



HAL
open science

An Implemented Theory of Mind to Improve Human-Robot Shared Plans Execution

Sandra Devin, Rachid Alami

► **To cite this version:**

Sandra Devin, Rachid Alami. An Implemented Theory of Mind to Improve Human-Robot Shared Plans Execution. The Eleventh ACM/IEEE International Conference on Human Robot Interaction, Mar 2016, Christchurch, New Zealand. pp.319-326, 10.1109/HRI.2016.7451768 . hal-01330339

HAL Id: hal-01330339

<https://hal.science/hal-01330339>

Submitted on 10 Jun 2016

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.

An Implemented Theory of Mind to Improve Human-Robot Shared Plans Execution.

Sandra Devin
CNRS, LAAS,
Univ de Toulouse, INP, LAAS,
7 avenue du colonel Roche,
F 31400 Toulouse, France

Rachid Alami
CNRS, LAAS,
Univ de Toulouse, LAAS,
7 avenue du colonel Roche,
F 31400 Toulouse, France

Index Terms—Theory of Mind; Joint Action; Shared Plan; HRI.

Abstract—When a robot has to execute a shared plan with a human, a number of unexpected situations and contingencies can happen due, essentially, to human initiative. For instance, a temporary absence or inattention of the human can entail a partial, and potentially not sufficient, knowledge about the current situation. To ensure a successful and fluent execution of the shared plan the robot might need to detect such situations and be able to provide the information to its human partner about what he missed without being annoying or intrusive. To do so, we have developed a framework which allows a robot to estimate the other agents mental states not only about the environment but also about the state of goals, plans and actions and to take them into account when executing human-robot shared plans.

I. INTRODUCTION

In robotics, one of the current research interests is to create robots able to work jointly with humans or to help them in everyday life. To do so, robots need to perform joint actions with humans. They are already capable of computing shared plans [1], not only for themselves, but also for other agents (humans or robots) involved in a joint task ([2], [3]). They can execute their part of the plan while monitoring the activity of the other agents. However, when working with humans, the global execution of these shared plans is not always simple. Indeed, unexpected situations and contingencies can happen due essentially to human initiative and robots need to correctly interpret them and act accordingly.

One of these unexpected situation is the temporary absence or inattention of a human. If it happens, when the human comes back, he can lack information about what happened during his absence. Some of the facts are observable while others are not, leading the human to have a partial, and potentially not sufficient, knowledge about the current situation. To ensure a successful and fluent execution of the shared plan, the robot needs to inform the human about what he missed. However, systematically informing him about all missing information can quickly lead to an intrusive and annoying behaviour of the robot. Consequently, the robot needs to be able to distinguish which information the human really needs and which he does not. To do so, the robot needs to estimate correctly and take into account the mental state of its human partner.

The ability to reason about other people perception, beliefs and goals and to take them into account is called Theory of

Mind (ToM) ([4], [5]). In robotics, previous work in ToM focused mainly on perspective taking and belief management: robots are able to reason about what humans can perceive or not, and then construct representations of the world from their point of view ([6], [7], [8]). They can use this knowledge to learn tasks, solve ambiguous situations or understand humans behaviour. However, there is still a gap between such representations and those necessary to permit an effective shared plan action execution.

In order to fill this gap, we present, as a first step, a framework that allows the estimation by the robot of its human partner mental state related to collaborative task achievement. Such mental states contain not only the state of the world but also the state of the goals, plans and actions. They are based on the capability of the robot to permanently compute the spatial perspective of its partners and to track their activity. As a second step, we present how the robot uses these mental states to perform joint actions with humans and more particularly to manage the execution of shared plans in a context of humans and robots performing collaborative objects manipulation. As a result, the robot is able to adapt to humans decisions and actions and to inform them when needed without being intrusive by giving (unnecessary) information that the human can observe or infer by himself.

We first present the global architecture in §III and then briefly situate, in §IV, our work relatively to the joint action literature. The formal definition of what we call a goal, a plan and an action are described in §V and we define more precisely the representation of the mental state used by our robot in §VI. These definitions are then used in §VII to explain how the mental states are estimated and updated. §VIII shows how the robot uses this information to manage shared plans execution. Finally, we evaluate our system in §IX and conclude.

II. BACKGROUND

Previous work on ToM in robotics mainly concern perspective taking and belief management. One pioneer work is [9] where two models from social sciences are analysed conducting to a first architecture integrating ToM for robots. Trafton et al., based on the *Polyscheme* architecture, compute spatial perspective taking to help the understanding of instructions in dialogue between astronauts [6]. Based on the ACT-R architecture, [10] models the mechanisms used to take

decisions during the Sally and Anne test ([11]). ACT-R is also used in [12] to explain, with the help of perspective taking, unexpected behaviours from a human partner. Breazeal et al. use perspective taking to learn tasks from ambiguous demonstrations [13]. Gray et al. simulate the outcomes of actions from other agents point of view to predict the actions effects in others' mental states [14]. Milliez et al. propose a belief management that enables a robot to pass the Sally and Anne test [8]. This reasoning is used in [15] to solve ambiguous situations in dialogue and in [3] to compute human-aware plans taking into account potential divergence of beliefs between the human and the robot or lack of information in the human knowledge about the current world state.

