Enabling Effective Information Sharing in Human-Robot Teams

Vaibhav V. Unhelkar and Julie A. Shah
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139
Email: {unhelkar, julie_a_shah}@csail.mit.edu

I. Motivation

Artificial agents are increasingly interacting with humans to support them in myriad tasks [21, 29, 30]. Gradually, the role of robots in these interactions is transitioning from that of a passive tool for the human user to that of a proactive collaborator. While noteworthy progress has been made in recent years [26], achieving seamless collaboration between robots and humans remains an active and challenging problem.

Realizing the full benefit of human-robot collaboration hinges on the resolution of several challenges, including the development of algorithms that enable robots to (a) model and predict the behavior of humans, and (b) utilize these models to make execution-time decisions. In several domains, such as disaster response and collaborative manufacturing, these algorithmic challenges are further exacerbated due to the adaptive and inter-dependent nature of human and machine behavior, and partial knowledge of the environment that the human-robot team is operating in.

We believe that effective sharing of information between the human and the robot is key to successful resolution of these challenges and achievement of seamless human-robot collaboration. This belief is founded in prior research on human teams [6, 8, 10] and human-machine teams [12, 14, 22], which indicate that *effective* communication is critical to the success of the collaboration. Thus, we are developing algorithms for human-machine collaboration that (a) leverage communication between humans and robots to learn models of human's decision-making and (b) enable robots to effectively share information during execution of human-robot collaborative tasks.

II. Challenges for Effective Information Sharing in Human-Robot Teams

In order to motivate the problem of effective information sharing, we consider a human-robot dyad performing a task with a known, shared objective (Fig. 1). During task execution, both the human and the artificial agent have autonomy over their actions and receive observations from the environment. The information that each agent has regarding the shared environment and the team can differ; this might happen due to decentralized nature of the multi-agent task, incomplete knowledge of the environment, an inaccurate model of the teammate's behavior, or differences in reasoning approaches of the agents.

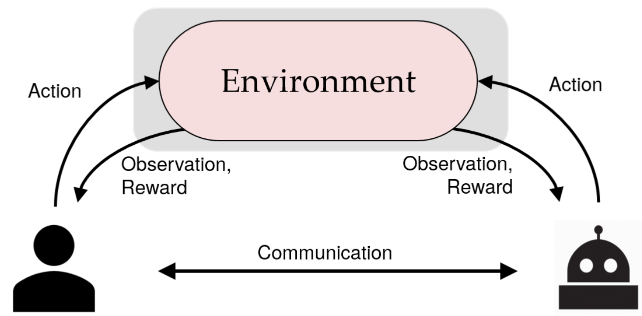

Fig. 1. An abstraction for a human-machine team performing a collaborative task in a shared environment. An (artificial or human) agent will typically have only partial information regarding its environment and teammate. Sharing information when it is beneficial allows agents to make more informed decisions, and the team to improve its task performance.

In such scenarios, sharing information *when it is beneficial* will allow agents to make better-informed decisions, and the team to improve its task performance. However, communicating ineffectively may lead to adverse effects on task performance and situation awareness, information overload or even humans ignoring communications from the robot altogether [4, 9, 18]. Hence, algorithms to decide when it is beneficial for the robot to share information with its human teammate are needed.

Decision-making for communication, however, is challenging due to a variety of reasons; including the reasons for which communication is needed (e.g., decentralized execution, limited knowledge of the team's environment, etc.) as well as the need to estimate value of information during execution time [15, 16]. Despite these challenges, human teams are capable of achieving coordination with the help of effective information sharing even within time-critical and safety-critical scenarios. Successful teams have been found to exhibit anticipatory communication strategies [6, 8, 10, 22]. This suggests that by developing algorithms that allow machines to anticipate and respond to communication needs of its human teammate, while respecting the associated cost and benefit trade-off, one can achieve effective communication in human-machine teams. This abstract summarizes our prior and on-going work towards this problem of effective information sharing in human-robot teams. Here, we assume the presence of a communication modality available to the robot (such as, [3, 17, 25, 24]) and focus on developing algorithms for judicious utilization of such capability in human-robot dyadic interactions.

