Understanding Intentions in Human Teaching to Design Interactive Task Learning Robots

Preeti Ramaraj, Matt Klenk and Shiwali Mohan

Palo Alto Research Center
Palo Alto, CA 94306 USA
Email: {pramaraj, klenk, smohan}@parc.com

***Abstract*—The goal of Interactive Task Learning (ITL) is to build robots that can be trained in new tasks by human instructors. In this paper, we approach the ITL research problem from a human instructor perspective. The research question that we address here is *how do we understand and leverage the intentionality of the instructors to enable natural and flexible ITL*. We propose a taxonomy based on *Collaborative Discourse Theory* that organizes human teaching intentions in a human robot teaching interaction. This taxonomy will provide guidance for ITL robot design that leverages a human's natural teaching skills, and reduces the cognitive burden of non-expert instructors. We propose human participant studies to validate this taxonomy and gain a comprehensive understanding of teaching interactions in ITL.**

I. Introduction

We envision a future where robots can help people with a myriad set of tasks in diverse and dynamic environments such as homes, offices, shopping centers, etc. The diversity in environments and tasks requires that people can teach robots new tasks and relevant information about their environments on the fly. Interactive Task Learning (ITL) [12] aims to contribute to this future. The research goal of ITL is to build robots that can be trained in new tasks by a human instructor through a combination of natural language and demonstration.

ITL relies on the fact that people naturally engage in interactive teaching and learning with other people, and therefore can apply this skill to teach a robot as well. A crucial aspect of teaching is the instructor's mental model of the learner - understanding what the learner knows and does not know and how they apply their knowledge to perform various tasks in the environment. An instructor may apply various interactive strategies to develop this understanding about the learner. A human parent (instructor) may ask the child (learner) to demonstrate a skill (bring your teddy bear), to identify concepts (where is your head?), to generate a concept (can you show me what angry looks like?), compare objects (what is bigger, an orca or a beluga?) etc. Responses to these prompts aid the instructor in adapting their instructions. Without reasonable estimates of a learner's capabilities, it is challenging for a human teacher to teach effectively. An ideal learner, then, must be able to provide appropriate responses to an instructor's prompts and integrate new information incrementally. yellow This instruction process is challenging to implement in an ITL scenario for two reasons. First, as ITL robots continuously acquire new knowledge and adapt their behavior, the instructor needs to appropriately update their mental model of the robot's capabilities. Second, human instructors are used to engaging in more flexible linguistic teaching strategies than current ITL robots are capable of performing.

Our research goal is to enable ITL robots to engage in teaching interactions that are natural for humans. In ITL, the instructor guides the robot using a task teaching process. This process includes providing examples of task-relevant concepts, evaluating the knowledge the robot already has, designing lessons to teach concepts and tasks of varying difficulty, etc. Each instruction that the instructor provides has a specific intention to which the robot must respond. The research question here is *how do we understand and leverage the intentionality of the instructors to enable flexible and natural ITL*. To answer this question, we look to collaborative discourse theory in which Grosz and Sidner [10] posit that human task-oriented, collaborative conversations can be understood as a sequence of intentional exchanges where each partner intends the other to update their beliefs based on the scenario, to act according to a goal etc. Rich et al. [23] demonstrate how collaborative discourse theory is used to design natural and flexible human-computer interaction. Previous work [19] leveraged the collaborative discourse theory to develop a computational model to manage ITL interactions. While this approach enabled flexible and mixed-initiative interactive behavior, it is robot-centred. Interactions are largely driven by the robot's learning needs with very little understanding of how humans teach.

