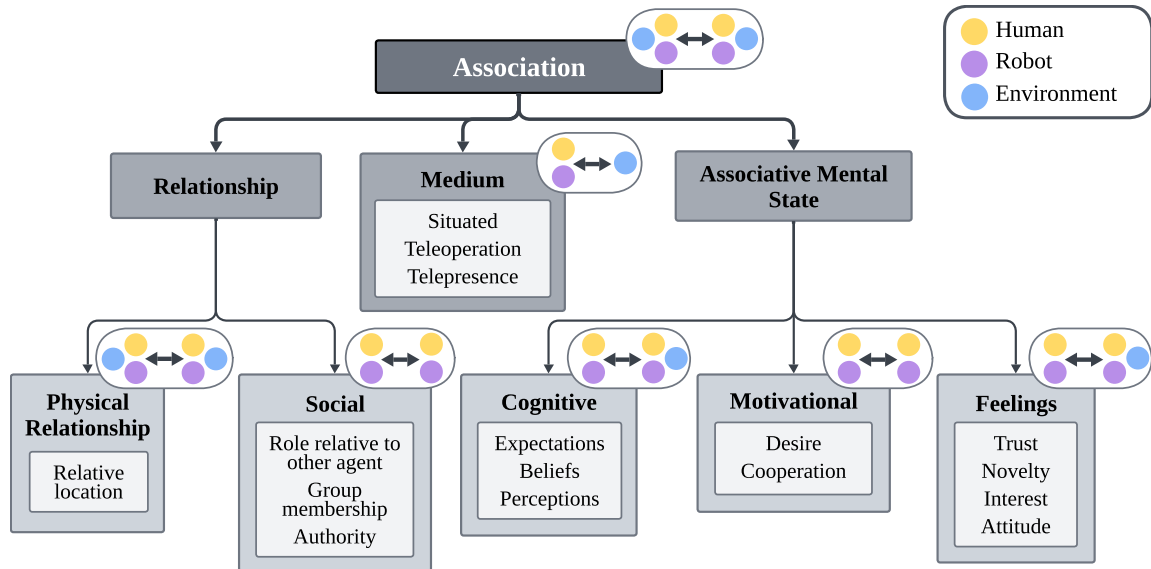


Figure 3.5: Taxonomy for socially contextual information of an association. Yellow ● and purple symbols ● convey agent-to-agent associations. Blue symbols ● convey associations that involve environments. These colors are consistent with Figure 3.2 and the prior taxonomies in Figure 3.3 and Figure 3.4.

[202, 282] or with robots in general [68]), an agent’s compliance with social norms [250], or other personality traits [106, 104] can also influence their actions in an interaction. Third, agents have *motivational* states that can influence their actions, which include attributes related to goals [46, 7, 362, 22], needs [67, 103], and preferences [282]. Finally, mental states can be about an agent’s *feelings*, including self-esteem [103], stress [282], or other emotions [250, 329].

3.4.4.3 Association attributes

As shown in Figure 3.5, attributes of associations generally fall into one of three categories:

Relationships. Attributes about relationship associations can be of two types: they can encode information about *physical relationships* between any entities or information about *social relationships* between agents. For example, Chang et al. described interactions in

different rooms (e.g., lobby and activity room) in the same building, so the two environments had a physical relationship [68]. Papers also described how the distance between agents or between an agent and an object can influence the agent’s actions [362, 345]. Attributes about social relationships include whether an agent has a human or robot partner [256], the agent’s group membership [362, 15], or the authority agents have over one another [15].

Medium. For agent-environment associations, the way in which the agent experiences the environment can influence their behavior and, hence, have an effect on human-robot interactions. For example, it is common for an agent to experience an environment by being situated in it, being co-located with other agents (as in Figure 3.2d). Additionally, agents could experience an environment through computer interfaces, such as a remote teleoperation interface for a robot (as in Figure 3.2a), a video call, virtual reality, etc.

Associative mental states. There are a variety of internal, mental states that involve more than one entity and can influence interactions. In the papers that we reviewed, we identified associative mental states that are *cognitive* factors, such as expectations [328, 15, 250], beliefs [46], or perceptions [47] about other agents or environments. Also, we identified *motivational* mental states, e.g., whether a human wants to interrupt a robot’s task [250] or touch a robot [15], or thoughts about whether two agents are cooperating versus competing [256]. Associative mental states can also be about *feelings*, such as trust in another agent [328, 15], novelty [250], interest [250], or attitudes [103, 104].

3.5 Discussion

Given the varied usage of the term “social context” in HRI, we proposed a conceptual model for the social context of human-robot interactions that can bridge different perspectives. Now, we discuss ways in which we foresee HRI practitioners and researchers leveraging this

model in the future and highlight open challenges and interesting future research directions.

3.5.1 Planning for human-robot interactions

The taxonomies provided in Section 3.4.4 can serve as an initial checklist for thinking about different kinds of socially contextual information that could influence a human-robot interaction of interest. First, outside of robotics, research has shown that interaction designers can have blind spots to novel conditions that come up during the deployment phase of a technology [262], requiring contingency planning in the interaction design process. Likewise, this could happen in HRI, where the situations that come up during a human-robot interaction can be novel and hard to predict [238, 152]. Second, there can be unanticipated human behavior around the adoption of new technology. For instance, while a growing body of research suggests that robots can help support educational efforts [30], research has also indicated that teachers can have negative attitudes towards education robots [272]. Attitudes towards robots can be socially contextual information, as described by the taxonomy from Figure 3.5, making it important to plan ahead for them (e.g., by working with stakeholders to facilitate the introduction of robotics technology). Overall, by thinking about a variety of potentially-relevant socially contextual information ahead of an interaction based on the proposed conceptual model, HRI practitioners and researchers can prepare for novel situations and challenges that may come up in practice during human-robot interactions.

A difficulty in planning for human-robot interactions is dealing with unknown socially contextual information. Prior work has explored understanding the social contexts of certain environments [317] and developing systems that consider the social context during planning [318, 311]. Our taxonomies of socially contextual information are not exhaustive as there

can be more attributes that matter for a given interaction of interest than those reported in the papers that we reviewed. Further, some attributes in our taxonomies may not be relevant to all interactions, and people may change over time, inducing changes in the factors that influence their interactions with robots. For these reasons, we advocate for iterative interaction design processes (e.g., as in [205, 151, 230]), where interactions are repeatedly prototyped, tested, and refined. Additionally, because user testing can be slow and expensive, it is important for the HRI community to continue innovating in design and evaluation methodologies which can accelerate understanding of social contexts. These methodologies may include the use of simulations [148, 335], virtual and augmented reality technologies [313, 351], front-end human-robot interfaces [330], crowdsourcing [206, 334], etc. While these methodologies may not fully replicate real-world results [217, 336, 107], they can accelerate the identification of socially contextual information and, hence, help develop better interaction paradigms and robust robotics technology.

3.5.2 Robot behavior generation during interactions

The proposed conceptual model for the social context of human-robot interactions can aid in developing autonomous robot behavior. Previous work has shown that adapting a robot's behavior to different contexts can improve user experience [267, 240]. Intuitively, imagine that a robot had a computational model of the social context of a human-robot interaction – that it understood what attributes of the relevant agents, environment(s) and of their associations influenced the actions of the agents of interest. If the robot could predict the outcome of these influencing effects, it could then generate suitable behavior by searching over its action space for the best actions that help it achieve a desired outcome. This behavior generation setup is generally intractable but approximations have found value in HRI, e.g., via receding horizon planning or optimization [249, 121, 273].

Computationally representing social contexts and, further, learning the dynamics of interactions in a way that captures all relevant socially contextual information is a difficult challenge. First, a particular interaction may contain extensive amounts of socially contextual information. Foundational work on computationally modeling the social context of an agent in robotics utilized symbolic representations [251, 254]. While effective at describing varied contexts, such representations can be potentially prohibitive from a space requirement due to the explicit nature of the symbolic abstractions. Given recent advances in representation learning with neural networks, there is opportunity for implicit representations of social contexts to be more scalable. But HRI data is scarce, making the use of inductive learning biases likely necessary for effective generalization [233]. In particular, we hypothesize that utilizing relational abstractions or graphs (as in Figure 3.2) and machine learning models designed specifically to reason about these abstractions, like Graph Neural Networks [28], could be beneficial for HRI given the importance of associations between agents and environments in the social context of an interaction. Indeed, recent work in HRI has utilized graphs as abstractions to encode attributes of the environment and team dynamics [8], for learning cost functions for motion planning [227], and as state abstractions for learned social robot behavior policies [125].

Second, it is not clear what the best level of specificity is for computationally abstracting socially contextual information. For instance, consider a person’s utterance. It could be computationally abstracted as a high-level intent, text, or a sound wave. Which abstraction is more useful in practice depends on what the robot is trying to achieve during an interaction. For example, high-level intent could be useful for coordinating the robot’s behavior with the user [156]. The sound wave could aid in synchronizing the robot’s speech with the user to build rapport [253]. Thus, we suspect that general and efficient computational abstractions for social contexts will ultimately need to be hierarchical. In alignment with

this hypothesis, the spatial reasoning community in robotics has advocated for hierarchical, metric-semantic environment maps to enable complex physical robot behavior [157].

Third, important socially contextual information is not directly observable by robots, such as internal mental states. While machine learning techniques are fueling a variety of approaches for inferring internal human states, e.g., from affective states [9, 303, 78] to perceptions of robot behavior [89, 308, 371], it remains difficult to measure these internal states in a scalable manner. This poses challenges for building datasets on which to train models and evaluating prediction performance in practice.

Given these challenges and recent advancements in generative Artificial Intelligence, it may seem natural to resort to large neural network models, like Vision-Language Models (VLMs) to generate socially-contextual robot behavior. For robot manipulation and navigation, VLMs built on large-scale Internet data are serving as effective backbones for improved scene understanding, and bridging the gap between high-level instructions and low-level control (e.g., see [349, 301]). In HRI, large models are becoming increasingly popular to create models of humans [368], generate more varied robot speech [300, 185], and implement a variety of functionality in the control system of a robot [358]. However, there is limited data capturing subjective human feedback that can be used in HRI to learn end-to-end robot policies with large models. Additionally, though large models seem to be ever-improving, their reasoning remains opaque, making it difficult to understand why, or even when, these models make mistakes. This motivates incremental learning approaches focused on continued improvement of robot autonomy [77] as well as utilizing a variety of human feedback [71, 70, 113] to adapt or steer robot behavior policies.

3.5.3 Post-interaction analyses

The proposed conceptual model for the social context of human-robot interactions can help to understand interactions after they have taken place. For example, one could analyze the appropriateness of robot behavior based on the relevant social norms that apply to them, which are part of the social context. This idea is in line with prior work in social robot navigation, which has categorized social situations to identify types of interactions where robot performance needs improvement [335], investigated how organizational factors affect the way people respond to delivery robots [239], or classified the environment to adapt robot behavior to different social norms [22]. Furthermore, our taxonomies for socially contextual information can help HRI researchers think about potential confounds that could lead to incorrect conclusions in experimental HRI work because, by definition, the attributes of the social context influence interactions. Finally, our conceptual model could also aid in defining the concept of “generalization” in HRI. Research in Human-Robot Interaction has called for building a generalizable theory of HRI from “in the wild” social encounters [178]: *“a principled understanding of what to expect with different types of robots, performing different types of tasks, in different types of social situations and cultures.”* Our conceptual model for the social context of human-robot interactions can serve to establish a broadly-applicable notion of generalization for such theories, where effective generalization entails predicting accurate outcomes in novel social contexts. These novel contexts are characterized by novel values for known socially contextual information, as well as by completely novel environment, agent, and association attributes that influence an interaction of interest.

It would be transformative if robots could reason about causal effects in social encounters based on their observations of interactions and new data that they collect. Recent work

demonstrates the feasibility of learning causal relationships within the HRI domain given a known set of relevant features [102, 62, 63], but more work is needed in this direction to capture a wider range of socially contextual information. Ultimately, causal competency, including understanding the social context of human-robot interactions, may be necessary for autonomous, social robots to behave ethically in novel situations [142].

Chapter 4

Soliciting Explicit Signals

A publication on this experiment appeared in the 2023 ACM/IEEE International Conference on Human-Robot Interaction [57]:

“Verbally soliciting human feedback in continuous human-robot collaboration: Effects of the framing and timing of reminders,” K. Candon, H. Zhou, S. Gillet, and M. Vázquez. In Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2023.

Now that we have a framework for social context, we explore how robots can actively shape this social context to solicit more explicit feedback for learning.

Specifically, in situations where humans are not providing enough feedback during an interaction, we propose that 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., [108, 150, 123, 197, 289, 82]). However, it then becomes essential that



Figure 4.1: Experimental setup for our study.

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. 4.1. The novelty of this interaction setup was twofold. First, the game was a continuous task, more similar to autonomous driving [283] than more typical turn-based interactions in collaborative robotics [97, 252, 309]. Second, the human interactant had additional objectives other than solely providing feedback to the robot (as is common in robot learning [162, 52, 13, 49, 196]). 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 work provides interesting insights on how to design robots that learn from humans in realistic, continuous collaboration scenarios.

4.1 Related Work

As discussed in Chapter 2, an active area of research concerns how to equip robots to adapt to individual preferences [320, 4], which are often learned through human feedback [88]. However, much of this research investigates turn-based tasks and scenarios where users are only focused on providing feedback (e.g., [49, 309]). 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 [135] or asking too many questions [356]. Thus, robots have the difficult task of ensuring requests or reminders for input are frequent enough to be useful, but not too incessant [34]. One approach is to identify opportune times for interruptions [1], such as by modeling user attention [183]. 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 [65, 180, 321, 109], we are interested in

understanding if robots can influence how much feedback humans provide. Close to our work, Rogers and Howard [278] 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 [192]. 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 [108, 150, 289, 269, 361] and in groups [297, 197, 82, 123, 104]). 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 [256]. Additionally, how a robot attributed blame amongst a group influenced how much humans trusted the robot [132, 179, 341]. Close to our work, Salomons et al. [285] 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 [79]. For instance, the time when a robot helps a human can impact the human’s perception of the robot [23]. The timing of robot actions can also affect the fluency of human-robot interactions [51, 165]. 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 [321, 186], we investigated the above question in relation to an important change in robot behavior during interactions.

4.2 Interaction Task: Space Invaders

Typically, when a robot learns from a human, the human's only objective is to teach the robot (e.g., [324, 288, 190]). Also, tasks are usually turn-based, where the robot takes an action and then waits for the human to provide feedback [352]. 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 [199] and unexpected help from a virtual agent [55].

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. 4.1). 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

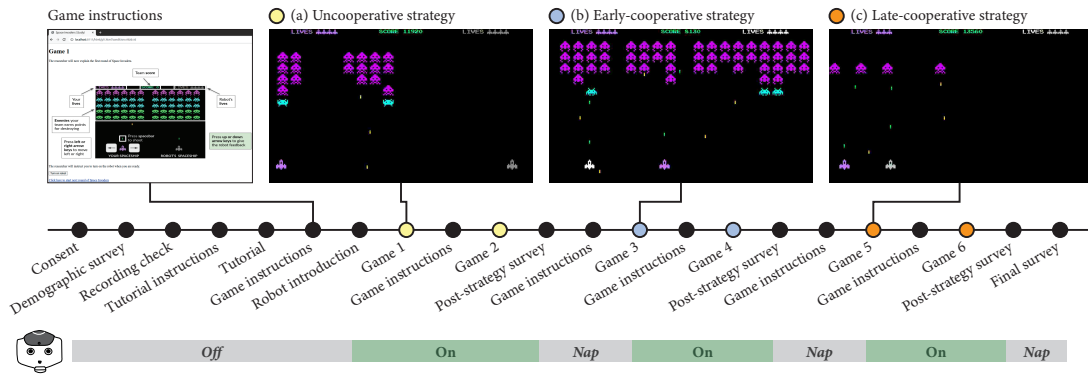


Figure 4.2: Experiment timeline. The robot reminded participants to provide feedback about the robot’s performance in games 3 and 4. The first image shows the game instructions that were shown to the participant before each round of Space Invaders (best viewed in digital form). The other images show the three gameplay strategies for the robot’s spaceship, as described in Sec. 4.2. 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.

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. 4.2(a).

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. 4.2(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. 4.2(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) [268] to provide game information to the robot. The supplementary material provides more implementation details.

4.3 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. 4.2. 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. 4.2). 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.

4.3.1 Study Design

To investigate the research questions outlined in Sec. 4.1, 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. 4.2, 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. 4.3 illustrate the difference between the Before and After reminders for the Team framing. See our supplementary video for examples of experimental conditions.

4.3.2 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 [209] and perceived responsibility for helping [210]. 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 [83].

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 [153, 366, 316]. Also, humans provide feedback both in response to past actions and to guide future behavior [321, 186].

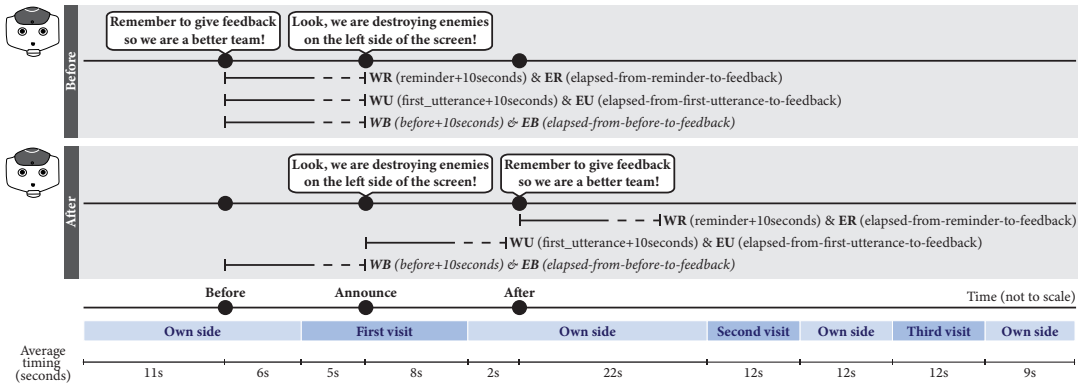


Figure 4.3: Robot behavior and related measures in the early-cooperative strategy (games 3 & 4) for the two timings of feedback (Before & After). Measures include ten-second windows (W) and elapsed time to next feedback press by participant (E) measures from game events. Italicized labels signify static 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.

4.3.3 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. 4.1.

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.

4.3.4 Procedure

Fig. 4.2 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) [128] and the Berkeley Expressivity Questionnaire (BEQ) [134].

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.

4.3.5 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. 4.3.

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. 4.3 shows the elapsed time measures (E) we analyzed and how they differ between Timing conditions.

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.

4.3.6 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 4.1 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

Table 4.1: Participant demographics by condition.

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

once a month.

4.3.7 Manipulation Checks

4.3.7.1 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.

4.3.7.2 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. 4.4 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.

4.4 Results

This section presents our results based on the measures described in Sec. 4.3.5. Unless otherwise noted, we used linear mixed models estimated with Restricted Maximum Likelihood (REML) analyses [325] via JMP Pro [164] 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. 4.3.7.2 as fixed effects. We conducted post-hoc Tukey Honestly Significant Difference (HSD) tests or post-hoc Student's t-tests as appropriate.

4.4.1 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. 4.3).

4.4.1.1 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




Figure 4.4: Rate of feedback in three windows described in Fig. 4.3 by Timing. Unit is feedback signals per minute (fspm).

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.

4.4.1.2 Specific windows

As discussed in Sec. 4.3.1, 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. 4.4. 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 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.

4.4.2 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.

4.4.3 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).

4.4.4 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.

4.4.5 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

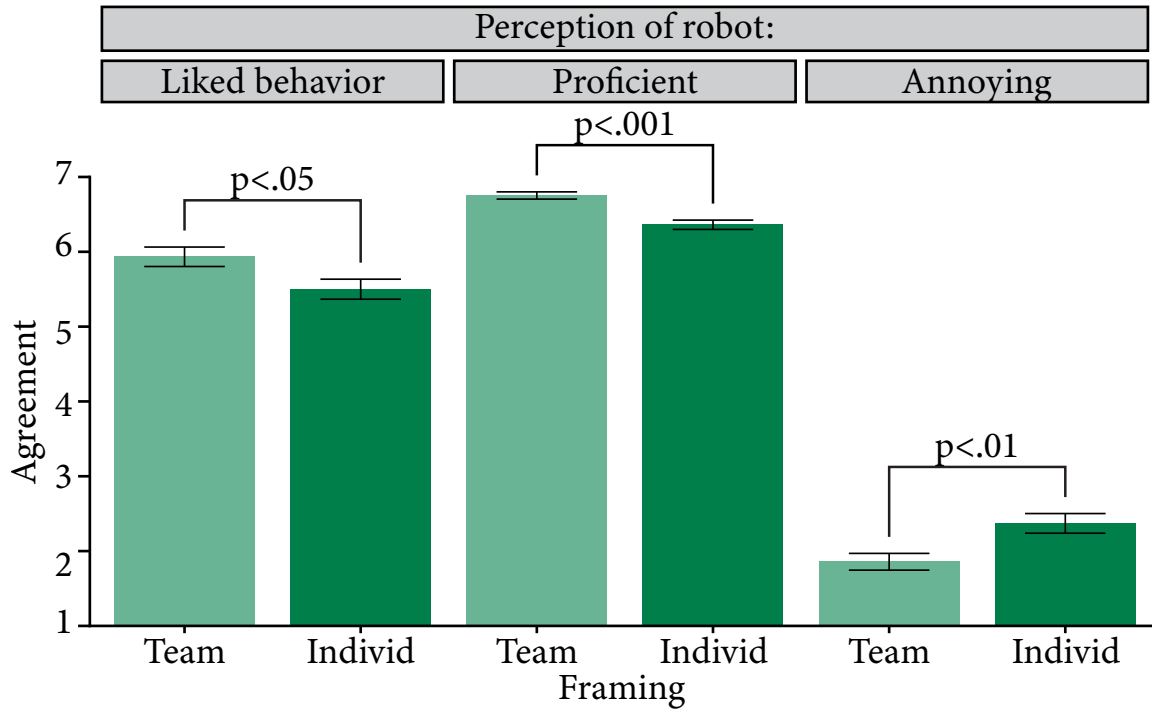


Figure 4.5: Participant agreement with “I liked the behavior of the robot in the game” (Liked), “The robot was proficient at the game” (Proficient) and “The robot was annoying” (Annoying) on a 7-point responding format.

strategy, robot Gameplay Strategy was included as a main effect. However, given that it is not the focus of this chapter, we do not include results for differences by robot Gameplay Strategy.

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.

4.5 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. 4.4.4). 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. 4.4.5). 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 [83, 285].

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. 4.4.1, 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. 4.4.1). 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 [214]).

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.

4.6 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 [281, 355], build physical objects [139, 6], or deliver parts in assembly lines [338]. 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.

4.7 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 con-

tinuous 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.

Chapter 5

Interpreting Facial Expressions as Implicit Signals

This chapter draws on a publication from the 2023 International Conference on Autonomous Agents and Multiagent Systems [56]:

“Nonverbal Human Signals Can Help Autonomous Agents Infer Human Preferences for Their Behavior,” K. Candon, J. Chen, Y. Kim, Z. Hsu, N. Tsoi, and M. Vázquez. In Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2023.

While the previous chapter demonstrated that soliciting extra explicit feedback is feasible, we should explore other avenues. As described in Chapter 2, when humans collaborate with one another, some learning happens implicitly. People observe others’ actions and infer preferences, without formal teaching. We propose robots should similarly learn from implicit signals alongside explicit feedback. Using this “leaked” feedback is another way

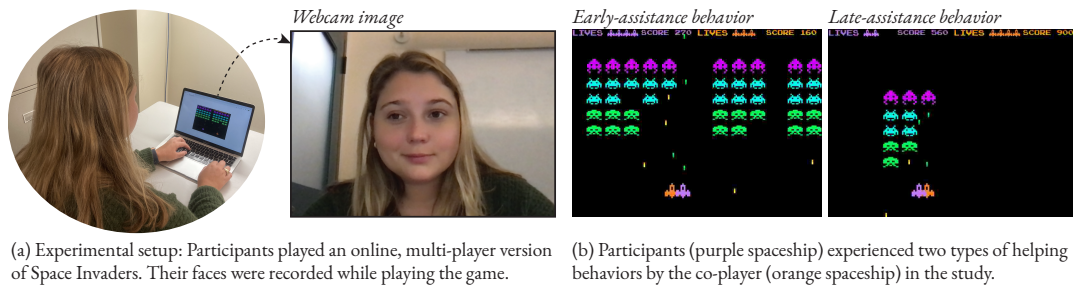


Figure 5.1: We collected data from an online human-agent interaction to investigate the usefulness of including human nonverbal signals into models for predicting which agent behavior participants preferred. In our study, the participants were recorded while playing the game with an autonomous agent via their webcams.

for robots to leverage the dynamic social context of human-robot interactions.

Despite the challenges when learning from implicit human feedback outlined in Chapter 2, we hypothesize that nonverbal human signals can provide clues about how people want computational agents to interact with them. To investigate this idea, we conducted an online study in which participants interacted with an autonomous agent in a fast-paced collaborative task. As shown in Figure 5.1, the participants played two games of a multi-player video game with different agent behaviors. We did not prime participants for cooperation (nor competition) in the game because we wanted to see how they would naturally react to an agent that tried to help them in the game when this help was not necessarily expected a priori. Also, we recorded participants via their webcams during interactions to later analyze their nonverbal signals.

We present results from multiple analyses to understand how different sources of information impact the prediction of user preferences for agent behaviors in the multi-player game scenario. First, we study the possibility of predicting preferences based on impressions of the interaction reported via surveys. These types of subjective impressions are commonly gathered in user studies [213, 201, 26] and survey data can help understand key factors that influence preferences. Then, we investigate whether the in-

clusion of nonverbal human signals improves preference predictions, per our hypothesis. Finally, we investigate whether additional context information, like information about the state of the game, further helps infer preferences over agent behaviors. To the best of our knowledge, this is the first study to investigate the usefulness of nonverbal human signals in predicting user preferences in a fast-paced and collaborative human-agent interaction scenario. The anonymized data that we used for our analyses can be found in https://github.com/yale-img/collabHAI_pref to facilitate future replication efforts and more complex preference modeling.

Overall, our findings provide insights for creating real-time adaptable autonomous agent behavior that leverages spontaneous nonverbal human reactions in the future. This could potentially help reduce the need to query users often in order to understand which agent behaviors work well in comparison to others.

5.1 Related Work

Unlike the previous corpus of work on learning from implicit human feedback described in Chapter 2, we are not focused on specific emotion or event recognition. Rather, our work investigates if nonverbal reactions provide insight into a user’s preferences for helping behaviors by an autonomous agent.

