Challenges for Communication Decision-Making in Sequential Human-Robot Collaborative Tasks

Vaibhav V. Unhelkar
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
unhelkar@csail.mit.edu

X. Jessie Yang
University of Michigan
Ann-Arbor, Michigan 48109
xijyang@umich.edu

Julie A. Shah
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
julie_a_shah@csail.mit.edu

***Abstract*—Effective communication between teammates is critical to the success of collaboration, including human-robot collaboration. For enabling human robot communication, several modalities are actively being researched — such as, text, speech, visual signals, and legible motion. The design of these modalities is necessary to achieve effective communication; however, it is not sufficient. Communication typically requires both the human and robot to expend resources; and too much or too little communication has the potential to adversely affect task performance. Here, we focus on the complementary but relatively less explored problem of decision-making for communication — i.e., deciding if and when to use an available communication modality during human-robot collaborative tasks. We briefly discuss the modeling and algorithmic challenges for communication decision-making in human-robot and human-agent teams. The article concludes with an overview of our past and on-going research in developing algorithms for communication decision-making in sequential human-robot collaborative tasks.**

I. Introduction

Communication is an important aspect of collaboration. Prior research has demonstrated the need and utility of effective communication for successful collaboration in human-robot teams [15, 17, 33]. For several human-robot collaborative scenarios, communication is essential — such as, answering a user query, or tutoring students [30]. In others (e.g., collaborative manufacturing and disaster response) it aids in sharing observations regarding the environment [39], and making inferences regarding teammates [40].

Indeed, there has been significant interest in communication for human-robot interaction (HRI) [1, 6, 20]. Several modalities for human robot communications are actively being researched — including natural language [23, 36], visual signals [3], and robot motion [10, 21, 35]. Design of these modalities is important for achieving human-robot communication. However, we posit that to achieve effective communication a robot teammate additionally needs the capability to decide if and when to use its communication modality. We refer to this complementary but relatively less explored problem as decision-making for communication in sequential tasks.

The key motivation for our focus on the problem of communication decision-making is the observation that communication incurs costs, i.e., it typically requires both the sender and the receiver to expend resources, which might otherwise be used for improving their collaborative task performance. This results in a trade-off between the cost and benefit of a communication. The challenge of communication decision-making is to resolve this trade-off during execution-time to achieve effective communication in a human-robot team.

Our vision is to develop general decision-making algorithms that work across tasks and communication modality — similar to existing decision-making paradigms in robotics such as motion and task planning. This would allow for using communication decision-making across robots, and allow for benchmarking different approaches. Here, we discuss the modeling and algorithmic challenges arising in the problem of decision-making for communication in human-robot teams. We also mention relevant research on communication in human teams, multi-robot teams and human-robot teams, while noting that our objective is not to provide an exhaustive survey. We conclude with an overview of our past and on-going research in developing algorithms for communication decision-making in sequential human-robot collaborative tasks.

II. Research Challenges

Decision-making for communication in sequential tasks for human-robot teams presents several modeling and algorithmic challenges, primarily arising due to

- challenges in quantifying cost of communication,
- challenges in modeling human teammates,
- difficulty in estimating benefit of communication,
- inherent decentralized nature of multi-agent tasks, and
- the need for execution-time communication decisions.

Next, we describe a few of these research challenges. We note that these challenges span across any communication modality available to the robot.

1) Communication cost: To weigh the cost and benefit of its communication, the robot will need to model communication costs. For multi-robot teams these costs are primarily restricted to physical quantities, such as, power and energy. In contrast, for human teams these costs also include difficult-to-model cognitive variables, such as, limited attention of humans which impacts their ability to process and incorporate communicated information. Communicating too often may lead to information overload [4, 28] or even humans ignoring communications from the robot [11, 26]. Research in interruption management and human team communication provides examples of challenges arising in modeling these costs and attempts to develop corresponding computational models [18, 22].

2) Estimating human state, plan and belief: To assess the benefits of its communications, the robot will need to estimate the state, plan and belief of its teammates [7, 13]. These inferences are challenging — since, the robot might not have full observability of teammate's states/actions. Further, the agents might have access to *only* an inaccurate model of their (potentially dynamic) environment prior to task execution.

