



HAL
open science

A Human-Aware Task Planner Explicitly Reasoning About Human and Robot Decision, Action and Reaction

Guilhem Buisan, Rachid Alami

► **To cite this version:**

Guilhem Buisan, Rachid Alami. A Human-Aware Task Planner Explicitly Reasoning About Human and Robot Decision, Action and Reaction. HRI '21: ACM/IEEE International Conference on Human-Robot Interaction, Mar 2021, Boulder CO USA (virtual), United States. pp.544-548, 10.1145/3434074.3447231 . hal-03192891

HAL Id: hal-03192891

<https://hal.science/hal-03192891v1>

Submitted on 8 Apr 2021

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.

A Human-Aware Task Planner Explicitly Reasoning About Human and Robot Decision, Action and Reaction

Guilhem Buisan, Rachid Alami

► **To cite this version:**

Guilhem Buisan, Rachid Alami. A Human-Aware Task Planner Explicitly Reasoning About Human and Robot Decision, Action and Reaction. HRI '21: ACM/IEEE International Conference on Human-Robot Interaction, Mar 2021, Boulder CO USA, United States. pp.544-548, 10.1145/3434074.3447231 . hal-03192891

HAL Id: hal-03192891

<https://hal.archives-ouvertes.fr/hal-03192891>

Submitted on 8 Apr 2021

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.

A Human-Aware Task Planner Explicitly Reasoning About Human and Robot Decision, Action and Reaction

Guilhem Buisan, Rachid Alami

name.surname@laas.fr

LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

ABSTRACT

The complexity of the tasks autonomous robots can tackle is constantly increasing, yet we seldom see robots interacting with humans to perform tasks. Indeed, humans are either requested for punctual help or given the lead on the whole task. We propose a human-aware task planning approach allowing the robot to plan for a task while also considering and emulating the human decision, action, and reaction processes. Our approach is based on the exploration of multiple hierarchical tasks networks albeit differently whether the agent is considered to be controllable (the robot) or uncontrollable (the human(s)). We present the rationale of our approach along with a formalization and show its potential on an illustrative example involving the assembly of a table by a robot and a human.

KEYWORDS

human-aware task planning, human robot collaboration, hierarchical task planning

ACM Reference Format:

Guilhem Buisan, Rachid Alami. 2021. A Human-Aware Task Planner Explicitly Reasoning About Human and Robot Decision, Action and Reaction. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction (HRI '21 Companion)*, March 8–11, 2021, Boulder, CO, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3434074.3447231>

1 INTRODUCTION

As technology progresses, humans are more and more assisted in their everyday tasks. Moreover, these autonomous assistants are becoming more and more complex and able to handle harder tasks. However, humans and robotic assistants are seldom working on a task at the same time and at the same place. Indeed, robots rarely include human action and even less the human decision process in their own decision process.

In this work we aim at proposing a robot task planning scheme for human robot collaboration. The goal of our approach is to integrate, during the robot task planning process, mechanisms allowing it to anticipate the human decision, action or reaction leading to plans where the robot may create situations calling, when needed or preferred, for human action.

To do so, we took inspiration from previous work on human-aware task modelling and planning to build a task planner focusing

on human robot collaboration. Our planner is able to explore in a distinct manner the deliberation and plan elaboration processes of the robot and the human in order to build robot plans and anticipate the decisions and actions of the human. Moreover, it maintains one belief base per agent (human or robot), and actions preconditions and effects can be expressed in any of these belief bases, allowing to represent situational or inherently non observable actions from agents, knowledge transfer actions, and to detect beliefs divergences and plan accordingly. Our scheme is designed to provide a suitable framework allowing to anticipate the beliefs and potential decision, reaction or contribution to a shared goal or to an interaction situation of both agents. This overall anticipation process takes place within the robot decisional activity.

In this paper, we first discuss briefly related work and present our approach and its intents and rationale. Then, we provide a formalization of our planner scheme. Finally, a proof of concept example will be presented before concluding.