III. Decision Making for Communication

In multi-agent systems research, a number of works have aimed to design communication strategies that support agents in communicating only when necessary, reducing communication overhead and potentially improving collaborative task performance [33, 23, 31]. Prior decision-theoretic approaches for generating online communication [20, 32, 2] have largely focused on tasks modeled using extensions of Dec-POMDP [5] that include communications [19, 11] and assume complete knowledge of action and sensing uncertainty present in the model. These approaches, however, do not explicitly consider a human teammate and are particularly suited for multi-agent settings where the associated uncertainty – namely, transition and observation probabilities – can be quantified a priori.

For human-robot teams, in contrast, the knowledge representation and decision making of the human and the robot differ. In several teaming scenarios, while (both human and artificial) agents can often achieve the desired outcome from a chosen action in a robust fashion (e.g., through the use of dynamic controllers), they do not have complete knowledge of their environment. For instance, in a disaster response scenario, the map of the team's environment might be inaccurate. This indicates the need of novel algorithms, which address aforementioned challenges of human-robot collaboration, for effective information sharing between humans and robots.

Towards this, we have developed Contact an algorithm that enables robots to make execution-time communication decisions during collaborative tasks [27]. Contact is designed for tasks with known, shared objective wherein the model of the team's environment (specified by transition function) is initially unknown but deterministic in nature. The algorithm allows for decision-making of each agent (i.e., planner used by each agent) to be different; however, the agents are assumed to have common knowledge of the planning behavior, initial state and goal state of their collaborators. In applications, this common knowledge can be derived from prior coordination. Briefly, the algorithm, detailed in [27], includes the following components:

- a model representation maintained by the robot (Fig. 2);
- procedures to update the model with and without communication during decentralized execution; and
- a method to generate communication decisions and trigger re-planning when warranted by gauging value of information using robot's model representation.

By maintaining an estimate of the belief maintained by the human collaborator, the robot can gauge the benefit of sharing novel information during decentralized execution.

To evaluate the algorithm, we have conducted multi-agent simulation evaluations motivated by rescue operations during disaster response scenarios. Our communication decision-making approach was compared against a baseline motivated by the algorithm Dec-Comm [20], where agents do not use local information without communicating it and communicate only if the expected reward is higher post communication. Through these experiments, we observed that while task

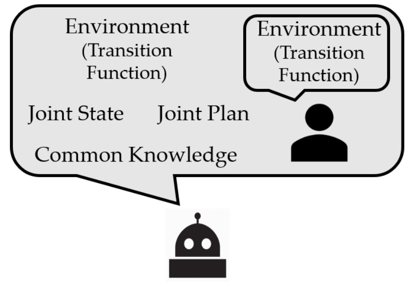

Fig. 2. A cartoon depicting the model representation for a robot performing execution-time communication decision making. The agent maintains an estimate of its environment, and state and plan of its (human) teammate. In addition, the agent also maintains an estimate of the belief of its teammate.

performance was comparable for both the algorithms, teams using Contact could achieve comparable performance with over 60% reduction in the number of communications.

This behavior is indicative of *effective* information sharing and suited for human-robot teams, where excessive communication is undesirable (Section II). In addition, we have conducted human-robot teaming experiments motivated by applications in collaborative manufacturing. An Amazon Alexa device and a template-based speech-to-text module were used for communication between the human and the robot (a Kuka youBot). Application of Contact requires a model of decision-making for the robot's teammate; for the experiments, a model of rational decision-making was ascribed to the human teammate and used to estimate her decisions. The experiments confirmed that our approach to communication decision-making despite making a fewer number of communications could result in comparable task performance.

IV. Learning Decision-Making Models without State Specification

To coordinate its actions with its teammate and to gauge the benefit of sharing information, a collaborative robot needs to estimate the state, plan and belief of its human teammate [7]. While in certain applications it is possible to handcraft a model of human's decision-making or approximate it assuming rational behavior, more generally it is desirable that robots learn these models from data and interaction. Prior approaches to learning models of sequential decision-making require a complete specification of the features that impact an agent's decisions [34]. A key challenge for human-robot collaboration, however, is that the behavior of the human teammate might depend on unknown and latent features, such as preference, trust or behavior of robot teammate [1].

