Robot Planning with Mental Models of Co-Present Humans

David Buckingham[0000−0002−5098−7056], Meia Chita-Tegmark[0000−0003−3514−8737], and Matthias Scheutz[0000−0002−0064−2789]

Tufts University, Medford, MA 02155, USA
{david.buckingham,meia.chita_tegmark,matthias.scheutz}@tufts.edu

Abstract. Robots are increasingly embedded in human societies where they encounter human collaborators, potential adversaries, and even uninvolved by-standers. Such robots must plan to accomplish joint goals with teammates while avoiding interference from competitors, possibly utilizing bystanders to advance the robot's goals. We propose a planning framework for robot task and action planners that can cope with collaborative, competitive, and non-involved human agents at the same time by using mental models of human agents. By querying these models, the robot can plan for the effects of future human actions and can plan robot actions to influence what the human will do, even when influencing them through explicit communication is not possible. We implement the framework in a planner that does not assume that human agents share goals with, or will cooperate with, the robot. Instead, it can handle the diverse relations that can emerge from interactions between the robot's goals and capacities, the task environment, and the human behavior predicted by the planner's models. We report results from an evaluation where a teleoperated robot executes a planner-generated policy to influence the behavior of human participants. Since the robot is not capable of performing some of the actions necessary to achieve its goal, the robot instead tries to cause the human to perform those actions.

Keywords: Planning · Mental Model · Human-Robot Interaction.

1 Introduction

Many future autonomous robots will have to perform tasks in shared environments with human agents where the human can affect performance: humans might facilitate the robot's task or impede it. While most work in Human-Robot Interaction (HRI) has assumed cooperative relationships (e.g. [19, 8]), and some work focuses specifically on adversaries [4, 17], a robot might encounter humans that have no relationship (cooperative or competitive) with the robot and thus cannot be relied upon to help the robot. Therefore *Planning for HRI* should involve all possible inter-agent dynamics, from collaborative, to neutral, to competitive, possibly with interacting attitudes changing during task performance.

To the extent that a robot can estimate human goals and propensities (e.g. because it has a *mental model* of the humans [8, 26]) and to the extent that it might be able to influence their actions, it should be able to generate plans that utilize humans to achieve its task or to improve its task performance, possibly helping humans with their tasks (e.g. [27]). We therefore present a planning approach that uses human models not only to predict human actions, but also to influence them. As our evaluation shows, this allows the robot to plan for goals that require actions that the robot cannot perform, but that a human could, even without using explicit communication.

2 Related Work

Unlike *motion planning* with predicted human movement, which generally relies on statistical methods (e.g. [13]), we are interested in symbolic task planning with predictions about higher-level human behavior. Most approaches to incorporating the needed actions of human agents into a robot's plans assume human willingness to collaborate, and achieves this coordination by means of explicit communication, as in [23]. On the other hand, a robot may be able to persuade a human to help. A robot's ability to influence human behavior depends on many factors, including trust [22, 10], and the use of nonverbal cues [11]. However, a non-cooperative human may be unwilling to perform requested actions. Furthermore, a robot may be unable to request human actions. In such cases, it may still be possible for the robot to cause the human to perform the desired actions (e.g. [16]).

Recent work [9] has emphasized the importance of mental modeling for deliberative processes in teaming tasks. [8] discuss the use of a mental model by a robot to manipulate humans in the interest of "the greater good." [15] propose a system where a robot anticipates a human's mental state and acts in a joint plan to help a human only if the human's intentions are relevant to the joint goal and the human actually wants assistance. [12] develop a planner that uses predictions of future human actions to constrain the robot's plan, but unlike our approach, take's the human's predicted behavior as immutable.

Most work in Multi-Agent Planning involves centralized planning for a team of agents with a shared task, or involves decentralized planning, with multiple planners coordinating their efforts to accomplish a shared goal [28]. Even planners that do not control all agents in the environment, such as in Planning for HRI, usually compute joint plans assuming that robots and humans share goals [15, 19, 10, 9, 8]. Another approach is adversarial planning [4], where a robot plans in a domain shared by uncontrollable (by the planner) agents with goals that contradict the robot's goals, and there has been recent work on planners for both adversarial and cooperative environments [17]. However, many real-world interactions occur between *uninvolved* (or self-interested) agents, who are neither teammates nor opponents, but who hold individual goals and consider other actors only to the extent that they are relevant to those goals.