In case of multi-agent systems, a common approach is to assume existence of a joint policy from prior coordination, or that all agents are employing similar (potentially decentralized) algorithms to generate local policies [25]. These assumptions are typically not valid when a robot is reasoning regarding its human teammate(s) [9, 27, 37]. Recent research in algorithmic HRI provides successful examples of employing (variants of) inverse reinforcement learning to estimate human's reward/cost function, and then use it to estimate human plans during collaboration [24, 32] in scenarios where accurate environment models are available.

In addition, the robot will also need to estimate the impact of the communication on the state, plan and belief of the human teammate. We refer to this as estimating the reception state of the communication. In multi-agent teams this estimation is often resolved by implementing a pre-specified behavior upon receipt of a communication — such as, a Bayesian update to the belief, or generating a novel plan [31, 42]. However, for human teammates the behavior upon receiving communication might vary based on their workload, perceived utility of robot communication, among other factors [2]. Further, in collaborative tasks, impact of communication on human's plan and belief will additionally depend on the human's estimate of robot's behavior - which might not necessarily be known and/or accurate.

3) Execution-time communication decisions: Lastly, communication decision-making will need to be performed in execution-time with limited computational time. In contrast to planning robot actions in collaborative scenarios, where robot action space is known a priori and action policies can be computed offline, the possible communications that an agent can make varies during execution. Further, in scenarios with inaccurate environment models an agent might receive observations which it cannot anticipate prior to execution but might benefit from sharing them with its teammate. This necessitates the design of novel representations and algorithms for communication decision-making.

In multi-agent systems, several variants of dec-POMDP models [25] and algorithms have been developed that explicitly model communication [14, 29, 34] and reason about communication decisions [31, 41, 42] in cases where accurate environment models are available. However, these approaches assume agents using identical (potentially decentralized) decision-making, which typically does not extend to human-robot teams. In human-agent teamwork, Kamar et al. [19] extend the SharedPlans formalism [16] using Probabilistic recipe trees to efficiently represent joint plans for communication decision-making ; however, this assumes full observability and resolution of challenges described in II-2.

III. Decision-Making for Communication

Despite the challenges described in Sec. II for communication decision-making, human teams are capable of achieving coordination via communications even within time critical and safety critical scenarios. Studies of human team coordination indicate significant differences between communication strategies employed by well performing teams vis-á-vis teams that achieve low task performance [5]. Further, successful teams exhibit anticipatory communication strategies [8, 12, 33]. This suggests that by developing algorithms that allow robots to anticipate and respond to communication needs of its human teammate(s), while respecting the cost and benefit trade-off of communications, one can achieve effective communication in human-robot teams. We conclude with a brief description of our effort towards developing such algorithms.

1) Multi-agent teams in deterministic, unknown environments: As a pre-cursor to developing algorithms for human-robot teams, we have explored communication decision-making in multi-agent teams motivated by challenges of human-robot teams. Specifically, we assume that each agent independently plans its actions and further this planning mechanism might differ across agents; this modeling choice is motivated by the fact that in human-robot teams, human and robot agents could use different action planning approaches. In this first step, we restrict our focus to static domains with deterministic action outcomes and assume common knowledge of the planning mechanism (but not the plan) of each agent. However, we consider that agents do not have complete knowledge of the environment prior to execution (i.e., unknown environment), thus making offline communication decision-making infeasible and highlighting the challenge II-3.

Our algorithm, ConTaCT, makes execution-time communication decisions, and as compared to a baseline approach results in comparable task performance while reducing the number of communications. The details of the task model, agent model and communication algorithm are available in [38]. The performance of ConTaCT is especially desirable for human-robot teams, in which excessive communications from a robot risk affecting human's task performance or leading to human teammates altogether ignoring robot communications.

2) Human-robot teams in unknown environments: We are currently extending the above approach for human-robot teams, by developing solutions to resolve uncertainty in human plans (see II-2), along with the need to make communication decisions in execution time, and with partial observability of human state (see II-3). This is motivated by the fact that not only human planner might be different from that of the robot, but also might not be accurately known to the robot. Towards this we have designed an experiment scenario motivated by applications in collaborative manufacturing, where human and robot need to operate in a shared environment to perform a sequential assembly task. We believe that by developing these solutions for communication decision-making in human-robot teams, there is potential to improve team fluency, task performance, and transparency of robot behavior, which is a precursor to explainable AI.

