

About Decisions During Human-Robot Shared Plan Achievement: Who Should Act and How?

Sandra Devin, Aurélie Clodic, and Rachid Alami

CNRS, LAAS, 7 avenue du colonel Roche, F-31400 Toulouse, France
Univ de Toulouse, LAAS, F-31400 Toulouse, France, `surname.name@laas.fr`

Abstract. It has been shown that, when a human and a robot have to perform a joint activity together, they need to structure their activity based on a so-called "shared plan". In this work, we present a scheme and an implemented system which allow the robot to elaborate and execute shared plans that are flexible enough to be achieved in collaboration with a human in a smooth and non-intrusive manner. We identify and analyze the decisions that should preferably be taken at planning time and those that should be better postponed. We also show in which conditions the robot can determine when it has to take the decision by itself or leave it to its human partner. As a consequence, the robot avoids useless communication by smoothly adapting its behavior to the human.

1 INTRODUCTION

One of the challenges in human-robot interaction is to devise robots able to work with humans to perform what is called Joint Action in social sciences [27]. In this context, robots first need to establish a joint goal with their human partners. Once the joint goal has been agreed upon, the participants need to agree on a way to achieve this goal: they need to have a shared plan [15]. Shared plans can be based on known procedures or can be elaborated on-line. In robotics, several planners allow to take into account the human [8, 20]. They allow the robot to reduce resource conflicts [5] or take divergent beliefs into account [32, 29]. To communicate about the shared plan, several possibilities have been studied: the robot computes a plan for both agents and shares/negotiates it [1, 23], the human imposes the shared plan to the robot [26, 25], or the robot adapts its actions to the human ones [22, 12]. In [14], it is shown that subjects prefer letting the robot plan when the task is too complex, prioritizing efficiency. In more simple tasks, it has been shown in [2] that a robot proactively helping the human is preferred to one waiting before proposing help. On another hand, the relevance of using a shared plan in HRI has been studied by [19]. They suggest that the shared plan should be fully communicated in order to sustain effective collaboration. However, when two humans share a plan, they usually do not communicate every detail at every step of the plan, but only the parts of their individual plans which

⁰ This work has been partly supported by the French National Research Agency (ANR) under grant agreement ANR-16-CE33-0017-01 (JointAction4HRI Project)

interact in space, time and/or resources [4]. In this work, we try to get rid of the entire shared plan verbalization by taking the right decision at the right time in order to come up with a robot which communicates only when necessary in order to promote efficiency, fluency and acceptability.

Then, when it comes to execute the shared plan, several contributions have been done to allow more flexibility. [7] proposes a method to plan only a few steps in advance and plan the actions further in an iterative way. Chaski, a task-level executive [28], allows to choose when to execute the robot actions adapting to a human partner. A system that mixes plan recognition and adaptation is described in [22]. It computes all possibilities for the plan and chooses an action based on the choice of the human and causal links. [16] proposes an adaptive action selection mechanism for a robotic teammate, making anticipatory decisions based on the confidence of their validity and their relative risk. [18] presents Pike, an online executive that unifies intent recognition and plan adaptation for temporally flexible plans with choices. Turn-taking abilities have also been studied, mainly to take a turn during a conversation [17] but also during physical interaction [6]. Finally, the HRI domain also took inspiration from cooperative multi-robot literature where task allocation and cooperative activity achievement have been thoroughly investigated [30, 3, 13]. However, the two domains are very different, since the human and the robot are not equal in any aspect and communication in HRI is essentially multi-modal, while most often robots communicate with each other through wireless network.

The work presented here is an extension of previous work in elaboration and achievement of human-robot shared plans [12, 10]. The contribution aims at improving flexibility and fluency during shared plan achievement by allowing the robot to identify which are the needed decisions and to choose the most pertinent moment to take them while smoothly adapting to the human behavior.

2 THE CONTEXT AND THE PROBLEM

We place ourselves in a context where a team-mate robot has to work jointly with a human. They share an environment and mutually observe each other. In order to focus on the paper contribution, we consider here that the joint goal has already been established; the robot and its human partner have a commitment to achieve the task [10] and none of them will abort it unless they know that it is not achievable any more. We also consider that the robot is working with only one human, even if the proposed scheme can be easily adaptable to several humans. The question now is to achieve the task.