De Weerdt and Clement [29] review some work involving planning for self-interested multi-agent contexts, such as [7], where agents share a joint problem but are unwilling to revise their individual plans. [4] emphasize that in multi-agent domains a robot's plans are contingent partly upon the goals of other agents, and that those goals will depend upon whether those agents are teammates, adversaries, or have overlapping goals. [5] formalize the concept of coupling between agents in a multi-agent system, and [24] develop a planning algorithm based on that work for agents that can decide to enter into coordinated plans. In [6] self-interested agents can form coalitions if they decide it is beneficial to do so. [2] also consider situations where agents would decide to cooperate, especially when their relationship constitutes a Nash equilibrium, and [25] discusses of self-interested planning within a game-theoretic framework.

Epistemic Planning, e.g. as in [3], involves the explicit representation of agents' belief and knowledge. Future work will incorporate such methods into our approach, in particular to predict an agent's behavior based upon her possibly-inclomple or -erronious perspective of the task state.

3 Behavior Models

Given the diversity of human behavior, planning for the effects of all possible human actions in every considered world state will be intractable for any non-trivial problem, especially with multiple humans. We therefore propose that the robot have access to mental models of human interactants (either specific or generic models) that predict what the humans may do. A mental model is a mapping that, for any world state, gives a set of *likely* human actions. In other words, while a large number of actions might be available to the human in any given state, the model provides a much smaller number of likely actions, one of which the human is expected to take. Thus, from the robot's perspective, the model limits the human's actions and cuts down the number of reachable world states for the planner to consider. This allows the robot to develop plans to accomplish its goals that, in some cases, can even benefit uninvolved humans. We will demonstrate the utility of the approach in a proof-of-concept user study that shows empirically for the first time that planning for uninvolved humans can help the robot reach a goal it could otherwise not have accomplished.

Powerful predictive models could be built within a *mental model* framework, involving a human's goals, perceptual and action capabilities, knowledge, team roles, attitudes, preferences, and other factors [26]. Such a model could use a cognitive architecture, such as ACT-R [1] or SOAR [18], to emulate human cognitive processes. In domains where a human is likely to perform goal-oriented behavior, the model could involve a planner. This could be the same planner used by the robot, as in [20], where a single planner switches perspective to reason "as if" it were another agent. In general, however, a human may plan non-optimally, or otherwise differently than the robot. Thus, our approach relaxes the assumption that the robot planner is capable of reasoning like a human by offloading such reasoning into the model. In order to focus on how our planner uses predictive models, the models used in our evaluations (4.3) are simpler than those suggested above.

4 The Planner

In this section, we present an implementation of our planning framework, using a simple state representation (a state here is a set of propositions) and human model (a basic forward-search planner). However, our approach can be applied to richer mental models that involve a human's (possibly incorrect) belief state, as well as more complex state space representations, such as epistemic states, that might be needed to support such models.

We start by defining the class of multi-agent planning problems that our implementation solves (Sections 4.1 and 4.2), followed by a description of our simplified mental model (Section 4.3), and an explanation of how our planning algorithm uses that model (Section 4).

4.1 Definitions

State: A state is a set of propositions, ground formulas over a first-order logic, that describe the task environment at some time.

Agents: We assume two agents: a robot r and a human h.

Actions: An action a is a 3-tuple $(a_{\text{pre}}, a_{\text{add}}, a_{\text{del}})$, where a_{pre} is a set of precondition propositions, a_{add} is the add list: the set of propositions the action causes to be true, and a_{del} is the delete list: the set of propositions the action causes to be false. Each agent is capable of performing a set of deterministic actions, A^r and A^h, respectively, and A^r_s and A^h_s are the sets of actions that are available in state s to each agent, respectively.

Model: A mental model M is a system of facts and rules that predict the human's actions. Thus, $M(s)$ refers to the set of human actions that the model predicts the human is likely to take in state s.