2 BACKGROUND

When a robot has to carry a complex task, it probably needs to plan before executing any action, otherwise risking to follow a sub-optimal path of actions or even engage itself in a non recoverable situation. Task planning aims at finding a feasible sequence of actions leading to the completion of a goal or to the fulfilment of a high level task [10]. Now, when acting in the presence of humans, the robot has not only to consider constraints such as pertinence and acceptance by the humans of its activity but also their course of actions. Most approaches to the so-called *human-aware task planning* problem assume a fully controllable and cooperative human, willing to participate in the accomplishment of a common goal [5, 17, 20]. Then, information and plan sharing are done in a post-processing step or at execution time [8, 16].

Other approaches are explicitly considering an external human model, which can be used to predict future human actions, and plan accordingly. Cirillo et al. present an approach where the robot considers an external human model allowing it to adapt its plan to multiple possible predicted human plans to avoid undesired situations [6]. However, the tasks considered are not shared, and the robot has no influence on the human actions. Talamadupula et al. propose to consider the human as a planning agent with a distinct goal and use plan recognition techniques to generate a robot plan coordinated with the human one [18]. Buckingham et al. integrate a separate human model taking into account a mental model of the human to generate plan for collaborative, competitive, and non-involved humans indifferently [4]. Finally, Unhelkar et al. propose a scheme integrating the human action model as an agent Markov model to predict the human action into a POMDP, allowing the robot to follow policies influencing the human actions

through communication [21]. However, these approaches consider the human model as an oracle on which reasoning is hardly possible during the planning process.

Planning problems for Human-Robot collaboration can also be seen as an instance of the general problem of multi-robot distributed task allocation [9] and multi-robot collaborative task achievement [3]. Indeed, there are similarities in terms of the identification of the situations which call for decision and the associated models for task representation [2, 7].

We propose an approach in which the human action models and beliefs can explicitly be reasoned on and managed during the task planning process, while also explored in order to elaborate robot plans taking into account human capabilities of planning, decision and reaction. We took inspiration for our human model from human task modeling described in [13]. In this work, the human task is modeled in order to evaluate an interactive system, ensuring that at any state, any possible human action is supported by the system. We want to enhance this approach by adding the ability to project agents actions into the future and find multi-agent plans by reasoning on both controllable and uncontrollable agent action model.

Since we would like a scheme allowing to anticipate goal-based as well situation-based decisions of each agent, we have naturally chosen to adopt HTN based planning and refinement mechanisms. Our planning scheme is inspired from HATP [1, 14, 17] where a single HTN is explored to produce a human robot joint plan. However, in our approach, we consider distinct action models for the humans and the robot resulting not in one but in several HTNs to be explored (one per considered agent).

3 DESCRIPTION AND NOTATION

We separate agents involved in a given task into two categories: the controllable agents (*i.e.* the robots) for which the planner needs to select the best course of actions to generate a policy; and the uncontrollable agents (*i.e.* the humans) on whom the planner has no direct control but, still, has a representation of their decision and action models. The two agent types are fundamentally different: (1) the robot is controllable since the process is run by the robot, (2) the human agent is not controllable since the process can only "speculate" on her/his decisions and actions, but can still influence them, (3) the two agents are not equivalent, the robot agent role is to help, assist and facilitate human and to synthesize pertinent, legible and acceptable behaviour. We want to devise a planner allowing controllable agents to plan for their actions while anticipating the decisions, actions and reactions of uncontrollable agents. Moreover, we want the planner to be able to generate plans where robots actions elicit situations calling for humans decision, action and reaction, thus creating and anticipating collaboration and interaction.

This problem may be seen as a classical non deterministic planning problem, but enriched with the ability of the robot to model the actions, beliefs and decision process of the human. Thus, we have to consider distinct action models, beliefs and execution streams for each of the agents involved. Doing so with classical STRIPS-style planning approaches would lead to an intractable search space. Therefore, we chose to use HTN planning for both the controllable

and uncontrollable agents. HTN planning aims at decomposing abstract tasks into atomic primitive tasks by choosing from a list of available context-dependent refinements for each abstract task, ensuring that preconditions and effects of refined primitives tasks are respected throughout the created plan. Similarly to HATP [17], our planner elaborates a plan with several streams of actions each associated to an agent involved in the task. But while in HATP, all the streams are built starting from an initial root node corresponding to a shared goal of all agents, our planner starts from multiple initial root nodes corresponding to the decision process of the different agents.

