Leveraging Implicit Human Feedback to Better Learn from Explicit Human Feedback in Human-Robot Interactions

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

ABSTRACT

My work aims to enable robots to more effectively learn how to help people. The way in which people want to be helped by robots can vary by task, person, or time, among other factors. Thus, it is important that robots can learn to tailor their behavior based on a person's evolving preferences during an interaction. Robots typically learn from humans via explicit feedback, such as evaluative feedback, preferences, demonstrations, or corrections. However, this type of feedback can interrupt the flow of an interaction and it places an additional cognitive burden on the human. We know that humans "leak" information through their non-verbal behavior that gives clues about their internal states during interactions–can this information be used to augment how a robot learns from humans? My research aims to explore how to incorporate feedback that humans provide implicitly into robot learning paradigms.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; • **Computing methodologies** → *Artificial intelligence.*

KEYWORDS

human-robot interaction, nonverbal behavior, human feedback

ACM Reference Format:
Kate Candon. 2024. Leveraging Implicit Human Feedback to Better Learn from Explicit Human Feedback in Human-Robot Interactions. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24 Companion), March 11–14, 2024, Boulder, CO, USA*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3610978.3638368

1 INTRODUCTION

Imagine you are making pizza with a robot and are limited to providing evaluative (i.e., "good" versus "bad") explicit feedback. You prefer to bring all of your ingredients to the workstation before you start anything else. If the robot starts chopping peppers before you are ready, you might consider providing negative feedback. You frown, look towards the "bad" button, but ultimately choose not to provide the feedback to the robot, since it was not a huge mistake. I believe it would lead to a better interaction if the robot were able to learn from the implicit feedback you provided, perhaps by asking for explicit feedback in that moment. Thus, my research aims to answer the question: **How can we leverage implicit human feedback so that robots can better learn from explicit human feedback in human-robot interactions?**

My goal is to enable robots to more effectively learn how to help people in a personalized manner during human-robot interactions. As the capabilities of robots continue to increase, humans will want to collaborate with robots in a wide range of tasks in a variety of settings. Thus, tasks will tend to be less objective and increasingly driven by personal preferences [4, 17], making it infeasible to pre-program robot behaviors. Rather, it is important to enable robots to better learn from non-expert human teachers [1].

Robots typically learn from humans via explicit feedback, such as evaluative feedback [15], preferences [4], demonstrations [11], or corrections [12], and recent work has explored how a robot learner can combine multiple types of feedback through a common framework [13]. However, relying solely on explicitly provided feedback poses challenges in human-robot interactions. For one, depending on the interface through which people can provide feedback to a robot learner, explicit feedback might not capture the subtlety of the person's preferences or goals. Additionally, explicit feedback can take time and attention away from the person's own tasks during a collaboration with a robot and interupt the flow.

Some recent works have explored how robots can learn from implicit human feedback (e.g., [10, 16]). Implicit human feedback encompasses a wide range of communicative signals inadvertently conveyed during interactions, without the explicit intent to convey information to the robot. For this work, I will focus on facial expressions, eye gaze, head position and orientation, and the actions a human takes in the task (e.g., moving an item). Interpreting implicit feedback presents challenges as reactions can be highly individualized and can vary across scenarios or cultures [3]. While it is challenging to interpret, we know that humans "leak" information through their nonverbal behavior that gives clues about how they perceive their social encounters [14, 20].

While my prior work has investigated explicit and implicit human feedback individually, I plan to focus my future work on reasoning about both explicit and implicit human feedback together to enhance robots' learning capabilities compared to existing approaches. Both implicit and explicit feedback could help make better sense of different cues or inputs, compared to considering only one type of feedback in isolation. By improving robots' understanding of how people perceive their actions and enabling adaptive behavior, robots would be able to more consistently act in alignment with user preferences when learning new skills. Thus, my hope is that considering implicit and explicit human feedback together will enable more seamless human-robot interactions.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
HRI '24 Companion, March 11–14, 2024, Boulder, CO, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0323-2/24/03.
https://doi.org/10.1145/3610978.3638368

2 PRIOR WORK

As a first step, I explored participants' perceptions of an agent's helpfulness. In an exploratory study, we had participants play a video game with an interactive agent (*co-player*) and we studied what factors influenced the perceived helpfulness of the agent in the interactions [8]. We found that even when participants were given a clear objective goal, they had different interpretations of whether or not assistive behaviors from an agent were in fact helpful. The perceived helpfulness of an agent was more correlated with emotional perceptions (e.g., how annoying the human reported they found the agent to be) than game objectives (e.g., how many points the co-player scored for the human). Our findings are in line with recent work in Artificial Intelligence suggesting that in human-AI teams, it is not always the most accurate model that ends up being the best [2]. It is challenging to know what to optimize for when determining a robot's behavior, so it is important that robots are capable of learning to adjust their behavior during interactions.