Concerning shared plans execution, to our knowledge, only a few contributions integrate the consideration of the human partner mental state. Human activities can be fairly well monitored ([16], [17]), robot actions can be executed taking into account the humans in the environment ([18], [19]) and human-aware plans can be computed ([2], [3]). However, very few architectures track and take into account the human partner state during the execution of a joint task. Clodic et al. present SHARY, a supervision system which enables to execute human-aware plans [20]. In this system, the execution can be stopped and a new plan computed in case of unexpected situations. Fiore et al. extended this work in a more robust way that includes new aspects of joint action like reactive action execution [21]. Fong et al. present HRI/OS, a system able to produce and schedule tasks for different agents based on their capacities [22]. However, the agents act mostly in a parallel and independent way. Shah et al. present *Chaski*, a task-level executive that allows to choose when to execute the robot actions adapting to the human partner [23]. However, none of these architectures explicitly takes into account the humans mental state when executing a shared plan with him.

Previous work has also been conducted in order to model shared plans [24] or agent intentions [25]. Researchers in Artificial Intelligence (AI) developed Beliefs, Desires and Intention (BDI) theories to model rational agents for multi-agents activities ([26], [27]). Frameworks have been devised to model multi-agents joint activities and shared plans execution in a robust way with regards to joint intention and beliefs about joint intention of agents but without taking into account other agents knowledge about the state of the plans and actions nor spatial perspective taking ([28], [29]).

III. OVERALL ARCHITECTURE

In this paper, we place ourselves in a context where an assistant or team-mate robot has to work jointly with other agents (humans or robots). They share an environment and mutually observe each other. Fig. 1 illustrates the implemented architecture on the robot. It is composed of:

- **A Situation Assessment module [8]:** which takes as input the sensor data and maintains the current world state from the point of view of all agents based on spatial perspective-taking. It also computes a set of symbolic facts that represent the observable world state from

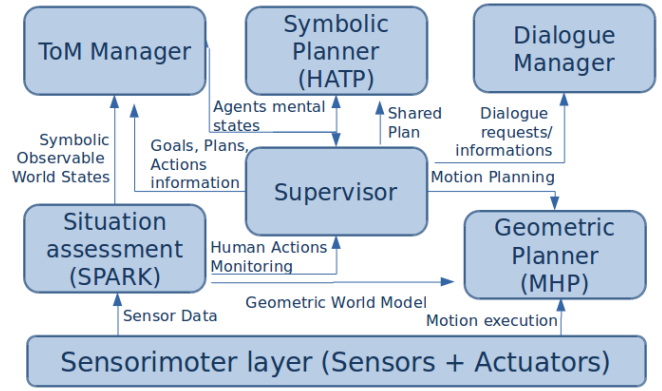


Fig. 1: The implemented architecture. The work presented here concerns the ToM Manager and the Supervisor.

the different agents point of view. These facts concern relations between objects (e.g. (mug isOn table)) or agents affordances (e.g. (mug isVisibleBy human), (book isReachableBy robot)).

- **A ToM Manager:** which takes the symbolic world models computed by the Situation Assessment module and information from the supervisor about the execution of goals, plans and actions in order to estimate and maintain the mental state of each agent involved in the cooperation. These mental states contain not only world state information but also an estimation of the agents beliefs about the tasks and the other agents capacities. More details concerning this module are given in §VII.
- **A high-level task planner [3], [30]:** which allows the robot to synthesize shared plans containing the actions of all agents involved in a given task. It is an HTN task planner which bases its reasoning on the symbolic representations of the world from the Situation Assessment module and on the agents abilities on the current context (for example which agent can reach which object). It searches for the best possible plan to achieve the shared goal by taking into account human-aware costs in order to come up with highly acceptable plans. An example of such a plan is given in Fig. 3(b).
- **A geometric action and motion planner [31]:** which allows to compute trajectories as well as objects placements and grasps in order to perform actions like Pick or Place while taking into account the human safety and comfort.
- **A dialogue manager [32], [33]:** which allows to verbalize information to the human and to recognize basic vocal commands.
- **A Supervisor [21]:** in charge of collaborative activity. To do so, it takes the human partner mental state into account to decide when to perform actions or to communicate with him. It also interprets the information coming from the Situation Assessment module in order to recognize human actions like Pick or Place. More details concerning

this module can be found in §VIII.

IV. OUR WORK IN THE JOINT ACTION CONTEXT

The first step during a joint action is to share a goal. This implies a number of issues related to agent commitment. In order to focus on the paper contribution, we consider here that the joint goal has already been established: we consider that the robot and its human partners have a commitment following the definition of *weak achievement goal* [34]. Consequently, we assume that none of the humans will abort the goal unless he knows that the goal is not achievable any more.

Once the joint goal had been established, the involved agents need to share a plan. This plan can come from multiple sources: it can be imposed by a human, negotiated through dialogue or the robot can compute it. To focus on the execution of the shared plan, we choose, for this paper, to let the robot compute the plan. However, the processes presented in this paper hold even if the plan comes from a different source. Moreover, as the symbolic planner used in the architecture takes into account human-aware costs, we assume that the computed plan will be close enough to human expectations to be accepted by him. This plan is automatically shared by the robot at the beginning of the interaction (by displaying it on a screen near the robot or by verbalization). In this paper, we will discuss about the processes used during the execution of shared plans, and, we will not focus on issues linked to the computation or the communication of these plans.

