



Towards a system that allows robots to use commitments in joint action with humans

Ely Repiso, Guillaume Sarthou, Aurélie Clodic

► To cite this version:

Ely Repiso, Guillaume Sarthou, Aurélie Clodic. Towards a system that allows robots to use commitments in joint action with humans. IEEE RO-MAN 2023, Aug 2023, Busan (Corée), South Korea. <10.1109/RO-MAN57019.2023.10309590>. <hal-04206081>

HAL Id: hal-04206081

<https://laas.hal.science/hal-04206081v1>

Submitted on 15 Sep 2023

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.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Towards a system that allows robots to use commitments in joint action with humans

Ely Repiso¹

Guillaume Sarthou¹

Aur lie Clodic¹

Abstract—In collaborative tasks, expectations for achieving shared goals arise at all hierarchical plan levels, including plans, tasks, subtasks, and actions. However, these expectations also generate uncertainties for individuals executing the joint plan. If left unresolved, these uncertainties can impede successful task completion. Uncertainties may relate to the agents’ motivation to initiate, continue, or complete their plan (motivational uncertainty), the best way to execute their shared plan (instrumental uncertainty), and their knowledge of other agents and the environment (common ground uncertainty). These expectations can be either normative or descriptive, but only normative expectations trigger reactions from agents to resolve the aforementioned types of uncertainties. Thus, this paper introduces a theoretical model that enables a robot to consider all agents’ expectations and take actions that reduce the uncertainties associated with their shared plan. By doing so, we aim to enhance the likelihood of success in joint plans between robots and humans. To demonstrate the effectiveness of our theoretical commitment model, we have implemented a proof of concept for a client service use case in a food shop.

I. INTRODUCTION

To enable the effective use of robots by the general public, incorporating philosophical and psychological insights from human-human joint actions into the robot’s behavior is crucial. As human actions are influenced by experience, human interactions serve as the closest model for humanoid-social robots. Consequently, social behaviors, such as those related to commitment theories in shared plans, must be developed by robots. Commitments play a vital role in establishing and reinforcing expectations regarding beliefs, intentions, or behaviors among agents, thereby reducing various sources of uncertainty in joint action (Michael et al., 2015; Castro et al., 2019). By addressing the uncertainties underlying agents’ expectations, smoother and more efficient collaboration can be fostered, enhancing agents’ predictability and increasing the likelihood of achieving their shared goal.

To address these concerns, we propose the integration of a Commitment Manager into the robot’s behavior. By applying the commitment theory, we can directly influence human behavior towards achieving shared goals. This approach has several advantages, including maintaining human motivation, reducing instrumental uncertainties related to task performance, and enhancing understanding of the robot’s behavior. These benefits are particularly important for the general population, as interacting with robots can be

challenging due to unfamiliarity with their behaviors. Also, incorporating commitments into the robot’s behavior serves a dual purpose. Firstly, it helps bridge the gap between people’s expectations and the robot’s perceived capabilities by reducing uncertainties regarding the robot’s abilities. Secondly, it reduces uncertainties for the robot concerning unexpected human behaviors or knowledge about the human’s abilities to accomplish the shared plan.

Incorporating psychological and philosophical insights from human behavior into robotic behaviors is challenging, as these tasks must be adapted to the robot’s capabilities. For instance, the robot needs to determine the responsible agent for each part of the plan, identifying the task performer, and anticipating potential uncertainties for the other agent during task execution. Furthermore, the robot must take actions that are understandable to humans and capable of resolving uncertainties for all agents involved, whether they are humans or robots.

This paper presents a robot behavior framework incorporating commitments to address uncertainties that arise among agents during collaborative tasks due to their expectations. Firstly, we define the expectations that can emerge for any shared plan involving both agents. Secondly, we outline how the robot can detect each type of expectation related to the plan, task, or actions to be performed jointly. Thirdly, we propose the robot’s actions to resolve uncertainties for both agents. Lastly, we provide a proof of concept implementation within the context of a client service in a food shop. This implementation serves as a demonstration of how our theoretical commitment manager can effectively mitigate uncertainties for agents during the execution of a shared plan.

In the remainder of the paper, Sec. II presents the related work. In Sec. III, we introduce our theoretical commitment manager, which connects commitment theories with their implementation in a robotic agent. Next, in Sec. IV, we showcase a proof of concept to demonstrate that this theoretical model can be implemented in a robot. It illustrates how a robot can take actions to address both agents’ uncertainties in a collaborative plan. Finally, we include a discussion of the method’s limitation and generalization in Sec. V, as well as discussed conclusions and future work in Sec. VI.