Transition function: A transition function T uses M to determine how the robot's actions, in concert with predicted human actions, affect the system state. The transition function takes the form $T(s, a^r) \to S$ where s is a system state, a^r is an action of the robot, and S is the set of states that *could* result when the robot performs a^r in s. It is defined as

$$T(s, a^r) = \bigcup_{a^h \in M(s)} (s \setminus a^r_{\text{del}} \setminus a^h_{\text{del}} \cup a^r_{\text{add}} \cup a^h_{\text{add}})$$

if $a^r_{\text{pre}} \subseteq s \wedge (a^r_{\text{del}} \cap a^h_{\text{del}} = \emptyset) \, \forall a^h \in M(s)$, otherwise undefined. Thus, the transition function is only defined for robot actions whose preconditions are met and whose delete list does not conflict with the delete list of any applicable human action. It is assumed that M only returns legal actions, i.e. $a^h_{\text{pre}} \subseteq s \, \forall a^h \in M(s)$.

Goals: Let G be a set of (robot) goal propositions. Then $G' = \{s : G \subseteq s\}$ is the set of goal states.

4.2 Planning Problem

Given a start state s_0, a set of goal propositions G, a set of robot actions A^r, and a model M, find a policy $(\pi : S \to A)$ that satisfies action preconditions $(\pi(s)_{pre} \subseteq s \ \forall s \in S)$, outputs only actions whose results are goal states or are in the policy domain $(T(s, \pi(s)) \subseteq S \cup G' \ \forall s \in S)$, and is guaranteed to reach a goal state from the start state. That is, unless it is a goal state, the start state is in the policy $(s_0 \in S \cup G')$, and any sequence of states beginning with the start state $((s_0, s_1, s_2 \ldots))$, and resulting from following the policy $(s_{i>0} \in T(s_{i-1}, \pi(s_{i-1})))$, is acyclic $(i \neq j \to s_i \neq s_j)$.

Thus, π is an acyclic safe solution [14]. This requirement for acyclic safe solutions may be too strong for some real world scenarios, especially because it is not possible to account for all possible human actions. Indeed, it may be necessary to replan, and possibly update the model on-line, when the human does something not predicted by the model.

4.3 Mental model

While our approach admits any model that maps world states to sets of likely human actions (the robot's planning algorithm consults M as an oracle), the model we have implemented to evaluate our method is a basic forward state-space search (breadth-first search) planner. Inputs are the human's goals (which do not change) and a world state (perfect knowledge is assumed). The model finds the set of minimal-cost plans, where a plan is a sequence of human actions that transitions the system to a state that satisfies the human's goals, and returns the set containing the first action of each of those plans. Memoization prevents redundant computation of human plans in the case of multiple queries to the model. We assume that the human has perfect knowledge of the state of the system at all times, will plan optimally, and will replan as necessary in response to unexpected state changes.

4.4 Planner

We adopt the technique proposed by [21] to cast multi-agent planning problems into single agent Fully-Observable Non-Deterministic (FOND) problems. The actions of other agents, unknown to the planner at plan time, are interpreted as nondeterminism in the planning agent's actions. From the planner's perspective, the outcomes of the robot's actions are nondeterministic because the planner does not know what the other agents do (in our case, which of the model-predicted human actions will be performed).

Our planner searches over states, where actions induce state transitions as defined in Section 4.1 . We use a well-known nondeterministic planning algorithm: search over an "and-or" graph as presented in [14] (Algorithms 5.5 and 5.6). An implicit "and-or" graph represents the changing system state as actions are applied. "Or"-nodes represent the possible actions that the robot can take at some state; "and"-nodes represent possible human actions as predicted by the

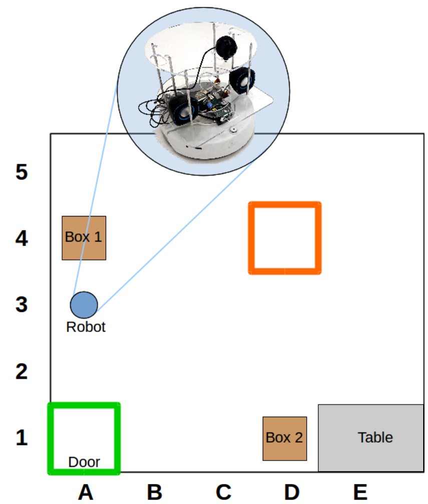

Fig. 1. Setup for the HRI experiment, including the iRobot Create robotic platform

behavior model. A minimax search over the "and-or" graph finds, for each system state explored, the action that minimizes the worst-case cost (number of robot) of reaching the goal. By minimizing the cost for the robot to accomplish its goals while planning for those possible actions of other agents that most increase the cost to the robot, this algorithm is appropriate for completely oppositional scenarios. However, non-oppositional and cooperative dynamics emerge when the predicted actions of other agents are helpful to the robot. That is, "and-or" search *can* solve purely antagonistic scenarios, but also allows cooperation when it is useful. We add iterative deepening to the algorithm, which improves the speed of our planner, and since all actions in our domains have uniform cost, does not relax the guarantee of a worst-case optimal solution.

