



HAL
open science

A Knowledge Rich Task Planning Framework for Human-Robot Collaboration

Shashank Shekhar, Anthony Favier, Rachid Alami, Madalina Croitoru

► **To cite this version:**

Shashank Shekhar, Anthony Favier, Rachid Alami, Madalina Croitoru. A Knowledge Rich Task Planning Framework for Human-Robot Collaboration. AI 2023 - 43rd SGAI International Conference on Artificial Intelligence, Dec 2023, Cambridge, United Kingdom. pp.259-265, 10.1007/978-3-031-47994-6_25 . hal-04385641

HAL Id: hal-04385641

<https://laas.hal.science/hal-04385641v1>

Submitted on 17 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Public Domain

A Knowledge Rich Task Planning Framework for Human-Robot Collaboration

Shashank Shekhar¹, Anthony Favier^{1,2}, Rachid Alami^{1,2}, and Madalina Croitoru³

¹ LAAS-CNRS, Universite de Toulouse, CNRS, INSA, Toulouse, France

² Artificial and Natural Intelligence Toulouse Institute (ANITI)

{shashank.shekhar, anthony.favier, rachid.alami}@laas.fr

³ LIRMM, University of Montpellier, Montpellier, France

madalina.croitoru@lirmm.fr

Abstract. In this paper, we position ourselves within the context of collaborative planning. Drawing upon our recent research, we introduce a novel, knowledge rich task planning framework that represents and reasons and effectively addresses agents’ *first-order* false beliefs to solving a task planning problem with a shared goal. Our contributions to this work are as follows: First, to enhance the reasoning abilities of an existing framework, an intuitive observability model for situation assessment is addressed, and for that, an improved knowledge modeling is considered. This effectively captures agents’ predictable false beliefs and motivates us to exploit the power of off-the-shelf knowledge reasoners. Second, a new planning approach that incorporates this improved encoding. And, to show the effectiveness of our planner and present our initial findings and proof of concept, we conduct a thorough use case analysis.

Keywords: Collaborative planning · Knowledge reasoning · Inference · Theory of Mind (ToM) · Human-aware task planning.

1 Introduction

Human-robot collaboration/interaction (HRC/I) is a current research focus due to the growing number of robot-assisted applications [13]. Collaborative robots add clear value to real-world domains like household [17], workshops [15], or medical facilities [12].

In this work, we address a task planning problem where humans and robots jointly achieve a shared goal. For seamless human-robot collaboration, we believe that it is essential not to restrict human behaviors and, hence, consider humans as uncontrollable agents [7]. Our focus is on developing a planning framework that generates collaborative behaviors for an autonomous robot while anticipating the possible human behaviors and being congruent to humans’ choices. Frameworks that support such requirements, in general, include HAT-P/EHDA [6], ADACORL [16] and CommPlan [17], to cite but a few.

Restricted observability available in the environment often leads agents to form false beliefs about the progress of the task and other agents’ beliefs. For

instance, an agent may rely on “only believing what they see or can infer” to update or form new beliefs. For the robot to anticipate such situations while planning its behavior to collaborate, it should consider Theory of Mind (ToM) to maintain and manage agents’ distinct beliefs.

Our recently proposed planner extends the HATP/EHDA framework and considers the first-order ToM, that is, managing (false) beliefs about the physical world [10]. Roughly, its main idea works as follows: First, the existing problem encoding is enhanced. To be specific, we categorize each environmental property as either *observable* – meaning its real value can be assessed directly from the environment – or *indirectly inferable*. In the latter case, the changes in the property’s status persist in the environment when an action is executed, but cannot be observed directly. However, an agent can assess this knowledge either as an observer or as an actor, or when they are communicated. Second, to consider restricted observability, we introduce the concept of *co-presence* and *co-location*. We formulate the environment such that a collection of discrete places where action manipulation can take place, permitting the enhanced encoding to enable precise action executions, or an agent assessing the value of an environmental property when they are physically situated within a specified place.

The encoding is used in HATP/EHDA as input, modifying the planner to use improved knowledge modeling to handle better the evolution of agents’ first-order (false) beliefs and estimate the actions humans are likely to perform. (Due to space restrictions, we will provide only a high-level description of it.)