Our motivation for studying nonverbal human reactions during human-agent interactions stems from recent work on adapting agent behavior based on implicit feedback with reinforcement learning [87, 215]. Different to these prior efforts, though, we do not bias humans’ internal goals or rewards in our study. Nor do we ask our participants to be expressive in a particular manner during human-agent interactions [364, 215] because this

can result in fatigue. Our goal is to instead understand if their natural reactions while playing a collaborative game with an agent are useful for predicting user preferences for agent behaviors.

Research has shown that agents can learn to adjust their behavior based on explicit user preferences [280, 32, 133, 173, 96]. Querying users for preferences involves interaction costs, so previous work has investigated how to ask for preferences in a way that minimizes annoyance [135] and reduces the number of queries needed for adaptation [35, 354]. To complement this line of work, we investigate:

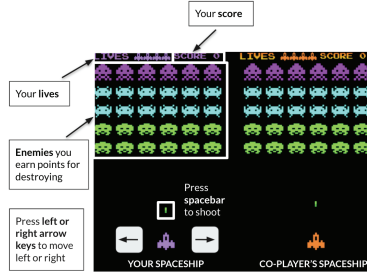
RQ1: *Can natural, nonverbal human reactions be leveraged to better predict user preferences for an agent's helping behaviors compared to only using explicit survey responses about game and agent attributes?*

Because recent work in psychology suggests that context is key when interpreting nonverbal human behavior [168, 24], we also ask:

RQ2: *Does the inclusion of additional context, beyond human nonverbal reactions, help predict human preferences over agent behaviors?*


Following Schilit and Theimer [290], we broadly consider context to be any element providing information about the environment in which an agent is situated. Nonverbal human reactions can be considered as context for an agent in our multi-player game. Furthermore, the background and personality traits of users as well as task (or activity) statistics are additional contextual factors relevant to the human-agent interactions studied in this work.

You are the **purple ship**, and you start with 4 lives. Use the **left and right arrow keys** to move, and press the **spacebar** to shoot. Your shooting speed is limited, so pressing the spacebar may not always result in a bullet being shot— that's part of the game. **You get points for enemies destroyed on the left half of the screen.**



(a) Introduction to the participant player

There will be an **orange co-player** in the game, automatically controlled by **artificial intelligence**. The co-player gets points for enemies destroyed on the right half of the screen.



(b) Introduction to the co-player (with AI identity)

Figure 5.2: Participants were introduced to the game as shown above. The explanation included explicit pointers to the scores and number of remaining lives for each player, where the enemies would appear, and how to move the spaceship and shoot bullets.

5.2 Method

In order to investigate the research questions outlined in the prior Section, we collected data through an exploratory online interaction. For the interaction, the participants completed a web survey through which they interacted with an agent in a two-player version of the Space Invaders game (Figure 5.1). The participant controlled one spaceship and the other spaceship, referred to as the *co-player*, was algorithmically controlled using simple heuristic behaviors. With their consent, we recorded participants during the games of Space Invaders via their webcams. The online study that we conducted for this data collection was approved by our local Institutional Review Board. We used the collected data to build models that predicted which co-player behavior was preferred by the participants, which is the focus of this chapter. For additional motivations, details, and results of the user study, please refer to the work by Candon et al. [55].

5.2.1 Data Collection

5.2.1.1 Participants

We recruited 360 participants for the study through Prolific [260]. The recruitment criteria required participants to be 18 years of age or older, be fluent in English, reside in the United States, and have normal or corrected-to-normal vision.

Out of the 360 participants, 194 participants were included in the final dataset, and 166 participants were excluded. Participants were excluded if they asked to withdraw, if their game logs indicated poor Internet connection, or if there was insufficient webcam data captured during the game. More details about our participant exclusion criteria are in Section A of the Supplementary Material.

Of the 194 final participants, 78% identified as female, 20% identified as male, 1% identified as nonbinary, and 1% preferred not to say. The participants' ages ranged from 18 to 66 years old, with an average age of 26.87 years ($SD = 9.30$). Participants indicated using computers often: 93% of participants used a computer daily, 6% used a computer 4-6 times a week, and 1% used a computer 2-3 times a week. A little more than half of the participants (53%) played video games once a week. When asked specifically about Space Invaders, 45% of the participants reported that they had played the game before, 46% reported they had not played it, and 8% were unsure.

5.2.1.2 Space Invaders Game

The Space Invaders game was introduced to participants as shown in Figure 5.2. The participant controlled a purple spaceship and the participant's co-player was an orange

spaceship. The participant's spaceship started on the left side of the screen, and the co-player's spaceship started on the right side of the screen, but both could move left and right within the full bounds of the game screen. Both spaceships could shoot upwards to destroy enemies, and each was assigned points individually for enemies destroyed on the side of the screen on which they originally started, regardless of whose bullet destroyed the enemy. Both the participant and co-player started the game with four lives. A player lost a life when hit by an enemy or a bullet from the enemies. Enemies moved left and right across the screen and slowly downwards, closer to the spaceships, until they were hit by a bullet or reached the bottom of the game screen. A game over screen appeared showing the final scores once all enemies were destroyed, both players lost all their lives, or an enemy reached the bottom of the game screen.

5.2.1.3 Co-player Behaviors

Because we were interested in predicting preferences between agent behaviors, we created two different co-player behaviors: early-assistance and late-assistance. In our game, the co-player could provide assistance to the participant by travelling to the participant's side of the game screen. Once on the participant's side, the co-player could help destroy enemies for which the co-player received points. We designed the two behaviors so they differed by the timing of when the co-player travelled to the participant's side of the screen to provide assistance. In the *early-assistance* behavior, the co-player went over to the participant's side of the screen on two occasions during the game while there were still enemies on the co-player's side of the screen (Figure 5.1b, left). In the *late-assistance* behavior, the co-player only went to the participant's side of the screen after all of the enemies on the co-player's side were already destroyed (Figure 5.1b, right).

5.2.1.4 Procedure

For the interaction, participants completed an online Qualtrics survey, which included two games of Space Invaders. The participant first consented to participate in the study, to have their video recorded, and to have their images and video clips shared publicly. The survey then included a video check to ensure that their webcam was working and that their face could be detected in the webcam images. The participants then completed a demographic section of the survey, as discussed in Section 5.2.1.1. This section also included personality data via the Revised Competitiveness Index [154] and the Ten Item Personality Measure (TIPI) [128]. The survey then introduced the Space Invaders game with a combination of text explanations and visual instructions, as illustrated in Figure 5.2. We purposefully did not prime participants for cooperation nor competition with the co-player. The participants experienced two games of Space Invaders, each followed by a post-game survey about their perceptions of the game and of the co-player. Each game involved a different co-player behavior, as described in Section 5.2.1.3. Participants were not informed of the order in which they experienced the two behaviors, and the order of behaviors was counterbalanced between participants. After playing two games of Space Invaders and answering both sets of post-game questions, participants were asked a final set of questions about the differences between the games and their preferences. Participants were presented with an optional Berkeley Expressivity Questionnaire (BEQ) [134]. The study took about 18 minutes to complete. Individuals were paid \$3.60 for participating.

5.2.2 Preference Prediction Task

In this work, we study the problem of predicting which co-player behavior a participant reported that they preferred at the end of their session. We consider this problem a multi-

class classification problem because the final section of our survey asked the participants to state whether they preferred the first co-player behavior, the second one, or did not prefer one over the other. Based on which order the participant experienced the two behaviors, we encoded the targets as either *Early*, *Late*, or *No Preference*.

5.2.2.1 Inputs to Preference Classifiers

We considered different combinations of four types of input features for preference classifiers. These feature types corresponded to 1) features derived from post-game survey responses, 2) nonverbal reaction data, 3) participant demographic data, and 4) game context. These four types of data were selected for different reasons. First, survey data is commonly used for understanding human preferences over agent behaviors (e.g., [339, 86, 100]). Second, nonverbal reaction data corresponded to implicit feedback that we hoped could be leveraged to better understand human preferences (per *RQ1*). Third, we suspected individual human factors would affect preferences, so we included demographic data (per *RQ2*). Finally, because our interaction domain was dynamic, it was important to consider game data to understand what the user and agent were doing and the state of the environment (per *RQ2*). While it is rare to consider all four data types in conjunction, we believe that there is value in systematically studying how they can all be used to predict user preferences.

For our analysis, we considered a variety of preference classifier models that differed in terms of the features they received as input. First, we considered models that took as input survey features only (Survey). These models were regarded as baselines because survey responses are commonly used to understand human perceptions of an interaction (e.g., [213, 201, 26]). Second, we added nonverbal reaction data (Survey+Nonverbal). Comparing Survey and Survey+Nonverbal allowed us to experimentally investigate whether nonverbal data helped predict user preferences (*RQ1*). Third, we considered models that

also took as input demographics and game data, to incorporate additional context. This led to three more sets of inputs: Survey+Nonverbal+Demo, Survey+Nonverbal+Game, and Survey+Nonverbal+Demo/Game. These feature combinations allowed us to experimentally investigate whether additional context further helped predict user preferences (*RQ2*). The next sections provide more detail about each specific feature type.

Post-Game Survey Responses: Our first set of input data was encoded from survey questions participants were asked after playing each game of Space Invaders. The questions were about the game experience, the perception of the co-player, and whether or not the participant thought they had helped the co-player. Game experience questions included agreement with statements from Large et al. [199] about whether the game was enjoyable, difficult, boring, and fun. Co-player perception questions included agreement with statements from Large et al. [199] about whether the co-player was helpful, proficient, intelligent, and annoying. Co-player perception questions also included if the participant liked the behavior of the co-player or if they thought anything about the behavior of the co-player seemed unusual. In addition, the participants evaluated the level of competence, warmth, and discomfort of the co-player using the 18 attributes from the Robotic Social Attributes Scale (RoSAS) [60]. More details about survey questions are included in Section B.1 of the Supplementary Material.

We considered both a Full set of survey features and a Selected set of survey features, which we experimentally found to be most important for preference prediction. For the Full set of survey features, we processed the raw data (e.g., via scaling and one-hot encoding) for 14 survey questions to arrive at 18 features encoding survey information for each game (as detailed in Section B.1 of the Supplementary Material). We additionally explored reducing the number of included survey features using the notion of Gini importance [44] from a trained Random Forest model, as implemented by the scikit-learn Python Library [264].

In particular, we selected the six post-game survey features with importance greater than 0.05: liked behavior, competence, annoyingness, helpfulness, warmth, and discomfort. Additionally, both sets of survey features included a feature encoding the order in which the participant experienced the two games, as well as a one-hot encoding for which of three co-player identities the participant experienced. We provided the survey data to a classifier in two ways: passing features for both games by concatenating them (Concat), and passing the difference between the value from the game with the early-assistance behavior and the game with the late-assistance behavior (Diff).

Nonverbal Reactions: We analyzed human facial and body reactions captured while the participants played Space Invaders. Our version of the Space Invaders game captured images of a participant via the participant’s own webcam at a framerate of 15 frames per second. We analyzed the images automatically using OpenFace 2.0 [21], a open-source toolkit for automatic behavior analysis. For each image, OpenFace 2.0 [21] extracted information about head pose, eye gaze, facial landmarks, and facial action units.

We explored different ways to incorporate nonverbal reaction data. Table 5.1 describes the features and the summary statistics that we used to arrive at 59 nonverbal reaction

Table 5.1: Description of Open Face 2.0 [21] attributes used and which summary statistics (standard deviation (stdev), mean, and/or maximum value) from the capture frames were included.

Attribute	Description	Summary Statistics
<i>Gaze</i>	Averaged gaze direction for both eyes in left-right and up-down directions	stdev
<i>Translational Pose</i>	Location of head with respect to camera	stdev
<i>Rotational Pose</i>	Pitch, yaw, roll for head	stdev
<i>Facial Action Unit Intensity</i>	Intensities (from 0 to 5) of 17 action units	stdev, mean, maximum

features for each game. For this set of features, we considered four approaches to provide them to a classifier: Full, Visit-Split, Visit-Diff, and Visit-Hadamard. In the Full case, we concatenated the summary features for both games and input them into the model. In the other cases, we computed summary features for when the co-player visited the participant’s side of the screen (“co-player visit”), and when it was on its side (“no co-player visit”).¹ In the Visit-Split case, both subsets of summary statistics for both games were concatenated and input to the model. In the Visit-Diff case, we computed the difference between the “co-player visit” and “no co-player visit” statistics within each game, concatenated the results, and then input them into to the model. Lastly, for Visit-Hadamard, we computed the element-wise product between the “co-player visit” and “no co-player visit” statistics within each game, concatenated the results, and passed the output to the classifier. Because the Nonverbal feature sets were much larger than the other feature sets, we considered a reduced set of nonverbal features derived by applying Principal Component Analysis (PCA) with 5, 10, 20, and 40 principal components in addition to the full set of original features.

Demographics: We preprocessed raw self-reported demographics information from the survey to arrive at 20 features describing the personality, age and gender of each participant as well as how often this person played video games, used a computer, and whether they had played Space Invaders before (as detailed in Section B.2 of the Supplementary Material). Similar to the survey data, we considered a Selected set of five demographic features, again with Gini importance from a trained RF model. The Selected set of demographic features included: how often the participant played video games, competitiveness index [154], negative expressivity from BEQ [134], and the extroversion and emotional stability dimensions of TIPI [128].

¹We grouped data based on the location of the co-player because we suspected that participants would react differently when it was on their side of the game screen. Also, in order to capture facial reactions as the co-player was traveling to or back from the participant’s side, we included 100 frames before and after a co-player visit when splitting the facial features into the “co-player visit” and “no co-player visit” sets.

Game Context: To enable classifier models to reason about contextual factors related to what occurred during Space Invaders, we included information from game logs. These game logs contained information about the state of the game, participant actions, and co-player actions for each rendered frame of Space Invaders. In particular, we extracted and analyzed: number of times the co-player and the participant moved to opposite sides of the game screen, the total number of frames in which the co-player and the participant stayed on opposite sides of the game screen, the number of participant enemies destroyed by the co-player, the number of co-player enemies destroyed by the participant, and the participant and co-player’s final scores and number of lives remaining. We provided the game data to a classifier in two ways: passing features for both games by concatenating them (Concat), and passing the difference between features for the two games (Diff).

5.2.2.2 Models for Predicting Participant Preferences

We considered various popular Machine Learning (ML) algorithms from the scikit-learn Python library [264]: Support Vector Machines (SVM), Random Forests (RF), K-Nearest Neighbors (KNN), and Multi-Layer Perceptron (MLP). We selected these simple, well-established classifiers because more complex models are likely to overfit on our small dataset and also tend to be more opaque. For each classifier, we performed a gridsearch over a range of suitable hyperparameters with survey data as inputs. We then used those hyperparameters for all models trained on different permutations of input feature sets. Details of the gridsearch procedure are presented in Section C of the Supplementary Material.

5.2.2.3 Evaluation of Preference Classifier Performance

We evaluated the classifiers using $F_1\text{-Score} = 2(\text{precision}^{-1} + \text{recall}^{-1})^{-1}$ because it balances different kinds of prediction errors and is less biased by class imbalance [141]. Precision = $|TP|/(|TP| + |FP|)$ is the proportion of positive predictions that are true positive targets, where $|TP|$ is the number of true positives and $|FP|$ is the number of false positives. Recall = $|TP|/(|TP| + |FN|)$ is the proportion of true positive targets that are correctly identified, where $|TP|$ is the number of true positives and $|FN|$ is the number of false negatives. We use the macro notion of $F_1\text{-Score}$. That is, we calculate the $F_1\text{-Score}$ for each of the three target classes and take the average.

In order to make the most use of our limited number of samples, we evaluated models using leave-one-out cross validation (LOOCV), rather than dividing the dataset into static training and testing sets. For each participant, we trained a model on the data from the other 193 participants and made a single test prediction. We then calculated $F_1\text{-Score}$ from the confusion matrix created from the 194 individual predictions (i.e., from the 194 folds of LOOCV).

5.3 Results

This section presents our evaluation of ML algorithms to predict which co-player behavior the participants preferred. We first discuss three naive baselines for predicting preferences. Then, we describe our analysis on whether or not nonverbal signals aid in modeling co-player behavior preferences (*RQ1*). Finally, we discuss results about incorporating additional context into our models (*RQ2*) and investigate feature importance values to gain insights into what kind of information is being leveraged by our preference classifiers.

5.3.1 Naive Baselines for Predicting Preferences

Of our 194 participants, 116 (60%) preferred the late-assistance co-player behavior, 50 (26%) preferred the early-assistance co-player behavior, and 28 (14%) did not prefer one over the other. With this distribution, always selecting the most dominant class results in an F_1 -Score of 0.25. Randomly sampling from the three classes and running the sampling $N = 10$ times results in an F1-score of $M = 0.29$ ($SD = 0.02$). Weighted sampling from the distribution of the true labels results in an F_1 -Score of $M = 0.33$ ($STD = 0.03, N = 10$).

5.3.2 Nonverbal Reactions Can Help Model Preferences

We first investigated whether nonverbal human reactions can help predict user preferences for an agent’s helping behaviors compared to only using explicit survey responses (*RQ1*). Table 5.2 shows the highest F_1 -Scores for the ML algorithms that were trained on both the Survey and Survey+Nonverbal input permutations described in Section 5.2.2.1. All F_1 -Scores are notably higher than the F_1 -Scores from naive baselines in Section 5.3.1. For each of the SVM, RF, and MLP algorithms, the highest F_1 -Score was from the set of models that incorporated nonverbal information. For KNN, the Survey only inputs resulted in a higher F_1 -Score than Survey+Nonverbal inputs. However, KNN was the lowest performing classifier among the ML algorithms considered in this work. We suspect that the low performance of KNN is related to the algorithm’s known trouble with the presence of outliers in training data [33] since nonverbal reactions can vary greatly. KNN’s low performance could also be due to the increase in dimensionality of input features with the addition of nonverbal data given our limited-sized dataset.

5.3.3 Incorporating Additional Context

Driven by the question of whether additional context information can help predict human preferences (*RQ2*), we analyzed the performance of preference classifiers with additional permutations of input data (Section 5.2.2.1). In particular, we considered demographic information, game context, and the combination of the two as additional context that could help reason about human perceptions of the agent. Figure 5.3 presents the highest F_1 -Score within each of the five permutations of input data for each ML algorithm considered in this analysis. Table 1 of the Supplementary Material includes information about the input features that resulted in the highest F_1 -Score for each type of classifier.

For each type of classifier, the highest F_1 -Score was from a set of models that included nonverbal reaction data and at least one other kind of additional context. Across all classifiers and all input permutations, SVM with Survey+Nonverbal+Game had the highest F_1 -Score (0.60). For both RF and MLP, the set of models with Survey+Nonverbal+Demographics/Game inputs had the highest F_1 -Scores of 0.59 and 0.57, respectively. The highest F_1 -Score for KNN (0.54) was obtained with the Survey+Nonverbal+Demographics data. Once again, KNN underperformed the other classifiers. Again, all F_1 -Scores are notably higher than the F_1 -Scores from naive baselines in Section 5.3.1.

Notably, including both kinds of additional contextual information did not always result in the highest F_1 -Score (see results for SVM and KNN in Figure 5.3). This means that one cannot take for granted that including more features will result in higher performance. Rather, it is important to further explore how to incorporate additional context when reasoning about internal human states.

In addition to classifier performance, we were interested in understanding which features

Table 5.2: Best F_1 -Scores for each classifier for Survey and Survey+Nonverbal input combinations. F_1 -Score was calculated over a confusion matrix derived from individual predictions of 194 folds of LOOCV.

Classifier	Survey	Survey+Nonverbal
Support Vector Machine	0.50	0.56
Random Forest	0.52	0.55
Multi-Layer Perceptron	0.51	0.54
K-Nearest Neighbors	0.50	0.47




Figure 5.3: Best F_1 -Scores for each input combination. F_1 -Score was calculated over a confusion matrix derived from individual predictions of 194 folds of LOOCV. Dots on left indicate which information was considered in the model. Machine learning algorithms (SVM, RF, MLP, KNN) are ordered left-to-right in decreasing order of highest F_1 -Score. The darkest bar highlights the highest F_1 -Score for each algorithm. Dotted lines separate input combinations into: Survey only (original baseline), Survey + Nonverbal (new baseline), Survey + Nonverbal + one type of additional interaction context, and Survey + Nonverbal + both types of additional interaction context.

played a key role in predicting participant preferences for agent helping behaviors. First, we considered doing this analysis with the SVM model because it had the top performance; but it employed a radial basis function kernel, making it hard to disentangle feature importances. Thus, we instead analyzed feature importance values with the RF algorithm. The RF models had good performance. Also, the high-level of interpretability of the underlying decision trees used for the preference predictions made it easy to understand what features mattered.

We compared the importance of features for a model considering Survey+Nonverbal inputs and Survey+Nonverbal+Demo/Game inputs. Instead of using LOOCV as for the prior results, we trained one RF model for each input combination with the data from all 194 participants for this analysis. This resulted in two models for which we then calculated

Gini feature importance values [44]. Gini importance indicates the total decrease in node impurity, weighted by the probability of reaching that node in the tree, averaged over all trees of the forest. Intuitively, this importance value can be thought of as a measure of a features' contribution to the homogeneity of the nodes and leaves in the forest.

Figure 5.4 illustrates the feature importance across the four types of input features considered by the RF with Survey+Nonverbal and Survey+Nonverbal+Demo/Game input combinations. In both cases, survey features were the most important to classifiers. Additionally, the results confirm our prior findings suggesting that the inclusion of additional context, via demographic and game information, contributes to improving classifier performance.

To investigate individual features, we examined feature importance with the Survey+Nonverbal+Demo/Game model. The three most important features were from survey responses: how much the participant liked the behavior of the agent (1st: 0.16), how annoying the participant found the agent (2nd: 0.15), and the competence of the agent (3rd: 0.08). The fourth (0.08) and fifth (0.05) most important features were principal components from the PCA analysis of nonverbal features. These principal components were driven by features describing AU45 (the “blink” action unit) for the game with the late-assistance co-player behavior as well as features describing AU2 (“outer brow raiser”), AU15 (“lip corner depressor”), and AU20 (“lip stretcher”) in the game with the early-assistance co-player behavior. The most important game feature was the number of frames the co-player was on the left side of the game screen (7th: 0.04). The most important demographic feature was how often the participant reported that they played video games (9th: 0.03).

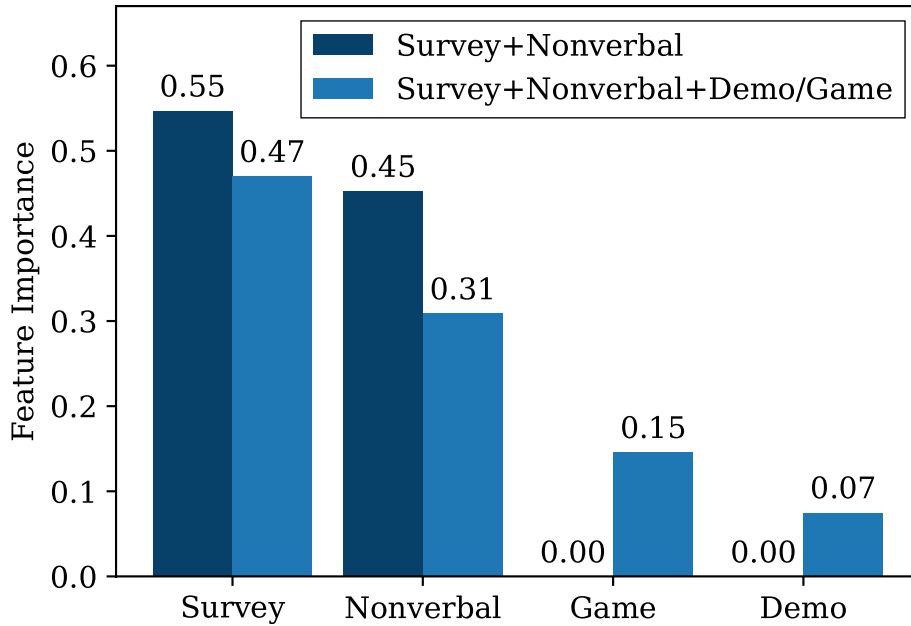


Figure 5.4: Feature importance by type of input feature.

5.4 Discussion

Our results support the idea that implicit feedback in the form of nonverbal behavior can help model user preferences. We trained a variety of ML classifiers to predict co-player behavior preferences and evaluated them based on F_1 -Score. Across four different types of ML algorithms, the three with the highest F_1 -Score all had better performance when implicit nonverbal clues were included as inputs in addition to explicitly provided survey responses. All four ML algorithms outperformed naive baselines based on the distribution of true labels. Repeatedly asking participants for their preferences can be annoying [135] or cause participants to lose interest in the interaction [277], so it is encouraging that analyzing nonverbal human signals can provide information about their preferences.

Figure 5.5 provides illustrative examples to highlight the value of implicit feedback, and why including it might help us better understand participant preferences. Participant P0354Z preferred the late-assistive behavior, which they experienced after the early-assistive

behavior. This participant was very expressive as can be seen in Figure 5.5(a) and 5.5(b). They moved their eyebrows, changed the expression of their mouth, and opened their eyes wide. A similar image highlighting implicit feedback provided by another, less expressive participant is included in Section F of the Supplementary Material. Additional examples are included in our supplementary video. We are excited about the potential for computational agents to leverage this kind of information to better render prosocial behavior and proactively cooperate with users in the future.

Additionally, our results suggest that considering additional context (participant and game information in our case) is important to interpret nonverbal behavior. For all four ML algorithms, the set of models with the highest F_1 -Score not only included nonverbal data, but also included additional context via demographic data, game data, or both. This aligns with research in psychology that contests the assumption that emotions are recognized and communicated universally with particular facial expressions and argues for the importance of considering the context of interactions [168, 24].

At first it may appear obvious that more features would lead to better predictions, but that is not necessarily the case. Jensen and Shen [172] argue against the notion that more features in datasets translates to better performance due to feature redundancy and relevancy [91, 198] as well as the curse of dimensionality [29]. Thus, it is important to carefully study what information predictive models should include when an autonomous agent is reasoning about how favorably a user is viewing their behavior.

Being able to predict user preferences over agent behaviors would open up doors for agents to reason about which behaviors are “better” for a given user and, thus, incorporate more of those behaviors into their interactions. Going forward, a better understanding of implicit human signals could also help human-agent interactions in other ways, e.g., enabling agents to learn how and when to ask a human collaborator for help [279].