In this paper, we take a human-centred approach to understand teaching interactions in ITL. We propose a taxonomy to organize human teaching intentions in a human-robot teaching interaction. The purpose of this taxonomy is to provide guidance for robot design that leverages a human's natural ability to teach. This human-centred design will reduce the cognitive burden of non-expert instructors. To validate and extend this taxonomy, we propose a set of human participant studies. The goal of these studies is to ensure this taxonomy is informed by actual non-expert instructor expectations and experiences. The taxonomy that we propose is by no means exhaustive. We expect the proposed studies to validate aspects of the taxonomy and uncover gaps that must be accounted for to create a comprehensive picture of teaching interactions.

II. ITL AGENTS

Our research studies and develops interaction mechanisms for ITL robots such as ROSIE [18] and AILEEN [20] that learn novel concepts and tasks through situated instruction. They are built in the Soar cognitive architecture [13] which has been augmented with computer vision and concrete action planning modules to enable operation in robotic domains. Research with Rosie has demonstrated that it can be taught over 60 games and puzzles [11] and several mobile delivery and navigation tasks [17] through a combination of situated natural language commands and demonstrations. AILEEN is a variation of ROSIE that focuses on learning novel *concepts* that provide grounding to embodied language processing, structure scene understanding, and organize action execution. For the purposes of this paper, we will be focusing on AILEEN.

One of the goals of research in AILEEN is to create an interaction model that enables a human teacher (who is not a robot expert) to teach diverse types of concepts relevant to tasks. The interaction patterns should be natural for human teachers. To focus our experiment design, data gathering, and analyses, we use concrete learning scenarios that AILEEN is expected to learn from. AILEEN lives in a simulated robotic world which has a table on which various simple objects can be placed (shown in Figure 1). AILEEN perceives its world through a simulated camera placed above the table providing a top-down perspective. It can act by picking up objects and placing them in desired locations on the table.

The goal of the instructor is to introduce AILEEN to novel **visual concepts** such as colors (*red*) and shapes (*box*); **spatial concepts** such as configurations (*yellow box is to the right of the red cylinder*); and **action concepts** (*move the red box to the right of the yellow box*). Visual and spatial concepts support scene understanding in AILEEN while spatial and action concepts support task execution. AILEEN learns through guided participation - conjoint stimuli with a demonstration accompanied with a linguistic description. The instructor generates static scenes by setting up objects to demonstrate visual and spatial concepts. A sequence of scenes that show temporal changes in the configuration of objects is used to demonstrate an action. AILEEN jointly learns concepts and the language to describe them. AILEEN employs analogical reasoning and generalization to learn concepts and we have demonstrated that learned concept definitions enable language understanding as well as task execution [20]. AILEEN implements an interaction model (composed of an interaction state and a heuristics-driven policy) based on collaborative discourse theory [19] that supports mixed-initiative interactions with an instructor and seamlessly integrates with learning and task execution. Findings from the human participant studies proposed below will be integrated into AILEEN's interaction model to enable flexible and robust ITL.

III. RELATED WORK

Previous research has looked at how a robot can provide information about itself, its understanding, and knowledge to help a human collaborator build a mental model. This research explores multiple modalities of interaction including language [7, 16, 21, 22, 24, 25, 27], gaze [5, 21, 26], gestures [5, 8, 21] and visualization [14, 21, 22]. Our research contributes to this line of research by analyzing interactions at the *intentional* level whereas previous work focuses on low-level, concrete interactions.

There has been some work in the domain of machine teaching where researchers have explored and incorporated human teaching perspective in a robot learning setting [2, 6, 15, 26]. Cakmak and Thomaz [6] study how non-experts naturally teach machine learners and propose providing guidance to them to leverage human intelligence and flexibility. Our work is along similar lines, where we want to look at this problem from a top-down human teaching perspective so that we can design robots that leverage natural human teaching behaviors. MacLellan et al. [15] propose a framework to describe cognitive system training interactions and how it can be used to enable people to *naturally* train cognitive systems using language. They describe a list of types of interactions that humans and robots can engage in, in the course of a teaching interaction. Our taxonomy shares some similarities with the list, however, we are proposing an approach to look at these interactions from a more abstract intentional level as it relates to the complete task.