The blocks building scenario: To illustrate the overall process, we use a task inspired from [9]: a human and a robot have to build a blocks construction (Fig. 1). At the beginning of the task, the robot and the human have several colored blocks they can access.

The problem: We identify three types of decisions to be taken during shared plans elaboration and execution:

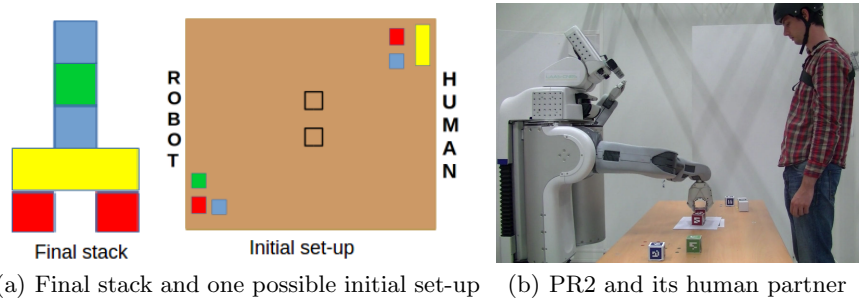


Fig. 1: Blocks building task needing collaboration between the robot and the human. There are two identical placements where to put the two red cubes.

Which action to perform in which order? This is done through an adaptation of HATP, a human-aware planner which has been demonstrated to be well suited to human-robot joint action [20].

Who will perform which action? We endow the robot with the ability to detect whenever the choice is trivial (only one agent can perform an action, or the efforts to perform it clearly point to one particular agent) or whenever a decision is needed. Only then, the robot communicates concerning the choices to be made and/or simply adapts its behavior to the human one, avoiding useless communication.

With which object? when an action has been allocated, some conflicts may remain concerning the action instantiation (which object, tool or placement to use). We endow the robot with the ability to postpone this action refinement and adapt to the human choices.

This work has been fully integrated into a human-aware architecture [21, 11] that is presented in Fig. 2. Two key components will be discussed: the planner and the supervisor. In section Sec. 3, we present the adaptation of HATP in order to elaborate more flexible shared plans and to promote decisions that can be better taken at the last moment. In section Sec. 4, we present the essential mechanisms that have been designed to build a supervisor that is able to control and refine on-line the shared plan in order to promote fluency and pertinence of robot decisions and actions including verbal interactions. In section Sec. 5, we present experiments we have ran with the system and exhibit quantitative and qualitative results that validate the approach.

3 AN IMPROVED SHARED PLAN ELABORATION

HATP is a Human-Aware Task Planner based on Hierarchical Task Network which has been specifically designed to integrate a number of features that are meant to promote the synthesis of shared plans that are acceptable by humans and easily if not trivially understandable by them, taking into account a number of social rules [20]. The resulting shared plans contain the actions of all agents involved in the task. In previous work, all the actions of the computed plans were

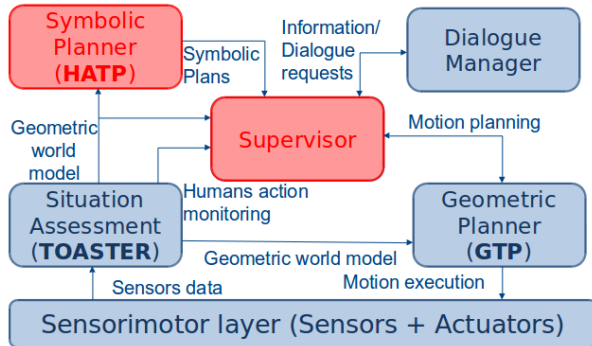


Fig. 2: The architecture. A Situation Assessment module [24] maintains the world state from the point of view of all agents based on perspective-taking. HATP[20] allows to compute a symbolic shared plan for the robot and its human partners. The plan is executed by the Supervisor with the help of the Dialogue Manager and a Geometric Planner [31] that computes trajectories as well as objects placements and grasps while taking into account human safety and comfort. The work presented here concerns HATP and the Supervisor (in red).

allocated and completely instantiated during plan elaboration. For example; the first action of the shared plan for the blocks building task could be: [ROBOT PICK_AND_PLACE RED_CUBE1 PLACEMENT_1]. In a collaborative setting, taking in advance all these decisions could be counter intuitive and it can be interesting to let the agents decide at execution time who will do what and how. To do so, we made some changes to the use of HATP resulting to the plan Fig. 3(a):

Similar objects: In order to allow to postpone the decisions of which exact objects to use, we use the notion of objects *similarity*: two *similar objects* will have the same role in the task. With this modification, the first action of the shared plan becomes [ROBOT PICK_AND_PLACE **RED_CUBE** PLACEMENT].