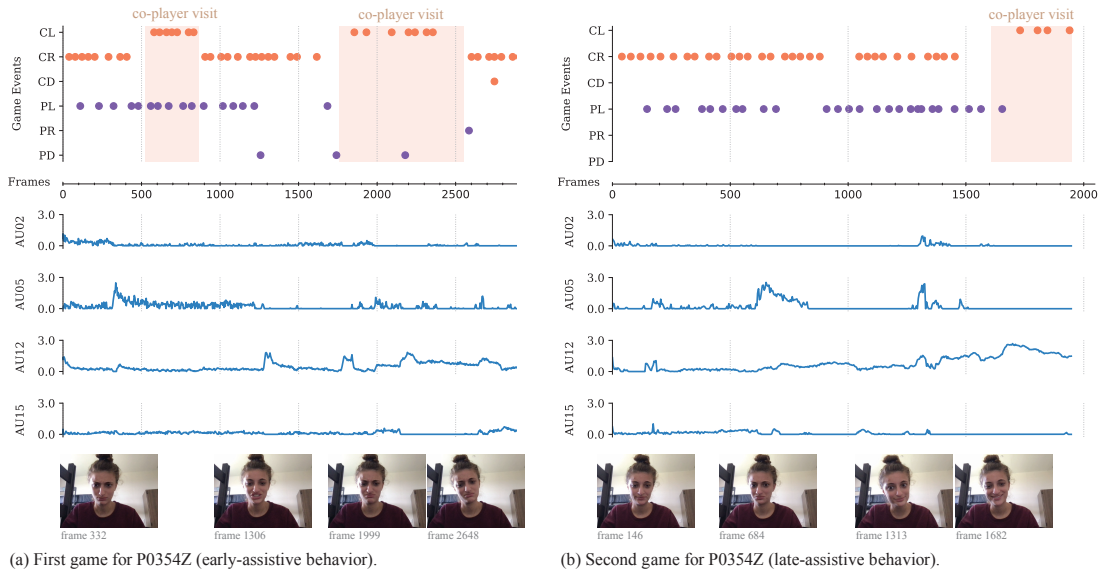


Figure 5.5: Visualization of data for participant P0354Z. Each plot shows game data (top) and webcam data (bottom). Game data includes six game events: co-player destroys enemy on the left side of the game screen (CL) or right side (CR), co-player dies (CD), participant destroys enemy on the left side (PL) or right side (PL), and participant dies (PD). The orange highlight identifies when the co-player is on the participants’ side (left side). Webcam data includes predicted action units (AU) by OpenFace: AU02 (outer brow raiser), AU05 (upper eyelid raiser), AU12 (lip corner puller), and AU15 (lip corner depressor). The x-axis corresponds to frame number as the game progresses from start to finish. All participants consented to having their images shared publicly.

5.5 Limitations and Future Work

Our work has limitations, which motivate interesting future research directions. First, our analysis is bound to the domain of Space Invaders. It would be interesting to investigate predicting preferences using human nonverbal reactions in other interactive scenarios, including human-agent interactions that occur in person (e.g., with embodied virtual agents or robots).

Second, future work could explore a richer utilization of implicit feedback. We found encouraging results from simple ML models using summary statistics of nonverbal reaction features over frames of a game, but stronger results may be discovered if we consider the

temporal nature of implicit feedback with more powerful ML algorithms (such as recurrent neural network models [146]).

Lastly, while this line of work is exciting, we must be cognisant of the ethical implications of designing autonomous agents able to analyze our nonverbal behavior and make inferences about our preferences. Going forward, it will be important to respect individual privacy and ensure individuals interacting with such autonomous agents are aware of the capabilities of the agents. Limitations on the social manipulation of autonomous agents will also be critical as agents become better able to understand humans in interactions.

5.6 Conclusion

This work investigated the usefulness of natural nonverbal human signals to predict preferences between two agent behaviors. We collected data via an online interaction in which participants played two games of Space Invaders with an autonomous co-player exhibiting different behaviors. We built models to predict which of the two games was preferred by the participant and analyzed results from different combinations of input data for the models. Without biasing humans to be expressive, we found that we could leverage “free” information that they provided via nonverbal reactions to improve our ability to predict their preferences for agent behaviors.

Based on our findings, we propose two key recommendations for future interactive agents. First, we recommend designing agents with the capability to reason about nonverbal human reactions. This capability can improve the speed at which interactive agents adapt to personal preferences because nonverbal signals are readily available during interactions. Second, it is important to incorporate additional types of context, such as user personality

or task statistics, into models that interpret nonverbal human signals. In our future work, we plan to take advantage of these insights to design better cooperative agents that reason about and adapt to the preferences of users with whom they interact.

Chapter 6

Two Datasets for Analyzing Implicit and Explicit Signals Together

The dataset and analyses of this chapter were presented as a short contribution at the 2024 ACM/IEEE International Conference on Human-Robot Interaction [58]:

“REACT: Two Datasets for Analyzing Both Human Reactions and Evaluative Feedback to Robots Over Time,” K. Candon, N. Georgiou, H. Zhou, S. Richardson, Q. Zhang, B. Scassellati and M. Vázquez. In Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2024.

In order to facilitate further research on how robots may interpret and integrate signals from the social context of human-robot interactions, we contribute the Reactions and Evaluative feedback over Time (REACT) database. REACT consists of two datasets that contain observations of humans, robots, and task-related data during human-robot




Figure 6.1: Overview of REACT. In REACT-Nao, people played a collaborative video game with a Nao robot (a). In REACT-Shutter, participants interacted with a Shutter robot during a photography task (d). For both datasets, we captured images of participants throughout the interaction (b,e) and provide facial analyses of the images (c,f).

interactions (as shown in Figure 6.1). The first dataset, **REACT-Nao**, consists of data from interactions from a user study [57] in which humans played a video game with a Nao robot while providing explicit feedback so that the Nao could learn to be a better teammate. REACT-Nao includes approximately 864 minutes of data collected across 72 participants. The second dataset, **REACT-Shutter**, consists of observations from interactions with a tabletop social robot during a photography task. REACT-Shutter includes approximately 160 minutes of data collected across 40 participants. Part of the latter data was used to investigate different annotation methods of robot performance during interactions [369]. In this work, we augmented this data with additional observations over the whole interaction to provide a more complete dataset to study human implicit signals in HRI. Together, the datasets provide a rich set of observations to analyze how human reactions are related to explicitly provided robot feedback. The datasets and documentation are available at: github.com/yale-img/react.

As a second contribution, we analyze the datasets to evaluate a common assumption in how machine learning models are used to make predictions about users from their nonverbal behavior in HRI. In particular, prior work often focuses on making predictions from short horizons of observations (e.g., [89, 369]). However, our analyses suggest that humans may become less reactive to robots over time. Thus, in the future, it is important for data-driven models to more explicitly account for interaction history in HRI. The data that we contribute

Table 6.1: Comparison of related available datasets. “Interactive task” indicates whether the human is actively interacting with the robot. “Additional task(s)” indicates if the participant had additional tasks other than just providing feedback to the robot (e.g., playing game in REACT-Nao). “Evaluative feedback” refers to if the dataset includes explicit, evaluative feedback about the robot from the participant throughout the interaction (either live or through annotations). The “Context” columns describe what additional context is provided in the dataset: Environment (e.g., location of enemies in REACT-Nao); Human (e.g., whether human spaceship moved left or right in REACT-Nao); Robot / agent (e.g., actual text of robot utterances in REACT-Shutter).

Dataset	Nonverbal Features					Task				Context			History
	Head pose	Gaze	Facial landmarks	Facial AUs	Raw images	Colocated robot	Interactive task	Additional task(s)	Evaluative feedback	Environment	Human	Robot	Spans interaction
EMPATHIC [89]	✓	✓	✓	✓	✓	X	X	X	X	✓	X	✓	✓
Errors in HRI [306]	X	X	X	✓	X	✓	✓	✓	X	X	X	X	✓
<i>REACT-Nao</i>	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
<i>REACT-Shutter</i>	✓	✓	✓	✓	X	✓	✓	✓	✓	✓	X	✓	✓

in this work opens up possibilities in this respect.

6.1 Related Work

6.1.1 Existing Datasets

There is a long history of open datasets with human nonverbal reactions (e.g., see [270] for a survey on human facial expression recognition); however, such datasets are still scarce within HRI. There exist some datasets of human nonverbal reactions to robots [45, 247, 171, 31, 89, 306]. Out of this set, the two publicly available datasets that are closest to REACT involve participants watching robots commit errors during an interactive task

[306] and watching agents perform a task sub-optimally [89], as detailed in Table 6.1. The other datasets [45, 247, 171, 31] provide great value to the field of HRI, but do not facilitate research examining both nonverbal human reactions and explicit evaluative feedback during a task in which both the human and robot play a key role. Our dataset includes both explicit, evaluative feedback and implicit, nonverbal reactions from participants that were actively interacting with a robot during a task. In comparison, the BAD Dataset [45] does not involve humans that are actively interacting with or explicitly evaluating a robot, but rather are reacting to videos that they observe online as bystanders. Similarly, the other datasets [247, 171, 31] do not include explicit feedback during the task. Rather, these datasets support other specific research avenues (e.g., modeling user engagement).

6.1.2 Interaction History

In prior work, models that reason about human nonverbal reactions to robots typically fail to account for a rich interaction history. It is a common approach to reason about nonverbal cues at the individual snapshot level (e.g., [306]), especially when inferring specific emotions or user states (e.g., [207]). Another approach is to examine changes in expressivity over fixed windows (e.g., [215]). While some models incorporate recurrence, they do not explicitly account for how human feedback may change over time (e.g., [369]). Our analyses suggest that as human-robot interactions evolve over time, human nonverbal signals may become more muted, requiring potentially different interpretations based on the interaction history. Going forward, it will be important to investigate algorithms that intelligently reason about feedback that is dependent on other factors, such as a longer interaction history or modeling of internal human states. This type of approach has been explored for reasoning about explicit human feedback, e.g., COACH learns from policy-dependent feedback [226].

6.2 The REACT-Nao Dataset

The first dataset, REACT-Nao, contains observations throughout a collaborative game between a Nao robot and humans, as presented in Chapter 4.

6.2.1 Data Collection

First, participants consented to take part in the data collection, be video recorded and have their data shared. Participants played six games of Space Invaders with a Nao robot (Figure 6.1a). They were instructed to provide feedback to the robot via their keyboard during the game so the robot could learn to be a better teammate. In the Space Invaders game, the goal was to destroy all enemies as a team. Each player generally took care of destroying enemies on one side of the game screen. However, the Nao employed different gameplay strategies across games which varied by when the robot’s spaceship crossed over to the human’s side of the gamescreen to help destroy enemies – we refer to these events as “visits”. During games 1 and 2, the robot did not crossover to the human’s side to provide assistance. During games 3 and 4, the robot crossed over to the human’s side for assistance on three separate “visits”. During games 5 and 6, the robot only crossed over for one “visit” at the end of the game, after it had destroyed all of the enemies on its own side. Participants were not prompted to speak during the interactions, but experimenters noted that some participants did speak at times.

Participants answered survey questions after each pair of games, and a final set of survey questions. The interaction lasted approximately 35 minutes, and the participants were compensated US\$10. The protocol was reviewed by our Institutional Review Board (IRB) and refined via pilots. For additional motivations and details of the user study, please

refer to Chapter 4.

6.2.2 Data Processing

The dataset consists of data collected for 72 participants during the six games of Space Invaders that they each played.

Facial Features. To analyze the images captured during the interaction, we used OpenFace 2.0 [21], a open-source toolkit for automatic facial behavior analysis. OpenFace 2.0 [21] uses computer vision algorithms to analyze each image and extract features about head pose, eye gaze, facial landmarks, and facial action units (AUs). Our data is organized in individual CSV files per game and participant. Each CSV file has one row per frame that includes a frame number and the output from running OpenFace on the image from that frame. A detailed description of individual features is included in the dataset documentation.

For our analyses, we first smoothed individual OpenFace features with a Gaussian filter (with a rolling window with a width of 30 data points and a Gaussian function with a standard deviation of 10). We then segmented the frames into “visits” by when the robot’s spaceship was on the participant’s side of the screen. We examined the mean of values of OpenFace activation values during various “visits” across the games of Space Invaders to see how participants reacted to a change in robot gameplay behavior. All post-processing scripts are included in github.com/yale-img/react.

Other features. Our dataset includes additional information that provides context about the interaction. For each game, we provide a json file that contains game state information, robot game actions, and participant game actions (including explicitly provided feedback via keyboard presses). We also provide a CSV that provides demographic information




Figure 6.2: Mean of sum of AU values during robot visits in REACT-NaO. Error bars are standard error. Letters (A,B,C) denote statistical significance. If visits do not share a letter, there is a statistically significant difference between values.

for each participant. Additionally, the raw images of the participant during the games is available at github.com/yale-img/react.

6.2.3 Results

We first analyzed how the robot’s visits affected human nonverbal signals as the data collection progressed. We used linear mixed models estimated with Restricted Maximum Likelihood (REML). The Game Number-Visit combination (e.g., Game3-First, Game4-Third, etc.) was a main effect and the participant ID was a random effect in the models. We conducted post-hoc Tukey Honestly Significant Difference (HSD) tests when appropriate.

We first examined the sum of AU activation values, as a proxy for participant expressiveness, during the robot visits in the interactions. Our analysis showed a significant difference by Game Number-Visit combination, $F(7, 7) = 16.54, p < 0.0001$. The post-hoc test revealed that the average of the sum of participant AU values during all three visits of both

Game 3 and Game 4 were significantly higher than the robot’s single visits in Games 5 and 6. Additionally, the average of the sum of participant AU values during the first visit of Game 3 was significantly higher than the third visit of Game 4. These differences between earlier and later visits show that participants reacted differently to similar stimuli based on when in the interaction they occurred. Figure 6.2 shows these results. A table of results is included in github.com/yale-img/react.

6.3 The REACT-Shutter Dataset

REACT-Shutter contains data from interactions with a robot photographer [369]. A subset of this data was previously published [369], but it only included observations during specific robot actions. REACT-Shutter provides the complete interaction history, enabling better analyses and modeling.

6.3.1 Data Collection

First, participants consented to take part in the data collection, be video recorded, and have their data shared. Each participant then sat in front of a small robot while the robot took six photographs of them (as in Figure 6.1d). The robot, called Shutter, is a social robot with a screen face mounted on a small arm [2, 211]. Shutter took photos of the participants via a camera mounted on its head.

Each photograph was preceded by a series of four robot actions. These actions consisted of a mix of robot dialogue (telling jokes, telling the person to smile, and telling the person to relax) and changes to the robot’s pose. The physical pose actions included aiming the robot’s face directly at the participant, orienting its face away from the participant, or

moving to one of four fixed poses. Actions were selected via weighted sampling, and an action could not be selected twice in a row – additional action details are included in the dataset documentation. Similar to subsection 6.2.1, participants were not prompted to speak during the interactions.

Throughout the data collection, participants annotated robot actions based on their impressions of the robot’s performance and answered survey questions. The whole interaction lasted between 45 minutes and one hour, and participants were compensated US\$20. The protocol was approved by the local IRB. For more details about the data collection, please refer to Zhang et al. [369].

6.3.2 Data Processing

The dataset consists of data collected for 40 participants, each of which completed six photography tasks.

Facial Features. The facial features were computed as in subsection 6.2.2, but the data is organized into CSVs by photography task.

For our analyses, we first smoothed individual OpenFace features with a Gaussian window function, using the same approach as in subsection 6.2.2. Additionally, we segmented the frames into action segments, splitting up the interaction based on when a new action began. We looked at the mean, median, maximum, and standard deviation of values of OpenFace features in each action segment. Post-processing scripts included in github.com/yale-img/react.

Other features. Our dataset includes additional information that provides context about the interaction. For each photography task, we include a CSV that provides the timestamps




Figure 6.3: Median of sum of AU values over the photography interaction. Error bars are standard error. Trend line is a linear regression model with a 95% confidence interval.

and details of robot actions throughout the task (e.g., specific utterance for a “joke” action). Additionally, we provide a summary CSV that provides additional information for each participant, including demographic information, the order of tasks, and the self-annotations. A full description of the features is available in the dataset documentation.

6.3.3 Results

We first explored how the expressiveness of participants changed over time as the interaction progressed. Considering all participants, we examined a variety of statistics (mean, median, max, standard deviation) over the sum of action unit activation values during the 24 actions that preceded the individual photos in order. For example, see Figure 6.3 for the median values over each action segment.

For each statistic calculated over the sum of AU activation values during action segments, we employed a linear regression model to predict the statistic considering action number as the independent variable. Table A of the dataset documentation displays the results computed with the `scipy.stats` Python library [295]. Across all four summary statistics, there was a statistically significant negative slope, suggesting that participants became less

expressive to robot actions over time. However, the slopes were just slightly negative, and the Pearson correlation coefficients were low suggesting that the model may not adequately capture the underlying relationships within the data. This is to be expected since expressivity likely depends on many other factors and warrants further study.

We fit another set of linear regression models, but this time considered whether the actions occurred first, second, third, or fourth in a mini-series before a photo as the independent variable. For these models, the slopes were positive for mean, median, and maximum values of the sum of action unit values over action segments (Table B of the dataset documentation). Taken with the previous results, this suggests that within a short photography task, participants got more expressive, but over time gradually became less expressive.

6.4 Discussion

The REACT database has the potential to influence HRI work by facilitating research that examines automated reasoning about human reactions. This could enable a deeper understanding of the dynamics of human-robot interactions, which is essential for designing more effective robots. As we work towards enabling robots to help with physical and social tasks in human environments, it will be important to consider how novelty effects diminish and people change their responses to robots during interactions. Failing to account for changes in user expressivity could cause robots to fail to adjust their behavior to muted reactions later on in interactions.

As with all human subject data, there are ethical considerations [158] for the use of the REACT database. Responsible use guidelines include ensuring that the data is not used for

purposes that would negatively manipulate or impact people.

Our database facilitates exciting research directions but it is not without limitations. The datasets showcase interactions for two different tasks, allowing users to explore model generalizability; however, it is unclear how analyses or models specific to these two tasks would translate to other interaction scenarios. Also, there are other forms of implicit communicative signals, such as the tone of verbal communications, that are not included in the datasets.

6.5 Conclusion

We contributed two datasets that can facilitate studying how robots can improve their behavior based on naturalistic human reactions. Additionally, we found preliminary evidence highlighting the importance of considering the interaction history when interpreting human reactions in HRI. We hope that the REACT database and initial findings encourage the HRI community to further explore how robots can learn from implicit human feedback over time.

Chapter 7

Combining Implicit Signals Via Physical Actions with Explicit Feedback

The approach discussed in this chapter was presented at the 2026 ACM/IEEE International Conference on Human-Robot Interaction [59]:

“Learning Human Preferences Over a Human-Robot Collaboration Based on Explicit and Implicit Human Feedback,” K. Candon, Q. Zhang, A. Lew, H. Claire, L. Qian, A. Quarles, C. Sarkar, and M. Vázquez. In Proceedings of the 2026 ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2026.

The culmination of this dissertation is a method for learning from various channels of the dynamic social context of a human-robot interaction during the collaboration. We contribute an approach for **P**reference learning from **I**mplicit and **E**xplicit feedback (**PIE**) in collaborative human-robot interactions with fixed roles. The key novelty of our




Figure 7.1: We study the problem of learning human preferences over human-robot collaborations. We propose to leverage human actions as implicit feedback that can help the robot learn preferences over the collaboration, alongside traditional explicit feedback. We evaluate this idea in a laboratory, where participants assemble pizzas with a robot and can give explicit feedback by pressing buttons.

work lies in framing human preferences over the human-robot collaboration that a person experiences, instead of over the robot’s behavior only, as is typical in prior interactive robot learning work. We formalize the human preferences as a shared set of parameters encoding alignment between the human’s behavior and the desired robot’s behavior. Then, by modeling human actions in a task as being (approximately) optimal with respect to the human’s preferences, we show how a robot can treat the human’s behavior as implicit feedback during the collaboration, effectively “listening” to what the human’s actions reveal about the human’s underlying preferences. Finally, we combine this implicit feedback with more traditional explicit human feedback to enable a robot to quickly estimate human preferences during an ongoing collaboration. We evaluate our proposed approach in simulations and with real users in a cooking collaboration, a popular setup for studying collaboration (e.g., see [310, 255, 43]). Our simulation results confirm that combining multiple modalities of human feedback improves a robot’s ability to estimate human preferences, with a similar trend observed in real-world evaluations.

In summary, this chapter has three main contributions. First, we propose a novel formulation for preference learning over human-robot collaborations, which leverages *both* implicit and explicit human feedback. Second, we systematically investigate the effectiveness of our approach for preference learning (PIE) in simulations and in the real-world. Our experiments consider varying cooking tasks, assumptions about the rationality of human behavior, and different preferences. Lastly, we open-source our implementation to facilitate future replication and benchmarking efforts.¹

7.1 Related Work

Our work is inspired by this body of research exploring subtle human signals, as introduced in Chapter 2, but takes a distinct approach by treating the human’s task-oriented actions as a rich source of implicit feedback. Our work leverages the INQUIRE framework [113], which frames human feedback in terms of inferences about accepted and rejected robot behavior. However, instead of investigating interactions where a human serves as a dedicated teacher, whose only task is to teach a robot, we investigate collaborations where the human teacher also takes task-relevant actions.

7.2 Problem Setup: Learning Preferences Over a Human-Robot Collaboration

We consider collaborative interactions in which a human H and a robot R work together to complete a physical task and where each has a specific role in the collaboration. For

¹https://github.com/yale-img/pie_preflearning.

example, in our pizza-making scenario of Figure 7.1, the robot may pass ingredients to the human from a storage area, and the human may use the ingredients to assemble a desired pizza.

At a time-step t , the human and robot observe a given state \mathbf{s} and take simultaneous high-level actions $a_H \in \mathcal{A}_H$ and $a_R \in \mathcal{A}_R$. We model these high-level actions as parameterized actions. For example, two high-level actions for the robot may be `pick(<ingredient>, <location>)` and `place(<ingredient>, <location>)` in Figure 7.1.

Collaborations are characterized by teammates trying to maximize a shared reward [25, 312, 18]. It is typical for the reward to be based on task success only [25, 69] or human preferences that encapsulate the desired task outcome [222, 113, 228]. However, in many interactions, it can be helpful to explicitly model both the task goal and human preferences, e.g., because not taking action toward the task goal is worse than violating preferences [245], or the task goal is public information while the human preferences are not [372]. We assume that both situations are true in our work, so, building off Cooperative Inverse Reinforcement Learning [138], we propose to define the shared reward as a combined reward:

$$R(\mathbf{s}, a_R, a_H) = R^{\text{goal}}(\mathbf{s}, a_R, a_H) + \gamma R^{\text{pref}}(\mathbf{s}, a_R, a_H) \quad (7.1)$$

where γ is a domain-specific parameter that controls the relative importance of the two rewards. This reward formulation is similar to Zhao et al. [372] preference learning setup, but we assume that human preferences are over the human-robot collaboration (not just the human’s contribution to the task) and the robot must follow these preferences (rather than having its own individual reward). Our framing enables the robot to leverage observed human actions as implicit feedback for preference learning.

Goal Reward: Motivated by the fixed-role assignments, we propose to decompose the goal

reward in eq. (7.1) into two components, one for the robot and one for the human:

$$R^{\text{goal}}(\mathbf{s}, a_R, a_H) = R_R^{\text{goal}}(\mathbf{s}, a_R) + R_H^{\text{goal}}(\mathbf{s}, a_H) \quad (7.2)$$

Preference Reward: We assume that the preference reward, R^{pref} in eq. (7.1), does not conflict with R^{goal} . In addition, we assume that the human preferences for the team members are aligned with each other, such that the preference reward can also be decomposed into two terms parameterized by the same weights \mathbf{w} :

$$\begin{aligned} R^{\text{pref}}(\mathbf{s}, a_R, a_H) &= R_R^{\text{pref}}(\mathbf{s}, a_R) + R_H^{\text{pref}}(\mathbf{s}, a_H) \\ &= \mathbf{w}^\top \phi_R(\mathbf{s}, a_R) + \mathbf{w}^\top \phi_H(\mathbf{s}, a_H) \\ &= \mathbf{w}^\top (\phi_R(\mathbf{s}, a_R) + \phi_H(\mathbf{s}, a_H)) \end{aligned} \quad (7.3)$$

We implement the preference reward as a linear function of features of the state (encoded via ϕ_R and ϕ_H) to keep the reward interpretable in this work. While this setup is common in the preference learning literature (e.g., [174, 113]), future work could investigate ways to relax this assumption (e.g., via more complex reward models implemented as neural networks [75, 161]).

Taken together, the above assumptions mean that if the robot and human act rationally, they maximize their individual rewards:

$$R_R(\mathbf{s}, a_R) = R_R^{\text{goal}}(\mathbf{s}, a_R) + \gamma R_R^{\text{pref}}(\mathbf{s}, a_R) \quad (7.4)$$

$$R_H(\mathbf{s}, a_H) = R_H^{\text{goal}}(\mathbf{s}, a_H) + \gamma R_H^{\text{pref}}(\mathbf{s}, a_H) \quad (7.5)$$

The Robot’s Learning Objective: In this work, we assume that the robot knows the goal

reward and the features of the state that may matter for the human preferences over the collaboration (ϕ_R and ϕ_H in eq. (7.3)); however, the robot does not know the weights \mathbf{w} that parameterize the preference reward functions R_H^{pref} and R_R^{pref} . Thus, the goal of the robot is to estimate the weights \mathbf{w} based on human feedback gathered *during* the human-robot collaboration.

7.3 Preference Learning from Implicit and Explicit Feedback (PIE)

We propose that robots estimate a human’s preferences for their collaboration per eq. (7.3) based on both *explicit* feedback provided by the human about the robot’s behavior as well as *implicit* feedback provided by the human’s own actions in the task. In our work, explicit feedback corresponds to binary evaluative feedback and, in our real-world evaluation, is implemented via physical button presses (e.g., similar to [276]). The main assumption for the implicit feedback is that human actions are driven by the reward in eq. (7.5).