IV. A TAXONOMY OF TEACHING INTENTIONS

We propose a taxonomy based on the CDT model to organize intentions that are observed in a human-robot teaching interaction. The CDT model assumes a goal-driven agent that makes decisions over a joint space of interaction, action, and learning. The information captured in the current state of this model influences the robot's decision making in this joint space. Each instruction that an instructor provides adds information to the current interaction state. This information is provided by the instructor in an *intentional* manner for the purposes of teaching, verifying, or revising the robot's knowledge. We propose a framework that organizes intentions inherent in teaching interactions. The framework will be used to extend AILEEN's interaction model so that it is faithful to typical human teaching interactions. Let us look at a grounded example in Aileen in order to develop this taxonomy.

Let us assume that AILEEN starts with no knowledge of any visual, spatial, or action concepts. We have a human instructor who wants to to teach AILEEN to *move the red box to the right of the blue cylinder* as seen in Figure 1e. Teaching this action concept is challenging as several visual and spatial concepts must be taught as well. We expect a human instructor to correctly assess AILEEN's ignorance and create a lesson to teach a sequence of individual concepts. The instructor would first teach AILEEN the visual concepts of color (*red*, *blue*) and shape (*box*, *cylinder*). Then they would teach the spatial concept *right of* using the configurations presented in Figures 1b and 1c. The instructor can inform the robot that the blue cylinder *is to the right-of* the red box in Figure 1c. Assuming that the instructor has successfully taught these visual and spatial concepts, they will teach the move action

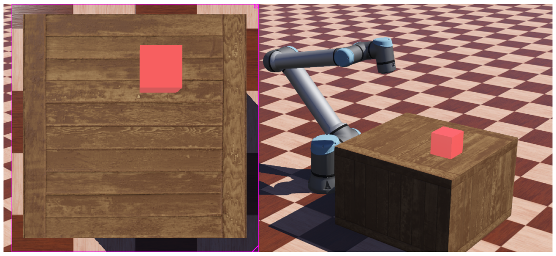

(a) A robot is shown a red box.

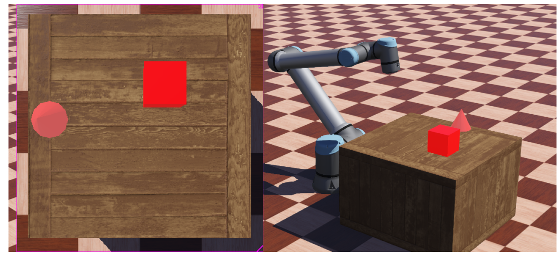

(b) A robot is shown that a red box is to the *right of* a red cone.

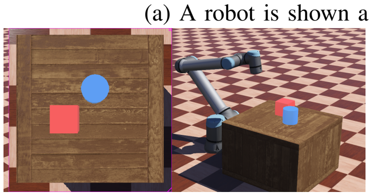

(c) Initial state: A blue cylinder is to the right of a red box.

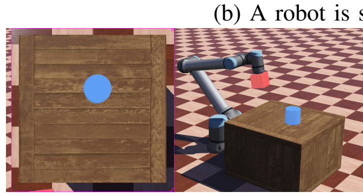

(d) The instructor guides the robot to move the red box.

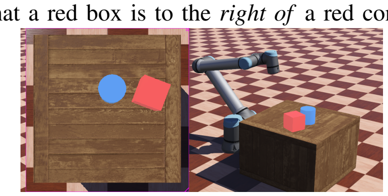

(e) Goal state: The red box is to the right of the blue cylinder.

Figure

Fig. 1: After learning concepts such as 'red', 'box', and 'right of' through configurations such as Figure 1a and 1b, the instructor demonstrates a move action as presented in Figure 1c, 1d, and 1e so that AILEEN can learn the *move* action.

by *demonstrating* a move action in the environment. After learning all the information required to follow the instruction, AILEEN can now move the red box to the right of the blue cylinder successfully in arbitrary initial situations.