The execution of shared plan is based on agent capacities. We distinguish between humans and robots capacities. We consider that a human has the capacities to:

- **Perform high level actions:** He is able to perform a set of high level actions like Pick or Place. The shared plan is then computed based on this assumption. In order to know what a human can do in a given situation, the Situation Assessment module computes, based on geometry, an estimation of his ability to reach the objects of the environment.
- **Perceive:** The Situation Assessment module computes an estimation of what the human can see in the environment. This information added to a belief management algorithm allows the robot to estimate, at any moment, what the human knows about the environment. Concerning actions, we make the assumption that a human will see and understand an action of another agent (mainly robot actions) when he is present and looking at the agent.
- **Communicate:** The dialogue manager allows the human, at any moment, to ask the robot to perform an action. We also assume that when he is present, the human is able to hear and understand the information verbalized by the robot.

We also consider that the robot has the following capacities:

- **Perform high level actions:** like Pick, Place, Drop, Handover, etc... Similarly to the human, we consider a set of high level actions that the robot is able to perform and that are taken into account to build the shared plan.

Likewise, reachabilities are computed for the robot in order to decide which action it is able to perform at each time.

- **Perceive:** The Situation Assessment module allows to keep an estimation of the current state of the environment: the robot is able to detect and localize objects and agents. The robot is also able to recognize simple high level actions performed by a human like Pick, Place or Drop.
- **Communicate:** Thanks to the Dialogue module, the robot is able to ask to a human to perform an action and to inform him about the state of the environment, the goal, a plan or an action. The robot is also able to share a plan through speech synthesis (plan verbalization) or by displaying it on a screen.

V. REPRESENTATIONS

Let's define the goal which the robot needs to perform with other agents, the plans computed to achieve this goal and the actions which compose these plans.

Let's G be the set of all goals. A goal $g \in G$ is defined as:

$$g = \langle name_g, AGG_g, O_g \rangle$$

Where $name_g$ is used to identify the goal, O_g is a set of facts representing the desired world state and AGG_g is a set of agents¹ involved in the goal

Let's P be the set of all plans. A plan $p \in P$ is defined as:

$$p = \langle id_p, g_p, ACP_p, L_p \rangle$$

Where id_p is used to identify the plan and g_p is the goal that the plan allows to achieve. ACP_p is a set of actions (see below) that compose the plan and L_p is a set of links defining actions order and causal links.

A link l is defined as $l = \langle previous_c, after_c \rangle$ where $previous_c$ is the id of the action which needs to be achieved before the action with the id $after_c$ is performed.

Let's ACT be the set of all actions. An action $ac \in ACT$ is defined as:

$$ac = \langle id_{ac}, name_{ac}, AGC_{ac}, PR_{ac}, PRE_{ac}, EFF_{ac} \rangle$$

Where id_{ac} is the action identifier and $name_{ac}$ represents its name (e.g. pick, place, etc...). AGC_{ac} is a set of agent names representing the actors of the actions and PR_{ac} a set of entities (objects or agents) which allows to define precisely the action (e.g. the name of the object to pick). PRE_{ac} and EFF_{ac} are sets of facts representing respectively the action preconditions and effects.

VI. MENTAL STATES

Let's also define how the robot represents other agents. Let's A be the set of all agents, an agent $ag \in A$ is defined as:

$$ag = \langle name_{ag}, type_{ag}, CAP_{ag}, MS_{ag}, AG_{ag} \rangle$$

Where $name_{ag}$ is used to identify the agent and $type_{ag}$ represents if the agent is a human or a robot. CAP_{ag} , the set of

¹Often only the robot and a human partner, but other robots and humans can be concerned