Beliefs: Let S be the set of all possible world states, we call beliefs of an agent a the state $s_a \in S$ in which this agent thinks the world is in. It is important to note that the state of the controllable agent is assumed to be the world state estimation for the planner, as we consider the planner as being part of the controllable agent.

Action models: We represent the action model of an agent a as $\Lambda_a = \langle Op_a, A_a, Me_a \rangle$ where Op_a are the primitive tasks (*i.e.* operators, actions) that the agent a can perform, A_a the set of abstract tasks and Me_a are the methods (*i.e.* decompositions) describing how an agent a can perform an abstract task through a refinement process.

Agents agendas and plans: An agenda d_a and a plan π_a (this agent only stream of actions) are defined for each agent a . The agenda d_a is a list of tasks (abstract or primitive) having to be performed by the agent. The plan π_a is a list of primitive tasks, built from the agenda, which the agent has to perform. Coordination between agent plans are represented by causal links between streams which correspond to effects of agents actions on the beliefs states of the other agents.

Agent triggers: We then define for each agent a a set of so-called *trigger functions* T_a . These trigger functions aim at representing reactions of agents to certain situations (subsets of worlds states).

Agents: Finally, we define an agent state as a tuple $\sigma_a = \langle d_a, \pi_a, s_a \rangle$, and an agent as being $a = \langle \text{name}_a, \sigma_a, \Lambda_a, T_a \rangle$. Then we represent the set of *controllable agents* as \widehat{Ag} and the set of the *uncontrollable agents* as \widetilde{Ag} . The set of all the agents considered by the planner is then represented as $Ag = \widehat{Ag} \cup \widetilde{Ag}$. Likewise, the set of controllable, uncontrollable and full agents states are represented as, respectively, $\widehat{\sigma} = \bigcup_{a \in \widehat{Ag}} \sigma_a$, $\widetilde{\sigma} = \bigcup_{a \in \widetilde{Ag}} \sigma_a$ and $\sigma = \bigcup_{a \in Ag} \sigma_a$. Let Σ be the set of all the possible agents states. For simplicity purposes, in the following, we will only consider one controllable agent, named *robot*, and one uncontrollable agent, named *human*.

4 PLANNING PROCESS

The cooperative agents planning problem consists in a set of agents Ag_{start} with their respective agenda filled with tasks to achieve and their beliefs about the current world. It is important to note that if for the controllable agent, the beliefs correspond to the planner ground truth, for the uncontrollable agent, their beliefs need to be estimated, through, for example, situation assessment component [12, 15].

Agent type	Readable	Writable
Controllable	$s_{self}, s_{other},$ π_{self}, π_{other} (1)	$s_{self}, s_{other},$ (2) d_{self}, d_{other} (3)
Uncontrollable	s_{self}, π_{self} (4)	$s_{self}, s_{other}, d_{self}$ (5)

Table 1: Readable and writable elements (belief states, agenda, plan) of the agents state by method, primitive task and trigger functions.

The result is a robot policy Π being a tree of alternating robot and human primitive tasks. Any path from the root to the leaves is a feasible sequence of primitive tasks (*i.e.* each primitive task application leads to a state where the following one is applicable) leading to a state where all the controllable agents agenda are empty.

To solve such a problem we need to augment the search space from world states S only to all the agents states considered by the planner σ , with their agenda, plan and beliefs. The exploration starts with σ^{start} and consecutively applies operators associated to the robot and to the human, leading to new agents states σ^i until all the controllable agents have an empty agenda: $\forall a \in \bar{A}g, d_a = ()$.