However, some limitations to the above line of work are: First, the context under which the reasoning works is limited. For example, associating an observable state property directly to a discrete place in the environment is not always intuitive and straightforward and enables sound reasoning. To understand this, consider the following scenario: A robot holding a coffee mug can be assessed from where the robot is currently situated. Moreover, if the robot moves from one place to another, then the *condition* under which the robot holding the coffee mug can be assessed, also changes implicitly. Or, if a table is moved, then the condition to assess the orientation of blocks on top of it would change implicitly, too. Second, the system’s ability to *refute* a false belief is limited.

In this regard, our contributions to this paper are the following.

1. We first extend and generalize the concepts introduced earlier for observing the value of an environmental property. We provide a new encoding that is much cleaner and intuitive and achieves a sound resolution and refutation.
2. Second, the planning algorithm is improved by incorporating the new encoding for better managing the evolution of agents’ predictable (false) beliefs. To show the efficacy of our planning approach and present our preliminary findings and proof of concept, we conduct a use case analysis.

2 The Illustrative Example

Consider an adapted version of the scenario appearing in [10]. The scenario includes a pot of water on the stove, while the pasta packet is in the adja-

cent room. One sub-task is to pour the pasta into the pot, which can only be done after turning on the stove (results in holding `StoveOn`) and adding salt to the water (results in `SaltIn`). Suppose, the robot is responsible for turning on the stove (`turn-stove-on`), while the human is tasked with fetching the pasta (`grab-pasta`, however, for that, they need to `move` to the other room where the *pasta* packet is kept) and pouring it into the pot (`pour-pasta`). Both the robot and the human have the capability to add salt (`add-salt`) to the pot. Humans can be *uncontrollable* but assumed to be *rational*, so they can choose to either first add salt, or first fetch the pasta. Next, the robot’s actions depend on the human’s decision, resulting in the generation of different false beliefs. Such variability in beliefs observed can be precisely attributed to the restricted observability and available human choices within the system.

Suppose, when the human is temporarily absent to bring pasta, the robot can act in *kitchen*. When they return, they can “*assess*” whether the robot turned on the stove, which is observable. However, the presence of salt (*SaltIn*) is not directly observable to her. To prevent misunderstandings, a proactive robot needs to anticipate false beliefs held by humans. Here, humans may mistakenly believe that salt is not added to the pasta or be uncertain about it.

3 The Basic Architecture and Recent Advancements

We provide a brief overview of the HATP/EHDA (*human-aware task planning by emulating human decisions and actions*) architecture [6], building upon previous research, e.g. [1, 2, 8, 5]. We also discuss recent advancements that have been made based on its planning framework. (For a better understanding of the basic terminologies, definitions, and planning models, we refer readers to [11].)

The HATP/EHDA architecture considers: (1.) distinct capabilities and characteristics of humans and robots, (2.) utilization of human task and/or action models by the robot to inform decision-making, considering human capabilities and environmental dynamics, and (3.) flexibility for humans to choose an action from the available choices, which is congruent to the shared goal.

The HATP/EHDA planning scheme utilizes a two-agent model consisting of a *human* and a *robot*, represented as two-agent-HTN (*hierarchical task networks*) [11], adapted as per the requirements. Roughly speaking, each agent has their own initial belief state, action model & methods, and task network. Our focus is on a sequential task to be solved, where both agents have a shared initial task network that needs to be decomposed. Considering perfectly aligned agents’ beliefs with the real world to begin with, the classical scheme assumes that the belief state of an agent consists of properties that are either *true* or *false*. Assuming the environment is fully observable, the planning scheme uses agents’ task/action models and beliefs to decompose their task network into valid primitive components, resulting in an implicitly coordinated plan for the robot within a joint solution policy [6]. This policy encompasses both the robot’s actions and an estimated policy for the human. For more details, see [10].

A new formalism is proposed to enhance the planning scheme’s robustness for estimating human initiatives when dealing with restricted observability. It considers the first-order Theory of Mind for managing agents’ (false) belief updates, resulting in improved management of their mental state and enhanced collaboration [9]. As we pointed earlier, this scheme addresses the need for a less abstract problem formulation, and conceptualizes place-based *situation-assessment* models, and hence capturing collaboration nuances more accurately.

An offline planning approach is proposed to detect predictable false beliefs of humans and address them through explicit communication actions if mandatory w.r.t. the goal. It promotes more accurate understanding and alignment between the robot and humans, allowing for minimal communication. Another way, as mentioned in [10], is to delay critical actions, enabling humans to reach a state s.t. the robot’s actions become perceivable.

