

Improved Task Planning through Failure Anticipation in Human-Robot Collaboration

Silvia Izquierdo-Badiola^{1,3}, Gerard Canal², Carlos Rizzo¹ and Guillem Alenyà³

Abstract—Human-Robot Collaboration (HRC) has become a major trend in robotics in recent years with the idea of combining the strengths from both humans and robots. In order to share the work to be done, many task planning approaches have been implemented. However, they don't fully satisfy the required adaptability in human-robot collaborative tasks, with most approaches not considering neither the state of the human partner nor the possibility of adapting the collaborative plan during execution or even anticipating failures.

In this paper, we present a planning system for human-robot collaborative plans that takes into account the agents' states and deals with unforeseen human behaviour, by replanning in anticipation when the human state changes to prevent action failure. The human state is defined in terms of *capacity*, *knowledge* and *motivation*. The system has been implemented in a standardised environment using the Planning Domain Definition Language (PDDL) and the modular ROSPlan framework, and we have validated the approach in multiple simulation settings. Our results show that using the human model fosters an appropriate task allocation while allowing failure anticipation, replanning in time to prevent it.

I. INTRODUCTION

Human-Robot Collaboration (HRC) has undoubtedly become a major trend in robotics in recent years. The idea of combining the comparative strengths from both humans and robots to improve the efficiency and productivity of plans is very appealing and promising in several applications. One of these application areas is the industrial scenario, which is characteristic for having teams of humans and robots cooperating. The main objective of HRC is for the human and the robot to achieve a shared goal through collaborating as a team, by executing tasks that contribute towards achieving the common goal. We refer to this as collaborative plans.

The main challenge in HRC is dealing with the stochastic nature of a human-centric dynamic environment. The human becomes an *uncontrollable* agent, making flexibility and adaptability key features when planning for shared goals in

these scenarios. In human-human collaboration, humans plan ahead to achieve a goal but generally manage to automatically adapt to the changes in the dynamic and uncertain real world by re-evaluating the current state of the world and the partner and replanning to reach the goal. Furthermore, humans are usually able to anticipate events based on the partner's behaviour.

Automated task planning for robots aims to achieve this, by generating a plan consisting in a sequence of actions to accomplish a goal. Nevertheless, existing task planning approaches don't satisfy the required adaptability in human-robot collaborative tasks, and we believe a major issue is the lack of modelling of the partner and its integration into the automated task planning system.

In this work, we intend to generate plans that solve a human-robot collaborative task whilst dealing with the unplanned behaviour of the human partner. Our hypothesis is that by modelling and taking into account the agents' states (physical and mental) in the plan generation and execution, the robot can better understand non-planned human behaviour and react accordingly, adding adaptability and anticipation to unforeseen situations through modelling and replanning.

This paper intends to contribute towards closing the gap between human state modelling and AI planning to anticipate and avoid failures due to human behaviour in human-robot collaborative tasks. We intend to answer the following research questions:

- **R1** - Can HRC plans be improved in terms of success rate and action failure prevention by integrating an agent's state model into the planning stage?
- **R2** - Can an appropriate and balanced task allocation between the agents still be achieved in this case?

The contribution of this paper consists in a planning system for human-robot collaborative plans that takes into account the agents' states and avoids failures due to unforeseen human behaviour, by replanning if the human state changes. The human state is defined in terms of *capacity*, *knowledge* and *motivation*. The system generates two parallel sequences of actions assigned to the robot and the human. By integrating the human state into the planning component, the system can anticipate the failure of certain actions assigned to the human and replan ahead. We implement this system in a standardised environment using the Planning Domain Definition Language (PDDL) [1] and the modular ROSPlan framework [2], so that it can be easily extended to a variety of scenarios and applications involving human-robot collaboration. We then validate the described system in

¹Eurecat, Centre Tecnològic de Catalunya, Robotics and Automation Unit, Barcelona, Spain {silvia.izquierdo, carlos.rizzo}@eurecat.org

²Department of Informatics, King's College London, United Kingdom gerard.canal@kcl.ac.uk

³Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain {sizquierdo, galenya}@iri.upc.edu