All the interactions described below occur within the context of situated interaction i.e., all concepts and knowledge are linked to the shared environment. While this framework focuses on the instructor's intentions and utterances, it operates under the assumption that the robot will provide positive evidence of processing the instructions. This evidence can take the form of acknowledgments, relevant next turns as well as continued attention towards the interaction [9].

1) **Inform:** Using the *inform* intention, the human instructor introduces new concepts to the robot. Teaching a new concept can include providing examples and demonstrations [1, 20] as well as more complicated instructions such as task descriptions [11].
   AILEEN learns through examples and demonstrations. For example, the instructor can use the scenario in Figure 1a to *inform* AILEEN that the object on the table is of shape *box* and whose color is *red*. The instructor can also use the *inform* intention to teach AILEEN task execution by demonstrating an action as shown in Figure 1.
2) **Evaluate:** Using the *evaluate* intention, the instructor learns to what extent the robot can correctly apply concepts. We can further divide the evaluate intention into two sub-intentions based on the specific knowledge the instructor would like to access.
   - **Instantiate:** When the instructor asks the robot to identify examples of known concepts in the environment, the instructor's intention is to learn the extent of its conceptual knowledge in understanding the scene and executing tasks. In Figure 1c, the instructor can ask AILEEN to *identify* the blue cylinder in the setup. AILEEN could respond by highlighting the blue cylinder or pointing at it. By introducing various shades of blue and asking AILEEN to identify blue objects, the instructor can develop an understanding of the generality of AILEEN's conceptual knowledge of blue. In order to verify the robot's task execution knowledge, the instructor can ask the robot to perform an action or a set of actions in the environment in various setups. For example, the instructor can ask AILEEN to *move the red box to the right of the blue cylinder* in various initial states (an example is shown in the configuration presented in Figure 1c). If AILEEN achieves the configuration in Figure 1e, it confirms that AILEEN learned the action and the related concepts successfully.
   - **Describe:** In the *describe* intention, the instructor evaluates if the robot can retrospect on its experience and use its conceptual knowledge to summarize it. For example, in Figure 1c, the instructor can ask AILEEN to describe the scenario, in order to learn whether it has accurately learned the visual and spatial concepts. The instructor can also request the robot to provide a linguistic description of the actions that it performed in the environment. This description illuminates AILEEN's understanding of the task structure to the instructor. For example, a move action done by AILEEN comprises primitive actions pick-up and place. AILEEN should look back at its last instruction to provide a complete description of picking the *red box* up and placing it on a location (identified by (x,y,z) coordinates) that is to the right of the blue cylinder.
3) **Elaborate:** The elaborate intention allows the instructor to gain a deeper understanding of the robot's knowledge when there is a failure. Failures can occur when the concept definitions are not appropriate (either are over-general or over-specific) to support correct scene understanding or task execution. In this intention, when the robot either stops

arbitrarily or indicates that it is unable to do any further, the instructor requests specifics about why and where the failure occurred. For example, AILEEN can fail when asked to *move the red box to the right of the blue cylinder* because it doesn't know how to recognize *box*. In this scenario, the instructor may ask for an elaboration about where the failure occurred. If AILEEN can identify that the failure occurred because it doesn't understand *box*, the instructor can help it learn by providing more examples.

4) **Revise:** The *revise* intention allows the instructor to provide feedback to the robot so that it can revise its scene understanding or task execution knowledge. The instructor would typically *revise* the robot's knowledge after they have *evaluated* that the robot has learned this knowledge. For example, let us assume that instructor believes that Figure 1e does not accurately reflect the relation *right of* after AILEEN moves the red box to that location. The instructor would *revise* the co-ordinates of the red box or provide more examples of *right of*, so that AILEEN can reformulate the plan for future move actions.
   The instructor should be able to provide both online and offline feedback. Thomaz and Breazeal [26] demonstrate that people tend to provide online feedback by providing guidance while the agent is doing the task. In case of offline feedback, the instructor should be able to provide feedback or suggest changes for specific steps, either once the task is completed or when the robot finishes formulating its plan for the task.