The X agent: In order to postpone the decision of who should perform an action, we define an *X agent* whose capabilities correspond to the intersection of capabilities of the human and the robot (with a lower cost). It allows the planner to identify for which action the decision is trivial and for which action a decision is needed. With this modification, our action becomes [**X-AGENT** PICK_AND_PLACE **RED_CUBE** PLACEMENT].

4 SHARED PLAN EXECUTION

We will now explain how the supervisor contributes to smooth the plan execution. Based on previous work [10], the supervisor is able to maintain the current shared plan and compute the set of feasible actions which need to be performed at each time, using causal links and actions preconditions. This previous work also allows the robot to elaborate a new plan whenever an unexpected event leads to a plan not valid anymore (and if no plan is found, the robot will abort the goal). We describe below the main steps of the plan execution.

1) Action selection: When, according to the plan, there are actions to perform by the robot or the *X agent*, the robot selects an action based on priorities. In our case, a higher priority is given to the actions allocated to the robot compared to those allocated to the *X agent*. Then, the priorities of the robot's actions are the same, so the robot will simply select one. However, it would be possible to integrate costs, e.g. to enable to select the action the farthest of what the human is currently doing. Finally, concerning the priorities of the actions allocated to the *X agent*, we choose to put a higher priority on what we call *analogous* actions. Two actions are *analogous* when they have exactly the same decomposition (same action name and same parameters). For example, in the blocks building scenario, the two first actions allocated to the *X agent* in Fig 3(a) are *analogous*: the first selected action will be [X-AGENT PICK_AND_PLACE RED-CUBE PLACEMENT]. All these choices are made with the idea to postpone as much as possible robot's decisions whenever it's possible, in order to leave the human the ability to take the initiative until the last moment.

2) Action allocation: If the action selected is attributed to *X agent*, the robot needs to decide whether it will execute it or let the human do it. To do so, the robot first looks for the possible actors of this action: agents which verify the preconditions of the action and which are not already performing another action. For example, if the human is currently doing something else and there is an action of the *X agent* to perform, the robot will be assigned to this action. Then, if there are still several possible actors for the action, the robot will compare the action cost for itself and the human and assign it to itself if it is less costly. In the current implementation, the only cost used concerns the *analogous* actions which will have lower cost for the robot (the robot can execute one and let to the human the possibility to perform the other(s)). Finally, as a last resort, we have developed two possible modes for action selection. In the **negotiation** mode, the robot directly asks its human partner if he wants to perform the action and then allocated the action according to this answer. In the **adaptation** mode, the robot waits a certain amount of time, and if the human does not take the initiative to perform the action, the robot does it. Allocating an action to an agent can lead to other actions being automatically allocated. For example, in our scenario, the choice of who places the first blue cube automatically leads to the other agent placing the second one. For this reason, after each allocation of an action of the *X agent*, HATP is called to get a new plan.

3) Actions Instantiation: Once the robot has decided to execute an action e.g. [Robot PICK_AND_PLACE RED-CUBE PLACEMENT], it needs to be able to deal with the *similar* objects. This choice will be instantiated as close as possible to the execution, i.e. in a PICK_AND_PLACE action which will be decomposed by a PICK(object) and then a PLACE(object, placement), the choice for the destination (PLACEMENT_1 or PLACEMENT_2) will be made when it comes to do the PLACE (leaving the time to the human to choose a placement for his red cube). Then, concerning the choice of the object, the robot will choose the less costly one. We choose to put a lower cost on objects accessible only by the robot. Then we use

a simple cost based on distance between the agents right-hands and objects that leads the robot to choose the object that are the closest to itself on one side and the farthest from the human on the other side. With these choices, we still want to let the maximum choices to the human while minimizing his effort.