*S. Izquierdo is a fellow of Eurecat's *Vicente López* PhD grant program. This work was financially supported by the Catalan Government through the funding grant ACCIÓ-Eurecat, and partially supported by MCIN/ AEI /10.13039/501100011033 under the project CHLOE-GRAPH (PID2020-118649RB-I00); by MCIN/ AEI /10.13039/501100011033 and by the European Union NextGenerationEU/PRTR under the project COHERENT (PCI2020-120718-2). Gerard Canal has been partially supported by EPSRC grant THuMP (EP/R033722/1) and by the Royal Academy of Engineering and the Office of the Chief Science Adviser for National Security under the UK Intelligence Community Postdoctoral Research Fellowship programme.

simulated scenarios where failures due to human behaviour occur, demonstrating the advantages of integrating a human model in the planner for the generation and execution of collaborative plans.

The remaining of the paper is structured as follows: Section II provides an overview of related work. Section III describes the methodology and proposed framework for improved task planning through failure anticipation. Then, we present the results of our evaluation in Section IV. Finally, we point out the conclusions and discuss future work in Section V.

II. RELATED WORK

When humans and robots work collaboratively towards a common goal, the planning framework is expected to generate joint plans that take into account the inferred or predicted human behaviour, assigning the appropriate complementary actions and adapting them to the capabilities of the participants. An example of a planning framework for human-robot collaboration is the Hierarchical Agent-based Task Planner (HATP) [3], a Hierarchical Task Network (HTN) planner adapted for robotics that can split a plan into multiple streams for multiple agents. The system incorporates user-supplied social rules and interleaves HATP planning with geometry planning. The system is further developed in [4], where the robot can infer the users' intentions and monitor their actions to replan if the initial plan fails. A demonstration of the system for both multiple cooperating UAVs and human-robot collaborative manipulation is described in [5]. The partner's willingness to contribute is not considered in this work. In [6], the human's flexibility during the collaboration is taken into account by selecting a pre-defined negotiation or an adaptation mode. The authors aim to adapt the generated shared plans to the human's preferences using simple mechanisms involving cost estimation to decide who should perform the actions during the plan execution. Canal *et al.* [7] adapt robot task planning to user preferences by means of modifying the plan's cost, learning from user feedback at the end of the execution. However, they use a predefined user model, not taking into account the user's state during execution. In [8], Bezrucav *et al.* aim to improve collaborative plans to adapt to unforeseen situations through replanning, focusing on timing in the plan generation and execution. Our work differs from the ones presented here in the following ways: unlike ours, these works adapt the plan based on the user's actions rather than the user state. Replanning is therefore triggered after a failure in the collaboration. In our work, integrating an agent model and replanning based on the user state allows for anticipation of failures, leading to a smoother collaboration. Some works dive into the idea of integrating the human mental state in the plan generation and execution. Devin *et al.* [9] explicitly use an estimation of the human mental state in the shared plan execution, focusing on producing a less intrusive behaviour of the robot. The human mental state concerns information about the environment and the state of the plan from the human perspective, but doesn't consider the

emotional state of the human, nor the human commitment towards an agreed shared goal. In [10], the authors propose a solution for human-robot collaborative plans where the robot acts in an adaptive way based on the human partner's level of knowledge about the task. The level of knowledge is obtained from action outcomes and interactions, and the focus of the adaptation lies on choosing between two policies (teaching or efficiency) which provide different levels of interactions. Again, the level of engagement of the human partner is not considered. Furthermore, the human mental state is not directly integrated into the plan generation. Instead, the plan is initially generated before being adapted using external algorithms. In this paper, we propose a direct integration of the agent state as action costs that will enable the direct generation of a feasible and optimal plan based on the current world's and partner's state.

Another line of works focuses more on modelling the state of the human partner during a collaboration, in order to anticipate their behaviour and react to it. Hoffman *et al.* [11] remove the assumptions that the human is always committed to the goal or willing to receive help during a plan. Instead of only considering the human's visual and positional perspective-taking, they develop an architecture that integrates human's intention and emotion into a robot's decision-making to decide whether to help the human or not. They focus on the estimation and modeling of human mental states and their effects on assisting the human with certain actions. In [12], [13], the authors attempt to model several types of human partners through MDPs and POMDPs in order to estimate correctly and non-intrusively when to intervene and assist the human. Their objective is to anticipate the human state in terms of stubbornness, tiredness, and distraction, and to be able to improve the collaboration when facing unexpected human behaviour. Nikolaidis *et al.* [14] model the concept of human adaptability level in order to update an MOMDP policy for a human-robot collaborative task. Compared to our work, these works only deal with decision making on previously defined policies, rather than automatically generating full collaborative plans. The applications and range of plans are therefore vastly limited in terms of variety and complexity.