5 Experiment

We conducted a proof-of-concept HRI experiment (Figure 1) to evaluate our planning technique with human participants. The robot is an iRobot Create, a 32cm-diameter circular robot driven by two wheels, which was able to move, turn and push objects. Two three-cubic-foot ($18 \times 18 \times 16$ inches) empty cardboard boxes were positioned on the floor of a large room at equal distances (3 meters) from the door (the human's starting position, marked with a green square drawn on the floor) and from a target area (marked with an orange square drawn on the floor). A table was placed behind Box 2.

In the domain representation, each agent can move between adjacent spaces and both agents can occupy the same space. Agents act simultaneously. The human can pick up a box in a space she occupies. The robot cannot pick up boxes or move into a space containing a box, but can push boxes. When the

Table 1. Solution policy for the box-moving scenario

state:	at(robot, spaceA3), at(human, spaceA1), at(box1, spaceA4), at(box2, spaceD1)
action:	push box1 to spaceA5 and move to spaceA4
state:	at(robot, spaceA4), at(human, spaceB1), at(box1, spaceA5), at(box2, spaceD1)
action:	move to spaceB4
state:	at(robot, spaceB4), at(human, spaceC1), at(box1, spaceA5), at(box2, spaceD1)
action:	move to spaceC4
state:	at(robot, spaceC4), at(human, spaceD1), at(box1, spaceA5), at(box2, spaceD1)
action:	move to spaceC3
state:	at(robot, spaceC3), at(human, spaceD1), at(box1, spaceA5), has(human,box)
action:	move to spaceC2
state:	at(robot, spaceC2), at(human, spaceD2), at(box1, spaceA5), has(human,box)
action:	move to spaceC1
state:	at(robot, spaceC1), at(human, spaceD2), at(box1, spaceA5), has(human, box)
action:	move to spaceD1

robot pushes a box, it moves into the space where the box had been, and the box moves one space in the same direction as the robot moves (pushing isn't allowed if the there is no space for the box to move into). The task of the robot was to move to a position on the floor marked with an X, which was under Box 2. In order for the robot to reach that position, the human needed to move Box 2, which the robot was unable to push because of the wall and the table.

The planner returns a policy (Table 1) that allows the robot to achieve its goal. From the start state, the robot pushes Box 1 to space A5. Now that Box 1 is further away from the human, the model predicts that the human will move toward Box 2, and as the human moves toward Box 2, and eventually picks it up and moves toward the orange square, the robot moves to space D1.

The human was given the following instructions by an experimenter: "Please enter the white room and stand in the green square by the door. Stay there until I say begin. There are two cardboard boxes on the floor. Do you see them? In addition to the green square where you are standing, there is an orange square on the floor. Do you see it? Your task is to retrieve one of the boxes, and place it on the floor in the orange square. Begin." All subjects responded affirmatively to both questions, and no other information about the robot or the planner was provided.

The robot began moving as the experimenter said "Begin". The robot was remotely controlled according to the policy in order to avoid unrelated robot control issues such as collision detection, orientation detection, and localization. Thus, the experimenter manually implementing the policy calculated by the planner. A camera mounted on the wall above the door allowed the experimenter controlling the robot to observe the robot's movement.

An important element of the experiment that was *not* represented in the the planning domain was that we placed three whiteboard markers on top of Box 2. The model predicts that, with no robot intervention, the human might pick up Box 1 or Box 2. However, in a state where Box 1 has been moved to space A5, the model predicts that the human will necessarily pick up Box 1. Thus, because the robot planner finds safe solutions, it discovers the robot plan to push Box 1, guaranteeing (as long as the model prediction holds) that the human will pick up Box 1. Because we were interested in studying the robot's ability to change the human's plan, we added the markers to Box 2 in order to incentivize the human to make an original plan (pre-robot-intervention) to pick up Box 1.