References

[1] Henny Admoni and Brian Scassellati. Robot Nonverbal Communication as an AI Problem and Solution. In *AAAI Fall Symposium Series*, 2015.

[2] Brian P. Bailey and Shamsi T. Iqbal. Understanding Changes in Mental Workload During Execution of Goal-directed Tasks and Its Application for Interruption Management. *Transactions on Computer-Human Interaction*, 14(4), Jan 2008. doi: 10.1145/1314683.1314689.

[3] Kim Baraka, Stephanie Rosenthal, and Manuela Veloso. Enhancing Human Understanding of a Mobile Robots State and Actions using Expressive Lights. In *RO-MAN*, Columbia University, NY, August 2016.

[4] Michael D. Basil. Multiple Resource Theory. In Norbert M. Seel, editor, *Encyclopedia of the Sciences of Learning*, pages 2384–2385, Boston, MA, 2012.

[5] Abhizna Butchibabu, Christopher Sparano-Huiban, Liz Sonenberg, and Julie Shah. Implicit coordination strategies for effective team communication. *Human Factors*, 58(4):595–610, 2016.

[6] Crystal Chao and Andrea Thomaz. Timing in multimodal turn-taking interactions: control and analysis using timed petri nets. *Journal of Human-Robot Interaction*, 1(1), 2012.

[7] Sharolyn Converse. Shared mental models in expert team decision making. *Individual and group decision making: Current issues*, 221, 1993.

[8] J. Michael Crant. Proactive behavior in organizations. *Journal of Management*, 26(3):435–462, 2000.

[9] Anca D Dragan and Siddhartha S Srinivasa. A policy-blending formalism for shared control. *The International Journal of Robotics Research*, 32(7):790–805, 2013.

[10] Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. Legibility and predictability of robot motion. In *HRI*, pages 301–308. IEEE, 2013.

[11] Mary T. Dzindolet, Scott A. Peterson, Regina A. Pomranky, Linda G. Pierce, and Hall P. Beck. The role of trust in automation reliance . *International Journal of Human-Computer Studies*, 58(6):697 – 718, 2003.

[12] Elliot E. Entin and Daniel Serfaty. Adaptive team coordination. *Human Factors*, 41(2):312–325, 1999.

[13] Michael Georgeff, Barney Pell, Martha Pollack, Milind Tambe, and Michael Wooldridge. The belief-desire-intention model of agency. In *International Workshop on Agent Theories, Architectures, and Languages*, pages 1–10. Springer, 1998.

[14] Claudia V. Goldman and Shlomo Zilberstein. Optimizing information exchange in cooperative multi-agent systems. In *AAMAS*, 2003.

[15] Michael A. Goodrich and Alan C. Schultz. Human-robot interaction: a survey. *Foundations and trends in human-computer interaction*, 1(3):203–275, 2007.

[16] Barbara J. Grosz and Sarit Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2): 269–357, 1996.

[17] Guy Hoffman and Cynthia Breazeal. Collaboration in human-robot teams. In *Intelligent Systems Technical Conference*, 2004.

[18] Eric Horvitz and Johnson Apacible. Learning and reasoning about interruption. In *Intl. Conf. on Multimodal Interfaces*, pages 20–27. ACM, 2003.

[19] Ece Kamar, Ya'akov Gal, and Barbara J. Grosz. Incorporating helpful behavior into collaborative planning. In *AAMAS*, pages 875–882, 2009.

[20] Tobias Kaupp, Alexei Makarenko, and Hugh Durrant-Whyte. Human-robot communication for collaborative decision making A probabilistic approach. *Robotics and Autonomous Systems*, 58(5):444 – 456, 2010.

[21] Thibault Kruse, Patrizia Basili, Stefan Glasauer, and Alexandra Kirsch. Legible robot navigation in the proximity of moving humans. In *Advanced Robotics and its Social Impacts (ARSO)*, pages 83–88. IEEE, 2012.

[22] Jean MacMillan, Elliot E. Entin, and Daniel Serfaty. Communication overhead: The hidden cost of team cognition. In *Team cognition: Process and performance at the inter and intra-individual level*. American Psychological Association, 2004.