Despite the research efforts made towards improving AI planning human-robot collaborative systems, several challenges remain open for obtaining practical and efficient versions of such systems. A number of works show the potential of using different agents' physical and mental state models to improve adaptability during the generation and execution of collaborative plans. However, there are still several limitations in the presented systems, and we believe the potential of merging models and planning has not been fully exploited. By targeting the gap between maintaining an agent state model and explicitly using it in shared plan generation and execution, we believe the uncertainty can be decreased, and unexpected events can be anticipated and better dealt with. The integration of the human model and AI planning into an implemented framework is not straightforward and goes beyond the state-of-the-art, but is an

essential element to achieve a system capable of generating and executing flexible human-robot collaborative plans.

To the best of our knowledge, there is no work that robustly integrates the agents' states and emotions directly into the automated generation of a shared plan, in order to anticipate failures rather than react to them. This paper aims to contribute towards these efforts, implementing and evaluating a framework integrating these components.

III. METHODOLOGY

The framework. We propose a system consisting in the architecture shown in Figure 1, with the aim of generating plans for the robot and the human to achieve a task collaboratively, taking into account any change in the state of the team members. The approach relies on AI planning, which is implemented using the ROSPlan framework [2], [15], PDDL language [16] and the POPF planner [17]. The agent state (Figure 2) is continually sensed and inferred by the sensing component and given to the plan controller, a node that is constantly running and monitoring these states - whenever a change in the agent state occurs, the knowledge base is updated and a new plan is generated and executed. The key element is that the agents' states are taken into account by the AI planning component, integrating them in the PDDL domain as action costs. The reaction to a change in an agent state, and therefore in the action costs, will involve the reassignment of the actions in the generated plan. In this way, a previous plan that would have most certainly failed is replaced by a feasible plan before this failure happens. With a proper representation of the agent states as action costs in the PDDL domain, the planner should be able to anticipate and automatically adapt to unexpected changes in the agents' state by replanning with the updated costs. We choose replanning over online probabilistic planning, as we do not expect to reach unexpected outcomes that cannot be recovered from by replanning. This makes our problem probabilistically uninteresting [18], where planning for a stochastic domain would add unnecessary delays and complexity to the problem.

The flow execution describing the process is as follows:

- 1) Update Knowledge Base (KB) and PDDL problem with current state of the agents (as action costs)
- 2) Generate and execute initial plan (ROSPlan)
- 3) Continuously monitor the agents' states (Plan Controller)
- 4) If a change is detected:
 - 4.1 Wait for robot to finish current action
 - 4.2 Cancel current plan execution
 - 4.3 Update KB and PDDL problem with new agent state costs
 - 4.4 Replan and execute new plan

Agent Modelling. The agent state is represented by three main components: *capacity*, *knowledge* and *motivation* (Figure 2). These three elements are believed to cover the main reasons why an agent might not be able to achieve their expected behaviour towards completing the goal of

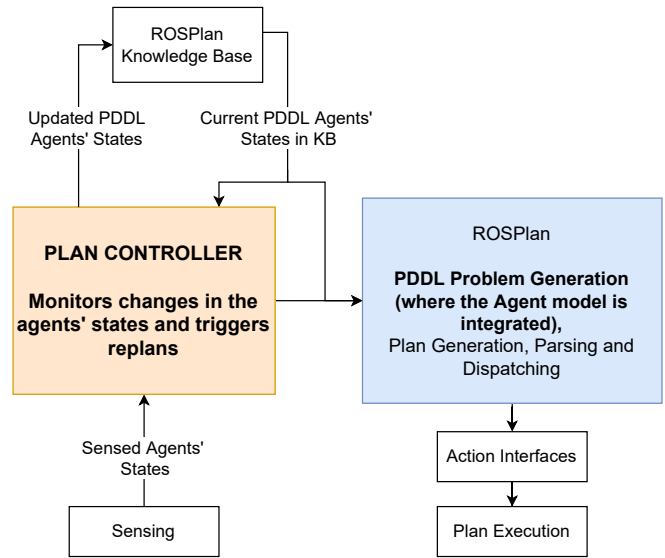


Fig. 1. System architecture - The plan controller node receives the agents' states and triggers a replan when needed to take these into account.

the collaboration. We have built this intuition on previous research: [9] where the knowledge component is used, and [13] which describes the capacity as well as the distracted and tired components, which we joined together into the motivation cost.