To allow for effective learning by robots, which typically relies on explicit feedback from humans, it is crucial that humans do, in fact, provide enough feedback during human-robot interactions. Thus, in prior work, we studied how a robot should remind people to provide explicit feedback in fast-paced, cooperative interactions [9]. We found that by reminding participants to provide feedback before a change in robot behavior, the robot could influence participants to provide feedback more quickly and more frequently.

While there are many opportunities to improve the explicit feedback humans provide in human-robot interactions, I believe there is potential in also studying the implicit feedback that humans provide. My prior work found that even without explicitly directing participants to be expressive, incorporating "free" nonverbal reactions improved our ability to predict their preferences between agent behaviors [6]. Furthermore, our results suggested that considering additional context is important when trying to interpret nonverbal human behavior effectively. In other work, we analyzed the data collected in two separate human-robot interactions to show that interaction history is an important factor that can influence human reactions to robots [7]. I plan to explicitly account for this history in future models for interpreting implicit feedback in HRI.

Challenges with reasoning about implicit human feedback have motivated two other prior works. First, we explored methods for self-annotation of implicit human feedback, so that we can collect informative data [21]. Then, to deal with class imbalance, which is often present when analyzing implicit human feedback (i.e., a person will have a neutral face the majority of the time and only a small number of frames might show a reaction to a robot), we introduced a new method for unifying the training and evaluation steps in binary classification [19]. These works enable us to collect better data and to train better models that can reason about implicit human feedback in more interesting ways going forward.

3 FUTURE WORK

We plan to investigate approaches to enable robots to better learn from explicit human feedback by considering implicit human feedback. Humans learn through a combination of explicit and implicit feedback, and it would be advantageous if robots were also able to do so [18]. The "freely provided" information from implicit feedback increases the amount of information from which a robot can learn, and the explicit feedback may help ground difficult to interpret implicit feedback. Specifically, I plan to investigate two research questions in a cooking setup (similar to Figure 1): **RQ1: How can a robot use implicit feedback to decide when to query a human for explicit feedback?** and **RQ2: How can a robot use implicit feedback to qualify explicit feedback?**

Figure 1: Collaborative cooking setup from prior work [5].

Regarding RQ1, when to ask for explicit feedback can be a challenging question. There are times when a person is distracted and would not be able to provide good feedback. There are other times when it is obvious how an interaction is going and not worth asking for feedback. I hypothesize that analyzing implicit human feedback could help identify appropriate times to query for explicit feedback. One potential approach is to predict a distribution of how people may respond to particular robot actions in context and compare their actual implicit signals to this distribution. If the human's actual actions are very different, it may be a good time to pose a query for explicit feedback. For example, if a robot brings something to the workstation and expected a smile, but the person keeps glancing at the recipe, it may be a good time to ask for feedback.

Regarding RQ2, I propose to use implicit human feedback to interpret explicit feedback in a more granular manner. Evaluative feedback ("good" vs. "bad") is one of the simplest ways a human teacher can provide explicit feedback to a robot learner. It is always possible to design a more complicated interface for people to provide more granular feedback to the robot, but that opens the possibility of confusing users, increasing their cognitive burden, or taking time away from other actions. I posit that rather than designing interfaces for humans to provide more granular feedback, we can leverage "freely provided" implicit feedback to qualify the explicit feedback they do provide. I plan to create a framework that uses implicit feedback to qualify explicitly provided feedback so that robots can learn better and faster. For example, if a person provides negative feedback via a button, but also rolls their eyes, that should be a stronger feeling than if a person hesitated before pressing the negative feedback button.

ACKNOWLEDGMENTS

Thank you to M. Vázquez for her consistent guidance and support as an advisor. This work was partly funded by the National Science Foundation (under Grant No. IIS-1924802 and IIS-2106690), Tata Sons Private Limited, Tata Consultancy Services Limited, and Titan.