Hence, currently, we are developing human-in-the-loop inference algorithms to characterize the sequential decision-making behavior of a human – without specification of its complete state features and reward optimization criteria – using observations (execution traces) of human's behavior and the ability to query the human [28]. We posit that by learning models that include latent features of human teammate's decision-making, the robot can better gauge the benefit of sharing information and consequently improve the fluency of collaboration [13].

References

[1] Stefano V. Albrecht and Peter Stone. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258:66 – 95, 2018.

[2] Christopher Amato, George D Konidaris, and Leslie P Kaelbling. Planning with macro-actions in decentralized POMDPs. In *Intl. Conf. on Autonomous agents and multi-agent systems (AAMAS)*, pages 1273–1280, 2014.

[3] Kim Baraka, Stephanie Rosenthal, and Manuela Veloso. Enhancing human understanding of a mobile robot's state and actions using expressive lights. In *RO-MAN*, pages 652–657. IEEE, 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] Daniel S Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. The complexity of decentralized control of Markov decision processes. *Mathematics of operations research*, 27(4):819–840, 2002.

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

[7] Sharolyn Converse, JA Cannon-Bowers, and E Salas. 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] 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.

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

[11] Claudia V. Goldman and Shlomo Zilberstein. Optimizing information exchange in cooperative multi-agent systems. In *Intl. Conf. on Autonomous agents and multi-agent systems (AAMAS)*, 2003.

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

[13] Guy Hoffman. Evaluating fluency in human-robot collaboration. In *Workshop on Human-Robot Collaboration at HRI*, 2013.

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

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

[16] 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*. APA, 2004.

[17] Cynthia Matuszek, Evan Herbst, Luke Zettlemoyer, and Dieter Fox. Learning to Parse Natural Language Commands to a Robot Control System. In *International Symposium on Experimental Robotics (ISER)*, pages 403–415, Heidelberg, 2013. Springer International Publishing.

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

[19] David V. Pynadath and Milind Tambe. Multiagent teamwork: Analyzing the optimality and complexity of key theories and models. In *Intl. Conf. on Autonomous agents and multi-agent systems (AAMAS)*, pages 873–880, 2002.

[20] Maayan Roth, Reid Simmons, and Manuela Veloso. Reasoning about joint beliefs for execution-time communication decisions. In *Intl. Conf. on Autonomous agents and multi-agent systems (AAMAS)*, 2005.

[21] Ruhi Sarikaya. The Technology Behind Personal Digital Assistants: An overview of the system architecture and key components. *IEEE Signal Processing Magazine*, 34 (1):67–81, 2017.

[22] 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.

[23] Matthijs T.J. Spaan, Geoffrey J. Gordon, and Nikos Vlassis. Decentralized planning under uncertainty for teams of communicating agents. In *Intl. Conf. on Autonomous agents and multi-agent systems (AAMAS)*, 2006.

[24] Daniel Szafir, Bilge Mutlu, and Terrence Fong. Communication of intent in assistive free flyers. In *Intl. Conf. on Human-Robot Interaction (HRI)*, pages 358–365. ACM, 2014.

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

[26] Andrea Thomaz, Guy Hoffman, Maya Cakmak, et al. Computational human-robot interaction. *Foundations and Trends® in Robotics*, 4(2-3):105–223, 2016.

[27] 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.

[28] Vaibhav V Unhelkar and Julie A Shah. Learning Models of Sequential Decision-Making without Complete State Specification using Bayesian Nonparametric Inference and Active Querying. CSAIL Technical Report, MIT, 2018.

[29] Manuela Veloso, Joydeep Biswas, Brian Coltin, Stephanie Rosenthal, Tom Kollar, Cetin Mericli, Mehdi Samadi, Susana Brandao, and Rodrigo Ventura. Cobots: Collaborative robots servicing multi-floor buildings. In *Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 5446–5447. IEEE, 2012.

[30] 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.

[31] Simon A. Williamson, Enrico H. Gerding, and Nicholas R. Jennings. Reward Shaping for Valuing Communications During Multi-Agent Coordination. In *Intl. Conf. on Autonomous agents and multi-agent systems (AAMAS)*, pages 641–648, 2009.

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

[33] Ping Xuan, Victor Lesser, and Shlomo Zilberstein. Modeling Cooperative Multiagent Problem Solving as Decentralized Decision Processes. In *AAMAS*, 2004.

[34] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum Entropy Inverse Reinforcement Learning. In *AAAI*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.