Algorithm 1 describes **PIE**, our proposed approach for preference learning, considering a given interaction step in a collaboration. First, the robot observes the current state of the world, takes action according to some policy π_R , and sees how the human behaves (lines 1-3 in Alg. 1). Every action that the human takes is then interpreted as implicit feedback for preference learning and the implications of the human’s behavior are stored in a feedback set (lines 4-6). Each element in this set includes a pair of accepted behavior (f^+) and rejected behavior (f^-) along with their associated state. When the human chooses to give explicit feedback, the implications of this feedback are also stored in the feedback set (lines 11-14). Finally, a non-parametric belief over preferences \mathbf{W} , implemented via M weight

Algorithm 1: Learning from Explicit & Implicit Feedback (PIE)

Input: Prior belief over pref. weights $\mathbf{W} = \{\mathbf{w}^i\}_{i=1}^M$, and prior feedback set \mathbf{F} with pairs of acceptable (f^+) and rejected (f^-) behaviors in prior states

Output: Updated belief \mathbf{W} , and updated feedback set \mathbf{F}

```
// Interact with the environment
1 Observe current state  $\mathbf{s}$ ;
2 Robot takes action  $a_R \leftarrow \pi_R(\mathbf{s})$ , where  $a_R \in \mathcal{A}_R$ ;
3 Observe current human action  $a_H \in \mathcal{A}_H$ ;

// Store implicit human feedback
4  $f_{imp}^+ \leftarrow a_H$ ; /* Observed  $a_H$  aligns with pref. */
5  $f_{imp}^- \leftarrow \mathcal{A}_H(\mathbf{s}) \setminus \{a_H\}$ ;
6  $\mathbf{F} \leftarrow \mathbf{F} \cup \{(f_{imp}^+, f_{imp}^-, \mathbf{s})\}$ ;

// Store explicit feedback on button press
7 if human indicates acceptable robot behavior then
8 |  $f_{exp}^+ \leftarrow a_R$ ; /* Robot action  $a_R$  aligns with pref. */
9 |  $f_{exp}^- \leftarrow \mathcal{A}_R(\mathbf{s}) \setminus \{a_R\}$ ;
10 else if human indicates unacceptable robot behavior then
11 |  $f_{exp}^+ \leftarrow \mathcal{A}_R(\mathbf{s}) \setminus \{a_R\}$ ; /* Other robot actions are better
    | aligned with pref. than  $a_R$  */
12 |  $f_{exp}^- \leftarrow a_R$ ;
13  $\mathbf{F} \leftarrow \mathbf{F} \cup \{(f_{exp}^+, f_{exp}^-, \mathbf{s})\}$ ;

// Update non-parametric belief over preference weights
14 foreach  $\mathbf{w}^i \in \mathbf{W}$  do
    | // MLE via gradient ascent, starting from prior  $\mathbf{w}^i$ 
15 |  $\mathbf{w}^i \leftarrow \text{optimize\_w\_to\_maximize\_likelihood}(\mathbf{F}, \mathbf{w}^i)$ ;
16 end
17 return  $\mathbf{W}, \mathbf{F}$ ;
```

samples, is computed using the feedback set (lines 15-17). PIE builds on the INQUIRE formalism [113] for combining various types of feedback during interactive robot learning.

One key difference between INQUIRE [113] and our approach, PIE, is that we consider implicit human feedback, not just explicit feedback. This results in different implications for preference learning. For explicit feedback directed intentionally from the human to the robot, the implications of the feedback are set in relation to the robot’s behavior. However, for the implicit feedback, we instead define the implications in relation to human behavior because this feedback is a direct consequence of the human’s own actions. Next, we describe in more detail how we interpret the feedback for preference learning and estimate the preference belief.

7.3.1 Implication of Explicit Human Feedback

We specifically consider explicit feedback in the form of binary feedback. When a robot takes a high-level action a_R in a state s (line 2 in Alg. 1), the human may choose to indicate whether the behavior is acceptable or not. If the human indicates that the robot behavior is acceptable (lines 7-9 in Alg. 1), the robot’s action is added to the accepted behavior set $f_{exp}^+ = \{a_R\}$, and all other viable robot actions in the state are added to the rejected behavior set $f_{exp}^- = \mathcal{A}_R(s) \setminus \{a_R\}$. However, if the human indicates unacceptable robot behavior (lines 10-12 in Alg. 1), the implication is the opposite. In this case, the accepted behavior set includes viable robot actions in the state that the robot did not take ($f_{exp}^+ = \mathcal{A}_R(s) \setminus \{a_R\}$), and the rejected behavior set includes only the action that was taken by the robot ($f_{exp}^- = \{a_R\}$).

Our formulation for the implication of binary feedback follows INQUIRE [113], with the exception that we do not study active preference learning in this work, so the robot does

not pose questions to the human for which feedback is received in return. Rather, the human may choose to give or not give explicit feedback at any point during the collaboration. Because explicit feedback is potentially sparse and prior findings show that people may reduce the amount of feedback that they give to a robot during collaborations [57], we propose to also consider implicit feedback.

7.3.2 Implication of Implicit Human Feedback

A key novelty of our work is framing the human’s actions as a source for implicit feedback about the human preferences over the collaboration. These preferences are encoded in the weight vector \mathbf{w} shared by $R_H^{\text{pref}}(\cdot)$ and $R_R^{\text{pref}}(\cdot)$, per eq. (7.3). If during the collaboration, the human takes actions that are approximately optimal with respect to the human’s reward $R_H(\cdot)$ in eq. (7.5), then the human’s behavior will leak information about \mathbf{w} through $R_H^{\text{pref}}(\cdot)$.

Formally, we assume that the human takes action on a given state \mathbf{s} following a Boltzmann rational policy:

$$P(a_H|\mathbf{s}) \propto \exp(\beta_H R_H(\mathbf{s}, a_H)) \tag{7.6}$$

where β_H controls how rational the human’s actions are. This model of human behavior is common in economics [299], psychology [20], and preference learning [174]. Importantly, eq. (7.6) results in myopic high-level decision-making because the human is said to take actions based on the reward of the current state. We find that this myopic assumption is reasonable for preference learning over high-level actions that span multiple time-steps during the collaboration and when the human’s reward is not sparse. However, for sparse rewards, this formulation would need to be adapted to a Boltzmann policy based on expected

future rewards (e.g., via Q-values [27]). We discuss this future work in Section 7.6.

Equipped with a model of human actions per eq. (7.6), we can formulate feedback based on these actions for preference learning. At a given state s , we consider the set of valid actions of the human, $\mathcal{A}_H(s)$, as the set of possible choices that the human has for implicit feedback in that state. Hence, when the human takes action $a_H \in \mathcal{A}_H(s)$ at a given point in the collaboration, the implication of that choice is that that action is accepted behavior ($f_{imp}^+ = \{a_H\}$) and that other viable human actions are rejected behavior ($f_{imp}^- = \mathcal{A}(s) \setminus \{a_H\}$). The implication of implicit feedback from human actions is outlined in lines 4-6 of Algorithm 1.

By reframing human actions as implicit feedback with the implication described previously, a robot can gain information about the preference weights w potentially *all throughout* a collaboration, without having to continuously query the human for feedback. Mathematically, the implication that we propose for implicit human feedback is equivalent to the implication for human demonstrations of robot behavior used in INQUIRE [113]. However, our implication defines sets of accepted and rejected *human* behavior, rather than sets of robot behavior, so the implications are conceptually different.

7.3.3 Estimating Belief Over Preference Weights

Whenever the robot receives human feedback, it stores the implications of the feedback (f_m^+, f_m^-) in a cumulative feedback set F , where m is the modality (explicit or implicit) of the feedback, alongside the current state s of the interaction when the feedback was received (see lines 6 and 13 in Algorithm 1). In PIE, the modality m of the feedback is critical because it dictates the perspective from which the robot should reason about the implications of the feedback.

The robot’s objective consists of estimating the preference weights that maximize the likelihood of the accepted behavior implied by the feedback in \mathbf{F} . Specifically, the likelihood is:

$$\mathcal{L}(\mathbf{w}) = \prod_{(f_m^+, f_m^-, \mathbf{s}) \in \mathbf{F}} P(f_m^+ | \mathbf{w}) = \prod_{(f_m^+, f_m^-, \mathbf{s}) \in \mathbf{F}} \frac{\sum_{a \in f_m^+} B_m(\mathbf{s}, a)}{\sum_{a \in f_m^+ \cup f_m^-} B_m(\mathbf{s}, a)} \quad (7.7)$$

where:

$$B_m(\mathbf{s}, a) = e^{\hat{\beta}_m \overbrace{(R_{\text{agent}}^{\text{goal}}(\mathbf{s}, a) + \gamma \mathbf{w}^\top \phi_{\text{agent}}(\mathbf{s}, a))}^{\text{agent's reward}}} \quad (7.8)$$

is the exponential component of the Boltzmann rationality model. The agent subscript in eq. (7.8) denotes the interactant associated with the modality m : when m is explicit, the agent is the robot R ; when m is implicit, the agent is the human H . Thus, the agent’s reward in eq. (7.8) is implemented per eq. (7.4) or eq. (7.5), respectively. Finally, the parameter $\hat{\beta}_m$ in eq. (7.8) models the robot’s assumptions about how rational the human is at providing explicit or implicit feedback (depending on m) as a function of the agent’s reward.

The goal of preference learning can then be expressed as: $\mathbf{w}^* = \arg \max_{\mathbf{w}} \mathcal{L}(\mathbf{w})$. We solve this Maximum Likelihood Estimation (MLE) problem using a belief distribution for the preference weights, which is implemented via a sample set $\mathbf{W} = \{\mathbf{w}_i\}_{i=1}^M$, as indicated in lines 14-16 of Algorithm 1. Specifically, we use gradient ascent on the log-likelihood $LL(\mathbf{w}) = \log \mathcal{L}(\mathbf{w})$ to find suitable preference weights. When the update takes place, gradient ascent is applied on each weight $\mathbf{w}_i \in \mathbf{W}$ using the gradient:

$$\nabla LL(\mathbf{w}) = \sum_{(f_m^+, f_m^-, \mathbf{s}) \in \mathbf{F}} \left(\frac{\sum_{a \in f_m^+} \hat{\beta}_m \gamma \phi_{\text{agent}}(\mathbf{s}, a) B_m(\mathbf{s}, a)}{\sum_{a \in f_m^+} B_m(\mathbf{s}, a)} - \frac{\sum_{a \in f_m^+ \cup f_m^-} \hat{\beta}_m \gamma \phi_{\text{agent}}(\mathbf{s}, a) B_m(\mathbf{s}, a)}{\sum_{a \in f_m^+ \cup f_m^-} B_m(\mathbf{s}, a)} \right) \quad (7.9)$$

The belief \mathbf{W} is randomly initialized when learning begins but, in subsequent time-steps,

we start gradient ascent using the belief estimated from the prior time-step. Reusing prior estimates of \mathbf{W} is essential for gradient ascent to converge quickly because the feedback set \mathbf{F} grows over time, making the objective more complex.

7.4 Evaluation in Simulation

We first evaluate our proposed approach for learning human preferences over a human-robot collaboration in simulation. We focus the evaluation on understanding the impact of the feedback modalities and key parameters of PIE. In particular, we consider different values for $\hat{\beta}_{imp}$ and $\hat{\beta}_{exp}$ in eq. (7.8). These parameters are used to find the weights \mathbf{w} that maximize the likelihood of the accepted behavior implied by the human’s feedback (see eq. (7.7)). For implicit feedback, $\hat{\beta}_{imp}$ indicates how rational the robot considers the human to be when choosing its own actions during the collaboration. For explicit feedback, $\hat{\beta}_{exp}$ indicates how rational the robot considers the human to be at deciding whether the robot’s behavior is acceptable or not. More specifically, our research questions are:

(RQ1) How does the type of feedback considered by the robot affect preference learning? A motivating hypothesis for this work is that combining explicit and implicit feedback will facilitate learning preferences over the collaboration. Thus, we compare three experimental conditions: 1) *explicit-only* feedback; 2) *implicit-only* feedback; and 3) *combined* feedback, where the robot learns from both explicit and implicit feedback with PIE.

We know that people can deviate from optimal decision making in varied ways [73, 204], so RQ1 considers different levels of rationality for the human’s actions in the collaboration (β_H in eq. (7.6)). Also, because RQ1 is focused on the effect of different feedback modalities, we assumed that the robot knows the level of rationality of the human’s actions,

so $\hat{\beta}_{imp} = \beta_H$. Lastly, we set $\hat{\beta}_{exp} = \hat{\beta}_{imp}$ for simplicity, as prior work often considers a single rationality coefficient β for integrating various types of feedback [113].

(RQ2) How does the correctness of the robot’s assumptions about the rationality of the human’s feedback affect preference learning? Our second experiment evaluates the performance of our PIE approach when we introduce the complication that the robot does not know how rational the human truly is. We systematically study in simulation how preference learning performance with PIE is affected by the alignment, or misalignment, between β_H and $\hat{\beta}_{imp}$. Specifically, we consider the actions taken by the human to be more ($\beta_H = 10$) or less rational ($\beta_H = 1$). Then, we consider two situations per β_H : the assumptions on the human’s rationality are aligned with the simulated human (e.g., $\hat{\beta}_{imp} = \beta_H = 1$), or they are misaligned (e.g., $\hat{\beta}_{imp} = 1$ but $\beta_H = 10$). Also, the robot reasons about human button presses in two ways. It assumes that the human’s explicit feedback is more rational with $\hat{\beta}_{exp} = 10$, or less rational with $\hat{\beta}_{exp} = 1$.

7.4.1 The Pizza-Making Task

We consider collaborations where the human and robot prepare pizzas together. The robot passes ingredients to the human, while the human is responsible for more complicated manipulation tasks involving assembling the pizza. This results in different action spaces for the human and the robot. The robot’s high level actions include: `pick(broccoli, storage)` or `place(broccoli, workstation)`, whereas the human’s high level actions include `add(pepperoni)` or `return(pepperoni)`. Both the robot and the human know the goal of the task, which is defined by the ingredients that comprise a given pizza. For example, they may work towards making a pizza with sauce, cheese, pepperoni, mushrooms, and olives. However, the robot does not know the true human preferences w^*

for how the team should reach the goal. Thus, the robot selects actions during a collaboration according to only the goal reward: $\pi_R(a_R|\mathbf{s}) = P(a_R|\mathbf{s}) \propto \exp(\beta_R R_R^{\text{goal}}(\mathbf{s}, a_R))$, where β_R controls how rational the robot’s actions are. We set $\beta_R = 1$ in our experiments.

We consider two types of preferences. First, the human can have *ordering preferences* for the ingredients, e.g., cheese should be placed on the pizza before the sauce. We considered a total of six ordering preferences, each of which is a feature in the preference weight vector \mathbf{w} . Second, the human can have *workspace preferences* over the number of ingredients that can be on the shared workspace at any given time, including one ingredient maximum, two ingredients maximum, or up to four ingredients. We describe these workspace preferences via two features in \mathbf{w} . Thus, the preference weight vector has a total of eight dimensions.

We set the components of the shared reward (eq. (7.1)) as follows. The goal reward is most positive when a topping that should be on the pizza is moved from the storage to the workstation or added to the pizza. The preference reward is defined as in eq. (7.3). The appendix includes a more detailed description of the reward, state, action, and preference space of the pizza-making task.

7.4.2 Simulating the Human in the Collaboration

Following common practice for evaluations in the preference learning literature [174, 113], we model human behavior in our simulations with a Boltzmann rationality model. We assume that the human tends to take rational actions per eq. (7.6). When a new time-step of the collaboration occurs in simulation, the simulated human always gives explicit feedback — in Section 7.5, we demonstrate PIE with real human feedback, which can vary in frequency over time [57].

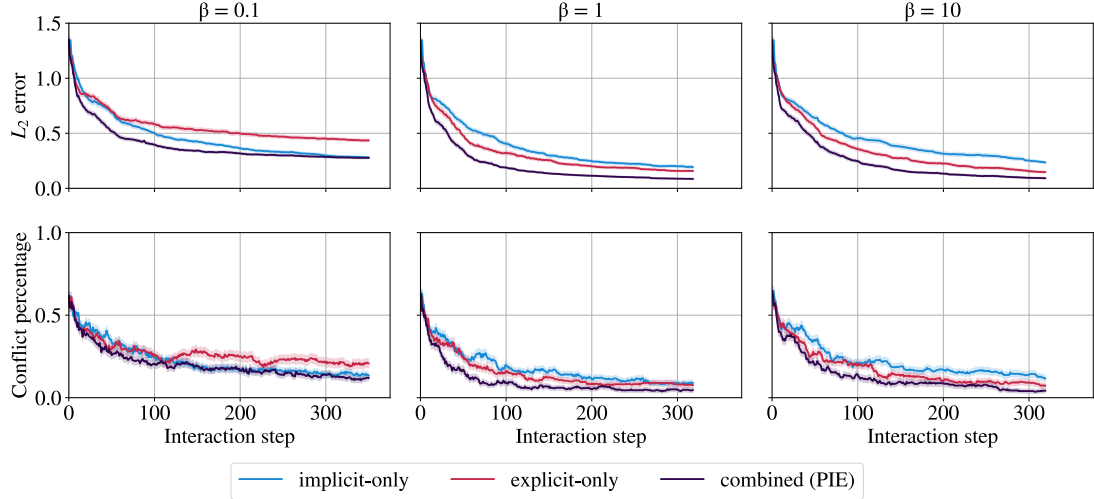


Figure 7.2: Results for RQ1: L_2 error with respect to the ground truth preference weights (top) and conflict percentage (bottom) as the interaction progresses. Lines represent average L_2 error and conflict percentage, and shading represents std. error across 100 interactions. Columns show results considering different β values, where $\beta_H = \hat{\beta}_{imp} = \hat{\beta}_{exp} = \beta$ for $\beta \in \{0.1, 1, 10\}$.

Prior work in interactive learning shows that binary human feedback tends to be “noise-reducing” in comparison to other types of explicit feedback [365]; thus, we simulated explicit feedback as rational feedback for studying RQ1 and RQ2. The simulated human decides which binary feedback signal to provide based on how well the robot performs relative to the best possible reward. To compute the best possible reward, we leverage the fact that the simulated human knows the reward of the robot R_R , as in eq. (7.4), because they know the goal reward and the true preference weights w^* . Then, when the robot takes action a_R at a given time-step with a state s , the simulated human compares the actual reward induced by the robot’s action, $R_R(s, a_R)$, with the highest possible reward $\max_{\hat{a}_R} R_R(s, \hat{a}_R)$ the robot could receive at state s , over all possible actions \hat{a}_R . If $R_R(s, a_R) == \max_{\hat{a}_R} R_R(s, \hat{a}_R)$, then the human gives positive binary feedback, indicating acceptable robot behavior. Otherwise, the simulated human gives negative feedback.

7.4.3 Evaluation Setup

We evaluate learning via two metrics:

L_2 error: As is common in preference learning, we evaluated learning in terms of the L_2 error with respect to the ground truth preference weights \mathbf{w}^* as the interaction progresses. At each timestep, we compute $\|\mathbf{w}^* - \sum_i \mathbf{w}_i/M\|$, with \mathbf{w}_i samples from the belief \mathbf{W} .

Conflict percentage: At each time step, we use the average weight estimate $\tilde{\mathbf{w}} = \sum_i \mathbf{w}_i/M$, with $\mathbf{w}_i \in \mathbf{W}$, to define a greedy robot policy $\pi_g(s) = \arg \max R_R(s, a_R; \tilde{\mathbf{w}})$. We then simulate a full interaction with a target pizza where the robot takes actions according to the greedy policy, and calculate the percentage of robot actions that deviate from the true optimal robot actions, based on \mathbf{w}^* . This metric allows us to quantify the practical effect of learning the preferences: a wrong estimate for the true weight \mathbf{w}^* may result in similar behavior to \mathbf{w}^* , or small deviations from \mathbf{w}^* could change the robot’s behavior in an undesired way.

7.4.4 Results

7.4.4.1 (RQ1) Type of Feedback

Figure 7.2 shows the results for 100 simulated human-robot interactions, each of which consisted of 10 pizzas and where the simulated human had a specific ground truth preference that was randomly sampled. We evaluated learning using different values of β , representing different (ir)rationality levels for the actions the simulated human took (β_H in eq. (7.6)) and for the robot’s assumptions when reasoning about explicit and implicit human feedback ($\hat{\beta}_{exp}$ and $\hat{\beta}_{imp}$ in eq. (7.8)).

We conducted a statistical analysis of the final L_2 error in Figure 7.2, after the 10 pizzas. Specifically, we used a linear mixed model analysis, estimated with REstricted Maximum Likelihood (REML), to evaluate the L_2 error. The model considered Interaction ID (100 levels) as random effect, Feedback Modality (*Explicit*, *Implicit*, or *Combined* with (PIE)) and Rationality ($\beta \in \{0.1, 1, 10\}$) as main effects, and the interaction effect of the latter two variables.

Feedback Modality had a significant effect on the L_2 error ($p < 0.0001$). As we hypothesized, a Tukey HSD post-hoc test indicated that learning preferences with *Combined* feedback (using PIE) led to significantly lower error than using a single feedback modality only. At the end of learning, the L_2 error for *Combined* feedback was $M = 0.15$ ($SE = 0.007$), for *Explicit* feedback was $M = 0.25$ ($SE = 0.01$), and for *Implicit* feedback was $M = 0.24$ ($SE = 0.007$).

The analysis also indicated a significant effect of the Rationality parameter (β) on the L_2 error ($p < 0.0001$). A Tukey HSD post-hoc test indicated that $\beta = 0.1$ ($M = 0.33$, $SE = 0.008$) led to significantly higher L_2 error than $\beta = 1$ ($M = 0.15$, $SE = 0.006$) and $\beta = 10$ ($M = 0.16$, $SE = 0.007$). As discussed later for RQ2, this finding can be due to an important mismatch between how the robot modeled the rationality of the human’s explicit feedback when $\beta = 0.1$ (which implied $\hat{\beta}_{exp} = 0.1$ for RQ1) and the perfectly-rational approach used by the simulated human to provide explicit feedback (as explained in Sec. 7.4.2).

Finally, we also found the Feedback Modality \times Rationality interaction to have a significant effect on the L_2 error at the end of the interactions ($p < 0.0001$). Significant pairwise differences from a Tukey HSD post-hoc test are shown in Fig. 7.3. Notably, using $\beta = 0.1$ and *Explicit* feedback only led to the highest error ($M = 0.44$; $SE = 0.01$) of all combinations of Modality and Rationality. Meanwhile, the *Combined* feedback led to

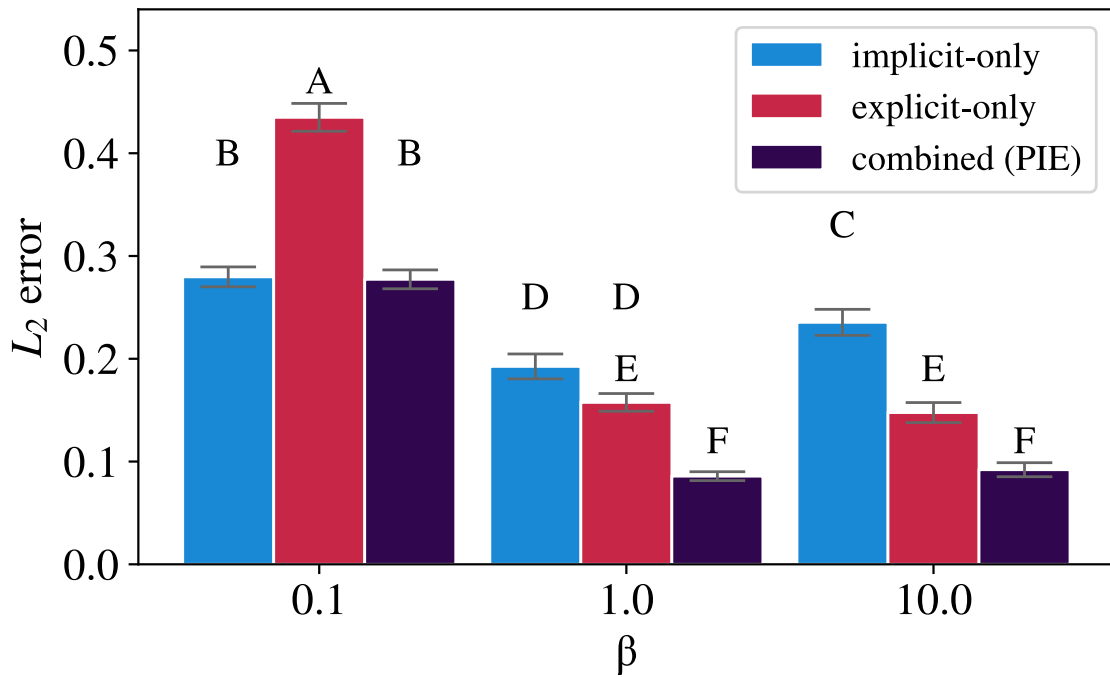


Figure 7.3: Results for RQ1: L_2 error at the end of the interactions based on Feedback Modality and Rationality ($\beta = \beta_H = \hat{\beta}_{imp} = \hat{\beta}_{exp}$). Error bars are std. error. Bars labeled with different letters (A-F) have significantly different error based on a Tukey HSD post-hoc test. See the text for more details.




Figure 7.4: Results for RQ2: Average L_2 error as the robot learns with PIE. The two lines in the plots indicate whether the robot’s assumption about the rationality of the human’s action ($\hat{\beta}_{imp}$) match the oracle’s behavior (β_H): solid lines correspond to $\hat{\beta}_{imp} = \beta_H$; dashed lines correspond to an erroneous assumption $\hat{\beta}_{imp} \neq \beta_H$. In the top row, when $\hat{\beta}_{imp} \neq \beta_H$, $\hat{\beta}_{imp} = 10$. In the bottom row, when $\hat{\beta}_{imp} \neq \beta_H$, $\hat{\beta}_{imp} = 1$. Shaded areas (visible when zoomed in) are the std. error in 100 interactions.

significantly smaller error with $\beta = 1$ ($M = 0.08$; $SE = 0.004$) and $\beta = 10$ ($M = 0.09$; $SE = 0.006$) compared to all other combinations.

7.4.4.2 (RQ2) Assumptions for $\hat{\beta}_{imp}$ and $\hat{\beta}_{exp}$

Figure 7.4 shows the L_2 error for PIE over 100 interactions, where each interaction consisted of 10 pizzas. The dashed lines show preference learning performance when there is a mismatch between how rational the human is at taking actions (β_H) and how the robot modeled this rationality ($\hat{\beta}_{imp}$); while the solid lines indicate performance when the robot’s

assumption was correct and $\hat{\beta}_{imp} = \beta_H$. In most cases, the dashed line leads to higher error, especially when $\beta_H = 1$. The results in Fig. 7.4 seemed less susceptible to the choice of $\hat{\beta}_{exp} \in \{1, 10\}$.

We conducted a linear mixed model analysis on the L_2 error at the end of the interactions, considering Interaction ID as random effect. The main effects were $\hat{\beta}_{imp}$ Alignment (which had a value of 1 when $\hat{\beta}_{imp} = \beta_H$, and 0 otherwise), and $\hat{\beta}_{exp}$ Alignment (which had a value of 1 when $\hat{\beta}_{exp} = 10$ and was 0 otherwise, because 10 better approximated the purely-rational feedback from the simulated human than $\hat{\beta}_{exp} = 1$). The analysis also considered the interaction effect between $\hat{\beta}_{imp}$ Alignment and $\hat{\beta}_{exp}$ Alignment.

The analysis indicated that $\hat{\beta}_{imp}$ Alignment had a significant effect on the L_2 error at the end of the interactions ($p < 0.0001$). As expected, a post-hoc t-test showed that $\hat{\beta}_{imp} = \beta_H$ led to significantly lower error than the misaligned $\hat{\beta}_{imp}$. Similarly, $\hat{\beta}_{exp}$ Alignment had a significant effect on the L_2 error ($p < 0.0001$). The post-hoc test indicated that $\hat{\beta}_{exp} = 10$ (more aligned) led to lower error than $\hat{\beta}_{exp} = 1$ (less aligned), although the difference was small ($M = 0.10$; $SE = 0.003$ vs. $M = 0.12$; $SE = 0.003$).