[23] Cynthia Matuszek, Evan Herbst, Luke Zettlemoyer, and Dieter Fox. Learning to Parse Natural Language Commands to a Robot Control System. In Jaydev P. Desai, Gregory Dudek, Oussama Khatib, and Vijay Kumar, editors, *Experimental Robotics: The 13th International Symposium on Experimental Robotics*, pages 403–415, Heidelberg, 2013. Springer International Publishing.

[24] Stefanos Nikolaidis, Ramya Ramakrishnan, Keren Gu, and Julie Shah. Efficient Model Learning from Joint-Action Demonstrations for Human-Robot Collaborative Tasks. In *HRI*, pages 189–196. ACM, 2015.

[25] Frans A. Oliehoek and Christopher Amato. The Decentralized POMDP Framework. In *A Concise Introduction to Decentralized POMDPs*, pages 11–32. Springer, 2016.

[26] Raja Parasuraman and Victor Riley. Humans and Automation: Use Misuse Disuse Abuse. *Human Factors*, 39(2):230–253, 1997.

[27] Stefania Pellegrinelli, Henny Admoni, Shervin Javdani, and Siddhartha Srinivasa. Human-robot shared workspace collaboration via hindsight optimization. In *IROS*, pages 831–838. IEEE, 2016.

[28] Robert W Proctor and Trisha Van Zandt. *Human Factors in Simple and Complex Systems*. CRC press, 2008.

[29] David V. Pynadath and Milind Tambe. Multiagent teamwork: Analyzing the optimality and complexity of key theories and models. In *AAMAS*, pages 873–880, 2002.

[30] Aditi Ramachandran, Chien-Ming Huang, and Brian Scassellati. Give Me a Break!: Personalized Timing Strategies to Promote Learning in Robot-Child Tutoring. In *HRI*, pages 146–155. ACM, 2017.

[31] Maayan Roth, Reid Simmons, and Manuela Veloso. Reasoning about joint beliefs for execution-time communication decisions. In *AAMAS*, 2005.

[32] Dorsa Sadigh, Shankar Sastry, Sanjit A. Seshia, and Anca D. Dragan. Planning for autonomous cars that leverages effects on human actions. In *R:SS*, 2016.

[33] Julie Shah and Cynthia Breazeal. An Empirical Analysis of Team Coordination Behaviors and Action Planning With Application to HumanRobot Teaming. *Human Factors*, 52(2):234–245, 2010.

[34] Matthijs T.J. Spaan, Geoffrey J. Gordon, and Nikos Vlassis. Decentralized planning under uncertainty for teams of communicating agents. In *AAMAS*, 2006.

[35] Daniel Szafir, Bilge Mutlu, and Terrence Fong. Communication of intent in assistive free flyers. In *HRI*, pages 358–365. ACM, 2014.

[36] Stefanie Tellex, Pratiksha Thaker, Robin Deits, Dimitar Simeonov, Thomas Kollar, and Nicholas Roy. Toward information theoretic human-robot dialog. In *R:SS*, 2012.

[37] Pete Trautman, Jeremy Ma, Richard M Murray, and Andreas Krause. Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation. *The International Journal of Robotics Research*, 34(3):335–356, 2015.

[38] Vaibhav V. Unhelkar and Julie A. Shah. ConTaCT: deciding to communicate during time-critical collaborative tasks in unknown deterministic domains. In *AAAI*, pages 2544–2550. AAAI Press, 2016.

[39] Sonia Waharte and Niki Trigoni. Supporting search and rescue operations with UAVs. In *Intl. Conf. on Emerging Security Technologies (EST)*, pages 142–147. IEEE, 2010.

[40] David Whitney, Eric Rosen, James MacGlashan, Lawson L.S. Wong, and Stefanie Tellex. Reducing Errors in Object-Fetching Interactions through Social Feedback. In *ICRA*. IEEE, 2017.

[41] Simon A. Williamson, Enrico H. Gerding, and Nicholas R. Jennings. Reward Shaping for Valuing Communications During Multi-Agent Coordination. In *AAMAS*, pages 641–648, 2009.

[42] Feng Wu, Shlomo Zilberstein, and Xiaoping Chen. Online planning for multi-agent systems with bounded communication. *Artificial Intelligence*, 175(2):487–511, 2011.