4.1 Action models restriction

Considering the definitions above, for any agent a the operators are defined as functions: $Op \ni o : \Sigma \rightarrow \Sigma \cup \perp$ which produce new agents state, being the effect of the application of the primitive task, or *false* if the task is not applicable. Methods are defined as tuple, containing an abstract task and a decomposition function: $Me \ni m = \langle \alpha, \delta \rangle$ with $\alpha \in A$ and $\delta : \Sigma \rightarrow (Op \cup A)^n \cup () \cup \perp$ with $n \in \mathbb{N}^*$, which, depending on agents states, decompose the abstract task returning a list of tasks (primitive or abstract), an empty list if the abstract task does not need to be decomposed, or *false* if the task cannot be decomposed in the current state. Multiple methods can address the same abstract task, the goal of the HTN planner is then to choose the right one to create a plan. Finally triggers function are defined as: $T \ni t : \Sigma \rightarrow (Op \cup A)^n \cup ()$ with $n \in \mathbb{N}^*$, returning a list of tasks to be inserted in an agent agenda as a reaction to specific agent states.

However, some constraints on these functions must be respected. Indeed, depending on whether the agent is controllable or not, their planning process will not take decisions based on the same information, and their action will not impact the world state in the same manner. We thus impose restrictions on what a function can read and write (writing means here having effects on agents states and is only in the case of primitive task functions) in the agents state. Then, the function constraints also depend on which agent is performing the action or making the decision (in method and trigger functions). When a function is applied we note *self* the agent which executes it and *other* the other agent. The rules for read and write access are given in Table 1.

(1). During robot planning, the decision and the action can depend on the beliefs of the robot and on the planned estimated beliefs of the human. Moreover, the current partial plan of the robot and the anticipated plan of human one can also be used to make decisions.

(2). The effects of robot actions obviously impact its own belief state (considered as the real world state by the planner), but also the beliefs of the human, for example, through their observation process and first order logic reasoning. More elaborate schemes to compute the effects can also be devised such as those described in [11].

(3). Besides, a robot action can add a new task to the agenda of the human. This is to account for communication actions requesting the human to do something.

(4). The human decisions and actions can only be done according to her own beliefs and partial plan. Indeed, we cannot add the robot ones as it is, or we would consider that the human estimation of the robot knowledge and past actions is always perfect. This would require a third type of agents, being the robot model as estimated by our estimation of the human. Here, we make the assumption that the human is a naive user, and thus, will not take their decision based on the estimated robot beliefs and past plan.

(5). The effects of the human actions obviously impact their beliefs and the robot (planner) ones. Moreover, the human agenda could also be updated through, for example, a positive answer of a task request.

4.2 Exploration algorithm

Our planner operates in a turn-taking scheme, based on the update of the agents beliefs states, the HTNs of the robot and the human are explored successively.

4.2.1 Controllable agent HTN exploration. The robot HTN exploration is a pretty standard depth first algorithm. The first task λ from its agenda d_{robot} its popped, then if it is an abstract task $\lambda \in A$, all the applicable methods are applied, and their result are prepended to the agenda, thus giving new agents states (with the same beliefs as the previous ones but with the robot agenda updated) and branching our search space. We iterate with the new task popped from the new robot agenda. Eventually, the popped task will be a primitive one $\lambda \in Op$, its function will then be applied to the currently explored agent states. If it returns *false*, the action is not applicable, and the exploration backtracks to another decomposition of an abstract task. However, if the action is applicable (returns a new agents state), the triggers are run for each agent, updating their agenda if necessary. Then, we question the human HTN to get their possible next actions from this new agents state, and, for each possible new agents state, we apply the triggers of each agent then we continue the robot HTN exploration. This exploration continues until the robot agenda is empty, or all the branches return false.

4.2.2 Uncontrollable agent HTN exploration. The human HTN exploration differs from classical HTN planner as the goal is not to produce a complete plan, but rather to list all the actions the human is likely to perform in a given agents state. To do so, we recursively decompose the first task of the human agenda d_{human} with every applicable methods, until we reach an applicable operator. All the agents states resulting from the application of the operators from all the applicable decompositions are then returned to the robot HTN exploration (containing both agents beliefs and human agenda updated, as stated in Table 1).

4.2.3 Default actions. Two special cases are handled during the exploration. If the human agenda is empty whereas the robot one is not, the exploration returns a default action *IDLE* – which does not modify agents beliefs nor agendas – for the human. This action represents the non-involvement of the human in a task. Besides, if for the human no applicable action is found a default action *WAIT* – which does not modify agents nor agendas – is returned. This action represents the impossibility of the human to act in the current situation, making them wait for the robot to proceed.