Fig. 2. Agent State Model represented in terms of the *capacity*, *knowledge* and *motivation* regarding the assigned actions in the plan.

We treat the *knowledge* and *motivation* elements as general for all actions, and therefore independent of any parameters. The *capacity* component, however, is dependent on a location parameter (an agent might be able to reach a location but not another one, affecting the cost of the action based on the location it refers to). All *knowledge* should be known and consistent for the *knowledge* cost to be low, and a higher value assigned to it will be interpreted as the agent lacking some knowledge related to a task assigned to them. For the *motivation* cost to be low, the agent should be focused and contributing to the plan. A low value for *capacity* of an agent at a location indicates the agent is capable of performing the tasks assigned at that location. The range of values for these costs was set to [0, 10]. The values of these elements are obtained from facts sensed in the environment. For this work, we assume the values are given, and the derivation of these from sensed facts is left for future work. As a proposal, methods such as human motion tracking or gaze tracking

could be used for the purpose. There also exists commercial software such as iMotions in order to sense human emotions.

Integration of the agent model into AI planning. The problem is defined as an AI planning problem using PDDL. The PDDL domain file includes the definition of the actions move, grasp and place with their respective parameters, preconditions and effects. The agent model is represented as PDDL functions, that are in turn integrated as action costs in the domain. The listing below contains a snippet of the domain showing how the agent model functions are integrated as costs into the move action. The PDDL problem file will include the current values for the elements in the agent model as PDDL functions (e.g. knowledge ?agent).

```
(: durative -action move
: parameters (?agent - agent ?from - wp ?to - wp)
: duration (= ?duration 50)
: condition (and ...)
: effect (and ...
(at start (increase (total-cost)
(+ (+ (knowledge ?agent)
(motivation ?agent))
(capacity ?agent ?to))))))
)
```

IV. EVALUATION

A. Simulation Setup

In order to present our approach, we define the scenario shown in Figure 3, consisting of a robot and a human that need to move boxes between different locations. Initially, all *motivation*, *capacity* and *knowledge* costs are set to a low value for both agents, meaning both agents are capable, focused, and have all the required knowledge to complete the shared task. The specified goal is for the green box to be moved from *wp3* to *wp2* and the blue box to be moved from *wp4* to *wp5*.

Simulating the agent states. As previously mentioned, the agent states are simulated and not inferred from sensed data. We specify a range of possible values for *motivation*, *knowledge* and *capacity*, and assign a value to them depending on the failure(s) simulated in each test case.

Simulating action failures. The human actions are given a success/failure outcome after a specified time. The failure causes are associated with one or several elements in the model: *motivation*, *knowledge*, *capacity*. We simulate the outcome of an action with a probability of success of $1 - \frac{\text{currentCostValue}}{\text{maxCostValue}}$. As an example, if there is a failure related to the motivation of an agent, the probability of success of an action assigned to the agent will be of $1 - \frac{\text{sensedAgentMotivation}}{\text{maxMotivationCost}}$. The robot actions are always successful, as we only evaluate failures due to the agent states modelled. We assume the robot is always willing and capable of contributing towards achieving the shared goal.

B. Illustrative Example

We present the result of a simulation that illustrates the advantage of integrating our agent model in Figure 4. In the optimal case, both agents are committed to the plan and have the *capacity*, *motivation* and *knowledge* required to complete




Fig. 3. Simulated scenario for evaluation consisting of two agents (robot and human) that have to transport several blocks to different locations.

it. All actions are therefore successful. We compare this case to the one where the human gets distracted after grasping the blue box, which causes the move action from *wp4* to *wp5* to fail. Without the model, the system is not aware of this lack of *motivation* and keeps reassigning the same move action to the human, that keeps failing. With the model, the lack of *motivation* is detected by an increase in the human *motivation* cost. The planner then assigns a place action to the human, assuming they will place the object so the robot can grasp it and take it to the correct location. A successful plan where the goal is reached is found after one replan. The same concept can be applied to the *knowledge* cost (e.g. the agent is lacking some important knowledge) and the *capacity* cost (e.g. the agent can't reach a location).