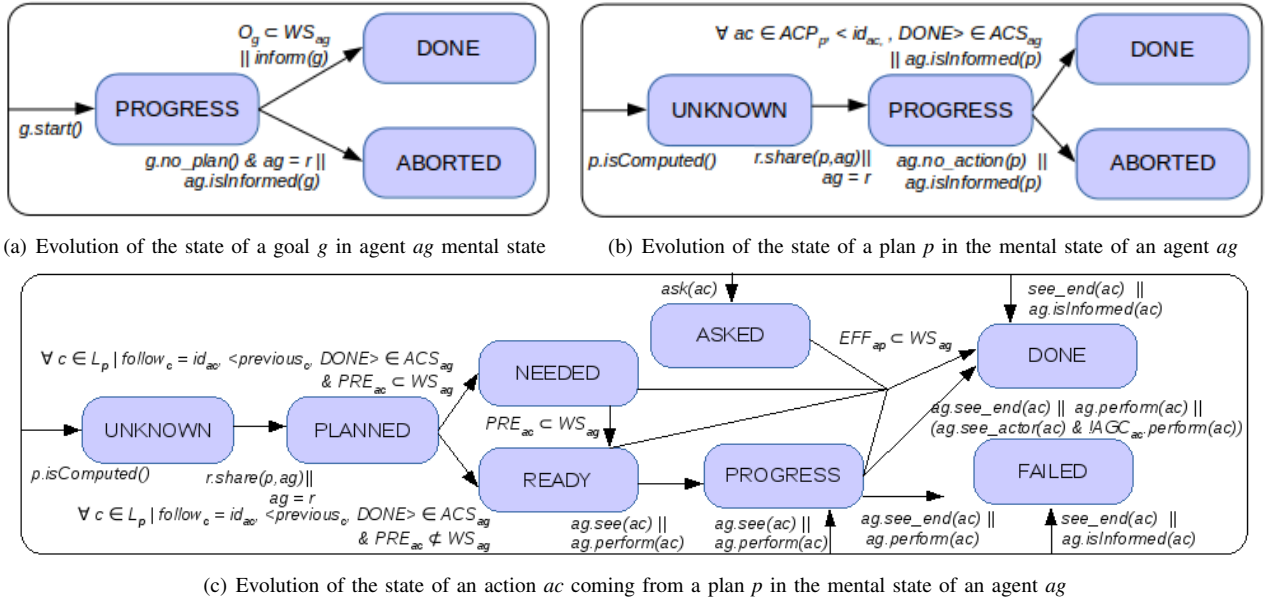


Fig. 2: Goal, plan and action evolutions.

high level action names, representing the actions that the agent is able to perform, is used by the symbolic planner to produce plans. MS_{ag} is the mental state of the agent (see below) and AG_{ag} is a set of agents containing all agents excepting ag : $AG_{ag} \subset A \mid AG_{ag} = \{A \setminus \{ag\}\}$. These agents are defined in the same way as ag and model how the given agent represents them. Accordingly, these agents will also contain other agents and so on in order to model high-order ToM [35].

The mental state MS_{ag} of ag is defined as:

$$MS_{ag} = \langle WS_{ag}, GS_{ag}, PS_{ag}, ACS_{ag} \rangle$$

Where WS_{ag} is a set of facts representing the current world state from the agent point of view.

GS_{ag} represents the state of the goals from the agent point of view. The state of a goal g is represented as $\langle name_g, state_g \rangle$ where $state_g$ can be either PROGRESS if the agent thinks the goal is in progress, DONE if the the agent thinks the goal has been achieved or ABORTED if the agent thinks the goal has been aborted.

PS_{ag} represents the state of the plans from the agent point of view. The state of a plan p is represented as $\langle id_p, state_p \rangle$ where $state_p$ can be either PROGRESS, DONE, ABORTED or UNKNOWN if the agent is not aware of the plan².

Finally, ACS_{ag} represents the state of the actions from the agent point of view. The state of an action ac is represented as $\langle id_{ac}, state_{ac} \rangle$ where $state_{ac}$ can be either PROGRESS, DONE, FAILED, ASKED (an agent asked for the action to be done), PLANNED (need to be done later according to the current plan), NEEDED (need to be done now according to the current plan but not possible), READY (need to be done now according to the current plan and possible) or UNKNOWN.

²An agent is not aware of a shared plan if it has not contributed to its synthesis or has not been informed by the agent(s) who has elaborated it

At each moment, goals, plans and actions can only have one state in each agent mental state.

VII. THE TOM MANAGER

The ToM Manager allows the robot to estimate and maintain the mental states of each agent interacting with the robot and is limited here to a first-order ToM the robot has its own knowledge and a representation of the other agents and their knowledge. However, these last agents do not have representations of other agents. In order to do that, we use an agent $r \in A$ which represents our robot. As defined previously, r contains a representation of the other agents AG_r . However, the agents in AG_r do not contain further other agents representation ($\forall ag \in AG_r, AG_{ag} = \emptyset$). We will present here how the robot estimates and maintains the mental state of an agent $ag \in AG_r$ but these processes are the same for all agents in AG_r and for the mental state of the robot r .

A. World state

The representation of the world state by an agent is composed of two types of facts:

- **Observable facts:** these facts are computed and maintained by the Situation Assessment module. They concern what the agent can observe about the world state. These facts represent the affordances of all agents (e.g. isVisibleBy, isReachableBy) and the relations between objects (e.g. isOn, isIn) visible to them.
- **Non-observable facts:** they can not be computed by the Situation Assessment module. They concern information that the agent can not observe (e.g. the fact that an object is into a closed box). There are two ways for an agent to be aware of a non-observable fact, it can be informed by

another agent or it can perform or see an action that has this fact in its side effects (when the robot estimates that an agent considers an action DONE, it considers that the agent is aware of all the effects of the action).

B. Goals

The evolution of the state of a goal $g \in G$ in the mental state of an agent ag is described in Fig. 2(a). As said before, we consider that the agent already commits to the goal. Consequently, when the robot starts to execute a goal ($g.start()$) the goal is considered in progress. The agent considers a goal achieved when all the facts belonging to the objective of the goal are known to it (accordingly to its knowledge, the desired world state has been reached). As we do not focus on issues relative to shared intention, we consider that the agent will not abort the goal on his own. The robot will abort the goal when no more plan can be found to achieve it ($g.no_plan()$). Finally, the agent can be informed about the result of a goal ($ag.isInformed(g)$).

C. Plans

The evolution of the state of a plan $p \in P$ in the mental state of an agent ag is described in Fig. 2(b). Each time the robot computes a plan ($p.isComputed()$) to achieve a goal, it shares it with the other agents ($r.share(p,ag)$). Accordingly, if the agent is present, it considers the plan in progress. The agent considers a plan achieved when he considers each action included in it performed. If, at any time, the agent considers that there is still actions from the plan to be performed but there are no action in progress or action to be done that are possible, the agent considers that the plan is aborted ($ag.no_action(p)$, Alg. 1). Finally, the agent can be informed about the result of a plan ($ag.isInformed(p)$).