Finally, the interaction effect between $\hat{\beta}_{imp}$ Alignment and $\hat{\beta}_{exp}$ Alignment was significant ($p = 0.036$). A Tukey HSD post-hoc test indicated that the error was significantly higher when $\hat{\beta}_{imp}$ and $\hat{\beta}_{exp}$ were both misaligned. Also, when $\hat{\beta}_{imp}$ was misaligned but $\hat{\beta}_{exp}$ was not, the error was significantly higher than when $\hat{\beta}_{imp}$ was aligned (whether $\hat{\beta}_{exp} = 1$ or $\hat{\beta}_{exp} = 10$).

Overall, these results reinforce findings for β with RQ1, and are consistent with prior work showing that having incorrect assumptions about rationality harms performance [66].

7.5 Real-World Evaluation

Having validated our approach in simulation, we conducted a real-world evaluation with 21 people. Each person collaborated on the pizza-making task with a robot, as illustrated in Fig. 7.1, while the robot tried to estimate their preferences for the collaboration. Through the real-world demonstration, we investigated:

(RQ3) *How well can the robot learn preferences over collaborations in real-world human-robot interactions?* The main challenge in this setup is learning from realistic human feedback, which may be noisy and sparse in more complicated ways than modeled in Section 7.4.

7.5.1 Experimental Protocol

The real-world evaluation was approved by our local Institutional Review Board. An experimental session typically lasted 45 min. Participants were compensated US\$15 and collaborated with the robot in the same pizza-making task from Sec. 7.4.1.

Experimental Setup. As shown in Fig. 7.1, the participants interacted with a robot system comprising two robots: a Franka Emika Panda arm, and a table-top robot called Shutter [326]. The Panda executed pick and place actions planned within the MoveIt Task Constructor framework [127]. During interactions, Shutter engaged with participants through its gaze and speech. Following Candon et al. [57], Shutter occasionally reminded participants to give explicit feedback. The reminders were framed as helping the robot to improve as a teammate (e.g., *“Remember that you can give me feedback so we can collaborate better in the future”*), and were only issued before the robot picked an object. High-level, multimodal robot behavior was controlled with behavior trees[81].

The table served as the collaborative workspace with defined storage and hand-off areas for the pizza ingredients. Each pizza ingredient was stored in a clear plastic container that could be grasped by the Panda hand. The table included two illuminated buttons that participants could press to give explicit feedback.

Procedure. The participants consented to participate in the interaction and to be audio- and video-recorded. The experimenter explained the goal of the interaction, introduced the robot, and started a tutorial. During the tutorial, the robot explained roles, the workstation, and how the person could provide explicit feedback. The participant then constructed a simple, practice pizza with a basic preference to see how preferences influenced the pizza-making interaction. The experimenter finally explained the set of preferences to choose from, had the participant select a preference, and went through hypothetical scenarios to ensure the participant understood their preferences. Each participant then worked with the robot to construct three different pizzas.

Participants. We recruited 21 participants 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. Sixteen of the evaluation participants (76%) were undergraduate or graduate students.

Evaluation. We conducted offline preference learning on the data collected from the 21 participants. We evaluated learning via L_2 error and conflict percentage, as in Sec. 7.4.3. For each participant, we first fit $\hat{\beta}_{imp}$ and $\hat{\beta}_{exp}$ using the data from the practice pizza, as an individual calibration step. We then used the fitted values for the three pizzas in the participant's interaction.

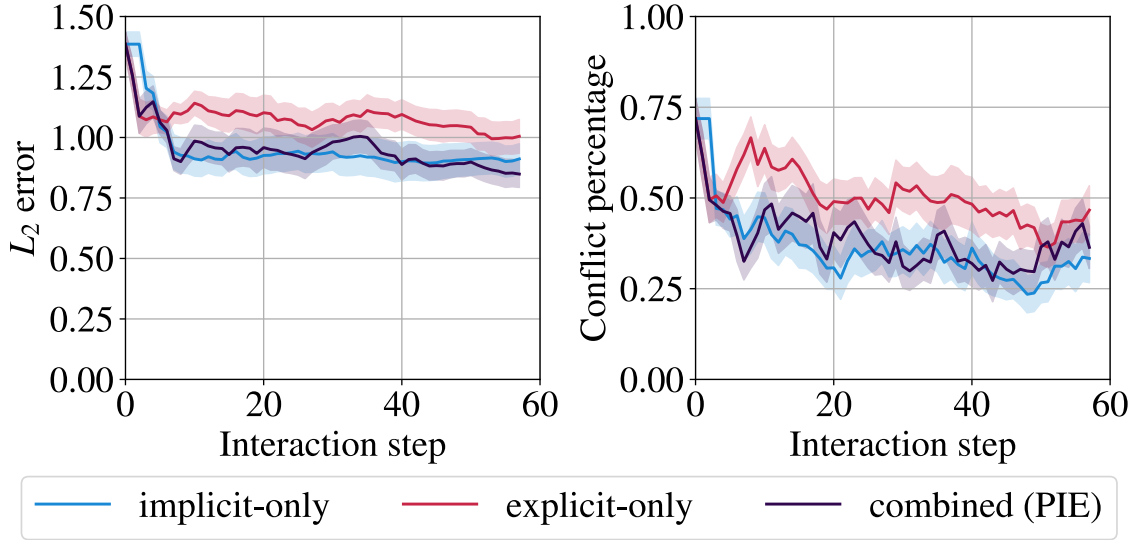


Figure 7.5: L_2 error with respect to the ground truth preference weights (left) and conflict percentage (right) as the interaction progresses. Lines represent average L_2 error and conflict percentage, and shading represents the standard error across the 21 participants in the real-world evaluation.

7.5.2 Results

Figure 7.5 shows results from 21 participants each making three pizzas with the robot. We analyzed the final L_2 error with a linear mixed model that considered Participant ID as random effect and Feedback Modality (*Explicit*, *Implicit*, or *Combined* with PIE) as main effect. We found a trend for Feedback Modality having an effect on the L_2 error: $F(2, 40) = 2.77, p = 0.07$. Examination of the model parameter estimates revealed that the *Explicit* condition was the only modality significantly different from the Grand Mean (Estimate = 0.083, $t(40) = 2.16, p = 0.037$, 95% CI [0.005, 0.161]), indicating significantly higher error than the overall average. Final L_2 errors were: *Implicit* ($M = 0.911, SE = 0.062$), *Explicit* ($M = 1.005, SE = 0.070$), and *Combined* feedback with PIE ($M = 0.849, SE = 0.055$). The conflict percentage results were similar, with $M = 0.333$ ($SE = 0.067$) for *Implicit* feedback, $M = 0.466$ ($SE = 0.067$) for *Explicit* feedback, and $M = 0.364$ ($SE = 0.058$) for *Combined* feedback.

7.6 Discussion

Our PIE approach outperforms single-modality baselines, achieving lower L_2 error and fewer conflicts in simulation. Even small reductions in conflict can meaningfully improve both task performance and human perception of the robot. Real-world results show a similar trend: leveraging multiple, naturally occurring feedback signals enhances a robot’s ability to infer human preferences, highlighting the value of richer feedback in human-robot collaboration. However, as the omnibus test did not reach conventional statistical significance and no post hoc comparisons were conducted to directly compare condition means, these findings should be interpreted with caution. We see our modest findings as an opportunity for future work, as they highlight the importance of incorporating accurate assumptions when reasoning about feedback. Future work could explore jointly learning preference weights and individualized betas [120] for assumptions about how rationally the human is providing feedback (β_m in eq. (7.8)) or addressing modality-specific effects on gradient updates in eq. (7.9).

We opted for a myopic objective so that we could reason about human feedback in relation to high-level actions with dense rewards. This limits applicability in sparse-reward settings, where reasoning over longer horizons would be beneficial. This direction would require reasoning about the gradients of value functions, as suggested in [174] for preference learning from multiple feedback.

Our work is also limited by implementation choices and assumptions made by PIE. For example, we only considered binary evaluative feedback via button presses and human task actions. Future efforts could integrate other feedback signals (e.g., corrections or language [174]). Also, future work could explore relaxing the assumption that the robot knows which features matter to humans, and instead aim to learn them [39].

By learning from naturally occurring, multimodal feedback, PIE moves toward more seamless, adaptive human-robot interactions, reducing the teaching burden on humans and fostering intuitive, productive collaboration.

Chapter 8

Conclusion

This dissertation advocates for a re-conceptualization of social context, moving it from a background variable to an active mechanism for robot learning. A distinguishing characteristic of robot learning in human-robot collaborations, as compared to traditional machine learning, is the dynamic social context. In a collaboration, the robot is not limited to learning from a fixed dataset. Rather, the robot can take an active role in interpreting and shaping the social context of the interaction. This approach facilitates the extraction of informative signals that may otherwise be ignored. By leveraging this feedback, a robot can adapt its behavior to the preferences of its collaborator in real-time.

We presented a novel definition of the social context of a human-robot interaction. We advance the field by demonstrating how a robot can use the social context to better learn during collaborations with non-expert users.

8.1 Contributions

The main contributions of this dissertation are as follows:

- **Formalization of the Social Context of a Human-Robot Interaction.** After surveying related literature in the field of HRI, we propose a novel definition of the “social context of a human-robot interaction” in Chapter 3. We provide a formal taxonomy and discuss how robot designers can use this tool before, during, and after interactions to facilitate better robot behavior.
- **Evidence of Robots Influencing People to Provide Feedback.** In Chapter 4, we studied the effects of framing and timing of reminders when a robot is asking a human to provide more feedback. To our knowledge, this was the first study to demonstrate that robots can influence both the feedback provided by their collaborators and their collaborators’ feelings about providing feedback.
- **Real World Examples of Robots Learning from Implicit Feedback Alongside Explicit Feedback in Human-Robot Collaborations.** We contributed two studies in which we demonstrated that robots can learn from implicit signals during collaborations without adding any additional burden to the user. First, we built models that could predict human preferences between agent behaviors considering nonverbal communicative signals via features derived from facial reactions, eye gaze, and head position (Chapter 5). Second, we show that a robot can learn human preferences from human task actions over a different action space than the robot’s actions (Chapter 7).
- **Dataset for the Analysis of Human Explicit and Implicit Feedback Together.** To facilitate further research on combining implicit and explicit human feedback, we developed the REACTdatabase (Chapter 6). This database contains more than 1,000 minutes of data on users’ natural reactions to and feedback provided to robots during a collaborative game and a photography scenario.
- **Method for Combining Implicit and Explicit Feedback Together for Human Preference Learning.** We propose Preference learning from Implicit and Explicit

feedback (PIE), a method for combining implicit and explicit human feedback to learn human preferences (Chapter 7). Our simulation results show that a robot can learn faster from both types of feedback together, without requiring anything extra of the user. Real world results show a similar trend, but further research is necessary.

8.2 Open Questions and Opportunities for Future Research

The work presented in this dissertation takes a step towards robots that are better able to learn during collaborations with non-expert users. In pursuit of this goal, there are many opportunities to advance this line of research.

8.2.1 Deeper Understanding of When Humans Choose to Give Feedback to Robots

While one of the motivating factors of this dissertation is that humans are hard to simulate, attempting to simulate a real person is inevitable in a robot learning dissertation. Trying to simulate *when* a person provides feedback highlighted this gap in current understanding. Many existing works facilitate simulating humans providing feedback. For example, there are ways to incorporate assumptions about how often a teacher makes a mistake or skips a query [204]. However, these works are typically in settings where the human is only teaching or it is a turn-taking task (i.e., there are clear “timesteps” when the human should be providing feedback).

In real-world collaborations, it is unrealistic for the default assumption to be that humans are constantly providing feedback during continuous tasks. A conceptual model to understand *how* and *when* a person chooses to provide explicit feedback could facilitate

better robot learning. In particular, it should be investigated how human expectations about the robot’s policy influence when they decide to provide feedback and what kind of feedback [225].

8.2.2 Evaluation of Feedback Quality

In a similar vein, it is not trivial to evaluate the quality of provided feedback. While Chapter 4 demonstrates that robots can influence the frequency and latency of human feedback, the *utility* or *correctness* of that feedback remains an open question. In a continuous collaboration, the lack of discrete “queries” makes it difficult to objectively measure the quality of provided feedback. Future work could explore a standardized definition of “feedback quality”. Frameworks grounded in information theory and cognitive science could help address this gap.

8.2.3 Broader Signals of Implicit Human Feedback

This dissertation explores facial reactions, eye gaze, head pose, and physical task actions as implicit signals from which a robot could learn. It would be fruitful to explore other modalities of implicit feedback. For example, gestures and proxemics could be rich signals for implicit human feedback.

Additionally, facial reactions, eye gaze, and head pose were only explored at the level of summary statistics in Chapter 5. A more nuanced view, especially considering the temporal nature of implicit feedback, could result in much stronger learning signals.

8.2.4 Verbal Feedback

With the success and prevalence of LLMs, we would be remiss to not encourage researchers to explore how to incorporate verbal feedback into the methods presented herein. It is a rich source of both explicit and implicit feedback. It is not clear how to segment verbal feedback into explicit versus implicit feedback categories. For example, a person could say “good job” and mean it, or they could say it sarcastically, requiring interpretation to realize this is actually negative feedback. Properties of verbal feedback such as prosody, volume, and speed are all part of the social context of the interaction.

8.2.5 Differentiated Interpretation of Different Types of Feedback

In the proposed PIE method (Chapter 7), implicit and explicit human feedback had implications, but the feedback instances were all treated the same. There is an opportunity to explore how to incorporate confidence in different feedback signals into the robot’s reasoning process. Related questions include how should a robot handle conflicting feedback, and can consistent instances of feedback reinforce each other.

8.2.6 Other Ways to Solicit Feedback

Chapter 4 validated that robots can influence their collaborators through targeted reminders to provide feedback. There is a rich opportunity for future work to explore other ways a robot can influence a person to provide more feedback. Can a social robot use eye gaze and proxies for attention (e.g., pointing) to solicit feedback in times of uncertainty? Would thanking a person for providing feedback lead to a person providing more feedback? Research has explored how transparency might influence robot learning – how can this be

viewed through the lens of social context.

8.2.7 Robot Learning During Group Human-Robot Collaborations

A critical extension of leveraging the social context in robot learning is moving from dyadic settings to multi-human group interactions. Robots have been shown to influence group dynamics in human-human interactions (e.g., [122, 332, 197, 296]). However, the interplay between multi-person explicit and implicit human feedback remains largely unexplored.

In multi-human group interactions, it is reasonable to expect that there could be instances of conflicting feedback. Consider a scenario where the robot hands Alice a pizza topping and Alice smiles. Bob then presses the “bad job” button, after which Alice looks saddened. How can a robot resolve these conflicting signals? Can a robot learn from signals between Alice and Bob?

8.2.8 Summary

In sum, this dissertation advances the field of HRI by highlighting the dynamic social context of human-robot interactions as a rich tool to facilitate better robot learning during collaborations from non-expert users.

Bibliography

- [1] Piotr D. Adamczyk and Brian P. Bailey. 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, page 271–278, New York, NY, USA, 2004. Association for Computing Machinery. ISBN 1581137028. doi: 10.1145/985692.985727. URL <https://doi.org/10.1145/985692.985727>. 49
- [2] Timothy Adamson, C Burton Lyng-Olsen, Kendrick Umstattd, and Marynel Vázquez. Designing social interactions with a humorous robot photographer. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pages 233–241, 2020. 104
- [3] Henny Admoni and Siddhartha Srinivasa. Predicting user intent through eye gaze for shared autonomy. In *2016 AAAI Fall Symposium Series*, 2016. 10
- [4] Muneeb Imtiaz Ahmad, Omar Mubin, and Joanne Orlando. A systematic review of adaptivity in human-robot interaction. *Multimodal Technologies and Interaction*, 1(3):14, 2017. 9, 49
- [5] Gopika Ajaykumar, Maureen Steele, and Chien-Ming Huang. A survey on end-user robot programming. *ACM Computing Surveys (CSUR)*, 54(8):1–36, 2021. 8
- [6] Gopika Ajaykumar, Maia Stiber, and Chien-Ming Huang. Designing user-centric programming aids for kinesthetic teaching of collaborative robots. *Robotics and Autonomous Systems*, 145:103845, 2021. 72
- [7] Merihan Alhafnawi, Edmund R Hunt, Severin Lemaignan, Paul O’Dowd, and Sabine Hauert. Mosaix: a swarm of robot tiles for social human-swarm interaction. *Int’l Conf. on Robotics and Automation (ICRA)*, 2022. 22, 37, 39
- [8] Batuhan Altundas, Zheyuan Wang, Joshua Bishop, and Matthew Gombolay. Learning coordination policies over heterogeneous graphs for human-robot teams via recurrent neural schedule propagation. *2022 IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems (IROS)*, pages 11679–11686, 2022. 43

- [9] Amir Aly and Adriana Tapus. Towards an online fuzzy modeling for human internal states detection. *2012 12th Int'l Conf. on Control Automation Robotics & Vision*, pages 1563–1570, 2012. 44
- [10] Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. Power to the people: The role of humans in interactive machine learning. *AI magazine*, 35(4):105–120, 2014. 9
- [11] Elisabeth André and Catherine Pelachaud. Interacting with embodied conversational agents. In *Speech technology*, pages 123–149. Springer, 2010. 10
- [12] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and autonomous systems*, 57(5):469–483, 2009. 9
- [13] Brenna D. Argall, Eric L. Sauser, and Aude G. Billard. Tactile guidance for policy adaptation. *Foundations and Trends® in Robotics*, 1(2):79–133, 2011. ISSN 1935-8253. doi: 10.1561/23000000012. URL <http://dx.doi.org/10.1561/23000000012>. 48
- [14] Michael Argyle, Adrian Furnham, and Jean Ann Graham. *Social situations*. Cambridge University Press, 1981. 16
- [15] Thomas Arnold and Matthias Scheutz. Beyond moral dilemmas: exploring the ethical landscape in hri. *Proc. of the 2017 ACM/IEEE Int'l Conf. on Human-Robot Interaction*, pages 445–452, 2017. 22, 35, 36, 38, 40
- [16] Thomas Arnold and Matthias Scheutz. Observing robot touch in context: How does touch and attitude affect perceptions of a robot’s social qualities? *Proc. of the 2018 ACM/IEEE Int'l Conf. on Human-Robot Interaction*, 2018. 22
- [17] Elliot Aronson, Timoty D Wilson, Robin M Akert, and SR Sommers. *Social psychology 9e*. England: Pearson, 2016. 15
- [18] Inbal Avraham and Reuth Mirsky. Shared control with black box agents using oracle queries. In *2025 IEEE International Conference on AI and Data Analytics (ICAD)*, pages 1–8. IEEE, 2025. 112
- [19] Andrea Bajcsy, Dylan P. Losey, Marcia K. O’Malley, and Anca D. Dragan. Learning robot objectives from physical human interaction. In *Proceedings of the 1st Annual Conference on Robot Learning*, volume 78 of *Proceedings of Machine Learning Research*, pages 217–226. PMLR, 13–15 Nov 2017. 9
- [20] Chris L Baker, Rebecca Saxe, and Joshua B Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009. 117

- [21] Tadas Baltrusaitis, Amir Zadeh, Yao Chong Lim, and Louis-Philippe Morency. Openface 2.0: Facial behavior analysis toolkit. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 59–66. IEEE, 2018. xxiii, 84, 102
- [22] Santosh Balajee Banisetty, Vineeth Rajamohan, Fausto Vega, and David Feil-Seifer. A deep learning approach to multi-context socially-aware navigation. *IEEE Int’l Conf. on Robot & Human Interactive Communication (RO-MAN)*, 2021. 15, 22, 36, 39, 45
- [23] Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh PN Rao, and Minoru Asada. Efficient human-robot collaboration: When should a robot take initiative? *The International Journal of Robotics Research*, 36(5-7):563–579, 2017. 50
- [24] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M Martinez, and Seth D Pollak. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psych. science in the public interest*, 20(1), 2019. 11, 77, 93
- [25] Samuel Barrett, Peter Stone, Sarit Kraus, and Avi Rosenfeld. Teamwork with limited knowledge of teammates. In *Proceedings of the AAAI conference on artificial intelligence*, volume 27:1, pages 102–108, 2013. 112
- [26] Christoph Bartneck, Tony Belpaeme, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Šabanović. *Human-robot interaction: An introduction*. Cambridge University Press, 2020. 15, 25, 75, 82
- [27] Andrew G Barto. Reinforcement learning: An introduction. by richard’s sutton. *SIAM Rev*, 6(2):423, 2021. 118
- [28] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018. 43
- [29] Richard E Bellman. *Dynamic programming*. Princeton university press, 2010. 93
- [30] Tony Belpaeme, James Kennedy, Aditi Ramachandran, Brian Scassellati, and Fumihide Tanaka. Social robots for education: A review. *Science robotics*, 3(21), 2018. 41
- [31] Atef Ben-Youssef, Chloé Clavel, Slim Essid, Miriam Bilac, Marine Chamoux, and Angelica Lim. Ue-hri: A new dataset for the study of user engagement in spontaneous human-robot interactions. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction, ICMI ’17*, page

- 464–472, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450355438. doi: 10.1145/3136755.3136814. URL <https://doi.org/10.1145/3136755.3136814>. 99, 100
- [32] Kush Bhatia, Ashwin Pananjady, Peter Bartlett, Anca Dragan, and Martin J Wainwright. Preference learning along multiple criteria: A game-theoretic perspective. *Advances in Neural Information Processing Systems*, 33:7413–7424, 2020. 77
- [33] Gautam Bhattacharya, Koushik Ghosh, and Ananda S Chowdhury. knn classification with an outlier informative distance measure. In *International Conference on Pattern Recognition and Machine Intelligence*, pages 21–27. Springer, 2017. 88
- [34] Adam Bignold, Francisco Cruz, Matthew E. Taylor, Tim Brys, Richard Dazeley, Peter Vamplew, and Cameron Foale. A conceptual framework for externally-influenced agents: An assisted reinforcement learning review. *Journal of Ambient Intelligence and Humanized Computing*, pages 1–24, 2021. 49
- [35] Erdem Biyik and Dorsa Sadigh. Batch active preference-based learning of reward functions. In Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto, editors, *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, pages 519–528. PMLR, 29–31 Oct 2018. URL <https://proceedings.mlr.press/v87/biyik18a.html>. 77
- [36] Erdem Biyik, Malayandi Palan, Nicholas C. Landolfi, Dylan P. Losey, and Dorsa Sadigh. Asking easy questions: A user-friendly approach to active reward learning. In *Proceedings of the 3rd Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pages 1177–1190. PMLR, 2019. 13
- [37] Erdem Biyik, Dylan P Losey, Malayandi Palan, Nicholas C Landolfi, Gleb Shevchuk, and Dorsa Sadigh. Learning reward functions from diverse sources of human feedback: Optimally integrating demonstrations and preferences. *The International Journal of Robotics Research*, 41(1):45–67, 2022. 13
- [38] Erdem Biyik, Aditi Talati, and Dorsa Sadigh. Aprel: A library for active preference-based reward learning algorithms. In *Proceedings of HRI*. IEEE, 2022. 2, 8
- [39] Andreea Bobu, Andi Peng, Pulkit Agrawal, Julie A Shah, and Anca D. Dragan. Aligning human and robot representations. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '24, pages 42–54, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400703225. doi: 10.1145/3610977.3634987. URL <https://doi.org/10.1145/3610977.3634987>. 132

- [40] Dan Bohus and Eric Horvitz. *Situated interaction*, page 105–143. Association for Computing Machinery and Morgan & Claypool, 2019. ISBN 9781970001754. URL <https://doi.org/10.1145/3233795.3233800>. 15
- [41] Adriana Bono, Agnese Augello, Manuel Gentile, and Salvatore Gaglio. Social practices based characters in a robotic storytelling system. *IEEE Int’l Conf. on Robot & Human Interactive Communic. (RO-MAN)*, 2020. 22
- [42] Jeffrey M Bradshaw, Paul J Feltovich, and Matthew Johnson. Human–agent interaction. In *The handbook of human-machine interaction*, pages 283–300. CRC Press, 2017. 8
- [43] Jake Brawer, Debasmita Ghose, Kate Candon, Meiyang Qin, Alessandro Roncone, Marynel Vázquez, and Brian Scassellati. Interactive policy shaping for human-robot collaboration with transparent matrix overlays. In *Proceedings of HRI*, 2023. 110
- [44] Leo Breiman, Jerome Friedman, Richard Olshen, and Charles Stone. Classification and regression trees. wadsworth int. *Group*, 37(15):237–251, 1984. 83, 91
- [45] Alexandra Bremers, Maria Teresa Parreira, Xuanyu Fang, Natalie Friedman, Adolfo Ramirez-Aristizabal, Alexandria Pabst, Mirjana Spasojevic, Michael Kuniavsky, and Wendy Ju. The bystander affect detection (bad) dataset for failure detection in hri. *arXiv preprint arXiv:2303.04835*, 2023. 99, 100
- [46] Gordon Briggs and Matthias Scheutz. Strategies and mechanisms to enable dialogue agents to respond appropriately to indirect speech acts. *2017 26th IEEE Int’l Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 323–328, 2017. 22, 35, 36, 39, 40
- [47] Gordon Briggs, Bryce Gessell, Matt Dunlap, and Matthias Scheutz. Actions speak louder than looks: Does robot appearance affect human reactions to robot protest and distress? *The 23rd IEEE international symposium on robot and human interactive communication*, pages 1122–1127, 2014. 22, 31, 38, 40
- [48] Andrew G Brooks and Ronald C Arkin. Behavioral overlays for non-verbal communication expression on a humanoid robot. *Autonomous robots*, 22:55–74, 2007. 12
- [49] Maya Cakmak and Andrea L. Thomaz. Designing robot learners that ask good questions. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI ’12*, page 17–24, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450310635. doi: 10.1145/2157689.2157693. URL <https://doi.org/10.1145/2157689.2157693>. 13, 48, 49

- [50] Maya Cakmak, Crystal Chao, and Andrea L. Thomaz. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development*, 2(2):108–118, 2010. doi: 10.1109/TAMD.2010.2051030. 13
- [51] Maya Cakmak, Siddhartha S. Srinivasa, Min Kyung Lee, Sara Kiesler, and Jodi Forlizzi. Using spatial and temporal contrast for fluent robot-human handovers. In *Proceedings of the 6th International Conference on Human-Robot Interaction*, HRI '11, page 489–496, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450305617. doi: 10.1145/1957656.1957823. URL <https://doi.org/10.1145/1957656.1957823>. 50
- [52] Sylvain Calinon and Aude Billard. Incremental learning of gestures by imitation in a humanoid robot. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, HRI '07, page 255–262, New York, NY, USA, 2007. Association for Computing Machinery. ISBN 9781595936172. doi: 10.1145/1228716.1228751. URL <https://doi.org/10.1145/1228716.1228751>. 48
- [53] Joseph Campbell and Katsu Yamane. Learning whole-body human-robot haptic interaction in social contexts. *2020 IEEE Int'l Conf. on Robotics and Automation (ICRA)*, pages 10177–10183, 2020. 22, 36, 37
- [54] Joana Campos, Patrícia Alves-Oliveira, and Ana Paiva. Looking for conflict: gaze dynamics in a dyadic mixed-motive game. *Autonomous Agents and Multi-Agent Systems*, 30(1):112–135, 2016. 26
- [55] Kate Candon, Zoe Hsu, Yoony Kim, Jesse Chen, Nathan Tsoi, and Marynel Vázquez. Perceptions of the helpfulness of unexpected agent assistance. In *Proceedings of the 10th International Conference on Human-Agent Interaction*, pages 41–50, 2022. 51, 78
- [56] Kate Candon, Jesse Chen, Yoony Kim, Zoe Hsu, Nathan Tsoi, and Marynel Vázquez. Nonverbal human signals can help autonomous agents infer human preferences for their behavior. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, pages 307–316, 2023. 6, 11, 74
- [57] Kate Candon, Helen Zhou, Sarah Gillet, and Marynel Vázquez. 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, page 290–300, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450399647. doi: 10.1145/3568162.3576980. URL <https://doi.org/10.1145/3568162.3576980>. 5, 47, 98, 117, 122, 129