V. STUDY PLAN

We are designing a set of human participant studies to understand how human teachers teach robots and what they expect from the learner robot. We have two goals: a) generate evidence that human teachers have intentions delineated in the taxonomy, and identify how they express these intentions, b) discover intentions expressed by the human teachers but are not covered in the taxonomy. We plan to use the recommendations specified in [3] to design the study to validate our proposed taxonomy. Our data collection will include self-assessments, behavioral observations as well as task performance metrics. We will be using examples from AILEEN [20] to structure our studies.

Given the qualitative nature of the taxonomy of intentions, we will begin with semi-structured interviews with non-expert participants. We will provide participants with basic knowledge about interacting with the robot, in terms of its current capabilities. AILEEN learns visual and spatial concepts as well as actions. The participants will be asked to begin by teaching simple concepts (such as shape, color or size), followed by actions and tasks that comprise multiple actions. We will present a typical robot learning environment (such as in Figure 1) with a limited set of objects. We will ask them to provide their instructions, and also ask them to detail how they would manipulate the environment during teaching. For the next iteration of this study, in addition to requesting these instructions, we will create a predefined list of counterfactual questions based on the limitations of AILEEN's abilities. We will also create a codebook to map these questions to expected predetermined instructions. Based on the participant's individual instructions, the interviewer will ask these counterfactual questions and collect responses to these questions. For example, let us take a scenario where the instruction is "Pick up the blue cylinder." An example of a counterfactual question would be "*Aileen responds with I don't know what a cylinder is. What will you do next?*"

We plan to conduct an inductive thematic analysis (which will include open and axial coding) [4] of these interactions, to build a taxonomy of intentions that were observed in the study. Our proposed taxonomy will act as an initial coding of turn-based interactions that we will expand upon, based on our findings. In the following study, we will provide participants with the updated list of intentions identified in the analysis. These intentions will have placeholders that participants can use to refer to specific concepts, objects or verbs. The goal of this study will be to validate the result of the explicit use of this intention framework in a teaching interaction. In order to focus on the instructor's teaching process, we will conduct this study using the Wizard-of-Oz method.

We have currently implemented reasoning mechanisms to process *inform* and *evaluate* intentions in AILEEN. We will use the results of this study to implement additional reasoning mechanisms that allow it to respond to new intentions that show up in these studies. In addition to identifying new intentions that are not presented here, we also want to learn the contexts under which we observe different intentions and the order in which they appear. For example, does the instructor intend to pick up or point to a block (*inform*) before providing the next instruction? When does the instructor use the *evaluate* intention? Is the *inform* intention always followed by an *evaluate* intention?

This study design is still in its initial stages. We expect to have a more concrete design once we conduct our first set of pilot interviews. We look forward to feedback from the community to help scope these studies to achieve our goals.

VI. CONCLUSION

In this paper, we approach the design of Interactive Task Learning robots from a human-centred perspective. Effective ITL robots must implement mechanisms that facilitate smoother interactions with non-expert instructors. Flexible interaction mechanisms enable people to teach robot tasks without the need of extensive reprogramming. In pursuit of this goal, we propose a taxonomy of intentions that are observed in a human-robot teaching scenario. This taxonomy will organize intentions that human teachers typically have when teaching a robot, and provide guidance for ITL robot design. Our goal is to validate and extend this taxonomy by conducting human participant studies with non-expert instructors and learn how these intentions appear in ITL scenarios. Our analyses will be useful in characterizing the state-action space of the interaction model and in developing heuristics to advance the interaction model.