REFERENCES

[1] Gopika Ajaykumar, Maureen Steele, and Chien-Ming Huang. 2021. A survey on end-user robot programming. *ACM Computing Surveys (CSUR)* 54, 8 (2021), 1–36.

[2] Gagan Bansal, Besmira Nushi, Ece Kamar, Eric Horvitz, and Daniel S Weld. 2021. Is the most accurate ai the best teammate? optimizing ai for teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 11405–11414.

[3] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M Martinez, and Seth D Pollak. 2019. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological science in the public interest* 20, 1 (2019), 1–68.

[4] Erdem Bıyık, Aditi Talati, and Dorsa Sadigh. 2022. Aprel: A library for active preference-based reward learning algorithms. In *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 613–617.

[5] Jake Brawer, Debasmita Ghose, Kate Candon, Meiying Qin, Alessandro Roncone, Marynel Vázquez, and Brian Scassellati. 2023. Interactive Policy Shaping for Human-Robot Collaboration with Transparent Matrix Overlays. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 525–533.

[6] Kate Candon, Jesse Chen, Yoony Kim, Zoe Hsu, Nathan Tsoi, and Marynel Vázquez. 2023. 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*. 307–316.

[7] Kate Candon, Nicholas C. Georgiou, Helen Zhou, Sidney Richardson, Qiping Zhang, Brian Scassellati, and Marynel Vázquez. 2024. REACT: Two Datasets for Analyzing Both Human Reactions and Evaluative Feedback to Robots Over Time. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*.

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

[9] Kate Candon, Helen Zhou, Sarah Gillet, and Marynel Vázquez. 2023. Verbally Soliciting Human Feedback in Continuous Human-Robot Collaboration: Effects of the Framing and Timing of Reminders. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 290–300.

[10] Yuchen Cui, Qiping Zhang, Alessandro Allievi, Peter Stone, Scott Niekum, and W. Bradley Knox. 2020. The EMPATHIC Framework for Task Learning from Implicit Human Feedback. In *CoRL*.

[11] Nadia Figueroa, Ana Lucia Pais Ureche, and Aude Billard. 2016. Learning complex sequential tasks from demonstration: A pizza dough rolling case study. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. Ieee, 611–612.

[12] Reymundo A Gutierrez, Elaine Schaertl Short, Scott Niekum, and Andrea L Thomaz. 2019. Learning from corrective demonstrations. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 712–714.

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

[14] Adam Kendon. 1990. *Conducting interaction: Patterns of behavior in focused encounters*. Vol. 7. CUP Archive.

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

[16] Guangliang Li, Hamdi Dibeklioğlu, Shimon Whiteson, and Hayley Hung. 2020. Facial feedback for reinforcement learning: a case study and offline analysis using the TAMER framework. *Autonomous Agents and Multi-Agent Systems* 34, 1 (2020), 1–29.

[17] Benjamin A Newman, Christopher Jason Paxton, Kris Kitani, and Henny Admoni. 2023. Towards Online Adaptation for Autonomous Household Assistants. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*. 506–510.

[18] Olivier Sigaud, Ahmed Akakzia, Hugo Caselles-Dupré, Cédric Colas, Pierre-Yves Oudeyer, and Mohamed Chetouani. 2022. Towards teachable autotelic agents. *IEEE Transactions on Cognitive and Developmental Systems* (2022).

[19] Nathan Tsoi, Kate Candon, Deyuan Li, Yofti Milkessa, and Marynel Vázquez. 2022. Bridging the Gap: Unifying the Training and Evaluation of Neural Network Binary Classifiers. 35 (2022), 23121–23134. https://proceedings.neurips.cc/paper_files/paper/2022/file/92440ec643f4e9f17409557b6516566e-Paper-Conference.pdf

[20] Morton Wiener, Shannon Devoe, Stuart Rubinow, and Jesse Geller. 1972. Nonverbal behavior and nonverbal communication. *Psychological review* 79, 3 (1972), 185.

[21] Qiping Zhang, Austin Narcomey, Kate Candon, and Marynel Vázquez. 2023. 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* (Stockholm, Sweden) *(HRI '23)*. Association for Computing Machinery, New York, NY, USA, 398–407. https://doi.org/10.1145/3568162.3576986