4 Improved Problem Encoding

We consider two types of agents. First is *GT*, which stands for *ground truth*, an omnipresent (virtual) agent that cannot achieve any task, however, it immediately knows all the effects when an action is applied. The second type refers to real agents, e.g., a robot. Real agents can have false beliefs, which means their beliefs about the world may diverge from the *GT*’s beliefs. Moreover, no real-world uncertainty is considered w.r.t. agents’ beliefs.

An environmental property is Boolean type, and is augmented by an argument, *?a - agent* (e.g., *Human/Robot*), which is also its first argument. It represents that the agent believes that this property holds true given it is contained in the agent’s beliefs. In this work, we restrict ourselves to those properties that are directly observable in the environment, while relying on the existing concepts with some minor updates required, for managing beliefs w.r.t. the facts that are indirectly inferable. To manage appropriate dynamics of the world, we augment the parameter list of *relevant* state properties which are observable with an argument *?p - Places* (at the end of its parameter list). However, for those that are observable but associating them directly to a place is not intuitive to the domain modeler, e.g., *StoveOn* is a property of a stove and can be assessed from where it is currently situated, we use a different strategy to cater to them.

With a lifted (primitive) action schema, we specify how it is only modeled to affect the beliefs of the *GT* agent, in principle, focusing on observable state properties.

```
:action turn-stove-on (?a - agent ?st - stove ?k - kitchen)
  :precondition (and (not(= GT ?a))(at ?a ?a ?k)(stoveAt ?a ?st ?k) ...)
  :effect (and (stoveOn GT ?st) ...)
```

Real agents assess available information systematically from the environment, and for that, we include the following in our problem specification:

Language & Interpretation: Informally, to assess the value of a property, an agent needs to respect certain *condition(s)*. Of course, a condition could be that

the agent is co-located, as supported by the existing system. However, inspired by the work in [4], we generalize this idea and specify *rules* for knowledge reasoning to capture more complex contexts. For one observable property (representing the *consequent*), conditions for assessing its value are “formulated” in the *antecedent* of a knowledge rule. Formally, a rule is of the form $p(a_1, x, y), q(x) \rightarrow s(y, b_1)$, which also means that $\forall x \forall y (p(a_1, x, y), q(x) \rightarrow s(y, b_1))$, where x and y are free variables, and p , q and s are names. a_1 and b_1 are constants.

For example, when an agent is situated at the stove’s location, it will acknowledge the current status of *StoveOn*. A corresponding rule to it is, R1: $\forall s \forall k \forall a (stoveAt(a, s, k), stoveon(GT, s), at(a, a, k) \rightarrow stoveon(a, s)$. Note that the GT’s beliefs play a key role here. Similarly, we can write other important rules for our example domain: R2: $\forall a_1 \forall k (at(GT, a_1, k) \rightarrow at(a_1, a_1, k))$, R3: $\forall a_1 \forall a_2 \forall k (at(GT, a_1, k), at(a_2, a_2, k) \rightarrow at(a_2, a_1, k))$, and R4: $\forall a \forall s (\neg stoveat(GT, s, kitchen), at(a, a, kitchen) \rightarrow \neg stoveat(a, s, kitchen))$. It treats the stove as a mobile object, which alters the context to also evaluate the status of *StoveOn* along with determining its current location.

Page restrictions allow us to offer a concise overview only, reserving detailed semantics for future research. In rules, it is important to note that the implications must not be contraposed. Anything that cannot be evaluated as true is considered false for agents. Our focus on *stratified* rule sets avoids potential issues related to unsafe negation usage and ambiguity in interpretation [3, 14].

5 Incorporating New Encoding for Planning

We adapt our planning mechanism presented in [10], to assess new encoding’s effectiveness. It assumes that the robot’s beliefs never deviate from GT’s beliefs, and hence essentially two distinct belief states are handled by the scheme.

Two questions arise at this point: First, how does the state transition occur during planning and how are the agents’ (*false*) beliefs updated? Second, how are the relevant false beliefs detected and addressed? We rely on the existing mechanism to address the second question and how beliefs are updated when agents do indirect inference. Our main focus will be to address how the new observability model is incorporated into the planning workflow.