References

[1] Sonya Alexandrova, Maya Cakmak, Kaijen Hsiao, and Leila Takayama. Robot programming by demonstration with interactive action visualizations. In *Robotics: science and systems*. Citeseer, 2014. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.650.6725&rep=rep1&type=pdf.

[2] Forough Arabshahi, Kathryn Mazaitis, Toby Jia-Jun Li, Brad A Myers, and Tom Mitchell. Conversational learning. Preprint on webpage at https://forougha.github.io/paperPDF/Conversational_Learning.pdf.

[3] Cindy L Bethel and Robin R Murphy. Review of human studies methods in HRI and recommendations. *International Journal of Social Robotics*, 2(4):347–359, 2010. URL https://link.springer.com/content/pdf/10.1007/s12369-010-0064-9.pdf.

[4] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative research in psychology*, 3(2):77–101, 2006. URL https://biotap.utk.edu/wp-content/uploads/2019/10/Using-thematic-analysis-in-psychology-1.pdf.pdf.

[5] Cynthia Breazeal, Cory D Kidd, Andrea Lockerd Thomaz, Guy Hoffman, and Matt Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 708–713. IEEE, 2005. URL http://alumni.media.mit.edu/~guy/publications/BreazealIROS05.pdf.

[6] Maya Cakmak and Andrea L Thomaz. Eliciting good teaching from humans for machine learners. *Artificial Intelligence*, 217:198–215, 2014. URL https://hcrlab.cs.washington.edu/assets/pdfs/2014/cakmak2014aij.pdf.

[7] Joyce Y Chai, Rui Fang, Changsong Liu, and Lanbo She. Collaborative Language Grounding Toward Situated Human-Robot Dialogue. *AI Magazine*, 37 (4), 2016. URL http://www.cse.msu.edu/~jchai/Papers/AIMagazine2016.pdf.

[8] Crystal Chao, Maya Cakmak, and Andrea L Thomaz. Transparent active learning for robots. In *Human-Robot Interaction (HRI), 2010 5th ACM/IEEE International Conference on*, pages 317–324. IEEE, 2010. URL https://www.cc.gatech.edu/~athomaz/papers/Chao-hri10.pdf.

[9] Herbert H. Clark and Susan E. Brennan. Grounding in communication. In *Perspectives on socially shared cognition.*, pages 127–149. 1991. ISBN 9781557981219. doi: 10.1037/10096-006. URL http://content.apa.org/books/10096-006.

[10] Barbara Grosz and Candace L Sidner. Attention, intentions, and the structure of discourse. *Computational linguistics*, 1986. URL https://www.aclweb.org/anthology/J86-3001.pdf.

[11] James R Kirk and John E Laird. Learning General and Efficient Representations of Novel Games Through Interactive Instruction. *Advances in Cognitive Systems*, 4, 2016. URL https://pdfs.semanticscholar.org/29ed/0fc57d9b79badfa249fb77b755b24693e6f1.pdf.

[12] John Laird, Kevin Gluck, John Anderson, Kenneth Forbus, Odest Chadwicke Jenkins, Christian Lebiere, Dario Salvucci, Matthias Scheutz, Andrea Thomaz, Greg Trafton, Robert Wray, Shiwali Mohan, and James Kirk. Interactive Task Learning. *IEEE Intelligent Systems*, 32(4):6–21, 2017. URL https://ieeexplore.ieee.org/document/8012335.

[13] John E Laird. *The Soar cognitive architecture*. MIT press, 2012. URL https://mitpress.mit.edu/books/soar-cognitive-architecture.

[14] Daniel A Lazewatsky and William D Smart. Context-sensitive in-the-world interfaces for mobile manipulation robots. In *RO-MAN, 2012 IEEE*, pages 989–994. IEEE, 2012. URL http://people.oregonstate.edu/~smartw/library/papers/2012/roman2012.pdf.