Once the robot agenda is emptied, the agents state is set as a success, the plan is added to the policy tree and the search can be continued until no decomposition is left for any task.

5 AN ILLUSTRATIVE EXAMPLE

We use here the same example as the one depicted in [19]: a robot needs to assemble a table, by mounting the four legs on the table upside down top, while a human is in its vicinity. The robot generates a plan where it assembles the table alone, but, during the execution, finds out that it cannot pick a table leg as it is not reachable to it. As a repair action, it asks the human to hand-over the table leg and then proceeds to complete the assembly.

We argue that if, during planning time, the robot was aware of the fact that the leg is not reachable to it, it could have planned to ask the human for it. Our planner can deal with such a reasoning ability.

The robot has to assemble four legs on a table top. Moreover, it is aware of a human presence in the same room who is doing nothing. Two table legs are not reachable by the robot, but the four legs are reachable by the human. Our planner is thus instantiated with one controllable agent (the robot) and one uncontrollable agent (the human). We did not set any belief differences between the robot and the human in this example. The two agents only differ by their name, action model and starting agenda. Indeed, the agenda of the robot is initiated with a high level task *assemble_table*, while the human agenda is empty. In a real scenario, the human agenda could be filled with their estimated current task via situation assessment and plan recognition components.

In the robot model the *assemble_table* abstract task has two decompositions. The first method tells how the robot can recursively assemble a leg until no leg is available, with the task *assemble_leg*. In the other method the robot asks the human to help it assemble the table with the two legs it cannot reach (operator *ask_help_for_assembly*), and then recursively adds the *assemble_leg* with the remaining legs, to its agenda. The operator *ask_help_for_assembly* fills the human agenda with a high level task to decompose. Then, the task *assemble_leg* has two decompositions, one where the robot *picks* and *assembles* the leg itself, and the other where the robot asks the human to give it the leg, waits until receiving it and assembles it. When the robot asks the human to give it the leg, operators are added to the human agenda: *pick*, *handover* and *wait_for_the_robot_to_take*.

The policy returned by the planner contains two main branches, splitting from the root node. Either the robot first action is asking the human for help, thus relying on planning capability of the human, as they have to decompose an abstract task; or the robot asks only the human to handle it one table leg, thus not making

them take any decision, and involving them only minimally in the task. Then, in each branch, when the robot needs to pick a leg, it can either ask the human for it, or take it itself (if it is reachable). In the policy, this is seen as two distinct branches when the robot retrieves a leg. However, in the branch where the robot asked the human for higher level help, the robot has to wait until the human completed the assembly of their two assigned table legs before giving a third to the robot.

With our approach we were able to successfully generate a policy allowing the robot to either elicit human planning capability by assigning them a high level task, allowing concurrent assembly, or to handle most of the assembly by itself, asking for simple help at the beginning of the task to avoid disturbing the human sparsely. Returning a policy and not only a selected plan allows a more robust execution, as the robot may switch between plan branches depending on the human action, and even change its course of actions if the human does not act as expected (e.g. is not willing to take part in the assembly).

6 CONCLUSION AND FUTURE WORK

We presented a novel approach for human-aware task planning allowing to explicitly include human decisions, actions and reactions. Our approach relies on a representation of each agent considered by the planner, with their own beliefs, agenda, stream of execution and action model. The action models are represented as HTNs which are explored consecutively yet differently if the agent is controllable (robots) or uncontrollable but rational (humans). This scheme aims at anticipating situations where the human may decide to act or wait for the robot to act, but also explicitly representing robot to human communication needs for beliefs alignment, goal sharing or action requests.

The policy returned is a tree containing human and robot streams of actions leading to a completion of the goal. However, how to choose a plan between the branches is an open problem. With the information retrieved during the exploration process, we want to allow for multiple plan selection criteria depending on a settable human hyper parameter. If a human is considered as busy, we could prefer plans which minimally involve them and avoid interrupted long sequences of *IDLE* actions. If a human is considered as in a hurry, we could choose shorter plans, whatever the human contribution may be. Moreover, this hyper parameter can bias the HTNs exploration to avoid plans which will not be selected in the policy.