- [58] Kate Candon, Nicholas C. Georgiou, Helen Zhou, Sidney Richardson, Qiping Zhang, Brian Scassellati, and Marynel Vázquez. React: Two datasets for analyzing both human reactions and evaluative feedback to robots over time. In *Proceedings of HRI, 2024*. 6, 97
- [59] Kate Candon, Qiping Zhang, Alexander Lew, Houston Claire, Lena Qian, Alyssa Quarles, Chayan Sarkar, and Marynel Vázquez. Learning human preferences over a human-robot collaboration based on explicit and implicit human feedback. In *Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction, HRI '26*, page 1040–1049, New York, NY, USA, 2026. Association for Computing Machinery. ISBN 9798400721281. doi: 10.1145/3757279.3785630. URL <https://doi.org/10.1145/3757279.3785630>. 7, 109
- [60] Colleen M Carpinella, Alisa B Wyman, Michael A Perez, and Steven J Stroessner. The robotic social attributes scale (rosas) development and validation. In *Proceedings of the 2017 ACM/IEEE International Conference on human-robot interaction*, pages 254–262, 2017. 83
- [61] Cristiano Castelfranchi. Modelling social action for ai agents. *Artificial intelligence*, 103(1-2):157–182, 1998. 8
- [62] Luca Castri, Sariah Mghames, Marc Hanheide, and Nicola Bellotto. Causal discovery of dynamic models for predicting human spatial interactions. *Int'l Conf. on Social Robotics*, pages 154–164, 2022. 46
- [63] Luca Castri, Sariah Mghames, Marc Hanheide, and Nicola Bellotto. Enhancing causal discovery from robot sensor data in dynamic scenarios. *Conf. on Causal Learning and Reasoning*, 2023. 46
- [64] Oya Celiktutan, Efstratios Skordos, and Hatice Gunes. Multimodal human-human-robot interactions (mhhri) dataset for studying personality and engagement. *IEEE Transactions on Affective Computing*, 10(4):484–497, 2017. 10
- [65] Joyce Y Chai, Maya Cakmak, Candace Sidner, and J Lupp. 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*, volume 26, 2017. 49
- [66] Lawrence Chan, Andrew Critch, and Anca Dragan. Human irrationality: both bad and good for reward inference. *arXiv preprint arXiv:2111.06956*, 2021. 128
- [67] Wan-Ling Chang and Selma Šabanović. Interaction expands function: Social shaping of the therapeutic robot paro in a nursing home. *Proc. of ACM/IEEE Int'l Conf. on Human-Robot Interaction*, 2015. 22, 23, 35, 36, 37, 39

- [68] Wan-Ling Chang, Selma Šabanovic, and Lesa Huber. Observational study of naturalistic interactions with the socially assistive robot paro in a nursing home. *The 23rd IEEE Int'l Symposium on Robot and Human Interactive Communication*, pages 294–299, 2014. 22, 35, 37, 39, 40
- [69] Min Chen, Stefanos Nikolaidis, Harold Soh, David Hsu, and Siddhartha Srinivasa. Planning with trust for human-robot collaboration. In *Proceedings of the 2018 ACM/IEEE international conference on human-robot interaction*, pages 307–315, 2018. 112
- [70] Sonia Chernova and Andrea L Thomaz. *Robot learning from human teachers*. Springer Nature, 2022. 44
- [71] Mohamed Chetouani. Interactive robot learning: an overview. *ECCAI Advanced Course on Artificial Intelligence*, pages 140–172, 2021. 44
- [72] Vivienne Bihe Chi and Bertram F. Malle. 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*, page 718–722. IEEE Press, 2022. 12
- [73] Susan EF Chipman. *The Oxford handbook of cognitive science*. Oxford University Press, 2016. 120
- [74] Meia Chita-Tegmark, Monika Lohani, and Matthias Scheutz. Gender effects in perceptions of robots and humans with varying emotional intelligence. *ACM/IEEE Int'l Conf. on Human-Robot Interaction*, 2019. 22
- [75] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30, 2017. 9, 113
- [76] Yun-Shiuan Chuang, Hsin-Yi Hung, Edwinn Gamborino, Joshua Oon Soo Goh, Tsung-Ren Huang, Yu-Ling Chang, Su-Ling Yeh, and Li-Chen Fu. Using machine theory of mind to learn agent social network structures from observed interactive behaviors with targets. *IEEE Int'l Conf. on Robot and Human Interactive Communication (RO-MAN)*, 2020. 22
- [77] Nikhil Churamani, Sinan Kalkan, and Hatice Gunes. Continual learning for affective robotics: Why, what and how? *2020 IEEE Int'l Conf. on Robot and Human Interactive Comm. (RO-MAN)*, 2020. 9, 44
- [78] Nikhil Churamani, Pablo Barros, Hatice Gunes, and Stefan Wermter. Affect-driven learning of robot behaviour for collaborative human-robot interactions. *Frontiers in Robotics and AI*, 9:717193, 2022. 44

- [79] F. Cini, T. Banfi, G. Ciuti, L. Craighero, and M. Controzzi. The relevance of signal timing in human-robot collaborative manipulation. *Science Robotics*, 6(58):eabg1308, 2021. doi: 10.1126/scirobotics.abg1308. URL <https://www.science.org/doi/abs/10.1126/scirobotics.abg1308>. 50
- [80] Caitlyn Clabaugh and Maja J. Matarić. 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*, pages 968–973. IEEE Press, 2016. doi: 10.1109/ROMAN.2016.7745226. 9
- [81] Michele Colledanchise and Petter Ögren. Behavior trees in robotics and ai, 2018. 129
- [82] Joe Connolly, Viola Mocz, Nicole Salomons, Joseph Valdez, Nathan Tsoi, Brian Scassellati, and Marynel Vázquez. Prompting prosocial human interventions in response to robot mistreatment. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, HRI '20*, page 211–220, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450367462. doi: 10.1145/3319502.3374781. URL <https://doi.org/10.1145/3319502.3374781>. 47, 50
- [83] Filipa Correia, Samuel Mascarenhas, Rui Prada, Francisco S. Melo, and Ana Paiva. Group-based emotions in teams of humans and robots. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, HRI '18*, page 261–269, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450349536. doi: 10.1145/3171221.3171252. URL <https://doi.org/10.1145/3171221.3171252>. 56, 70
- [84] Adam K Coyne, Keshav Sapkota, and Conor McGinn. “who said that?” applying the situation awareness global assessment technique to social telepresence. *ACM Trans. on Human-Robot Interaction*, 12(4):1–23, 2023. 22
- [85] Andrea Cuadra, Hansol Lee, Jason Cho, and Wendy Ju. Look at me when i talk to you: A video dataset to enable voice assistants to recognize errors. *arXiv preprint arXiv:2104.07153*, 2021. 10
- [86] Andrea Cuadra, Shuran Li, Hansol Lee, Jason Cho, and Wendy Ju. My bad! repairing intelligent voice assistant errors improves interaction. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1):1–24, 2021. 82
- [87] Yuchen Cui, Qiping Zhang, Alessandro Allievi, Peter Stone, Scott Niekum, and W. Bradley Knox. The empathic framework for task learning from implicit human feedback. In *Proceedings of the 4th Conference on Robot Learning (CoRL 2020)*, Cambridge MA, USA, November 2020. URL <http://www.cs.utexas.edu/users/ai-lab?CORL20-Cui>. 76

- [88] Yuchen Cui, Pallavi Koppol, Henny Admoni, Scott Niekum, Reid Simmons, Aaron Steinfeld, and Tesca Fitzgerald. Understanding the relationship between interactions and outcomes in human-in-the-loop machine learning. In *Proceedings of IJCAI*, 2021. doi: 10.24963/ijcai.2021/599. URL <https://doi.org/10.24963/ijcai.2021/599>. 9, 49
- [89] Yuchen Cui, Qiping Zhang, Brad Knox, Alessandro Allievi, Peter Stone, and Scott Niekum. The empathic framework for task learning from implicit human feedback. *Conf. on Robot Learning*, pages 604–626, 2021. 10, 44, 98, 99, 100
- [90] Sanobar Dar and Ulysses Bernardet. When agents become partners: A review of the role the implicit plays in the interaction with artificial social agents. *Multimodal Technologies and Interaction*, 4(4), 2020. ISSN 2414-4088. doi: 10.3390/mti4040081. URL <https://www.mdpi.com/2414-4088/4/4/81>. 10
- [91] Manoranjan Dash and Huan Liu. Feature selection for classification. *Intelligent data analysis*, 1(1-4):131–156, 1997. 93
- [92] Celso M de Melo, Peter J Carnevale, and Jonathan Gratch. Peoples biased decisions to trust and cooperate with agents that express emotions. Technical report, UNIVERSITY OF SOUTHERN CALIFORNIA LOS ANGELES, 2018. 10
- [93] Celso M. de Melo, Peter Khooshabeh, Ori Amir, and Jonathan Gratch. Shaping cooperation between humans and agents with emotion expressions and framing. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '18, page 2224–2226, Richland, SC, 2018. International Foundation for Autonomous Agents and Multiagent Systems. 10
- [94] Anind K Dey. Understanding and using context. *Personal and ubiquitous computing*, 5:4–7, 2001. 17, 26
- [95] Anna Dobrosovostnova, Jesse de Pagter, and Astrid Weiss. Borrowing, poking and entangling. in search of shared spaces between science and technology studies and human-robot interaction. *Companion of the 2023 ACM/IEEE Int'l Conf. on Human-Robot Interaction*, pages 21–29, 2023. 22
- [96] Carmel Domshlak, Eyke Hüllermeier, Souhila Kaci, and Henri Prade. Preferences in ai: An overview. *Artificial Intelligence*, 175(7-8):1037–1052, 2011. 77
- [97] Anca D. Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S. Srinivasa. Effects of robot motion on human-robot collaboration. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '15, page 51–58, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450328838. doi: 10.1145/2696454.2696473. URL <https://doi.org/10.1145/2696454.2696473>. 48

- [98] Paolo Dragone, Stefano Teso, and Andrea Passerini. Constructive preference elicitation. *Frontiers in Robotics and AI*, 4:71, 2018. 9
- [99] Jiali Duan, Qian Wang, Lerrel Pinto, C-C Jay Kuo, and Stefanos Nikolaidis. Robot learning via human adversarial games. *2019 IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, pages 1056–1063, 2019. 26
- [100] Wen Duan, Naomi Yamashita, Yoshinari Shirai, and Susan R Fussell. Bridging fluency disparity between native and nonnative speakers in multilingual multiparty collaboration using a clarification agent. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2):1–31, 2021. 82
- [101] Brian R Duffy and Karolina Zawieska. Suspension of disbelief in social robotics. *2012 IEEE RO-MAN: The 21st IEEE Int'l Symposium on Robot and Human Interactive Communication*, 2012. 22
- [102] Filip Edström, Thomas Hellström, and Xavier De Luna. Robot causal discovery aided by human interaction. *IEEE Int'l Conf. on Robot and Human Interactive Communication (RO-MAN)*, 2023. 46
- [103] Hadas Erel, Yoav Cohen, Klil Shafir, Sara Daniela Levy, Idan Dov Vidra, Tzachi Shem Tov, and Oren Zuckerman. Excluded by robots: Can robot-robot-human interaction lead to ostracism? *Proc. of the 2021 ACM/IEEE Int'l Conf. on Human-Robot Interaction*, pages 312–321, 2021. 22, 36, 37, 39, 40
- [104] Hadas Erel, Elier Carsenti, and Oren Zuckerman. A carryover effect in hri: Beyond direct social effects in human-robot interaction. *ACM/IEEE Int'l Conf. on Human-Robot Interaction (HRI)*, 2022. 22, 37, 39, 40, 50
- [105] Hadas Erel, Marynel Vázquez, Sarah Sebo, Nicole Salomons, Sarah Gillet, and Brian Scassellati. Rosi: A model for predicting robot social influence. *J. Hum.-Robot Interact.*, 13(2), June 2024. doi: 10.1145/3641515. URL <https://doi.org/10.1145/3641515>. 15, 25
- [106] Connor Esterwood, Kyle Essenmacher, Han Yang, Fanpan Zeng, and Lionel P. Robert. Birds of a feather flock together: But do humans and robots? a meta-analysis of human and robot personality matching. *IEEE Int'l Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2021. 22, 37, 39
- [107] Connor Esterwood, Ruijian Hannah Guan, Xin Ye, and Lionel P Robert. Virtually the same or realistically different?: A meta-analysis of real vs.'not so real'robots. *2025 20th ACM/IEEE Int'l Conf. on Human-Robot Interaction (HRI)*, pages 559–568, 2025. 42
- [108] Juan Fasola and Maja J Matarić. A socially assistive robot exercise coach for the elderly. *Journal of Human-Robot Interaction*, 2(2):3–32, June 2013.

doi: 10.5898/JHRI.2.2.Fasola. URL <https://doi.org/10.5898/JHRI.2.2.Fasola>. 47, 50

- [109] Taylor A. Kessler Faulkner, Elaine Schaertl Short, and Andrea L. Thomaz. Interactive reinforcement learning with inaccurate feedback. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7498–7504. IEEE, 2020. doi: 10.1109/ICRA40945.2020.9197219. 49
- [110] David Feil-Seifer and Maja J Mataric. Defining socially assistive robotics. In *9th International Conference on Rehabilitation Robotics, 2005. ICORR 2005.*, pages 465–468. IEEE, 2005. 8
- [111] David Feil-Seifer and Maja J. Matarić. Socially assistive robotics. *IEEE Robotics & Automation Magazine*, 18(1):24–31, 2011. doi: 10.1109/MRA.2010.940150. 9
- [112] Nadia Figueroa, Ana Lucia Pais Ureche, and Aude Billard. Learning complex sequential tasks from demonstration: A pizza dough rolling case study. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 611–612. Ieee, 2016. 2
- [113] Tesca Fitzgerald, Pallavi Koppol, Patrick Callaghan, Russell Quinlan Jun Hei Wong, Reid Simmons, Oliver Kroemer, and Henny Admoni. Inquire: Interactive querying for user-aware informative reasoning. *6th Annual Conf. on Robot Learning*, 2022. 9, 11, 44, 111, 112, 113, 116, 118, 121, 122
- [114] Tesca Fitzgerald, Pallavi Koppol, Patrick Callaghan, Russell Quinlan Jun Hei Wong, Reid Simmons, Oliver Kroemer, and Henny Admoni. Inquire: Interactive querying for user-aware informative reasoning. In *6th Annual Conference on Robot Learning*, Proceedings of Machine Learning Research. PMLR, 2022. 13
- [115] Terrence Fong, Illah Nourbakhsh, and Kerstin Dautenhahn. A survey of socially interactive robots. *Robotics and autonomous systems*, 42(3-4):143–166, 2003. 10
- [116] Mary Ellen Foster and Jane Stuart-Smith. Social robotics meets sociolinguistics: investigating accent bias and social context in hri. *Comp. of ACM/IEEE Int'l Conf. on Human-Robot Interaction*, 2023. 22
- [117] Mark G Frank, Paul Ekman, and Wallace V Friesen. Behavioral markers and recognizability of the smile of enjoyment. *Journal of personality and social psychology*, 64(1):83, 1993. 11
- [118] Helen L Gallagher, Anthony I Jack, Andreas Roepstorff, and Christopher D Frith. Imaging the intentional stance in a competitive game. *Neuroimage*, 16(3):814–821, 2002. 10

- [119] Nicholas C Georgiou, Rebecca Ramnauth, Emmanuel Adeniran, Michael Lee, Lila Selin, and Brian Scassellati. Is someone there or is that the tv? detecting social presence using sound. *ACM Trans. on Human-Robot Interaction*, 12(4): 1–33, 2023. 22
- [120] Gaurav R Ghosal, Matthew Zurek, Daniel S Brown, and Anca D Dragan. The effect of modeling human rationality level on learning rewards from multiple feedback types. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37:5, pages 5983–5992, 2023. 132
- [121] Debasmita Ghose, Michal Lewkowicz, David Dong, Andy Cheng, Tran Doan, Emma Adams, Marynel Vázquez, and Brian Scassellati. Planning with critical decision points: Robots that influence humans to infer their strategy. *2024 33rd IEEE Int’l Conf. on Robot and Human Interactive Communication (ROMAN)*, pages 777–782, 2024. 42
- [122] Sarah Gillet, Ronald Cumbal, André Pereira, José Lopes, Olov Engwall, and Iolanda Leite. Robot gaze can mediate participation imbalance in groups with different skill levels. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction, HRI ’21*, page 303–311, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450382892. doi: 10.1145/3434073.3444670. URL <https://doi.org/10.1145/3434073.3444670>. 139
- [123] Sarah Gillet, Maria Teresa Parreira, Marynel Vázquez, and Iolanda Leite. 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*, pages 265–274. IEEE Press, 2022. 47, 50
- [124] Sarah Gillet, Marynel Vázquez, Sean Andrist, Iolanda Leite, and Sarah Sebo. Interaction-Shaping Robotics: Robots that Influence Interactions Between Other Agents. *ACM Trans. on Human-Robot Interaction*, 13(1), 2024. 26
- [125] Sarah Gillet, Sydney Thompson, Iolanda Leite, and Marynel Vázquez. Templates and graph neural networks for social robots interacting in small groups of varying sizes. *2025 20th ACM/IEEE Int’l Conf. on Human-Robot Interaction (HRI)*, pages 458–467, 2025. 43
- [126] Manuel Giuliani, Nicole Mirnig, Gerald Stollnberger, Susanne Stadler, Roland Buchner, and Manfred Tscheligi. Systematic analysis of video data from different human-robot interaction studies: a categorization of social signals during error situations. *Frontiers in psychology*, 6:931, 2015. 10, 11
- [127] Michael Görner, Robert Haschke, Helge Ritter, and Jianwei Zhang. Moveit! task constructor for task-level motion planning. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 190–196, 2019. 129

- [128] Samuel D. Gosling, Peter J. Rentfrow, and William B. Swann Jr. A very brief measure of the big-five personality domains. *Journal of Research in Personality*, 37(6):504–528, 2003. ISSN 0092-6566. doi: [https://doi.org/10.1016/S0092-6566\(03\)00046-1](https://doi.org/10.1016/S0092-6566(03)00046-1). 58, 81, 85
- [129] Cedric Goubard and Yiannis Demiris. Cooking up trust: Eye gaze and posture for trust-aware action selection in human-robot collaboration. In *Proceedings of the First International Symposium on Trustworthy Autonomous Systems*, pages 1–5, 2023. 12
- [130] Simon Greipl, Katharina Bernecker, and Manuel Ninaus. Facial and bodily expressions of emotional engagement: How dynamic measures reflect the use of game elements and subjective experience of emotions and effort. *Proceedings of the ACM on Human-Computer Interaction*, 5(CHI PLAY):1–25, 2021. 10
- [131] Daniel Grimm, Philip Schörner, Moritz Dreßler, and J-Marius Zöllner. Holistic graph-based motion prediction. *2023 IEEE Int’l Conf. on Robotics and Automation (ICRA)*, pages 2965–2972, 2023. 22
- [132] Victoria Groom, Jimmy Chen, Theresa Johnson, F. Arda Kara, and Clifford Nass. Critic, compatriot, or chump? responses to robot blame attribution. In *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction*, HRI ’10, page 211–218. IEEE Press, 2010. ISBN 9781424448937. 50
- [133] Jasmin Grosinger, Federico Pecora, and Alessandro Saffiotti. Robots that maintain equilibrium: Proactivity by reasoning about user intentions and preferences. *Pattern Recognition Letters*, 118:85–93, 2019. ISSN 0167-8655. doi: <https://doi.org/10.1016/j.patrec.2018.05.014>. URL <https://www.sciencedirect.com/science/article/pii/S0167865518301843>. 77
- [134] James J. Gross and O.P. John. Revealing feelings: Facets of emotional expressivity in self-reports, peer ratings, and behavior. *Journal of Personality and Social Psychology*, 72:435–448, 1997. 58, 81, 85
- [135] Balint Gucsi, Danesh S. Tarapore, William Yeoh, Christopher Amato, and Long Tran-Thanh. 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*, page 7935–7940. IEEE Press, 2020. doi: 10.1109/IROS45743.2020.9341607. URL <https://doi.org/10.1109/IROS45743.2020.9341607>. 49, 77, 92
- [136] Luke Guerdan, Alex Raymond, and Hatice Gunes. Toward affective xai: facial affect analysis for understanding explainable human-ai interactions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3796–3805, 2021. 10

- [137] Reymundo A Gutierrez, Elaine Schaertl Short, Scott Niekum, and Andrea L Thomaz. Learning from corrective demonstrations. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 712–714. IEEE, 2019. 2
- [138] Dylan Hadfield-Menell, Stuart J Russell, Pieter Abbeel, and Anca Dragan. Cooperative inverse reinforcement learning. In *Advances in neural information processing systems*, volume 29, 2016. 112
- [139] Bradley Hayes and Brian Scassellati. Autonomously constructing hierarchical task networks for planning and human-robot collaboration. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5469–5476. IEEE, 2016. 12, 72
- [140] Cory J Hayes, Maria F O’Connor, and Laurel D Riek. Avoiding robot faux pas: Using social context to teach robots behavioral propriety. *ACM/IEEE Int’l Conf. on Human-Robot Interaction*, 2014. 22
- [141] Haibo He and Eduardo A Garcia. Learning from imbalanced data. *TKDE*, 21(9):1263–1284, 2009. 87
- [142] Thomas Hellström. The relevance of causation in robotics: A review, categorization, and analysis. *Paladyn, Journal of Behavioral Robotics*, 12(1):238–255, 2021. 46
- [143] Shannon Hennig and Ryad Chellali. Expressive synthetic voices: Considerations for human robot interaction. *IEEE Int’l Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2012. 22
- [144] Hui-Ru Ho, Edward M Hubbard, and Bilge Mutlu. “It’s Not a Replacement”: Enabling Parent-Robot Collaboration to Support In-Home Learning Experiences of Young Children. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2024. 12
- [145] Mark K. Ho, James MacGlashan, Michael L. Littman, and Fiery Cushman. Social is special: A normative framework for teaching with and learning from evaluative feedback. *Cognition*, 167:91–106, 2017. ISSN 0010-0277. doi: <https://doi.org/10.1016/j.cognition.2017.03.006>. 13
- [146] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 95
- [147] Guy Hoffman and Cynthia Breazeal. Collaboration in human-robot teams. *AIAA 1st intelligent systems technical conference*, page 6434, 2004. 26
- [148] Guy Hoffman and Wendy Ju. Designing robots with movement in mind. *Journal of Human-Robot Interaction*, 3(1):91–122, 2014. 42

- [149] Guy Hoffman and Xuan Zhao. A primer for conducting experiments in human–robot interaction. *J. Hum.-Robot Interact.*, 10(1), October 2020. doi: 10.1145/3412374. URL <https://doi.org/10.1145/3412374>. 15
- [150] Guy Hoffman, Jodi Forlizzi, Shahar Ayal, Aaron Steinfeld, John Antanitis, Guy Hochman, Eric Hochendoner, and Justin Finkenaur. Robot presence and human honesty: Experimental evidence. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '15, page 181–188, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450328838. doi: 10.1145/2696454.2696487. URL <https://doi.org/10.1145/2696454.2696487>. 47, 50
- [151] Guy Hoffman, Oren Zuckerman, Gilad Hirschberger, Michal Luria, and Tal Shani Sherman. Design and evaluation of a peripheral robotic conversation companion. *Proc. of the tenth annual ACM/IEEE Int'l Conf. on Human-Robot Interaction*, pages 3–10, 2015. 42
- [152] Shanee Honig and Tal Oron-Gilad. Expect the unexpected: Leveraging the human-robot ecosystem to handle unexpected robot failures. *Frontiers in Robotics and AI*, 8:656385, 2021. 41
- [153] Mohammed Moshiul Hoque, Tomomi Onuki, Dipankar Das, Yoshinori Kobayashi, and Yoshinori Kuno. 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, page 149–150, New York, NY, USA, 2012. Association for Computing Machinery. ISBN 9781450310635. doi: 10.1145/2157689.2157729. URL <https://doi.org/10.1145/2157689.2157729>. 56
- [154] John Houston, Paul Harris, Sandra McIntire, and Dientje Francis. Revising the competitiveness index using factor analysis. *Psychological Reports*, 90(1): 31–34, 2002. 81, 85
- [155] Chien-Ming Huang and Bilge Mutlu. The repertoire of robot behavior: Enabling robots to achieve interaction goals through social behavior. *Journal of Human-Robot Interaction*, 2(2):80–102, 2013. 20, 21
- [156] Matthew Huggins, Sharifa Alghowinem, Sooyeon Jeong, Pedro Colon-Hernandez, Cynthia Breazeal, and Hae Won Park. Practical guidelines for intent recognition: Bert with minimal training data evaluated in real-world hri application. *Proc. of the 2021 ACM/IEEE Int'l Conf. on Human-Robot Interaction*, pages 341–350, 2021. 43
- [157] Nathan Hughes, Yun Chang, Siyi Hu, Rajat Talak, Rumaia Abdulhai, Jared Strader, and Luca Carlone. Foundations of spatial perception for robotics: Hierarchical representations and real-time systems. *The Int'l Journal of Robotics Research*, 43(10):1457–1505, 2024. 44