Algorithm 1 $ag.no_action(p)$

```

if ( $\nexists \langle id, state \rangle \in ACS_{ag} \mid state = PROGRESS$ 
   $\parallel state = READY$ )
  & ( $\exists \langle id_2, state_2 \rangle \in ACS_{ag} \mid state_2 = PLANNED$ 
   $\parallel state_2 = NEEDED$ ) then
   $\langle id_p, ABORTED \rangle \in PS_{ag}$ 
end if

```

D. Actions

The evolution of the state of an action $ac \in AC$ belonging to a plan $p \in P$ in the mental state of an agent ag is described in Fig. 2(c). When a new plan is shared by the robot ($r.share(p,ag)$), all the actions in this plan are considered planned by the agent. The agent considers an action READY if all previous actions in the plan (based on the plan links) are considered done by the agent and if the agent considers all the preconditions of the action true. If the agent does not consider all the preconditions of the action true, it considers the action NEEDED. When an action is executed, from an agent perspective, if it performs it ($ag.perform(ac)$) or observes another agent performing it ($ag.see(ac)$), it knows then that the

action is in progress. In a same way, we consider that, when an action is over, if the agent performed it ($ag.perform(ac)$) or observed the end of its execution ($ag.see_end(ac)$), it knows the result of the action. When the agent is informed that an action has been done, it also infers the effects of the action. But an agent can also infer that an action has been done if it knows that the action was in progress or on its way to be done and it can see the effects of the action or if it knows that the action was in progress and can see the actors of the action ($ag.see_actor(ac)$) and that there are not performing the action any more ($!AGC_{ac}.perform(ac)$). An agent can also be asked (or ask to somebody else) to perform an action ($ask(ac)$) and be informed about the result of an action ($ag.isInformed(ac)$).

VIII. THE SUPERVISOR

The supervisor manages shared plans execution. To do so, it makes the robot execute its actions when it considers they are needed and possible (action READY in the robot knowledge) and, in parallel, the robot monitors the humans activities in order to detect their actions. If the supervisor estimates that the current plan is not feasible any more (plan ABORTED in its knowledge) it tries to compute a new plan. If it succeeds, it shares the plan and starts to execute it. If it fails, it aborts the goal. Thus, when a human performs an unexpected action or if an action fails, the supervisor is able to quickly produce a new plan and adapt to the new situation.

In this paper, we will focus on the activity of the supervisor which allows to manage divergent belief during the execution of a shared plan. Indeed, when two humans share a plan, they usually do not communicate all along the plan execution. Only the *meshing subplans* of the plan need to be shared [36]. Consequently, the robot should inform humans about elements of the shared plan only when it considers that the divergent belief might have an impact on the joint activity in order to not be intrusive by giving them information which they do not need or which they can observe or infer by themselves.

A. Weak achievement goal

If we follow the definition of *weak achievement goal* in [34], if the robot knows that the current goal has been achieved or is not possible anymore, it has to inform its partners. Accordingly, we consider that, when, in the robot knowledge, the state of a goal is DONE (resp. ABORTED) and the robot estimates that a human does not consider it DONE (resp. ABORTED), the robot informs him about the achievement (resp. abandoning) of the goal (if the agent is not here or is busy with something else, the robot will do it as soon as the agent is available). We extend this reasoning to plans: the robot informs in the same way about the achievement (resp. abandoning) of a plan.

B. Before human action

A divergent belief of a human partner can be an issue when it is related to an action that he has to do. To avoid that a human misses information to execute his part of the plan, each time the robot estimates that a human has to do an action

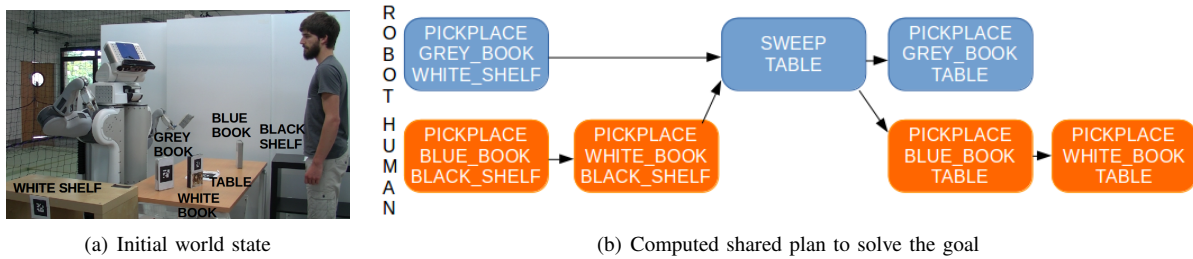


Fig. 3: Initial set-up of the Clean the table scenario. The human and the robot have to clean the table together. The computed shared plan is to remove the three books lying on the table, sweep it and place back the books.