C. Evaluation

We have performed an extensive analysis by running a series of simulations to evaluate the effects of integrating the agent's model in the planning stage. We want to evaluate the following hypotheses:


- **H1 - Failure anticipation.** The integration of an agent model increases the success rate in reaching the goal, whilst reducing action failure and the number of replans.
- **H2 - Workload allocation.** The integration of an agent model fosters an appropriate task allocation, avoiding failures whilst maximising the contribution of both agents.

Impact of the model on failure anticipation. In order to evaluate the advantage of integrating the agent model into our planning system, we have run a number of simulations for the scenario described in the example presented in Section III, introducing failures that might occur due to the human's motivation, knowledge, capacity, or a combination of these. We evaluate the results in terms of the success rate and number of replans in achieving the goal, as well as the percentage of failed actions over the total executed actions.

TABLE I
EVALUATION METRICS FOR DIFFERENT SIMULATED EXECUTION CASES

	Success rate [%]	Number of replans	Failed actions [%]	Assigned actions (human) [%]	Assigned actions (robot) [%]	Successfully executed actions (human) [%]	Successfully executed actions (robot) [%]
Case 1a	25.00	3.75	44.50	65.60	34.40	14.17	85.83
Case 1b	100.00	0.93	8.50	31.92	68.08	19.03	80.97
Case 2a	37.50	3.13	36.75	63.10	36.90	21.25	78.75
Case 2b	95.00	1.48	10.98	41.18	58.82	27.65	72.35

A) Optimal plan



B) Human gets distracted after grasping the blue box

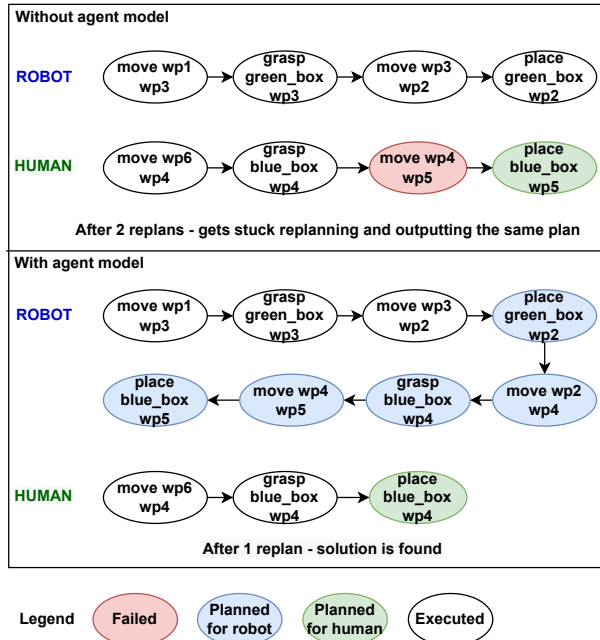


Fig. 4. Comparison of the resulting plans with and without an agent model for the case where the human gets distracted. Without using the model, the system keeps replanning and the goal is never reached. With the model, a feasible plan is found after one replan.

We compare the difference in workload between the human and the robot in terms of actions assigned and actions successfully executed by each agent. For these tests, we set the maximum number of replans to 5, as past this point the planner will get stuck outputting the same plan that will keep failing. Table I presents the mean values resulting from 40 simulations for the different test cases, with and without the agent model integrated into planning.

Case 1 - Non-contributing human partner. In this case, we evaluate the edge case where the human partner is not willing, not capable, or lacks some knowledge to perform certain actions assigned to them in the plan. Every action assigned to the human will fail. By comparing the results

(Table I) of integrating the model in the planning system (Case 1b) and not doing it (Case 1a), it becomes evident that the integration of the model greatly increases the success rate, always succeeding in reaching the goal, and reduces the number of replans as well as the percentage of failed actions. This proves that the system is able to predict when an action is likely to fail based on the partner's state, and avoids assigning this action to this agent. This is reflected in the percentage of assigned and executed actions by the human and the robot. In the case of not integrating the model, the planner is not aware of the reason why the human action keeps failing, and keeps reassigning the same actions to the human - the percentage of assigned actions to the human is higher than the robot's one, whilst the successfully executed actions of the human compared to the robot's one is extremely low.