- [158] Isabelle Hupont, Songül Tolan, Hatice Gunes, and Emilia Gómez. The landscape of facial processing applications in the context of the european ai act and the development of trustworthy systems. *Scientific Reports*, 12(1):10688, 2022. 107
- [159] Ahmed Hussein, Mohamed Medhat Gaber, Eyad Elyan, and Chrisina Jayne. Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35, 2017. 10, 11
- [160] Jungsik Hwang, Min-Gyu Kim, and Kwangsu Cho. Developing interaction scenarios of robot-mediated imitation training for children with autism spectrum disorders. *2013 IEEE RO-MAN*, pages 616–620, 2013. 22, 36, 37
- [161] Borja Ibarz, Jan Leike, Tobias Pohlen, Geoffrey Irving, Shane Legg, and Dario Amodei. Reward learning from human preferences and demonstrations in atari. *Advances in neural information processing systems*, 31, 2018. 113
- [162] Auke Jan Ijspeert, Jun Nakanishi, and Stefan Schaal. Movement imitation with nonlinear dynamical systems in humanoid robots. In *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, volume 2, pages 1398–1403. IEEE, 2002. 48
- [163] Anil Kumar Inkulu, MVA Raju Bahubalendruni, Ashok Dara, and SankaranarayanaSamy K. Challenges and opportunities in human robot collaboration context of industry 4.0-a state of the art review. *Industrial Robot: the international journal of robotics research and application*, 49(2):226–239, 2022. 12
- [164] SAS Institute. Jmp pro version 15.0.0, 2021. 63
- [165] Tariq Iqbal and Laurel D. Riek. Human coordination dynamics with heterogeneous robots in a team. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction, HRI '16*, page 619–620. IEEE Press, 2016. ISBN 9781467383707. 50
- [166] Charles Isbell and Christian Shelton. Cobot: A social reinforcement learning agent. In *Advances in Neural Information Processing Systems*, volume 14. MIT Press, 2001. URL <https://proceedings.neurips.cc/paper/2001/file/92bbd31f8e0e43a7da8a6295b251725f-Paper.pdf>. 12
- [167] I. Iturrate, L. Montesano, and J. Minguez. Robot reinforcement learning using eeg-based reward signals. In *2010 IEEE International Conference on Robotics and Automation*, pages 4822–4829, 2010. doi: 10.1109/ROBOT.2010.5509734. 10
- [168] Rachael E Jack, Oliver GB Garrod, Hui Yu, Roberto Caldara, and Philippe G Schyns. Facial expressions of emotion are not culturally universal. *Proceedings of the National Academy of Sciences*, 109(19):7241–7244, 2012. 11, 77, 93

- [169] Ryan Blake Jackson. Toward morally sensitive robotic communication. *2019 14th ACM/IEEE Int'l Conf. on Human-Robot Interaction (HRI)*, pages 715–717, 2019. 22
- [170] Natasha Jaques, Jennifer McCleary, Jesse Engel, David Ha, Fred Bertsch, Douglas Eck, and Rosalind Picard. Learning via social awareness: Improving a deep generative sketching model with facial feedback. In *Workshop on Artificial Intelligence in Affective Computing*, pages 1–9. PMLR, 2020. 10
- [171] Dinesh Babu Jayagopi, Samira Sheiki, David Klotz, Johannes Wienke, Jean-Marc Odobez, Sebastien Wrede, Vasil Khalidov, Laurent Nyugen, Britta Wrede, and Daniel Gatica-Perez. The vernissage corpus: A conversational human-robot-interaction dataset. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 149–150, 2013. doi: 10.1109/HRI.2013.6483545. 99, 100
- [172] Richard Jensen and Qiang Shen. Are more features better? a response to attributes reduction using fuzzy rough sets. *IEEE Transactions on Fuzzy Systems*, 17(6):1456–1458, 2009. doi: 10.1109/TFUZZ.2009.2026639. 93
- [173] Hong Jun Jeon, Smitha Milli, and Anca Dragan. Reward-rational (implicit) choice: A unifying formalism for reward learning. In *NEURIPS*, volume 33, 2020. URL <https://proceedings.neurips.cc/paper/2020/file/2f10c1578a0706e06b6d7db6f0b4a6af-Paper.pdf>. 13, 77
- [174] Hong Jun Jeon, Smitha Milli, and Anca Dragan. Reward-rational (implicit) choice: A unifying formalism for reward learning. *Advances in Neural Information Processing Systems*, 33:4415–4426, 2020. 9, 113, 117, 122, 132
- [175] Pengju Jin, Eshed Ohn-Bar, Kris Kitani, and Chieko Asakawa. A-exp4: Online social policy learning for adaptive robot-pedestrian interaction. *IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems (IROS)*, 2019. 22, 35, 37
- [176] Swapna Joshi and Selma Šabanović. A communal perspective on shared robots as social catalyts. *2017 26th IEEE Int'l Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2017. 22
- [177] Wendy Ju. *The design of implicit interactions*. Springer Nature, 2022. 25
- [178] Malte Jung and Pamela Hinds. Robots in the wild. *ACM Trans. on Human-Robot Interaction*, 7(1):1–5, 2018. 45
- [179] Poornima Kaniarasu and Aaron M. Steinfeld. Effects of blame on trust in human robot interaction. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication, RO-MAN*, pages 850–855. IEEE Press, 2014. doi: 10.1109/ROMAN.2014.6926359. 50

- [180] Tasneem Kaochar, Raquel Torres Peralta, Clayton T. Morrison, Ian R. Fasel, Thomas J. Walsh, and Paul R. Cohen. Towards understanding how humans teach robots. In *Proceedings of the 19th International Conference on User Modeling, Adaption, and Personalization, UMAP'11*, page 347–352, Berlin, Heidelberg, 2011. Springer-Verlag. ISBN 9783642223617. 49
- [181] Aditya Kapoor, Sushant Swamy, Pilar Bachiller, and Luis J Manso. Socnavgym: a reinforcement learning gym for social navigation. *IEEE Int'l Conf. on Robot and Human Interactive Communication (RO-MAN)*, 2023. 22
- [182] Adam Kendon. *Conducting interaction: Patterns of behavior in focused encounters*, volume 7. CUP Archive, 1990. 9
- [183] Taylor Kessler Faulkner, Reymundo A. Gutierrez, Elaine Schaertl Short, Guy Hoffman, and Andrea L. Thomaz. Active attention-modified policy shaping. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19*, page 728–736. IFAAMAS, 2019. ISBN 9781450363099. 49
- [184] Sara Kiesler and Pamela Hinds. Introduction to this special issue on human-robot interaction. *Human-Computer Interaction*, 19(1-2):1–8, 2004. 15
- [185] Callie Y Kim, Christine P Lee, and Bilge Mutlu. Understanding large-language model (llm)-powered human-robot interaction. *Proceedings of ACM/IEEE Int'l Conf. on Human-Robot Interaction*, 2024. 44
- [186] Elizabeth S. Kim, Dan Leyzberg, Katherine M. Tsui, and Brian Scasselati. How people talk when teaching a robot. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction, HRI '09*, page 23–30, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605584041. doi: 10.1145/1514095.1514102. URL <https://doi.org/10.1145/1514095.1514102>. 51, 56
- [187] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan P Foster, Pannag R Sanketi, Quan Vuong, et al. Openvla: An open-source vision-language-action model. *Conference on Robot Learning*, 2025. 15
- [188] Sangmin Kim, Jongsuk Choi, Yoonseob Lim, and Sonya S Kwak. Should a robot follow social norms? human-robot interaction design for social relations in mixed age group. *2022 IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems (IROS)*, pages 8234–8239, 2022. 22, 23
- [189] Yunkyung Kim and Bilge Mutlu. How social distance shapes human-robot interaction. *Int'l Journal of Human-Computer Studies*, 72(12):783–795, 2014. 15

- [190] W. Bradley Knox and Peter Stone. Tamer: Training an agent manually via evaluative reinforcement. In *2008 7th IEEE International Conference on Development and Learning*, pages 292–297. IEEE, 2008. doi: 10.1109/DEVLRN.2008.4640845. 2, 9, 51
- [191] W. Bradley Knox and Peter Stone. Interactively shaping agents via human reinforcement: the tamer framework. In *Proceedings of the Fifth International Conference on Knowledge Capture, K-CAP '09*, pages 9–16, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605586588. doi: 10.1145/1597735.1597738. URL <https://doi.org/10.1145/1597735.1597738>. 11
- [192] W. Bradley Knox, Brian D. Glass, Bradley C. Love, W. Todd Maddox, and Peter Stone. How humans teach agents. *International Journal of Social Robotics*, 4(4):409–421, 2012. URL <https://doi.org/10.1007/s12369-012-0163-x>. 50
- [193] Kheng Lee Koay, Dag Sverre Syrdal, Michael L Walters, and Kerstin Dautenhahn. Five weeks in the robot house—exploratory human-robot interaction trials in a domestic setting. In *2009 second international conferences on advances in computer-human interactions*, pages 219–226. IEEE, 2009. 12
- [194] Fernando Koch and Iyad Rahwan. Classification of agents-based mobile assistants. In *AAMAS workshop on agents for ubiquitous computing*. AAAMAS, 2004. 8
- [195] Dimosthenis Kontogiorgos, Minh Tran, Joakim Gustafson, and Mohammad Soleymani. A systematic cross-corpus analysis of human reactions to robot conversational failures. In *Proceedings of the 2021 International Conference on Multimodal Interaction, ICMI '21*, page 112–120, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450384810. doi: 10.1145/3462244.3479887. URL <https://doi.org/10.1145/3462244.3479887>. 11
- [196] Martin Kraft and Markus Rickert. 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*, page 942–949. IEEE Press, 2017. doi: 10.1109/ROMAN.2017.8172416. URL <https://doi.org/10.1109/ROMAN.2017.8172416>. 48
- [197] Minae Kwon, Mengxi Li, Alexandre Bucquet, and Dorsa Sadigh. Influencing leading and following in human-robot teams. In *Proceedings of Robotics: Science and Systems, RSS*, 2019. 47, 50, 139
- [198] Pat Langley et al. Selection of relevant features in machine learning. In *Proceedings of the AAAI Fall symposium on relevance*, volume 184, pages 245–271, 1994. 93

- [199] Jamie Large, Graham Stodolski, and Marynel Vázquez. Studying human-agent interactions in space invaders. In *Proceedings of the 8th International Conference on Human-Agent Interaction*, HAI '20, page 245–247, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450380546. doi: 10.1145/3406499.3418747. URL <https://doi.org/10.1145/3406499.3418747>. 51, 83
- [200] Przemyslaw A Lasota and Julie A Shah. Analyzing the effects of human-aware motion planning on close-proximity human–robot collaboration. *Human factors*, 57(1):21–33, 2015. 12
- [201] J. Lazar, J.H. Feng, and H. Hochheiser. *Research Methods in Human-Computer Interaction*. Elsevier Science, 2017. ISBN 9780128093436. URL <https://books.google.com/books?id=hbkxDQAAQBAJ>. 75, 82
- [202] Hee Rin Lee, Selma Šabanović, Wan-Ling Chang, Shinichi Nagata, Jennifer Piatt, Casey Bennett, and David Hakken. Steps toward participatory design of social robots: mutual learning with older adults with depression. *Proc. of the 2017 ACM/IEEE Int’l Conf. on Human-Robot Interaction*, pages 244–253, 2017. 22, 35, 37, 38, 39
- [203] Jin Joo Lee, Fei Sha, and Cynthia Breazeal. A bayesian theory of mind approach to nonverbal communication. *2019 14th ACM/IEEE Int’l Conf. on Human-Robot Interaction (HRI)*, pages 487–496, 2019. 22, 32, 36, 37
- [204] Kimin Lee, Laura Smith, Anca Dragan, and Pieter Abbeel. B-pref: Benchmarking preference-based reinforcement learning. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. 120, 136
- [205] Min Kyung Lee, Jodi Forlizzi, Paul E Rybski, Frederick Crabbe, Wayne Chung, Josh Finkle, Eric Glaser, and Sara Kiesler. The snackbot: documenting the design of a robot for long-term human-robot interaction. *Proc. of the 4th ACM/IEEE Int’l Conf. on Human-Robot Interaction*, pages 7–14, 2009. 42
- [206] Wen-Ying Lee, Mose Sakashita, Elizabeth Ricci, Houston Claire, François Guimbretière, and Malte Jung. Interactive vignettes: Enabling large-scale interactive hri research. *2021 30th IEEE Int’l Conf. on Robot & Human Interactive Communication (RO-MAN)*, pages 1289–1296, 2021. 42
- [207] Iolanda Leite, André Pereira, Carlos Martinho, Ana Paiva, and Ginevra Castellano. Towards an empathic chess companion. In *Proceedings of the 8th International Conference on Autonomous agents and Multiagent Systems (AAMAS 2009)*, pages 33–36, 2009. 11, 100

- [208] Mark R Lepper, Maria Woolverton, Donna L Mumme, and J Gurtner. Motivational techniques of expert human tutors: Lessons for the design of computer-based tutors. *Computers as cognitive tools*, 1993:75–105, 1993. 10
- [209] Mark Levine and Kirstien Thompson. Identity, place, and bystander intervention: Social categories and helping after natural disasters. *The Journal of Social Psychology*, 144(3):229–245, 2004. doi: 10.3200/SOCP.144.3.229-245. URL <https://doi.org/10.3200/SOCP.144.3.229-245>. PMID: 15168427. 56
- [210] Mark Levine, Amy Prosser, David Evans, and Stephen Reicher. Identity and emergency intervention: How social group membership and inclusiveness of group boundaries shape helping behavior. *Personality and Social Psychology Bulletin*, 31(4):443–453, 2005. doi: 10.1177/0146167204271651. URL <https://doi.org/10.1177/0146167204271651>. PMID: 15743980. 56
- [211] Alexander Lew, Sydney Thompson, Nathan Tsoi, and Marynel Vázquez. Shutter, the robot photographer: Leveraging behavior trees for public, in-the-wild human-robot interactions. *arXiv preprint arXiv:2302.00191*, 2023. 104
- [212] Kurt Lewin. *Principles of topological psychology*. Read Books Ltd, 2013. 16
- [213] Michael Lewis. Designing for human-agent interaction. *AI Magazine*, 19(2): 67–67, 1998. 75, 82
- [214] Guangliang Li, Hayley Hung, Shimon Whiteson, and W. Bradley Knox. Using informative behavior to increase engagement in the tamer framework. In *Proceedings of AAMAS*, 2013. 12, 71
- [215] Guangliang Li, Hamdi Dibeklioglu, Shimon Whiteson, and Hayley Hung. Facial feedback for reinforcement learning: a case study and offline analysis using the tamer framework. *Autonomous Agents and Multi-Agent Systems*, 34(1):1–29, 2020. 76, 100
- [216] Jia-Min Li, Ke-Xi Liu, Ji-Fei Xie, and Hao-Yu Wu. How does human-robot collaboration affect hotel employees’ proactive behavior? *International Journal of Human-Computer Interaction*, pages 1–15, 2024. 12
- [217] Rui Li, Marc van Almkerk, Sanne van Waveren, Elizabeth Carter, and Iolanda Leite. Comparing human-robot proxemics between virtual reality and the real world. *2019 14th ACM/IEEE Int’l Conf. on Human-Robot Interaction (HRI)*, pages 431–439, 2019. 42
- [218] Congcong Liu, Yuying Chen, Ming Liu, and Bertram E Shi. Avgcn: Trajectory prediction using graph convolutional networks guided by human attention. *IEEE Int’l Conf. on Robotics and Automation (ICRA)*, 2021. 22

- [219] Phoebe Liu, Dylan F Glas, Takayuki Kanda, and Hiroshi Ishiguro. Data-driven hri: Learning social behaviors by example from human–human interaction. *IEEE Trans. on Robotics*, 32(4):988–1008, 2016. 16
- [220] Robert Loftin, James MacGlashan, Bei Peng, Matthew E. Taylor, Michael L. Littman, Jeff Huang, and David L. Roberts. A strategy-aware technique for learning behaviors from discrete human feedback. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, AAAI’14, page 937–943. AAAI Press, 2014. 12
- [221] Dylan P Losey and Marcia K O’Malley. Including uncertainty when learning from human corrections. In *Conference on Robot Learning*, pages 123–132. PMLR, 2018. 9
- [222] Dylan P Losey, Andrea Bajcsy, Marcia K O’Malley, and Anca D Dragan. Physical interaction as communication: Learning robot objectives online from human corrections. *The International Journal of Robotics Research*, 41(1):20–44, 2022. 112
- [223] Matthias Luber, Luciano Spinello, Jens Silva, and Kai O Arras. Socially-aware robot navigation: A learning approach. *2012 IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems*, pages 902–907, 2012. 21, 22
- [224] Gale M Lucas, Jonathan Gratch, Aisha King, and Louis-Philippe Morency. It’s only a computer: Virtual humans increase willingness to disclose. *Computers in Human Behavior*, 37:94–100, 2014. 10
- [225] James MacGlashan, Mark K Ho, Robert Loftin, Bei Peng, Guan Wang, David L Roberts, Matthew E Taylor, and Michael L Littman. Interactive learning from policy-dependent human feedback. In *International Conference on Machine Learning*, pages 2285–2294. PMLR, 2017. 137
- [226] James MacGlashan, Mark K. Ho, Robert Loftin, Bei Peng, Guan Wang, David L. Roberts, Matthew E. Taylor, and Michael L. Littman. Interactive learning from policy-dependent human feedback. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2285–2294. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/macglashan17a.html>. 100
- [227] Luis J Manso, Ronit R Jorvekar, Diego R Faria, Pablo Bustos, and Pilar Bachiller. Graph neural networks for human-aware social navigation. *Advances in Physical Agents II: Proc. of the 21st Int’l Workshop of Physical Agents, November 19-20, 2020*, pages 167–179, 2021. 43
- [228] Daniel Marta, Simon Holk, Christian Pék, Jana Tumova, and Iolanda Leite. Aligning human preferences with baseline objectives in reinforcement learning.

- In *IEEE International Conference on Robotics and Automation (ICRA)*, MAY 29-JUN 02, 2023, London, ENGLAND. Institute of Electrical and Electronics Engineers (IEEE), 2023. 112
- [229] Eloise Matheson, Riccardo Minto, Emanuele GG Zampieri, Maurizio Faccio, and Giulio Rosati. Human–robot collaboration in manufacturing applications: a review. *Robotics*, 8(4):100, 2019. 12
- [230] Kayla Matheus, Marynel Vázquez, and Brian Scassellati. Ommie: The design and development of a social robot for anxiety reduction. *ACM Trans. on Human-Robot Interaction*, 14(2):1–34, 2025. 42
- [231] Christoforos Mavrogiannis, Alena M Hutchinson, John Macdonald, Patrícia Alves-Oliveira, and Ross A Knepper. Effects of distinct robot navigation strategies on human behavior in a crowded environment. *2019 14th ACM/IEEE Int’l Conf. on Human-Robot Interaction (HRI)*, pages 421–430, 2019. 26
- [232] Donald McMillan, Barry Brown, Ikkaku Kawaguchi, Razan Jaber, Jordi Solsona Belenguer, and Hideaki Kuzuoka. Designing with gaze: Tama – a gaze activated smart-speaker. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW), nov 2019. doi: 10.1145/3359278. URL <https://doi.org/10.1145/3359278>. 10
- [233] Tom M Mitchell. The need for biases in learning generalizations. Technical report, Computer Science Department, Rutgers University, 1980. 43
- [234] Martin Moder and Josef Pauli. Proactive robot movements in a crowd by predicting and considering the social influence. *IEEE Int’l Conf. on Robot and Human Interactive Commun. (RO-MAN)*, 2022. 22
- [235] Debasmita Mukherjee, Kashish Gupta, Li Hsin Chang, and Homayoun Najjaran. A survey of robot learning strategies for human-robot collaboration in industrial settings. *Robotics and Computer-Integrated Manufacturing*, 73: 102231, 2022. 9
- [236] Thibaut Munzer, Marc Toussaint, and Manuel Lopes. Preference learning on the execution of collaborative human-robot tasks. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 879–885. IEEE, 2017. 9
- [237] Robin R Murphy. *Introduction to AI robotics*. MIT press, 2019. 1
- [238] Robin R Murphy and Jennifer L Burke. From remote tool to shared roles. *IEEE Robotics & Automation Magazine*, 15(4):39–49, 2008. 41
- [239] Bilge Mutlu and Jodi Forlizzi. Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. *Proceedings of*

- the 3rd ACM/IEEE Int'l Conf. on Human Robot Interaction*, page 287–294, 2008. doi: 10.1145/1349822.1349860. URL <https://doi.org/10.1145/1349822.1349860>. 45
- [240] Bilge Mutlu, Steven Osman, Jodi Forlizzi, Jessica Hodgins, and Sara Kiesler. Task structure and user attributes as elements of human-robot interaction design. *IEEE Int'l Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2006. 15, 42
- [241] Bilge Mutlu, Toshiyuki Shiwa, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. Footing in human-robot conversations: how robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 61–68, 2009. 10
- [242] Vivek Myers, Erdem Biyik, Nima Anari, and Dorsa Sadigh. Learning multimodal rewards from rankings. In *Conference on robot learning*, pages 342–352. PMLR, 2022. 9
- [243] Amal Nanavati, Patricia Alves-Oliveira, Tyler Schrenk, Ethan K Gordon, Maya Cakmak, and Siddhartha S Srinivasa. Design principles for robot-assisted feeding in social contexts. *Proc. of the 2023 ACM/IEEE Int'l Conf. on Human-Robot Interaction*, pages 24–33, 2023. 22
- [244] Amal Nanavati, Patrícia Alves-Oliveira, Tyler Schrenk, Ethan K Gordon, Maya Cakmak, and Siddhartha S Srinivasa. Unintended failures of robot-assisted feeding in social contexts. *Companion of the 2023 ACM/IEEE Int'l Conf. on Human-Robot Interaction*, pages 884–886, 2023. 22
- [245] Austin Narcomey, Nathan Tsoi, Ruta Desai, and Marynel Vázquez. Learning human preferences over robot behavior as soft planning constraints. *arXiv preprint arXiv:2403.19795*, 2024. 112
- [246] Ingrid Navarro and Jean Oh. Social-patternn: Socially-aware trajectory prediction guided by motion patterns. *IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, pages 9859–9864, 2022. 22
- [247] Benjamin A Newman, Reuben M Aronson, Siddhartha S Srinivasa, Kris Kitani, and Henny Admoni. Harmonic: A multimodal dataset of assistive human-robot collaboration. *The International Journal of Robotics Research*, 41(1): 3–11, 2022. 99, 100
- [248] Benjamin A Newman, Christopher Jason Paxton, Kris Kitani, and Henny Admoni. Towards online adaptation for autonomous household assistants. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pages 506–510, 2023. 8

- [249] Eley Ng, Ziang Liu, and Monroe Kennedy. It takes two: Learning to plan for human-robot cooperative carrying. *2023 IEEE Int'l Conf. on Robotics and Automation*, pages 7526–7532, 2023. 42
- [250] Sara Nielsen, Mikael B Skov, Karl Damkjær Hansen, and Aleksandra Kaszowska. Using user-generated youtube videos to understand unguided interactions with robots in public places. *ACM Trans. on Human-Robot Interaction*, 12(1):1–40, 2023. 22, 35, 36, 37, 38, 39, 40
- [251] Aastha Nigam and Laurel D Riek. Social context perception for mobile robots. *2015 IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems (IROS)*, pages 3621–3627, 2015. 18, 21, 22, 27, 35, 43
- [252] Stefanos Nikolaidis, Anton Kuznetsov, David Hsu, and Siddharta Srinivasa. Formalizing human-robot mutual adaptation: A bounded memory model. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, HRI '16, page 75–82. IEEE Press, 2016. ISBN 9781467383707. URL <https://doi.org/10.1109/HRI.2016.7451736>. 48
- [253] Shogo Nishimura, Takuya Nakamura, Wataru Sato, Masayuki Kanbara, Yuichiro Fujimoto, Hirokazu Kato, and Norihiro Hagita. Vocal synchrony of robots boosts positive affective empathy. *Applied Sciences*, 11(6):2502, 2021. 43
- [254] Maria Francesca O'Connor and Laurel D Riek. Detecting social context: A method for social event classification using naturalistic multimodal data. *2015 11th IEEE Int'l Conf. and Workshops on Automatic Face and Gesture Recognition (FG)*, 3:1–7, 2015. 17, 27, 43
- [255] Carl Oechsner, Sven Mayer, and Andreas Butz. Challenges and opportunities of cooperative robots as cooking appliances. *AutomationXP@ CHI*, 2022. 110
- [256] Raquel Oliveira, Patrícia Arriaga, Patrícia Alves-Oliveira, Filipa Correia, Sofia Petisca, and Ana Paiva. Friends or foes? socioemotional support and gaze behaviors in mixed groups of humans and robots. *Proc. of the 2018 ACM/IEEE Int'l Conf. on Human-Robot Interaction*, pages 279–288, 2018. 22, 36, 40, 50
- [257] Anastasia K Ostrowski and Cynthia Breazeal. Design justice for robot design and policy making. *2022 17th ACM/IEEE Int'l Conf. on Human-Robot Interaction (HRI)*, pages 1170–1172, 2022. 22
- [258] Anastasia K Ostrowski, Cynthia Breazeal, and Hae Won Park. Long-term co-design guidelines: empowering older adults as co-designers of social robots. *2021 30th IEEE Int'l Conf. on Robot & Human Interactive Communication (RO-MAN)*, pages 1165–1172, 2021. 22

- [259] Anastasia K Ostrowski, Raechel Walker, Madhurima Das, Maria Yang, Cynthia Breazea, Hae Won Park, and Aditi Verma. Ethics, equity, & justice in human-robot interaction: A review and future directions. *2022 31st IEEE Int'l Conf. on Robot & Human Interactive Communication (RO-MAN)*, pages 969–976, 2022. 22
- [260] Stefan Palan and Christian Schitter. Prolific. ac—a subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, 17:22–27, 2018. 79
- [261] Panagiotis Papadakis, Anne Spalanzani, and Christian Laugier. Social mapping of human-populated environments by implicit function learning. *IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, 2013. 22
- [262] Donald J Patterson, Susan Elliott Sim, and Tosin Aiyelokun. Overcoming blind spots in interaction design: A case study in designing for african aids orphan care communities. *Information Technologies & Int'l Development*, 5(4):pp–75, 2009. 41
- [263] Bjarke Kristian Maigaard Kjær Pedersen, Bente Charlotte Weigelin, Jørgen Christian Larsen, and Jacob Nielsen. Using educational robotics to foster girls' interest in stem: A systematic review. *2021 30th IEEE Int'l Conf. on Robot & Human Interactive Communication (RO-MAN)*, pages 865–872, 2021. 22, 37
- [264] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. 83, 86
- [265] Jairo Pérez-Osorio, Davide De Tommaso, Ebru Baykara, and Agnieszka Wykowska. Joint action with icub: A successful adaptation of a paradigm of cognitive neuroscience in hri. *2018 27th IEEE Int'l Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 152–157, 2018. 22, 36, 37
- [266] Rosalind W Picard. *Affective computing*. MIT press, 2000. 10
- [267] David Porfirio, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. Transforming robot programs based on social context. *Proceedings of CHI Conf. on Human Factors in Computing Systems*, page 1–12, 2020. doi: 10.1145/3313831.3376355. URL <https://doi.org/10.1145/3313831.3376355>. 42
- [268] Morgan Quigley, Brian Gerkey, Ken Conley, Josh Faust, Tully Foote, Jeremy Leibs, Eric Berger, Rob Wheeler, and Andrew Ng. Ros: an open-source robot