Finally, as the exploration of all the human possible actions may lead to a large branching factor, we want to bias the exploration by stopping it at a certain level of depth, balancing between more or less probable courses of actions, and adding robot actions guiding the human to more probable actions.

Acknowledgements

This work has been partially funded by the French Agence Nationale de la Recherche through the grant NR-19-PI3A-0004, Artificial and Natural Intelligence Toulouse Institute - Institut 3iA (ANITI)

REFERENCES

- [1] Samir Alili, Rachid Alami, and Vincent Montreuil. 2009. A Task Planner for an Autonomous Social Robot. In *In: Asama H., Kurokawa H., Ota J., Sekiyama K. (eds)*

- Distributed Autonomous Robotic Systems 8*. Springer, Berlin, Heidelberg. Springer Berlin Heidelberg, Berlin, 335–344.
- [2] Bashira Akter Anima, Janelle Blankenburg, Mariya Zagainova, S. Pourya Hoseini A., Muhammed Tawfiq Chowdhury, David Feil-Seifer, Monica N. Nicolescu, and Mircea Nicolescu. 2019. Collaborative Human-Robot Hierarchical Task Execution with an Activation Spreading Architecture. In *Social Robotics - 11th International Conference, ICSR 2019, Madrid, Spain, November 26-29, 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11876)*, Miguel A. Salichs, Shuzhi Sam Ge, Emilia Ivanova Barakova, John-John Cabibihan, Alan R. Wagner, Álvaro Castro González, and Hongsheng He (Eds.). Springer, Cham, 301–310.
 - [3] Silvia C. Botelho and Rachid Alami. 1999. M+: a scheme for multi-robot cooperation through negotiated task allocation and achievement. In *IEEE International Conference on Robotics and Automation*. IEEE, Detroit, United States, 1234–1239 vol.2.
 - [4] David Buckingham, Meia Chita-Tegmark, and Matthias Scheutz. 2020. Robot Planning with Mental Models of Co-present Humans. In *Social Robotics (Lecture Notes in Computer Science)*, Alan R. Wagner, David Feil-Seifer, Kerstin S. Haring, Silvia Rossi, Thomas Williams, Hongsheng He, and Shuzhi Sam Ge (Eds.). Springer International Publishing, Cham, 566–577. https://doi.org/10.1007/978-3-030-62056-1_47
 - [5] T. Chakraborti, G. Briggs, K. Talamadupula, Y. Zhang, M. Scheutz, D. Smith, and S. Kambhampati. 2015. Planning for serendipity. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Chicago, IL, USA, 5300–5306. <https://doi.org/10.1109/IROS.2015.7354125>
 - [6] Marcello Cirillo, Lars Karlsson, and Alessandro Saffiotti. 2010. Human-aware task planning: An application to mobile robots. *ACM Transactions on Intelligent Systems and Technology (TIST)* 1, 2 (2010), 1–26.
 - [7] Silvia Silva da Costa Botelho and Rachid Alami. 2000. Robots that cooperatively enhance their plans. In *In: Parker L.E., Bekey G., Barhen J. (eds) Distributed Autonomous Robotic Systems 4*. Springer Japan, Cham, 55–65. <https://hal.laas.fr/hal-01979720>
 - [8] Sandra Devin and Rachid Alami. 2016. An Implemented Theory of Mind to Improve Human-Robot Shared Plans Execution. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction, HRI 2016, Christchurch, New Zealand, March 7-10, 2016*. IEEE Press, Christchurch, New Zealand, 319–326. <https://doi.org/10.1109/HRI.2016.7451768> bibtex: Devin-hri16 bibtex[bibsource=dblp computer science bibliography, <https://dblp.org>].
 - [9] Brian P. Gerkey and Maja J. Mataric. 2004. A Formal Analysis and Taxonomy of Task Allocation in Multi-Robot Systems. *Int. J. Robotics Res.* 23, 9 (2004), 939–954. <https://doi.org/10.1177/0278364904045564>