Case 2 - Varying level of of human contribution. We have performed the same comparison with a more realistic case where varying levels of contribution from the human are modelled, corresponding to a different level of contribution randomly assigned in each test. An action assigned to the human might fail with a variable action failure probability determined by the sensed human state. In this case, we see a better balance distribution of tasks between the human and the robot compared to Case 1, with the planner risking the allocation of tasks to the human in some occasions where there is some cost assigned to it. This is desirable, as the end objective in the collaboration is for both agents to contribute as much as possible. When comparing the results to the case of not using the model (Case 2a), we see that the model integration (Case 2b) still significantly improves the success rate, whilst reducing the number of replans and failures. In terms of the distribution of successfully executed actions, the use of the model slightly increases the balance. As expected, the robot is still performing the majority of successful tasks.

Impact of the model on the workload assignment. We have assessed the impact that a single modelled element has on the allocated workload, in terms of percentage of actions assigned to each agent. Figure 5 represents the results of workload assignment for different simulations. Each simulation involves the same initial situation and goal, causing the same plan to be generated, but we vary the level of human distraction. As the human distraction increases (or human motivation cost increases), the number of allocated actions to the robot increases, whilst the number of allocated actions to the human decreases. After a value of 10 in the motivation cost, the planner safely allocates all actions to the robot after the first replan, which is triggered by the sensing

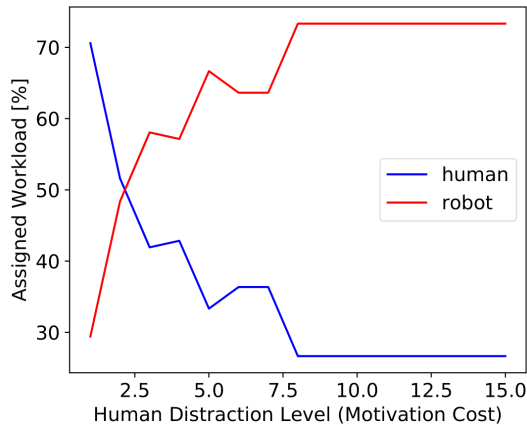


Fig. 5. Human and robot workload distribution by the planner as the human motivation cost increases.

of an increase in the motivation cost in the human agent. This also constitutes an initial validation of the range of values for the elements in our model.

V. CONCLUSIONS

In this paper we have presented an improved human-robot task planning system that anticipates failures in order to plan ahead and reassign tasks, taking into account any change in the state of the team members. The agent state is represented by three main components: *capacity*, *knowledge*, and *motivation*. When the agent state changes, the planner automatically replans to take this new state into account.

The agents' states are given to the plan controller, which is constantly running and monitoring them. Whenever a change in the agent state occurs, the knowledge base is updated and a new plan is generated and executed, cancelling the previously running plan. The agents' states are taken into account by the AI planning component, which integrates them in the PDDL domain as action costs. The reaction to a change in an agent state and therefore in the action costs will involve the reassignment of the actions in the generated plan. In this way, a previous plan that would have most certainly failed is replaced by a feasible plan before this failure happens. The system was implemented in a standardised environment using PDDL and the modular ROSPlan framework, so that it can be easily extended to a variety of scenarios and applications involving Human-Robot Collaboration.

Furthermore, the effect of the use of the human state model on the collaboration was evaluated by means of simulations. Our results show how promising the integration of the agent model into AI planning is in terms of anticipating failures in human-robot collaborative plans, and we have shown how anticipating replanning allows the planner to reallocate the tasks before they fail. Our two hypotheses regarding the integration of an agent model in the planning stage have been confirmed. Firstly, the success rate in reaching the goal is increased, whilst the action failure and the number of replans are reduced. Secondly, an appropriate workload allocation

between the agents is still achieved, avoiding failures whilst maximising the contribution of both agents.

As future work, we envisage to introduce an action type to trigger interactions for the cases where an interaction to solve an issue identified by the model might be preferred over the reallocation of actions. We will also look into the automatic inference and learning of the agent model values and action costs, which have been simulated in this work.