- operating system. In *ICRA Workshop on Open Source Software*, volume 3.2, page 5, 2009. 53
- [269] P.L. Patrick Rau, Ye Li, and Dingjun Li. Effects of communication style and culture on ability to accept recommendations from robots. *Computers in Human Behavior*, 25(2):587–595, March 2009. ISSN 0747-5632. doi: 10.1016/j.chb.2008.12.025. URL <https://doi.org/10.1016/j.chb.2008.12.025>. 50
- [270] Niyati Rawal and Ruth Maria Stock-Homburg. Facial emotion expressions in human–robot interaction: A survey. *International Journal of Social Robotics*, 14(7):1583–1604, 2022. 99
- [271] Byron Reeves and Clifford Nass. The media equation: How people treat computers, television, and new media like real people. *Cambridge, UK*, 10(10): 19–36, 1996. 25
- [272] Natalia Reich-Stiebert and Friederike Eyssel. Robots in the classroom: What teachers think about teaching and learning with education robots. *Int’l Conf. on Social Robotics*, pages 671–680, 2016. 41
- [273] Heiko Renz, Maximilian Krämer, and Torsten Bertram. Moving horizon planning for human-robot interaction. *Proc. of the 2024 ACM/IEEE Int’l Conf. on Human-Robot Interaction*, pages 939–943, 2024. 42
- [274] Laurel D Riek and Peter Robinson. Challenges and opportunities in building socially intelligent machines [social sciences]. *IEEE Signal Processing Magazine*, 28(3):146–149, 2011. 17
- [275] James K Rilling, David A Gutman, Thorsten R Zeh, Giuseppe Pagnoni, Gregory S Berns, and Clinton D Kilts. A neural basis for social cooperation. *Neuron*, 35(2):395–405, 2002. 10
- [276] Hannes Ritschel, Andreas Seiderer, Kathrin Janowski, Stefan Wagner, and Elisabeth André. Adaptive linguistic style for an assistive robotic health companion based on explicit human feedback. In *Proceedings of the 12th ACM international conference on PErvasive technologies related to assistive environments*, pages 247–255, 2019. 114
- [277] Claire Rivoire and Angelica Lim. Habit detection within a long-term interaction with a social robot: an exploratory study. In *Proceedings of the International Workshop on Social Learning and Multimodal Interaction for Designing Artificial Agents*, pages 1–6, 2016. 92
- [278] Kantwon Rogers and Ayanna Howard. 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 In-*

- teraction*, HRI '21 Companion, page 57–61, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450382908. doi: 10.1145/3434074.3447128. URL <https://doi.org/10.1145/3434074.3447128>. 50
- [279] Stephanie Rosenthal, Joydeep Biswas, and Manuela M Veloso. An effective personal mobile robot agent through symbiotic human-robot interaction. In *AAMAS*, volume 10, pages 915–922, 2010. 93
- [280] Aaron Roth, Jonathan Ullman, and Zhiwei Steven Wu. Watch and learn: Optimizing from revealed preferences feedback. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 949–962, 2016. 77
- [281] Radu Bogdan Rusu, Brian Gerkey, and Michael Beetz. Robots in the kitchen: Exploiting ubiquitous sensing and actuation. *Robotics and Autonomous Systems*, 56(10):844–856, 2008. ISSN 0921-8890. doi: 10.1016/j.robot.2008.06.010. URL <https://doi.org/10.1016/j.robot.2008.06.010>. 72
- [282] Selma Šabanović, Sarah M Reeder, and Bobak Kechavarzi. Designing robots in the wild: In situ prototype evaluation for a break management robot. *Journal of Human-Robot Interaction*, 3(1):70–88, 2014. 22, 36, 37, 38, 39
- [283] Dorsa Sadigh, Shankar Sastry, Sanjit A. Seshia, and Anca D Dragan. Planning for autonomous cars that leverage effects on human actions. In *Proceedings of Robotics: Science and Systems*, volume 2, pages 1–9, 2016. 48
- [284] Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. Active preference-based learning of reward functions. In *Proceedings of Robotics: Science and Systems*, RSS, 2017. 9
- [285] Nicole Salomons, Kaitlynn Taylor Pineda, Adérónké Adéjàre, and Brian Scassellati. "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, page 176–184. IEEE Press, 2022. 50, 70
- [286] Alan G Sanfey, James K Rilling, Jessica A Aronson, Leigh E Nystrom, and Jonathan D Cohen. The neural basis of economic decision-making in the ultimatum game. *Science*, 300(5626):1755–1758, 2003. 10
- [287] Lindsay Sanneman. Preference elicitation and explanation in iterative planning. In *IJCAI*, pages 6456–6457, 2019. 9
- [288] Joe Saunders, Chrystopher L. Nehaniv, and Kerstin Dautenhahn. Teaching robots by moulding behavior and scaffolding the environment. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction*, HRI '06, page 118–125, New York, NY, USA, 2006. Association for Computing

- Machinery. ISBN 1595932941. doi: 10.1145/1121241.1121263. URL <https://doi.org/10.1145/1121241.1121263>. 51
- [289] Shane Saunderson and Goldie Nejat. How robots influence humans: A survey of nonverbal communication in social human–robot interaction. *International Journal of Social Robotics*, 11(4):575–608, 2019. URL <https://doi.org/10.1007/s12369-019-00523-0>. 47, 50
- [290] Bill N Schilit and Marvin M Theimer. Disseminating active map information to mobile hosts. *IEEE network*, 8(5):22–32, 1994. 77
- [291] Albrecht Schmidt. Implicit human computer interaction through context. *Personal technologies*, 4(2):191–199, 2000. 9
- [292] David C Schrader. Judgments of acceptable partners for social goals based on perceptions of nonverbal behavior. *Journal of Social Behavior and Personality*, 9(2):353, 1994. 10
- [293] Mariah L. Schrum, Erin Hedlund-Botti, Nina Moorman, and Matthew C. Gombolay. Mind meld: Personalized meta-learning for robot-centric imitation learning. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 157–165, 2022. doi: 10.1109/HRI53351.2022.9889616. 11
- [294] Mariah L. Schrum, Erin Hedlund-Botti, Nina Moorman, and Matthew C. Gombolay. 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, page 157–165. IEEE Press, 2022. 9
- [295] SciPy Developers. Scipy linregress function documentation. <https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.linregress.html>, 2023. Accessed October 1, 2023. 106
- [296] Sarah Sebo, Brett Stoll, Brian Scassellati, and Malte F. Jung. Robots in groups and teams: A literature review. *Proc. ACM Human-Computer Interaction*, 4 (CSCW2), oct 2020. doi: 10.1145/3415247. URL <https://doi.org/10.1145/3415247>. 26, 139
- [297] Sarah Strohkorb Sebo, Priyanka Krishnamurthi, and Brian Scassellati. ”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, page 57–65. IEEE Press, 2020. ISBN 9781538685556. 50
- [298] Thomas B Sheridan. Human–robot interaction: status and challenges. *Human factors*, 58(4):525–532, 2016. 12
- [299] Herbert A Simon. A behavioral model of rational choice. *The quarterly journal of economics*, pages 99–118, 1955. 117

- [300] Gabriel Skantze and Bahar Irfan. Applying general turn-taking models to conversational human-robot interaction. In *ACM/IEEE Int'l Conf. on Human-Robot Interaction (HRI)*, pages 859–868. IEEE, 2025. 44
- [301] Daeun Song, Jing Liang, Amirreza Payandeh, Amir Hossain Raj, Xuesu Xiao, and Dinesh Manocha. Vlm-social-nav: Socially aware robot navigation through scoring using vision-language models. *IEEE Robotics and Automation Letters*, 2024. 44
- [302] Yue Song, Mengying Zhang, Jiajing Hu, and Xingping Cao. Dancing with service robots: The impacts of employee-robot collaboration on hotel employees' job crafting. *International Journal of Hospitality Management*, 103:103220, 2022. 12
- [303] Mariacarla Staffa and Silvia Rossi. Enhancing affective robotics via human internal state monitoring. *2022 IEEE Int'l Conf. on Robot and Human Interactive Communication (RO-MAN)*, 2022. 44
- [304] Charles Stangor, Rajiv Jhangiani, and Hammond Tarry. Principles of social psychology-1st int'l edition. *Victoria, BC: BCcampus*. <https://opentextbc.ca/socialpsychology>, 2014. 16
- [305] Maia Stiber, Russell Taylor, and Chien-Ming Huang. Modeling human response to robot errors for timely error detection. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 676–683. IEEE, 2022. 11
- [306] Maia Stiber, Russell H. Taylor, and Chien-Ming Huang. On using social signals to enable flexible error-aware hri. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction, HRI '23*, page 222–230, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450399647. doi: 10.1145/3568162.3576990. URL <https://doi.org/10.1145/3568162.3576990>. 99, 100
- [307] Maia Stiber, Russell H Taylor, and Chien-Ming Huang. On using social signals to enable flexible error-aware hri. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pages 222–230, 2023. 12
- [308] Maia Stiber, Dan Bohus, and Sean Andrist. “uh, this one?”: Leveraging behavioral signals for detecting confusion during physical tasks. *Proc. of the 26th Int'l Conf. on Multimodal Interaction*, 2024. 44
- [309] Sarah Strohkorb Sebo, Margaret Traeger, Malte Jung, and Brian Scasselati. 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*,

- page 178–186, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450349536. doi: 10.1145/3171221.3171275. URL <https://doi.org/10.1145/3171221.3171275>. 48, 49
- [310] Yuta Sugiura, Daisuke Sakamoto, Anusha Withana, Masahiko Inami, and Takeo Igarashi. Cooking with robots: designing a household system working in open environments. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 2427–2430, 2010. 110
- [311] Aamodh Suresh, Angelique Taylor, Laurel D Riek, and Sonia Martinez. Robot navigation in risky, crowded environments: Understanding human preferences. *IEEE Robotics and Automation Letters*, 8(9), 2023. 41
- [312] Prasanth Sengadu Suresh, Siddarth Jain, Prashant Doshi, and Diego Romeres. Open human-robot collaboration using decentralized inverse reinforcement learning. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7092–7098. IEEE, 2024. 112
- [313] Ryo Suzuki, Adnan Karim, Tian Xia, Hooman Hedayati, and Nicolai Marquardt. Augmented reality and robotics: A survey and taxonomy for ar-enhanced human-robot interaction and robotic interfaces. *Proc. of the 2022 CHI Conf. on Human Factors in Computing Systems*, pages 1–33, 2022. 42
- [314] Katia Sycara and Gita Sukthankar. Literature review of teamwork models. *Robotics Institute, Carnegie Mellon University*, 31:31, 2006. 8
- [315] Gilles Tagne, Patrick Henaff, and Nicolas Gregori. Measurement and analysis of physical parameters of the handshake between two persons according to simple social contexts. *2016 IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems (IROS)*, pages 674–679, 2016. 22
- [316] Xiang Zhi Tan, Sean Andrist, Dan Bohus, and Eric Horvitz. 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*, page 468–470, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450370578. doi: 10.1145/3371382.3378363. URL <https://doi.org/10.1145/3371382.3378363>. 56
- [317] Angelique Taylor, Sachiko Matsumoto, and Laurel D Riek. Situating robots in the emergency department. *AAAI Spring Symp. on Applied AI in Healthcare: Safety, Community, and the Environment*, 2020. 41
- [318] Angelique M Taylor, Sachiko Matsumoto, Wesley Xiao, and Laurel D Riek. Social navigation for mobile robots in the emergency department. *2021 IEEE Int’l Conf. on Robotics and Automation (ICRA)*, 2021. 41

- [319] Hamish Tennent, Dylan Moore, Malte Jung, and Wendy Ju. Good vibrations: How consequential sounds affect perception of robotic arms. *IEEE Int'l Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2017. 35, 36
- [320] Andrea Thomaz, Guy Hoffman, and Maya Cakmak. Computational human-robot interaction. *Foundations and Trends in Robotics*, 4(2), 2016. URL <https://doi.org/10.1561/23000000049>. 9, 15, 49
- [321] Andrea L. Thomaz and Cynthia Breazeal. 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, page 1000–1005. AAAI Press, 2006. ISBN 9781577352815. 49, 51, 56
- [322] Andrea L. Thomaz and Cynthia Breazeal. 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, pages 720–725, 2007. doi: 10.1109/ROMAN.2007.4415180. 12
- [323] Andrea L Thomaz and Cynthia Breazeal. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, 172(6):716–737, 2008. 12
- [324] Andrea L. Thomaz and Maya Cakmak. Learning about objects with human teachers. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction*, HRI '09, page 15–22, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605584041. doi: 10.1145/1514095.1514101. URL <https://doi.org/10.1145/1514095.1514101>. 51
- [325] Robin Thompson. Maximum likelihood estimation of variance components. *Statistics: A Journal of Theoretical and Applied Statistics*, 11(4):545–561, 1980. URL <https://doi.org/10.1080/02331888008801560>. 63
- [326] Sydney Thompson, Austin Narcomey, Alexander Lew, and Marynel Vázquez. Shutter: A low-cost and flexible social robot platform for in-the-wild deployments. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '24, pages 94–96, 2024. 129
- [327] Sydney Thompson, Kate Candon, and Marynel Vázquez. The social context of human–robot interactions. *Annual Review of Control, Robotics, and Autonomous Systems*, 9, 2025. 5, 14, 15

- [328] Leimin Tian and Sharon Oviatt. A taxonomy of social errors in human-robot interaction. *ACM Trans. on Human-Robot Interaction (THRI)*, 10(2):1–32, 2021. 22, 36, 37, 38, 40
- [329] Meg Tonkin, Jonathan Vitale, Sarita Herse, Mary-Anne Williams, William Judge, and Xun Wang. Design methodology for the ux of hri: A field study of a commercial social robot at an airport. *Proc. of the 2018 ACM/IEEE Int’l Conf. on Human-Robot Interaction*, pages 407–415, 2018. 22, 33, 34, 35, 36, 37, 38, 39
- [330] Russell Toris, Julius Kammerl, David V Lu, Jihoon Lee, Odest Chadwicke Jenkins, Sarah Osentoski, Mitchell Wills, and Sonia Chernova. Robot web tools: Efficient messaging for cloud robotics. *IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems (IROS)*, 2015. 42
- [331] Kristina Tornbjerg, Anne Marie Kanstrup, Mikael B Skov, and Matthias Rehm. Investigating human-robot cooperation in a hospital environment: Scrutinising visions and actual realisation of mobile robots in service work. In *Proceedings of the 2021 ACM Designing Interactive Systems Conference*, pages 381–391, 2021. 12
- [332] Margaret L. Traeger, Sarah Strohkorb Sebo, Malte Jung, Brian Scassellati, and Nicholas A. Christakis. Vulnerable robots positively shape human conversational dynamics in a human–robot team. *Proceedings of the National Academy of Sciences*, 117(12):6370–6375, 2020. doi: 10.1073/pnas.1910402117. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1910402117>. 139
- [333] Konstantinos Tsiakas, Maher Abujelala, and Fillia Makedon. Task engagement as personalization feedback for socially-assistive robots and cognitive training. *Technologies*, 6(2):49, 2018. 11
- [334] Nathan Tsoi, Mohamed Hussein, Olivia Fugikawa, JD Zhao, and Marynel Vázquez. An approach to deploy interactive robotic simulators on the web for hri experiments: Results in social robot navigation. *2021 IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems*, pages 7528–7535, 2021. 42
- [335] Nathan Tsoi, Alec Xiang, Peter Yu, Samuel S Sohn, Greg Schwartz, Subashri Ramesh, Mohamed Hussein, Anjali W Gupta, Mubbasir Kapadia, and Marynel Vázquez. Sean 2.0: Formalizing and generating social situations for robot navigation. *IEEE RA-L*, 7(4):11047–11054, 2022. 17, 42, 45
- [336] Nathan Tsoi, Rachel Sterneck, Xuan Zhao, and Marynel Vázquez. Influence of simulation and interactivity on human perceptions of a robot during navigation tasks. *ACM Trans. on HRI*, 13(4):1–19, 2024. 42

- [337] Paige Tuttosi, Emma Hughson, Akihiro Matsufuji, Chuxuan Zhang, and Angelica Lim. Read the room: Adapting a robot’s voice to ambient and social contexts. *IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems (IROS)*, 2023. doi: 10.1109/IROS55552.2023.10341925. 22
- [338] Vaibhav V. Unhelkar, Przemyslaw A. Lasota, Quirin Tyroller, Rares-Darius Buhai, Laurie Marceau, Barbara Deml, and Julie A. Shah. Human-aware robotic assistant for collaborative assembly: Integrating human motion prediction with planning in time. *IEEE Robotics and Automation Letters*, 3(3): 2394–2401, 2018. doi: 10.1109/LRA.2018.2812906. 72
- [339] Vaibhav V Unhelkar, Shen Li, and Julie A Shah. Decision-making for bidirectional communication in sequential human-robot collaborative tasks. In *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 329–341. IEEE, 2020. 82
- [340] Mikel Val-Calvo, José Ramón Álvarez-Sánchez, José Manuel Ferrández-Vicente, and Eduardo Fernández. Affective robot story-telling human-robot interaction: exploratory real-time emotion estimation analysis using facial expressions and physiological signals. *IEEE Access*, 8:134051–134066, 2020. 10
- [341] Diede P.M. Van der Hoorn, Anouk Neerincx, and Maartje M.A. de Graaf. ”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*, page 140–148, New York, NY, USA, 2021. Association for Computing Machinery. ISBN 9781450382892. doi: 10.1145/3434073.3444681. URL <https://doi.org/10.1145/3434073.3444681>. 50
- [342] Marieke Van Otterdijk, Diana Saplacan, Adel Baselizadeh, Bruno Laeng, and Jim Torresen. To shake or not to shake: Intuitive reactions of senior adults to a robot handshake in a western culture. *2023 32nd IEEE Int’l Conf. on Robot and Human Interactive Communication*, pages 890–896, 2023. 22, 36
- [343] Gary R VandenBos. *APA dictionary of psychology*. American Psychological Association, 2007. 2
- [344] Mascha Van’t Wout, René S Kahn, Alan G Sanfey, and André Aleman. Affective state and decision-making in the ultimatum game. *Experimental brain research*, 169(4):564–568, 2006. 10
- [345] Marynel Vázquez, Elizabeth J Carter, Braden McDorman, Jodi Forlizzi, Aaron Steinfeld, and Scott E Hudson. Towards robot autonomy in group conversations: Understanding the effects of body orientation and gaze. *Proc. of the 2017 ACM/IEEE Int’l Conf. on Human-Robot Interaction*, pages 42–52, 2017. 22, 36, 37, 40

- [346] Marynel Vázquez, Alexander Lew, Eden Gorevoy, and Joe Connolly. Pose generation for social robots in conversational group formations. *Frontiers in Robotics and AI*, 8:703807, 2022. 16
- [347] Alessandro Vinciarelli, Maja Pantic, and Hervé Bourlard. Social signal processing: Survey of an emerging domain. *Image and vision computing*, 27(12): 1743–1759, 2009. 10, 15
- [348] Hana Vrzakova, Mary Jean Amon, McKenzie Rees, Myrthe Faber, and Sidney D’Mello. Looking for a deal? visual social attention during negotiations via mixed media videoconferencing. *Proc. ACM Hum.-Comput. Interact.*, 4 (CSCW3), jan 2021. doi: 10.1145/3434169. URL <https://doi.org/10.1145/3434169>. 10
- [349] Quan Vuong, Sergey Levine, Homer Rich Walke, Karl Pertsch, Anikait Singh, Ria Doshi, Charles Xu, Jianlan Luo, Liam Tan, Dhruv Shah, et al. Open x-embodiment: Robotic learning datasets and rt-x models. *Towards Generalist Robots: Learning Paradigms for Scalable Skill Acquisition@ CoRL2023*, 2023. 44
- [350] Lennart Wachowiak, Peter Tisnikar, Gerard Canal, Andrew Coles, Matteo Leonetti, and Oya Celiktutan. Analysing eye gaze patterns during confusion and errors in human-agent collaborations. In *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 224–229, 2022. doi: 10.1109/RO-MAN53752.2022.9900589. 11
- [351] Michael Walker, Thao Phung, Tathagata Chakraborti, Tom Williams, and Daniel Szafir. Virtual, augmented, and mixed reality for human-robot interaction: A survey and virtual design element taxonomy. *ACM Trans. on Human-Robot Interaction*, 12(4):1–39, 2023. 42
- [352] Guan Wang, Carl Trimbach, Jun Ki Lee, Mark K. Ho, and Michael L. Littman. 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, page 649–657, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450367462. doi: 10.1145/3319502.3374824. URL <https://doi.org/10.1145/3319502.3374824>. 51
- [353] Lun Wang, Andrea Marrella, and Daniele Nardi. Investigating user perceptions of hri in social contexts. *2019 14th ACM/IEEE Int’l Conf. on Human-Robot Interaction (HRI)*, pages 544–545, 2019. 22
- [354] Ruiqi Wang, Weizheng Wang, and Byung-Cheol Min. Feedback-efficient active preference learning for socially aware robot navigation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11336–11343, 2022. doi: 10.1109/IROS47612.2022.9981616. 77

- [355] Xiaofei Wang, Kimin Lee, Kourosh Hakhmaneshi, Pieter Abbeel, and Michael Laskin. Skill preferences: Learning to extract and execute robotic skills from human feedback. In *6th Annual Conference on Robot Learning*, Proceedings of Machine Learning Research, pages 1259–1268. PMLR, 2022. 72
- [356] David Whitney, Eric Rosen, James MacGlashan, Lawson L. S. Wong, and Stefanie Tellex. Reducing errors in object-fetching interactions through social feedback. In *2017 IEEE International Conference on Robotics and Automation*, ICRA, pages 1006–1013. IEEE Press, 2017. URL <https://doi.org/10.1109/ICRA.2017.7989121>. 49
- [357] Morton Wiener, Shannon Devoe, Stuart Rubinow, and Jesse Geller. Nonverbal behavior and nonverbal communication. *Psychological review*, 79(3):185, 1972. 9
- [358] Tom Williams, Cynthia Matuszek, Ross Mead, and Nick Depalma. Scarecrows in oz: the use of large language models in hri. *ACM Transactions on Human-Robot Interaction*, 13(1):1–11, 2024. 44
- [359] Katie Winkle, Donald McMillan, Maria Arnelid, Katherine Harrison, Madeline Balaam, Ericka Johnson, and Iolanda Leite. Feminist human-robot interaction: Disentangling power, principles and practice for better, more ethical hri. *Proc. of the 2023 ACM/IEEE Int’l Conf. on Human-Robot Interaction*, pages 72–82, 2023. 22
- [360] Bryce Woodworth, Francesco Ferrari, Teofilo E. Zosa, and Laurel D. Riek. Preference learning in assistive robotics: Observational repeated inverse reinforcement learning. In *Proceedings of the 3rd Machine Learning for Healthcare Conference*, volume 85 of *Proceedings of Machine Learning Research*, pages 420–439. PMLR, 17–18 Aug 2018. URL <https://proceedings.mlr.press/v85/woodworth18a.html>. 9
- [361] Julia L. Wright, Shan G. Lakhmani, and Jessie Y.C. Chen. Bidirectional communications in human-agent teaming: The effects of communication style and feedback. *International Journal of Human-Computer Interaction*, 38(18-20):1972–1985, 2022. doi: 10.1080/10447318.2022.2068744. URL <https://doi.org/10.1080/10447318.2022.2068744>. 50
- [362] Yifan Xu, Theodor Chakhachiro, Tribhi Kathuria, and Maani Ghaffari. Solo t-dirl: Socially-aware dynamic local planner based on trajectory-ranked deep inverse reinforcement learning. *2023 IEEE Int’l Conf. on Robotics and Automation (ICRA)*, pages 12045–12051, 2023. 22, 35, 36, 39, 40
- [363] Keiichi Yamazaki, Michie Kawashima, Yoshinori Kuno, Naonori Akiya, Matthew Burdelski, Akiko Yamazaki, and Hideaki Kuzuoka. Prior-to-request and request behaviors within elderly day care: Implications for developing

- service robots for use in multiparty settings. In *ECSCW 2007*, pages 61–78. Springer, 2007. 10
- [364] Yukang Yan, Chun Yu, Wengrui Zheng, Ruining Tang, Xuhai Xu, and Yuanchun Shi. Frownonerror: Interrupting responses from smart speakers by facial expressions. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2020. 10, 76
- [365] Hang Yu, Reuben M Aronson, Katherine H Allen, and Elaine Schaertl Short. From “Thumbs Up” to “10 out of 10”: Reconsidering Scalar Feedback in Interactive Reinforcement Learning. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4121–4128. IEEE, 2023. 123
- [366] Zhou Yu, Dan Bohus, and Eric Horvitz. 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*, pages 402–406, 2015. 56
- [367] Jan-Nico Zaech, Dengxin Dai, Alexander Liniger, and Luc Van Gool. Action sequence predictions of vehicles in urban environments using map and social context. *2020 IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems (IROS)*, pages 8982–8989, 2020. doi: 10.1109/IROS45743.2020.9340643. 21, 22
- [368] Bowen Zhang and Harold Soh. Large language models as zero-shot human models for human-robot interaction. *IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems (IROS)*, 2023. 2, 44
- [369] Qiping Zhang, Austin Narcomey, Kate Candon, and Marynel Vázquez. Self-annotation methods for aligning implicit and explicit human feedback in human-robot interaction. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction, HRI ’23*, page 398–407, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450399647. doi: 10.1145/3568162.3576986. URL <https://doi.org/10.1145/3568162.3576986>. 98, 100, 104, 105
- [370] Qiping Zhang, Austin Narcomey, Kate Candon, and Marynel Vázquez. Self-annotation methods for aligning implicit and explicit human feedback in human-robot interaction. In *Proceedings of HRI*, 2023. doi: 10.1145/3568162.3576986. URL <https://doi.org/10.1145/3568162.3576986>. 10
- [371] Qiping Zhang, Nathan Tsoi, Mofeed Nagib, Booyeon Choi, Jie Tan, Hao-Tien Lewis Chiang, and Marynel Vázquez. Predicting human perceptions of robot performance during navigation tasks. *ACM Transactions on Human-Robot Interaction*, 14(3):1–27, 2025. 11, 44

- [372] Michelle D Zhao, Reid Simmons, and Henny Admoni. Learning human contribution preferences in collaborative human-robot tasks. In *Conference on Robot Learning*, pages 3597–3618. PMLR, 2023. 112
- [373] Zuyuan Zhu and Huosheng Hu. Robot learning from demonstration in robotic assembly: A survey. *Robotics*, 7(2):17, 2018. 12
- [374] Andreas Zimmermann, Andreas Lorenz, and Reinhard Oppermann. An operational definition of context. *Int'l and interdisciplinary conference on modeling and using context*, pages 558–571, 2007. 27