4) Adaptation to human initiative: The robot needs to be able to properly react to human behavior. The robot is, in parallel to the others steps described previously, constantly monitoring the human actions (it detects all actions performed by the human which involve objects from the task). If the action was previously allocated to the *X agent*, the robot asks to HATP a new plan to evaluate the possible allocations of the remaining actions. If the human performs an unexpected action with respect to the shared plan, the robot re-plan from the new situation induced by the human action. Moreover, if the unexpected human action is *analogous* to the one the robot is currently performing, the robot also stops its action (its actions may not be needed anymore). Finally, if the human approaches an object which is involved in the current robot action (e.g. in the blocks building scenario, the human places his red cube in the placement the robot has chosen), the robot first halts its action. Then, the robot looks if it can find another *similar* object. If it finds one, it continues its action with this object. If not, it waits the human to retreat from the object, and if the human actions did not lead to a new plan it continues its action if possible.

5 RESULTS

Illustrative examples on a real robot We ran experiments with several setups of the blocks building task described in Sec. 2 and with different kinds of human behaviors. Due to space constraint, we will present here only one execution which includes several interesting points. However more scenarios are available in the attached video¹.

The presented scenario starts with the setup in Fig. 1(a). The plan produced can be found in Fig. 3(a). It starts with two *analogous* actions for the *X agent* (place a red cube), the supervisor selects one and starts to execute it. So, the robot picks the red cube and, at the same time, evaluates the consequences of its choice. Since each agent owns only one red cube, the new plan computed (Fig. 3(d)) states that the human needs to place the second red cube. After picking its red cube the robot starts to place it on the placement the farthest to the human. However, the human picks his red cube and places it in the very same placement (Fig. 3(b)). So, the robot halts and adapts by placing its cube in the other placement (Fig. 3(c)). Then, the human places the stick on the red cubes. In this scenario, we have chosen the **negotiation** mode. Since the next action is allocated to the *X agent*, the robot asks the human if he wants to do it ("*Do you want to place the blue cube?*"). The human answers yes, leading the robot to compute a new plan in Fig. 3(e) where the human have to place the first blue cube and the robot the second one. Finally, the human and the robot perform their last actions and achieve the goal.

¹ Video available at <https://youtu.be/MZFG9GV3OD8>




Fig. 3: A situation where the robot adapts smoothly its behavior to the human actions. In (b), the human places his cube in the very same placement chosen by the robot. In (c), the robot adapts by placing its cube in the other placement

Quantitative results In order to evaluate our system, we ran it in simulation using the same scenario. Different set-ups were used as initial states and the robot was confronted to a simulated human. This simulated human performs all actions that are feasible only by him and answers robot questions. When confronted to an *X agent* action, he either chooses to perform it with 50% chance (50%-case), systematically chooses to perform it (hurry-case) or systematically chooses not to perform it (lazy-case). Then, we settled two different human behaviors:

- **the "kind" human (case=K)** who adapts his behavior to what the robot verbalizes (ie does an action if the robot asks him and stops an action if the robot says it will perform it)
- **the "stubborn" human (case=S)** who does not react nor comply to robot verbalization (he will not change his decision whatever the robot says).

A second independent variable corresponds to the 4 different modes:

- using the original system, called Reference System (RS), with all decisions and instantiations performed at planning time:
 - **RS-none mode:** the robot verbalizes nothing (unless it is strictly necessary)
 - **RS-all mode:** the robot informs the human when he has to perform an action and when it will act,
- using the proposed system, called New System (NS):
 - **NS-N:** the robot uses the **Negotiation** mode previously defined when a decision need to be made concerning *X agent* action,
 - **NS-A:** the robot uses the **Adaptation** mode.