II. STATE OF THE ART

A. Related works

In early works of HRI, we can find researchers that focus on the engagement as Sidner et al. [1] that found that people’s engagement was higher when the robot used social gestures during their interaction compared with the case without. In

This work has been supported by the Artificial Intelligence for Human-Robot Interaction (AI4HRI) project ANR-20-IADJ-0006.

¹The authors are from LAAS-CNRS, Universit  de Toulouse, CNRS, Mathematics and Computer Science, 7 Av. du Colonel Roche, Toulouse, France, 31400. {ely.repiso-polo@laas.fr, guillaume.sarthou@laas.fr, aurelie.clodic@laas.fr}

addition, Huang and Mutlu [2] presented an approach to enable robots to incorporate social behaviors that humans use in joint action tasks. They evaluate that the robot social behaviors increase their participants' engagement, and this engagement was highly dependent on gender. Furthermore, Garrell et al. [3] used the robot's social behaviors to engage people in a collaborative task that allows the robot to recognize human faces for future interactions.

Some works took inspiration from the joint intention theory of Bratman [4] and Cohen [5] to allow a robot to learn and perform a task collaboratively with a human partner. For example, Breazeal et al. [6], [7] include mechanisms to ensure both partners' commitment to fulfill the goals of their collaborative task. Strabala et al. [8] used the theories of Clark [9] to include social cues during handovers to maintain both partners committed to their joint task. Mioch et al. [10] is based on Singh's commitment model [11] to allow their participants to agree on which agent is committed to do a task for their joint plan. Vignolo et al. [12] explored the commitment of the agents to achieve a task related to high or low space-temporal coordination, where they found that high coordination was related to a higher degree of commitment. Finally, Chang et al. [13] presented an algorithm for Shared Cooperation that facilitates collaboration by promoting mutual responsiveness through three mechanisms: understanding the agents' intent, aligning their sub-plans, and providing assistants to the partner as needed.

B. Theoretical Foundation

The almost exclusive emphasis on agents' motivation in Human-Robot interaction research has resulted in a loss of critical information required for the success of shared plans. To address this issue, we turn our attention to the studies of Castro et al. [14] and Michael et al. [15], which offer a comprehensive perspective on all the expectations that may arise during collaborative interactions between agents. In addition to the theoretical ideas presented by Castro et al. [16], [17] and Belhassein et al. [18] on applying an extended theory of commitments in HRI, our work takes a further step by providing comprehensive and general definitions of how to implement a commitment manager in the robot's behavior to address all types of agents' expectations.

Commitments can be understood as "a triadic relation among two agents and an action, where one of the agents is obligated to perform the action as a result of having given an assurance to the other agent that she would do so, and of the other agent's having acknowledged that assurance under conditions of common knowledge" [19].

Two major theories co-exist regarding commitments in joint action.

1) *The Functionnal Approach* [15]: This approach emphasizes the major importance of using commitments in joint action because of its ability to produce reliable expectations. These expectations facilitate predictions and help reduce uncertainties that can have different causes.

- Motivation (are the participants still motivated/engaged to do the task?)

- Instrumentation (are the participants still able to handle the task, or does it exist an instrumental issue that prevent the execution?)
- Common ground (do the participants have all the necessary information to handle the task?)

2) *The Normative Approach* [14]: This approach also emphasizes the importance of expectations and describes two types.

- Descriptive expectations: expectations whose violation or frustration does not necessary triggers reactive attitudes;
- Normative expectations: expectations whose violation or frustration triggers reactive attitudes.

Our framework is build using jointly these two theoretical contributions.

C. Our contribution

Our contributions regarding the state-of-the-art are three-fold. First, we have analyzed all the expectations that can arise in a shared plan regarding both agents' points of view (the robot and the human). Second, we give ways for a robot to detect these expectations in any part of their shared plan. Third, we have developed ways for a robot to act to reduce the agents' uncertainties generated by these expectations. Finally, we have implemented a proof of concept to demonstrate that these theories can be applied in a real-life human-robot shared plan to reduce all the agents' uncertainties regarding their collaboration.

III. THEORETICAL COMMITMENT MANAGER

In this section, we propose a general description of a commitment manager who enables the robot to perform actions aimed at reducing uncertainties between agents during their collaborative efforts towards achieving a common goal. By reducing uncertainties, we enhance the likelihood of successfully accomplishing the shared plan. We consider the perspectives of both partners, the human and the robot. However, we can only control the robot and modify its behavior to consider the expectations of both partners during their interaction. This can be challenging, especially given the context or the general population's familiarity with robots. On one hand, the robot needs to manage the execution of its actions while also considering and managing human expectations about its behavior. On the other hand, the robot needs to monitor human actions and signal when its expectations regarding human actions are not met.