Once an action is executed, such as the human moving to the kitchen, the process of situation assessment is initiated to evaluate the environment from the perspectives of individual agents. Here, our approach diverges from current methods. Instead of directly matching the *co-location* of agents with state properties and assessing the latter’s values, we invoke an external reasoner to deduce such knowledge. Using the agent’s beliefs, the reasoner verifies the formula that captures the context specified (in the rules) to deduce new knowledge or to refute a fact that humans believe but is no longer true in the robot’s beliefs. This departure from the existing encoding allows us to remove the hardcoded context within an action. For instance, if the human moves from one room to another, carrying the pasta packet that was initially in the room, the new location for assessing the packet switches to the kitchen. Moreover, we needed to hardcode

it within each action that updates the place of an agent in the environment. However, with the new encoding, if the robot sees the human in the kitchen, then given the domain rules, it can assess whether the human is holding the pasta packet, by their presence in the kitchen.

Acknowledgments. This work has been partially financed by Défi Clé “Robotique Centrée sur l’Humain” supported by Région Occitanie and the Artificial and Natural Intelligence Toulouse Institute (ANITI).

References

1. Alami, R., Clodic, A., Montreuil, V., Sisbot, E.A., Chatila, R.: Toward Human-Aware Robot Task Planning. In: AAAI spring symposium 2006: to boldly go where no human-robot team has gone before (2006)
2. Alili, S., Warnier, M., Ali, M., Alami, R.: Planning and plan-execution for human-robot cooperative task achievement. In: Proc. of ICAPS (2009)
3. Apt, K. R., Blair, H. A., Walker, A.: Towards a Theory of Declarative Knowledge. In: Minker J. (ed.) Foundations of Deductive Databases and Logic Programming, pp. 89-148, Morgan Kaufmann (1988)
4. Baget, J. F., Leclère, M., Mugnier, M. L., Rocher, S., Sipieter, C.: Graal: A toolkit for query answering with existential rules. In: Rule Technologies: Foundations, Tools, and Applications: 9th International Symposium, RuleML (2015)
5. Buisan G., Alami, R.: A Human-Aware Task Planner Explicitly Reasoning About Human and Robot Decision, Action and Reaction. In: Proc. of HRI (2021)
6. Buisan G., Favier, A., Mayima, A., Alami, R.: HATP/EHDA: A Robot Task Planner Anticipating and Eliciting Human Decisions and Actions. In: Proc. of ICRA (2022)
7. Cirillo, M., Karlsson, L., Saffiotti, A.: A Human-Aware Robot Task Planner. In: Proc. of ICAPS (2009)
8. De Silva, L., Lallement, R., Alami, R.: The HATP hierarchical planner: Formalisation and an initial study of its usability and practicality. In: Proc. of IROS (2015)
9. Favier, A., Shekhar, S., Alami, R.: Robust Planning for Human-Robot Joint Tasks with Explicit Reasoning on Human Mental State. In: AI-HRI, AAAI Symposium (2022)
10. Favier, A., Shekhar, S., Alami, R.: Anticipating False Beliefs and Planning Pertinent Reactions in Human-Aware Task Planning with Models of Theory of Mind. In: PlanRob 2023, the ICAPS Workshop (2023)
11. Ghallab, M., Nau, D. S., Traverso, P.: Automated Planning - Theory and Practice. Elsevier (2004)
12. Jacob, M. G., Li, Y. T., Akingba, G. A., Wachs, J. P.: Collaboration with a robotic scrub nurse. Communications of the ACM, vol. 56(5), pp. 68–75, (2013)
13. Selvaggio, M., Cognetti, M., Nikolaidis, S., Ivaldi, S., and Siciliano, B.: Autonomy in physical human-robot interaction: A brief survey. IEEE Rob. Autom. Lett. (2021)
14. Thiébaux, S., Hoffmann, J., Nebel, B.: In Defense of PDDL Axioms. In: Proc. of IJCAI (2003)
15. Unhelkar, V. V., Lasota, P. A., Tyroller, Q., Buhai, R. D., and Marceau, L., Deml, B., Shah, J. A.: Human-aware robotic assistant for collaborative assembly: Integrating human motion prediction with planning in time. IEEE Rob. Autom. Lett., vol. 3(3), pp. 2394–2401, (2018)

16. Unhelkar, V. V., Li, S., Shah, J. A.: Semi-Supervised Learning of Decision-Making Models for Human-Robot Collaboration. In: Proc. of CoRL (2019)
17. Unhelkar, V. V., Li, S., Shah, J. A.: Decision-Making for Bidirectional Communication in Sequential Human-Robot Collaborative Tasks. In: Proc. of HRI (2020)