REFERENCES

- [1] M. Fox and D. Long, "PDDL2.1: An extension to PDDL for expressing temporal planning domains," *Journal of artificial intelligence research*, vol. 20, pp. 61–124, 2003.
- [2] M. Cashmore, M. Fox, D. Long, D. Magazzeni, B. Ridder, A. Carrera, N. Palomeras, N. Hurtós, and M. Carreras, "ROSPlan: Planning in the robot operating system," *Proceedings International Conference on Automated Planning and Scheduling, ICAPS*, pp. 333–341, 2015.
- [3] R. Lallement, L. de Silva, and R. Alami, "HATP: an HTN planner for robotics," *CoRR*, vol. abs/1405.5345, 2014.
- [4] M. Fiore, A. Clodic, and R. Alami, "On Planning and Task achievement Modalities for Human-Robot Collaboration," in *International Symposium on Experimental Robotics (ISER 2014)*, Marrakech, Morocco, Jun. 2014, p. 15p.
- [5] R. Lallement, L. De Silva, and R. Alami, "HATP: Hierarchical Agent-based Task Planner," in *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, Stockholm, Sweden, Jul. 2018. [Online]. Available: <https://hal.laas.fr/hal-01944178>
- [6] S. Devin, A. Clodic, and R. Alami, "About decisions during human-robot shared plan achievement: Who should act and how?" in *Social Robotics*, A. Kheddar, E. Yoshida, S. S. Ge, K. Suzuki, J.-J. Cabibihan, F. Eyssele, and H. He, Eds. Cham: Springer International Publishing, 2017, pp. 453–463.
- [7] G. Canal, G. Alenyà, and C. Torras, "Adapting robot task planning to user preferences: an assistive shoe dressing example," *Autonomous Robots*, vol. 43, no. 6, pp. 1343–1356, Aug 2019.
- [8] S. Bezrucav and B. Corves, "Improved AI planning for cooperating teams of humans and robots," in *Proceedings of the 8th ICAPS Workshop on Planning and Robotics (PlanRob)*, 2020.
- [9] S. Devin and R. Alami, "An implemented theory of mind to improve human-robot shared plans execution," *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 319–326, 2016.
- [10] G. Milliez, R. Lallement, M. Fiore, and R. Alami, "Using human knowledge awareness to adapt collaborative plan generation, explanation and monitoring," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2016, pp. 43–50.
- [11] O. Görür, B. Rosman, G. Hoffman, and S. Albayrak, "Toward integrating theory of mind into adaptive decision-making of social robots to understand human intention," in *HRI Workshop on the Role of Intentions*, 2017.
- [12] O. C. Görür, B. Rosman, F. Sivrikaya, and S. Albayrak, "Social cobots: Anticipatory decision-making for collaborative robots incorporating unexpected human behaviors," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2018, p. 398–406.
- [13] O. C. Görür, B. Rosman, and S. Albayrak, "Anticipatory bayesian policy selection for online adaptation of collaborative robots to unknown human types," in *International Conference on Autonomous Agents and MultiAgent Systems*, 2019, p. 77–85.
- [14] S. Nikolaidis, D. Hsu, and S. Srinivasa, "Human-robot mutual adaptation in collaborative tasks: Models and experiments," *The International Journal of Robotics Research*, vol. 36, no. 5-7, pp. 618–634, 2017.
- [15] G. Canal, M. Cashmore, S. Krivić, G. Alenyà, D. Magazzeni, and C. Torras, "Probabilistic Planning for Robotics with ROSPlan," in *Towards Autonomous Robotic Systems*. Springer International Publishing, 2019, pp. 236–250.
- [16] M. Ghallab, A. Howe, C. Knoblock, D. McDermott, A. Ram, M. Veloso, D. Weld, and D. Wilkins, "PDDL—The Planning Domain Definition Language," 1998.
- [17] A. Coles, A. Coles, M. Fox, and D. Long, "Forward-chaining partial-order planning," in *International Conference on Automated Planning and Scheduling*, 2010, pp. 42–49.
- [18] I. Little and S. Thiébaux, "Probabilistic planning vs. replanning," in *In ICAPS Workshop on IPC: Past, Present and Future*, 2007.