 - [10] Malik Ghallab, Dana Nau, and Paolo Traverso. 2016. *Automated Planning and Acting*. Cambridge University Press, Cambridge, England.
 - [11] Mamoun Gharbi, Raphaël Lallement, and Rachid Alami. 2015. Combining symbolic and geometric planning to synthesize human-aware plans: toward more efficient combined search. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Hamburg, Germany, 6360–6365.
 - [12] S. Lemaignan, Y. Sallami, C. Wallhridge, A. Clodic, T. Belpaeme, and R. Alami. 2018. UNDERWORLDS: Cascading Situation Assessment for Robots. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Chicago, IL, USA, 7750–7757. <https://doi.org/10.1109/IROS.2018.8594094> ISSN: 2153-0866.
 - [13] Célia Martinie, Philippe Palanque, and Marco Winckler. 2011. Structuring and Composition Mechanisms to Address Scalability Issues in Task Models. In *Human-Computer Interaction – INTERACT 2011 (Lecture Notes in Computer Science)*, Pedro Campos, Nicholas Graham, Joaquim Jorge, Nuno Nunes, Philippe Palanque, and Marco Winckler (Eds.). Springer, Berlin, Heidelberg, 589–609. https://doi.org/10.1007/978-3-642-23765-2_40
 - [14] Grégoire Milliez, Raphaël Lallement, Michelangelo Fiore, and Rachid Alami. 2016. Using human knowledge awareness to adapt collaborative plan generation, explanation and monitoring. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, IEEE, Chicago, IL, USA, 43–50.
 - [15] Grégoire Milliez, Matthieu Warmier, Aurélie Clodic, and Rachid Alami. 2014. A framework for endowing an interactive robot with reasoning capabilities about perspective-taking and belief management. In *IEEE RO-MAN 2014*. IEEE, Edinburgh, United Kingdom, 1103–1109.
 - [16] Stefanos Nikolaidis, Minae Kwon, Jodi Forlizzi, and Siddhartha Srinivasa. 2018. Planning with Verbal Communication for Human-Robot Collaboration. *J. Hum.-Robot Interact.* 7, 3, Article 22 (Nov. 2018), 21 pages. <https://doi.org/10.1145/3203305>
 - [17] Eugenio Sebastiani, Raphaël Lallement, Rachid Alami, and Luca Iocchi. 2017. Dealing with On-line Human-Robot Negotiations in Hierarchical Agent-Based Task Planner. In *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling, ICAPS 2017, Pittsburgh, Pennsylvania, USA, June 18-23, 2017*, Laura Barbulescu, Jeremy Frank, Mausam, and Stephen F. Smith (Eds.). AAAI Press, Menlo Park, California, USA, 549–557. <https://aaai.org/ocs/index.php/ICAPS/ICAPS17/paper/view/15742>
 - [18] Kartik Talamadupula, Gordon Briggs, Tathagata Chakraborti, Matthias Scheutz, and Subbarao Kambhampati. 2014. Coordination in human-robot teams using mental modeling and plan recognition. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, Chicago, IL, USA, 2957–2962. <https://doi.org/10.1109/IROS.2014.6942970>
 - [19] Stefanie Tellex, Ross Knepper, Adrian Li, Daniela Rus, and Nicholas Roy. 2014. Asking for Help Using Inverse Semantics. In *Proceedings of Robotics: Science and Systems, RSS* (Ed.). Berkeley, USA. <https://doi.org/10.15607/RSS.2014.X.024>
 - [20] Alejandro Torreño, Eva Onaindia, Antonin Komenda, and Michal Stolba. 2017. Cooperative Multi-Agent Planning: A Survey. *Comput. Surveys* 50, 6 (Nov. 2017), 84:1–84:32. <https://doi.org/10.1145/3128584>
 - [21] Vaibhav V. Unhelkar, Shen Li, and Julie A. Shah. 2020. Decision-Making for Bidirectional Communication in Sequential Human-Robot Collaborative Tasks. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, Cambridge United Kingdom, 329–341. <https://doi.org/10.1145/3319502.3374779>