Results can be found in Tab. 1. The dependent variables we measured are:

the number of verbal interactions between the human and the robot (either an information given by the robot or question asked).

the number of human/robot incompatible decisions: either both decide to perform the same action (and the robot stops its own action to avoid the conflict) or both decide not to perform the action (the robot first asks the human to

	Verbal interactions				Incompatible decisions			
	RS-none	RS-all	NS-N	NS-A	RS-none	RS-all	NS-N	NS-A
50%-K	0.4 (0.52)	6.0 (0.0)	1.2 (0.63)	0.0	0.6 (0.52)	0.0	0.0²	0.0
hurry-K	0.0				0.3 (0.48)			
lazy-K	0.9 (0.32)				0.9 (0.32)			
50%-S	0.4 (0.52)	6.4 (0.52)			0.5(0.53)	0.6 (0.52)		
hurry-S	0.0	6.0 (0.0)			0.3 (0.48)	0.3 (0.48)		
lazy-S	0.9 (0.32)	6.9 (0.32)			0.9 (0.32)	0.9 (0.32)		

Table 1: Results for the Reference System (RS) and the proposed New System (NS): NS-N for the Negotiation mode and NS-A for the Adaptation mode. The measures are the average number of verbal interactions and the average number of incompatible decisions between the human and the robot. The numbers correspond to means in 10 runs and their associated standard deviations.

perform the action after a predefined time and, if after another period the human has still not executed the action, the robot looks for a new plan where it can proceed).

Reference System performance: The verbalizations in the RS-none mode corresponds to the case where the human and the robot both choose not to execute the action: the robot tries to solve the conflict by asking the human to execute the action. Because the robot does not inform about its decisions, the number of verbal interactions is low in this mode. However, due to the same reason, there is several incompatible decisions in each conditions. The RS-all mode avoids incompatible decisions with the "kind" human. However, the number of verbal interaction is high (6 as the number of actions to execute in the task and so to verbalize). With the "stubborn" human, even if the robot informs the human, incompatible decisions remains. The number of verbal interaction also increases as, when the human does not want to perform the action, as it is stubborn, the robot needs to compute a new plan where it executes the action, and so inform about the new action.

New System performance: We can see that the robot is able to avoid conflicts in all cases without being too talkative (or without being talkative at all for the adaptation mode). Moreover, the efficiency of the system is not degraded with the "stubborn" human: the system allows the human to execute the actions he wants without an increase of verbal interaction. Finally, here the adaptation mode performs better than the negotiation one since the human is simulated and always performs his actions in time. However, in a real context, the negotiation

² These results have been obtained with a simulated human which always takes his decision in a time shorter than the time the robot waits. Consequently there are 0 incompatible decisions. In a more realistic context, there can be incompatible decisions if the robot waiting time is not well configured, or if the human changes his mind. However, there will always be less incompatible decisions than in the case where the robot does not inform at all.

mode would certainly have the benefit to ensure the absence of conflicts even if the robot is a little more talkative. Moreover, a human would surely be more comfortable with a robot which directly asks when (and only when) there is a decision to take compared to a robot which has unnecessary waiting time. Such measure of "satisfaction" cannot be easily simulated and further experiments will be done with real humans.

6 CONCLUSION

In this paper we have presented a system which allows the robot to adapt its behavior and let the maximum latitude and initiative to the human thanks to its ability to determine the decisions that should better be postponed. This also helps to reduce communication avoiding to have a too talkative robot while still avoiding conflicts. Furthermore, this work has been implemented in a complete architecture and is based on algorithms which are generic enough to be used for other tasks and contexts. A recent user study with real subjects has been performed. The first results of this study suggest that the proposed system is better perceived by the users and that the negotiation mode is preferred to the adaptation one for naive users. Finally, A number of the mechanisms described in the paper involve cost estimation in order to decide between options. In the current system, simple costs are used but could be easily replaced by more elaborate ones. For instance, a finer estimation of action costs based on geometric reasoning and human efforts or an estimation of accumulated costs of all actions remaining in the plan.