(action with the state `READY` in the robot knowledge and with a human in its actors) it checks if the human is aware that he has to and can perform the action (the state of the same action should be `READY` in the estimation of the knowledge of the human too). If it is not the case, there are three possibilities:

- The state of the current plan is not in `PROGRESS` in the estimation of the agent knowledge: the human misses information about the current plan, so, the robot shares it with him.
- The state of the action is `PLANNED` in the estimation of the human knowledge: the human misses information about previous achieved actions to know that his action has to be performed now according to the plan. The robot checks the state of all actions linked to the first one with the plan links and informs about the achievement of all actions with a state different of `DONE` in the estimation of the human knowledge.
- The state of the action is `NEEDED` in the estimation of the human knowledge: the human misses information about the world state to know that his action is possible. In such case, the robot looks into the preconditions of the actions and informs the human about all those the human is not aware of.

C. Preventing mistakes

A divergent belief of a human partner can also be an issue if it leads him to perform an action that is not planned or not to do now according to the plan. To prevent this, for each action that the robot estimates the human thinks `READY`, the robot checks if the action really needs to be done (action `READY` in its knowledge too). If it is not the case, the robot corrects the human divergent belief by two different ways:

- If the action is `PLANNED` in the robot knowledge: the human thinks that a previous action has been achieved successfully while it is not the case leading him to think he has to perform his action. The robot looks in all actions linked to the first one by the plan links and informs about their state if it is different in the estimation of the human knowledge and in the robot one.
- If the action is `NEEDED` in the robot knowledge: the human has a divergent belief concerning the world state that leads him to think that his action is possible while

it is not the case. The robot looks into the preconditions of the action and informs about divergent beliefs.

D. Signal robot actions

When the robot has to perform an action, it looks if it estimates that the humans are aware that it will act (the action should be `READY` in the humans knowledge). If it is not the case, the robot signals its action before performing it.

E. Inaction and uncertainty

Even if the robot estimates that the human is aware that he has to act (the state of the action which he must perform is `READY` in the estimation of his knowledge), it is possible that the human still does not perform this action. If the human is already busy by something else (there is an action in the robot knowledge with the state `PROGRESS` and with the human in its actors), the robot waits for the human to be available. If the human is not considered busy by the robot, the robot first considers that its estimation of the human mental state can be wrong, and that, in reality, the human is not aware that he should act. Consequently, the robot asks the human specifically to do the action. If the human still does not act while the action has been asked, the robot considers the action failed, aborts the current plan and tries to find an alternative plan excluding that action.

In order to avoid considering that the human is available, and so to disturb him while he is busy doing something that the robot can not recognize, we have added an action named *busy* used when the robot estimates that the human is doing something without knowing what. The action *busy* when executed by a human h can be defined as $\langle id, busy, \{h\}, \emptyset, \emptyset, \emptyset \rangle$.

IX. EVALUATION

A. Scenarios

1) *"Clean the table" scenario*: In this example, a PR2 robot and a human have to clean a table together. To do so, they need to remove all items from this table, sweep it, and re-place all previous items. The initial world state is the one in Fig. 3(a). We consider that the grey book is reachable only by the robot, the blue book only by the human and the white book by both agents. The actions considered during this task are *pick-and-place* and *sweep*. To pick and place an object on a support, the object and the support need to be reachable by the




Fig. 4: Initial set-up for the Inventory scenario. The coloured objects need to be scanned by the robot and then, put into a box of the same colour.

agent, and, to sweep a surface, it should not have any objects on it and it should be reachable by the agent. The initial plan produced to achieve the goal is shown in Fig. 3(b)³.

2) "Inventory" scenario: In this example, a human and a PR2 robot have to make an inventory together. At the beginning of the task, both agents have coloured objects near them as well as a coloured box (initial world state in Fig. 4). These coloured objects need to be scanned and then, stored in the box of the same colour. The actions considered during this task are *placeReachable*, *pickanddrop* and *scan*. We call *placeReachable* the action to pick an object and to place it such that it is reachable by the other agent. For an agent to perform this action, the object needs to be reachable by it. We call *pickanddrop* the action to pick an object and to drop it on a box. To perform such action, the object and the box need to be reachable by the agent. The *scan* action can only be performed by the robot and consists of scanning an object by orienting the head of the robot (assumed to be equipped by a scanner) in the direction of the object placed such that it is reachable.

B. Criteria and results

One objective of our contribution is to reduce unnecessary communication acts from the robot during the execution of a shared plan aiming at a more friendly and less intrusive behaviour of the robot. Consequently, in order to evaluate our system, we have chosen to measure the amount of information shared by the robot during a shared plan execution where the human misses some information. As it is not trivial to create and control a lack of knowledge from a human subject, we decided to evaluate our system in simulation. We ran experiments in the two scenarios described previously where a human and a robot have to perform a joint task together⁴. When the interaction starts, we consider that the joint goal is already established and that both human and robot already agreed on a shared plan. The robot executes the plan as described in the presented work and the simulated human executes the actions planned for him. We randomly sample

³This example has been fully implemented in a real robot. A detailed execution can be found in the attached video

⁴More details about these two scenarios can be found in the attached video

a time when the human leaves the scene and another time when the human comes back. While absent, the human does not execute actions and cannot see anything nor communicate.