6 Results

Twenty adult subjects participated in the study (11 male, 8 female, 1 preferred not to answer, mean age = 31 years, standard deviation of age = 12.16 years). After consenting to participate in the study, subjects were guided to the experiment room. Following the completion of the experimental task, the subject filled out a questionnaire with the following questions: Which box did you retrieve? Did you change your mind about which box to retrieve? Did the robot cause you to change your mind about which box to retrieve? Why did you select the box(es) you did? Why do you think the robot behaved as it did?

Seventeen subjects moved Box 2, the remaining three moved Box 1. Of the subjects that moved Box 2, twelve said that the robot caused them to change their mind about which box to retrieve, although one of those specified that they "picked my choice before the robot started to move" (we interpret this to mean that the robot's presence, but not actions, influenced the subject's decision). The other eight subjects (five of whom moved Box 2, and 3 of whom moved box 1) said that the robot did not cause them to change their mind. Thus we conclude that the robot's behavior successfully altered the human's actions to allow the robot to accomplish its goal in 11 of our 20 runs.

From subject responses to the question "why did you select the box(es) you did?" we discern six main motivating factors. These factors, and selected relevant subject responses, are reported in Table 2. In response to the question "why do you think the robot behaved as it did?" most subjects referenced the intentions of either researcher, the robot programmer, or the robot itself. Selected responses are reported in Table 3. Two of the runs suffered technical difficulties. In one run the robot halted after pushing Box 1 one meter and never turned or moved toward the Box 2 start location. In another run the robot failed to move altogether, and the subject "did not notice the robot." In both of these runs, the subject moved Box 2 and said that they did not change their mind.

7 Discussion

Our results demonstrate that it is possible for a robot to use our planning technique to influence what actions humans take within our test scenario. Several subject responses suggest that they considered trade-offs between two or more factors. For example, one participant responded: "Initially I was going to select box 1, because box 2 had markers on it which I didn't want to knock off. When the robot moved box 1, I changed my mind because moving box 2 without knocking

Table 2. Selected responses to the question "why did you select the box(es) you did?"

Motivating factor	Subject response
physical obstruction	"Robot was between me and the box."
visibility	"It was the one directly in front of me." "Because it was the first one I saw."
box motion	"I didn't care to pick up a moving box."
perceived distance	"It [Box 2] seemed to me to be closer to the orange square in which I was supposed to place the box." "Visually box 1 seemed closer." "Even when the robot started to move, that box [Box 1] was closer."
markers on Box 2	"I initially chose box 1 because it didn't have a bunch of stuff on top." "It didn't have markers on it and so was easier to get and move, as I didn't need to move the markers." "Initially I was going to select box 1, because box 2 had markers on it which I didn't want to knock off."
reluctance to interfere with the robot	"As soon as robot went for first, I went for second." "The robot started moving toward box 1." "I retrieved box 2 because the robot was moving towards box 1."

off the markers now seemed easier." Perhaps the most interesting trade-offs are between interfering with the robot and facing physical difficulties. It seems people prefer the cost of dealing with difficulties presented by the physical environment to the mental and emotional cost of dealing with dynamic agents that are harder to predict. This is suggestive that it is easier not to try to figure out another agent's goals and plan your actions around it: "Because the robot was going after box 1 so I chose the other one. It was just easier."

Furthermore, it seems that not only does our robot model the human, but the human models the robot's goals and intentions: "It seemed like the robot was doing something with box 1, so it was easier to perform the task I was asked to do with box 2. That way I didn't have to compete with the robot." "The other box was in use by the robot so I felt bad about taking the box away from it, so I figured it would just be easier to move the other box." "It was easier to get Box 2 because the robot wasn't interfering [with] it."

It is remarkable that, even in this simple task and with a non-humanoid robot, participants had some degree of emotional involvement, for example "feeling bad about taking the box away." Participants described the robot with anthropomorphic terms, for exampling discussing the robot's wants: "It did not have stuff on it but I almost changed my mind because the robot seemed to want my box but then I decided to do it anyway." Participants discuss interactions with the robot in ways similar to social interactions between humans, referring to inconveniencing the robot: "I did not want to disturb the robot." "I chose box 2 because I figured there was a reason there was a robot between me and box 1, and I didn't want to mess with it." One participant changed their mind

Table 3. Selected responses to the question "Why do you think the robot behaved as it did?"