References

1. Allen, J., Ferguson, G.: Human-machine collaborative planning. Third International NASA Workshop on Planning and Scheduling for Space (2002)
2. Baraglia, J., Cakmak, M., Nagai, Y., Rao, R., Asada, M.: Initiative in robot assistance during collaborative task execution. HRI, ACM/IEEE (2016)
3. Botelho, S.C., Alami, R.: M+: a scheme for multi-robot cooperation through negotiated task allocation and achievement. Robotics and Automation, IEEE (1999)
4. Bratman, M.E.: Shared intention. Ethics (1993)
5. Chakraborti, T., Zhang, Y., Smith, D.E., Kambhampati, S.: Planning with resource conflicts in human-robot cohabitation. AAMAS (2016)
6. Chao, C., Thomaz, A.: Timed petri nets for fluent turn-taking over multimodal interaction resources in human-robot collaboration. IJRR (2016)
7. Chien, S.A., Knight, R., Stechert, A., Sherwood, R., Rabideau, G.: Using iterative repair to improve the responsiveness of planning and scheduling. AIPS (2000)
8. Cirillo, M., Karlsson, L., Saffiotti, A.: Human-aware task planning: an application to mobile robots. IST (2010)
9. Clodic, A., Alami, R., Chatila, R.: Key elements for human-robot joint action. Frontiers in AI and Applications (2014)
10. Devin, S., Alami, R.: An implemented theory of mind to improve human-robot shared plans execution. HRI, ACM/IEEE (2016)
11. Devin, S., Milliez, G., Fiore, M., Clodic, A., Alami, R.: Some essential skills and their combination in an architecture for a cognitive and interactive robot. Workshop in HRI, ACM/IEEE (2016)

12. Fiore, M., Clodic, A., Alami, R.: On planning and task achievement modalities for human-robot collaboration. *ISER* (2014)
13. Gerkey, B.P., Mataric, M.J.: A formal analysis and taxonomy of task allocation in multi-robot systems. *The IJRR* (2004)
14. Gombolay, M.C., Gutierrez, R.A., Clarke, S.G., Sturla, G.F., Shah, J.A.: Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. *Autonomous Robots* (2015)
15. Grosz, B.J., Kraus, S.: Collaborative plans for complex group action. *Artificial Intelligence* (1996)
16. Hoffman, G., Breazeal, C.: Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. *HRI, ACM/IEEE* (2007)
17. Johansson, M., Skantze, G.: Opportunities and obligations to take turns in collaborative multi-party human-robot interaction. *SIGdial* (2015)
18. Karpas, E., Levine, S.J., Yu, P., Williams, B.C.: Robust execution of plans for human-robot teams. *ICAPS* (2015)
19. Lallée, S., Hamann, K., et al.: Cooperative human robot interaction systems: Iv. communication of shared plans with naïve humans using gaze and speech. *IROS, IEEE/RSJ* (2013)
20. Lallement, R., de Silva, L., Alami, R.: HATP: an HTN planner for robotics. *CoRR* (2014)
21. Lemaignan, S., Warnier, M., et al.: Artificial cognition for social human-robot interaction: An implementation. *Artificial Intelligence* (2016)
22. Levine, S.J., Williams, B.C.: Concurrent plan recognition and execution for human-robot teams. *ICAPS* (2014)
23. Milliez, G., Lallement, R., Fiore, M., Alami, R.: Using human knowledge awareness to adapt collaborative plan generation, explanation and monitoring. *HRI ACM/IEEE* (2016)
24. Milliez, G., Warnier, M., Clodic, A., Alami, R.: A framework for endowing an interactive robot with reasoning capabilities about perspective-taking and belief management. *RO-MAN* (2014)
25. Mohseni-Kabir, A., Rich, C., Chernova, S., Sidner, C.L., Miller, D.: Interactive hierarchical task learning from a single demonstration. *HRI, ACM/IEEE* (2015)
26. Petit, M., Lallée, S., et al.: The coordinating role of language in real-time multimodal learning of cooperative tasks. *Autonomous Mental Development, IEEE* (2013)
27. Sebanz, N., Bekkering, H., Knoblich, G.: Joint action: bodies and minds moving together. *Trends in cognitive sciences* (2006)
28. Shah, J., Wiken, J., Williams, B., Breazeal, C.: Improved human-robot team performance using chaski, a human-inspired plan execution system. *HRI* (2011)
29. Talamadupula, K., Briggs, G., et al.: Coordination in human-robot teams using mental modeling and plan recognition. *IROS, IEEE/RSJ* (2014)
30. Tambe, M.: Agent architectures for flexible, practical teamwork. *14th National Conf. on AI* (1997)
31. Waldhart, J., Gharbi, M., Alami, R.: A novel software combining task and motion planning for human-robot interaction. *AAAI Fall Symposium Series* (2016)
32. Warnier, M., Guitton, J., Lemaignan, S., Alami, R.: When the robot puts itself in your shoes. managing and exploiting human and robot beliefs. *RO-MAN* (2012)