During the interaction, we logged the number of facts (information chunks) given by the robot to the human. An information concerns either a change in the environment, the state of a previous action, the abortion of a previous plan or the sharing of a plan. We compared our system (called *ToM system*) to:

- a system which verbalizes any plan abortion, shares any new plan and informs about each action missed by the human (called *Missed system*).
- a system which verbalizes any plan abortion, shares any new plan and informs about each action performed by the robot even if the human sees it (called *Performed system*).

The obtained results are given in Table I.

Scenario System	Clean the table		Inventory	
	Average	Std Dev	Average	Std Dev
ToM	1.32	0.98	0.41	0.48
Missed	2.86	1.33	2.61	1.36
Performed	4.44	1.85	10.0	0.0

TABLE I: Number of information given by the robot during the two presented scenarios for the three systems (*ToM*, *Missed* and *Performed*).

We can see that our system allows to reduce significantly the amount of information given by the robot. In the "Clean the table" scenario, depending on when the human leaves, the robot might change the initial plan and take care of the book reachable by both agents instead of the human. This explains why the average number for the *Performed system* is higher than the number of actions initially planned for the robot: the robot performs more actions in the new plan and can communicate about the new plan and the abortion of the previous one. In this scenario, our system allows to communicate about the plan abortion only if the human can not infer it by himself and to not communicate about missed *pickandplace* actions as the human can infer them by looking at the objects placements. However, the robot will inform the human if he missed the fact that the robot has swept the table as it is not observable and it is a necessary information for the human to know before he can put back objects on the table.

In the inventory scenario, as all objects and boxes are reachable only by one agent, the robot does not change the plan when the human leaves. This explains the fact that the standard deviation is null for the *Performed system*: the number of actions performed by the robot never changes and there is no change in the plan. In this scenario, the *pickanddrop* and *scan* actions have non-observable effects (the human can not see an object in a box). However, we can see that our system still verbalizes less information than the *Missed system*: the robot communicates only the information which the human really needs (as the fact that an object the human should drop in a box has been scanned) and does not give information

which are not linked to the human part of the plan (as the fact that the robot scanned an object it have to drop in its box).

X. CONCLUSION

In this paper, we have presented a system that allows to estimate and maintain mental states of other agents concerning not only the environment but also about the state of goals, plans and actions in the context of human-robot shared plan execution. This system takes these mental states into account when executing human-robot shared plans by allowing the robot to manage human-robot joint activities taking into account the human perspective and its knowledge about the task. We have shown that this system allows to reduce the number of unnecessary information given to the human while giving the needed information to ensure a proper task achievement.

The novelty of this work is twofold: the estimated mental states concern not only observable information about the environment but also the state of current and previous goals, plans and actions, and, these mental states are taken into account during the execution of a shared plan in order to reduce unnecessary communication to produce a less intrusive behaviour of the robot. Moreover, this work has been implemented and run on a complete human-aware architecture, enabling the robot to fluently perform joint tasks with a human.

We have presented here one relevant use of the computed mental states. However, we strongly believe that they can also be used to tackle other challenges during human-robot joint actions. For example, another use could be to estimate the possible lack of information of the robot in order to allow it to ask for help or to ask information when needed. Another aspect, that we plan to explore in the future, is to reason about such mental states to better understand human “unexpected” behaviours.

Acknowledgments: This work has been funded by the French Agence Nationale de la Recherche ROBOERGOSUM project ANR-12-CORD-0030.