[15] Christopher J MacLellan, Erik Harpstead, Robert P Marinier III, and Kenneth R Koedinger. A Framework for Natural Cognitive System Training Interactions. *Advances in Cognitive Systems*, 6:1–16, 2018. URL http://www.cogsys.org/papers/ACSvol6/article14.pdf.

[16] Matthew Marge and Alexander I Rudnicky. Towards Overcoming Miscommunication in Situated Dialogue by Asking Questions. In *AAAI Fall Symposium: Building Representations of Common Ground with Intelligent Agents*, 2011. URL http://www.cs.cmu.edu/~mrmarge/MargeAAAIfss11.pdf.

[17] Aaron Mininger and John Laird. Interactively Learning Strategies for Handling References to Unseen or Unknown Objects. *Adv. Cogn. Syst*, 5, 2016.

[18] Shiwali Mohan. *From Verbs to Tasks: An Integrated Account of Learning Tasks from Situated Interactive Instruction*. PhD thesis, Department of Computer Science and Engineering,University of Michigan, Ann Arbor, 2015. URL http://web.eecs.umich.edu/~soar/sitemaker/docs/pubs/mohan_thesis_2015.pdf.

[19] Shiwali Mohan, Aaron H Mininger, James R Kirk, and John E Laird. Acquiring Grounded Representations of Words with Situated Interactive Instruction. In *Advances in Cognitive Systems*, volume 2, pages 113–130. Citeseer, 2012. URL http://www.cogsys.org/pdf/paper-3-2-136.pdf.

[20] Shiwali Mohan, Matt Klenk, Matthew Shreve, Kent Evans, Aaron Ang, and John Maxwell. Characterizing an Analogical Concept Memory for Newellian Cognitive Architectures. *arXiv preprint arXiv:2006.01962*, 2020. URL https://arxiv.org/pdf/2006.01962.pdf.

[21] Leah Perlmutter, Eric Kernfeld, and Maya Cakmak. Situated Language Understanding with Human-like and Visualization-Based Transparency. In *Robotics: Science and Systems*, 2016. URL http://www.roboticsproceedings.org/rss12/p40.pdf.

[22] Preeti Ramaraj, Saurav Sahay, Shachi H. Kumar, Walter S. Lasecki, and John E. Laird. Towards using transparency mechanisms to build better mental models. In *Advances in Cognitive Systems: 7th Goal

Reasoning Workshop, volume 7, pages 1–6, 2019. URL https://preetiramaraj.github.io/acs_hri_2019/ACS_HRI_2019_PreetiRamaraj.pdf.

[23] Charles Rich, Candace L Sidner, and Neal Lesh. Collagen: Applying collaborative discourse theory to human-computer interaction. *AI magazine*, 22(4):15–15, 2001. URL https://www.aaai.org/ojs/index.php/aimagazine/article/download/1589/1488/.

[24] Stephanie Rosenthal, Manuela Veloso, and Anind K Dey. Acquiring accurate human responses to robots' questions. *International journal of social robotics*, 4 (2):117–129, 2012. URL http://www.cs.cmu.edu/~mmv/papers/12ijsr-RosenthalVelosoDey.pdf.

[25] Stefanie Tellex, Ross A Knepper, Adrian Li, Daniela Rus, and Nicholas Roy. Asking for Help Using Inverse Semantics. Robotics: Science and Systems Foundation, 2014. URL http://cs.brown.edu/courses/csci2951-k/papers/tellex14.pdf.

[26] Andrea L Thomaz and Cynthia Breazeal. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence*, 172 (6-7):716–737, 2008. URL https://www.sciencedirect.com/science/article/pii/S000437020700135X.

[27] Emily Wu, Nakul Gopalan, James MacGlashan, Stefanie Tellex, and Lawson LS Wong. Social Feedback For Robotic Collaboration. 2016. URL https://h2r.cs.brown.edu/wp-content/uploads/2016/07/wu16thesis.pdf.