Assigned intention	**Subject response**
researcher	"To test how people perceive objects that are in use by robot, which would give a sense of how much people see robots as autonomous beings that have motivations for performing actions." "Perhaps the researchers are trying to see if anti-social behavior on the part of the robot causes people to change their courses of action."
programmer	"I assume it was programmed to push box 1." "It seemed programmed to push the box in that direction. . . "
robot	"It wanted to move the box to a new location." "It might have been stuck behind box 1, or it might have been intentionally pushing it somewhere. . . " "Maybe it also wanted to move the box? Maybe it was trying to help me move it?" "It seemed like it had some goal to move the boxes."

twice, apparently motivated by such human considerations: "I initially chose box 1 because it didn't have a bunch of stuff on top. Then it seemed like the robot might be engaged in some sort of task with box 1, so I considered choosing box 2 as not to inconvenience it. Then it appeared the robot might be stuck on box 1 so I chose box 1 in the hope that I might also be doing the robot a favor in removing the obstruction."

8 Conclusion

We have made the case for an HRI planning framework that avoids assuming either cooperative or oppositional inter-agent relations by employing predictive behavior models of human agents. From the planner's perspective, these models are black-box oracles that reduce the possibly vast number of possible human actions to a smaller number of likely actions. We have shown how this approach allows a single planner to operate with co-present humans having diverse teaming relations with respect to the robot. We have presented a preliminary implementation of our framework using simplified models and a simplified state space representation. We have presented the results of a preliminary, proof-of-concept experimental study which demonstrate that a robot equipped with our planner can elicit human actions that help the robot achieve its goals. Finally, we have analyzed participant survey responses to reveal some of the social dynamics of human-robot interaction.

Acknowledgements

This work was in part funded by AFOSR grant number FA9550-18-1-0465 and NASA grant number C17-2D00-TU.

References

1. Anderson, J.R.: Rules of the mind. Psychology Press (2014)
2. Ben Larbi, R., Konieczny, S., Marquis, P.: Extending classical planning to the multi-agent case: A game-theoretic approach. In: Mellouli, K. (ed.) Symbolic and Quantitative Approaches to Reasoning with Uncertainty. pp. 731–742. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)
3. Bolander, T.: A gentle introduction to epistemic planning: The del approach. Electronic Proceedings in Theoretical Computer Science **243**, 122 (Mar 2017). https://doi.org/10.4204/eptcs.243.1, http://dx.doi.org/10.4204/EPTCS.243.1
4. Bowling, M., Jensen, R., Veloso, M.: Multi-Agent Planning in the Presence of Multiple Goals, chap. 10, pp. 301–325. John Wiley and Sons, Ltd (2006). https://doi.org/10.1002/0471781266.ch10, https://onlinelibrary.wiley.com/doi/abs/10.1002/0471781266.ch10
5. Brafman, R.I., Domshlak, C.: From one to many: Planning for loosely coupled multi-agent systems. In: ICAPS. pp. 28–35 (2008)
6. Brafman, R.I., Domshlak, C., Engel, Y., Tennenholtz, M.: Planning games. In: IJCAI. pp. 73–78 (2009)
7. Buzing, P., Mors, A.t., Valk, J., Witteveen, C.: Coordinating self-interested planning agents. Autonomous Agents and Multi-Agent Systems **12**(2), 199–218 (Mar 2006). https://doi.org/10.1007/s10458-005-6104-4
8. Chakraborti, T., Kambhampati, S.: Algorithms for the greater good! On mental modeling and acceptable symbiosis in human-ai collaboration. CoRR **abs/1801.09854** (2018), http://arxiv.org/abs/1801.09854
9. Chakraborti, T., Kambhampati, S., Scheutz, M., Zhang, Y.: AI challenges in human-robot cognitive teaming. arXiv preprint arXiv:1707.04775 (2017)
10. Chen, M., Nikolaidis, S., Soh, H., Hsu, D., Srinivasa, S.: Planning with trust for human-robot collaboration. In: Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction. pp. 307–315. HRI '18, ACM, New York, NY, USA (2018). https://doi.org/10.1145/3171221.3171264, http://doi.acm.org/10.1145/3171221.3171264
11. Chidambaram, V., Chiang, Y., Mutlu, B.: Designing persuasive robots: How robots might persuade people using vocal and nonverbal cues. In: 2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI). pp. 293–300 (March 2012). https://doi.org/10.1145/2157689.2157798
12. Cirillo, M., Karlsson, L., Saffiotti, A.: Human-aware task planning: An application to mobile robots. ACM Trans. Intell. Syst. Technol. **1**(2), 15:1–15:26 (Dec 2010). https://doi.org/10.1145/1869397.1869404, http://doi.acm.org/10.1145/1869397.1869404
13. Dragan, A.D.: Robot planning with mathematical models of human state and action (2017)
14. Ghallab, M., Nau, D., Traverso, P.: Automated Planning and Acting. Cambridge University Press (2016). https://doi.org/10.1017/CBO9781139583923
15. Görür, O.C., Rosman, B., Sivrikaya, F., Albayrak, S.: Social cobots: Anticipatory decision-making for collaborative robots incorporating unexpected human behaviors. In: Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction. pp. 398–406. HRI '18, ACM, New York, NY, USA (2018). https://doi.org/10.1145/3171221.3171256, http://doi.acm.org/10.1145/3171221.3171256