REFERENCES

- [1] B. J. Grosz and C. L. Sidner, “Plans for discourse,” tech. rep., DTIC Document, 1988.
- [2] M. Cirillo, L. Karlsson, and A. Saffiotti, “Human-aware task planning: an application to mobile robots,” *Intelligent Systems and Technology*, vol. 1, no. 2, p. 15, 2010.
- [3] J. Guitton, M. Warnier, and R. Alami, “Belief management for hri planning,” *European Conf. on Artificial Intelligence*, p. 27, 2012.
- [4] S. Baron-Cohen, A. M. Leslie, and U. Frith, “Does the autistic child have a theory of mind?,” *Cognition*, vol. 21, no. 1, pp. 37–46, 1985.
- [5] D. Premack and G. Woodruff, “Does the chimpanzee have a theory of mind?,” *Behavioral and brain sciences*, vol. 1, no. 04, pp. 515–526, 1978.
- [6] J. G. Trafton, N. L. Cassimatis, M. D. Bugajska, D. P. Brock, F. E. Mintz, and A. C. Schultz, “Enabling effective human-robot interaction using perspective-taking in robots,” *Systems, Man and Cybernetics*, vol. 35, no. 4, pp. 460–470, 2005.
- [7] M. Berlin, J. Gray, A. L. Thomaz, and C. Breazeal, “Perspective taking: An organizing principle for learning in human-robot interaction,” in *Nat. Conf. on Artificial Intelligence*, vol. 21, AAAI Press, MIT Press, 2006.
- [8] G. Milliez, M. Warnier, A. Clodic, and R. Alami, “A framework for endowing an interactive robot with reasoning capabilities about perspective-taking and belief management,” in *Int. Symp. on Robot and Human Interactive Communication*, pp. 1103–1109, IEEE, 2014.
- [9] B. Scassellati, “Theory of mind for a humanoid robot,” *Autonomous Robots*, vol. 12, no. 1, pp. 13–24, 2002.
- [10] L. M. Hiatt and J. G. Trafton, “A cognitive model of theory of mind,” in *10th Int. Conf. on Cognitive Modeling*, pp. 91–96, Citeseer, 2010.
- [11] H. Wimmer and J. Perner, “Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children’s understanding of deception,” *Cognition*, vol. 13, no. 1, pp. 103–128, 1983.
- [12] L. M. Hiatt, A. M. Harrison, and J. G. Trafton, “Accommodating human variability in human-robot teams through theory of mind,” in *Int. Joint Conf. on Artificial Intelligence*, vol. 22, p. 2066, 2011.
- [13] C. Breazeal, M. Berlin, A. Brooks, J. Gray, and A. L. Thomaz, “Using perspective taking to learn from ambiguous demonstrations,” *Robotics and Autonomous Systems*, vol. 54, no. 5, 2006.
- [14] J. Gray and C. Breazeal, “Manipulating mental states through physical action,” *Int. Journal of Social Robotics*, vol. 6, no. 3, 2014.
- [15] R. Ros, S. Lemaignan, E. A. Sisbot, R. Alami, J. Steinwender, K. Hamann, and F. Warneken, “Which one? grounding the referent based on efficient human-robot interaction,” in *Int. Symp. on Robot and Human Interactive Communication*, pp. 570–575, IEEE, 2010.
- [16] J. Gray, C. Breazeal, M. Berlin, A. Brooks, and J. Lieberman, “Action parsing and goal inference using self as simulator,” in *Int. Workshop on Robot and Human Interactive Communication*, IEEE, 2005.
- [17] Y. Demiris and B. Khadhoury, “Hierarchical attentive multiple models for execution and recognition of actions,” *Robotics and autonomous systems*, vol. 54, no. 5, pp. 361–369, 2006.
- [18] J. Mainprice, E. A. Sisbot, L. Jaillet, J. Cortés, R. Alami, and T. Siméon, “Planning human-aware motions using a sampling-based costmap planner,” in *Int. Conf. on Robotics and Automation*, IEEE, 2011.
- [19] A. D. Dragan, S. Bauman, J. Forlizzi, and S. S. Srinivasa, “Effects of robot motion on human-robot collaboration,” in *Int. Conf. on Human-Robot Interaction*, pp. 51–58, ACM/IEEE, 2015.
- [20] A. Clodic, H. Cao, S. Alili, V. Montreuil, R. Alami, and R. Chatila, “Shary: a supervision system adapted to human-robot interaction,” in *Experimental Robotics*, pp. 229–238, Springer, 2009.
- [21] M. Fiore, A. Clodic, and R. Alami, “On planning and task achievement modalities for human-robot collaboration,” in *Int. Symposium on Experimental Robotics*, 2014.
- [22] T. Fong, C. Kunz, L. M. Hiatt, and M. Bugajska, “The human-robot interaction operating system,” in *Int. Conf. on Human-robot interaction*, pp. 41–48, ACM/IEEE, 2006.
- [23] J. Shah, J. Wiken, B. Williams, and C. Breazeal, “Improved human-robot team performance using chaski, a human-inspired plan execution system,” in *Int. Conf. on Human-Robot Interaction, HRI’11*, ACM, 2011.
- [24] B. J. Grosz and S. Kraus, “Collaborative plans for complex group action,” *Artificial Intelligence*, vol. 86, no. 2, pp. 269–357, 1996.
- [25] P. R. Cohen and H. J. Levesque, “Intention is choice with commitment,” *Artificial intelligence*, vol. 42, no. 2, pp. 213–261, 1990.
- [26] A. S. Rao and M. P. Georgeff, “Modeling rational agents within a bdi-architecture,” *Readings in agents*, vol. 91, pp. 473–484, 1991.
- [27] M. J. Wooldridge, *Reasoning about rational agents*. MIT press, 2000.
- [28] N. R. Jennings, “Controlling cooperative problem solving in industrial multi-agent systems using joint intentions,” *Artificial intelligence*, vol. 75, no. 2, pp. 195–240, 1995.
- [29] M. Tambe, “Towards flexible teamwork,” *Journal of artificial intelligence research*, pp. 83–124, 1997.
- [30] R. Lallement, L. de Silva, and R. Alami, “HATP: an HTN planner for robotics,” *CoRR*, vol. abs/1405.5345, 2014.
- [31] E. A. Sisbot and R. Alami, “A human-aware manipulation planner,” *Robotics, IEEE Transactions on*, vol. 28, no. 5, pp. 1045–1057, 2012.
- [32] E. Ferreira, G. Milliez, F. Lefevre, and R. Alami, “Users belief awareness in reinforcement learning-based situated human-robot dialogue management,” in *IWSDS*, 2015.
- [33] G. Milliez, R. Lallement, M. Fiore, and R. Alami, “Using human knowledge awareness to adapt collaborative plan generation, explanation and monitoring,” in *ACM/IEEE International Conference on Human-Robot Interaction, HRI’16, New Zealand, March 7-10, 2016*.
- [34] P. R. Cohen and H. J. Levesque, “Teamwork,” *Nous*, pp. 487–512, 1991.
- [35] R. Verbrugge and L. Mol, “Learning to apply theory of mind,” *Journal of Logic, Language and Information*, vol. 17, no. 4, pp. 489–511, 2008.
- [36] M. E. Bratman, “Shared intention,” *Ethics*, pp. 97–113, 1993.