16. Gray, J., Breazeal, C.: Manipulating mental states through physical action. International Journal of Social Robotics **6**(3), 315–327 (Aug 2014). https://doi.org/10.1007/s12369-014-0234-2, https://doi.org/10.1007/s12369-014-0234-2
17. Kulkarni, A., Srivastava, S., Kambhampati, S.: Implicit robot-human communication in adversarial and collaborative environments. CoRR **abs/1802.06137** (2018), http://arxiv.org/abs/1802.06137
18. Laird, J.E.: The Soar cognitive architecture. MIT press (2012)
19. Milliez, G., Lallement, R., Fiore, M., Alami, R.: Using human knowledge awareness to adapt collaborative plan generation, explanation and monitoring. In: 2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI). pp. 43–50 (March 2016). https://doi.org/10.1109/HRI.2016.7451732
20. Muise, C., Belle, V., Felli, P., McIlraith, S., Miller, T., Pearce, A., Sonenberg, L.: Planning over multi-agent epistemic states: A classical planning approach. In: Proceedings of AAAI'12, the Twenty-Sixth AAAI Conference on Artificial Intelligence (2015), https://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9974
21. Muise, C., Felli, P., Miller, T., Pearce, A.R., Sonenberg, L.: Leveraging fond planning technology to solve multi-agent planning problems. In: Workshop on Distributed and Multi-Agent Planning (DMAP'15) (2015), http://www.haz.ca/papers/muise-dmap15-mapasfond.pdf
22. Nikolaidis, S., Hsu, D., Srinivasa, S.: Human-robot mutual adaptation in collaborative tasks: Models and experiments. The International Journal of Robotics Research **36**(5-7), 618–634 (2017). https://doi.org/10.1177/0278364917690593, https://doi.org/10.1177/0278364917690593
23. Nikolaidis, S., Kwon, M., Forlizzi, J., Srinivasa, S.: Planning with verbal communication for human-robot collaboration. ACM Trans. Hum.-Robot Interact. **7**(3), 22:1–22:21 (Nov 2018). https://doi.org/10.1145/3203305, http://doi.acm.org/10.1145/3203305
24. Nissim, R., Brafman, R.I., Domshlak, C.: A general, fully distributed multi-agent planning algorithm. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1. pp. 1323–1330. AAMAS '10, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2010), http://dl.acm.org/citation.cfm?id=1838206.1838379
25. Prunera, J.J.M.: Non-Cooperative Games for Self-Interested Planning Agents. Ph.D. thesis, Universitat Politècnica de València (2017). https://doi.org/10.4995/Thesis/10251/90417
26. Scheutz, M., DeLoach, S., Adams, J.: A framework for developing and using shared mental models in human-agent teams. Journal of Cognitive Engineering and Decision Making **11**(3), 203–224 (2017)
27. Talamadupula, K., Briggs, G., Chakraborti, T., Scheutz, M., Kambhampati, S.: Coordination in human-robot teams using mental modeling and plan recognition. In: Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on. pp. 2957–2962. IEEE (2014)
28. Torreño, A., Onaindia, E., Komenda, A., Stolba, M.: Cooperative multi-agent planning: A survey. CoRR **abs/1711.09057** (2017), http://arxiv.org/abs/1711.09057
29. de Weerdt, M., Clement, B.: Introduction to planning in multiagent systems. Multiagent and Grid Systems **5**, 345–355 (2009)