This theoretical model assumes that we have knowledge of the agents' shared plan and common goal. This required information includes the agents' roles, which determine the part of the plan to be performed by each agent; the hierarchy of tasks and actions of their shared plan, which specify their behavior; and the location and timing of plan execution, which constrains the possible shared plans to perform and when they collaborate. Next, our method associates the shared plans/tasks/actions of the agents with their expectations. Where these expectations give rise to three types of uncertainties: motivational, instrumental, and

common ground, as described in Sec. III-A. Finally, the method considers these uncertainties to perform concrete robot's actions that help to solve them, in Sec. III-B. An overview of the theoretical commitment manager is provided in Fig. 1.

A. Redefine the expectations to use them in a Robot

To enable the robot to effectively manage these expectations, we need a definition of those expectations that links the theory with how the robot can detect each type of expectation.

- Motivational uncertainty appears when there is an uncertainty about the other agent's level of engagement in the task. Then, this can be related to not knowing if the other agent started her action or to spend more time than expected to execute one action. In such situations, doubts may arise concerning the other agent's motivation to initiate, continue, or complete the joint task. Then, the robot considers that there is an uncertainty about the agent motivation in two situations: when the start of that action is not easily recognizable by the other agent, for example, an action of thinking; and when there is a noticeable discrepancy between the expected and actual execution times of an action.
- Instrumental uncertainty arises when there are discrepancies in the plans, tasks, or actions that the agents can choose or the objects they can utilize to achieve their shared goal. One approach to detect this type of uncertainty is by identifying multiple parallel plans or tasks that can be pursued, or different objects that can be utilized to accomplish the same objective. Whenever the robot encounters such a situation, it recognizes it as an instance of instrumental uncertainty.
- Common ground uncertainty arises when there are differences in the agent's knowledge about other agents and their environment. For example, in the case of the general population, people are not used to interact with robots. Therefore, they need to know the time that the robot dedicates to do an action or if the robot understands that it needs to start the following action of the plan. Furthermore, the robot should have mental models of people to reduce its uncertainties about them. Finally, the agents can have uncertainties about their environment, such as the objects they can find inside. Then, the robot can detect this type of uncertainty by identifying missing information in an agent's knowledge about others or their environment.

These definitions may evolve in the future to include subtypes within each type of uncertainty.

B. Robot actions to solve the agents' uncertainties about their expectations

To develop a comprehensive robot behavior capable of resolving agents' uncertainties, it is necessary to define general types of robot actions to address each specific type of uncertainty. As with the previous section, these definitions

can be redefined in the future to include additional actions for resolving subtypes of uncertainties.

1) *Motivational Uncertainties*: In this initial theoretical model, we simplify the robot's behavior to recognize the person's motivation. The assumption is made that the robot lacks a perception system to detect the person's presence or social cues. Instead, the robot relies on the person's response to its questions as an indication of motivation to continue. Then, the robot expects that the person is always there but can take more than the expected time to interact with it. If the person exceeds the expected time to answer a question, the robot may interpret it as a lack of motivation to continue with their shared plan. In such cases, it can only ask the person if she wants to continue with the joint action to solve the robot's uncertainty about the person's motivation. Then, the robot should finish their interaction if the person does not answer in a reasonable amount of time.

Currently, we have implemented robots that are always motivated to interact with people, as it is their primary purpose. However, people may still experience uncertainty if the robot temporarily pauses their interaction to attend to a more urgent task. In such cases, the robot needs to inform the person about the situation while simultaneously inquiring if they would like to cancel the interaction or proceed with their shared plan afterward.

2) *Instrumental Uncertainties*: We associate these uncertainties with the possibility of having parallel plans or tasks, where different actions can be taken to achieve the same goal. In such cases, the robot can autonomously select and execute a particular plan. However, if the person disagrees with the chosen plan, she has the option to leave the interaction or attempt to correct the robot during the middle of its execution. These actions may cause errors. Then, the robot might handle this uncertainty differently. A better way to handle this uncertainty from the starting point is to ask the person's preference regarding the available parallel plans. Humans also sometimes ask their partner how she prefers to perform a task to avoid this type of uncertainty.

3) *Common Ground Uncertainties*: The robot must be prepared to interact with a general population with limited knowledge about its behavior. To solve common ground uncertainties regarding the robot's behavior, the robot's actions need to be as transparent as possible to the person. The robot needs to inform when it starts any action and also about the expected time it will take to complete the action. This helps avoid uncertainty regarding the progress of the joint task because humans only know how long it takes a person to do that action, but that time can be different for the robot. By providing information about the execution time, the robot ensures that the person comprehends the time required for the action, which helps maintain their engagement in the joint task. In parallel with the robot's action execution, it should control the time it is using to do its action for two reasons. On the one hand, it should notify the person if the action is progressing well but exceeding the expected execution time. On the other hand, it should inform the person if the action is not going well to restart or abort the action.

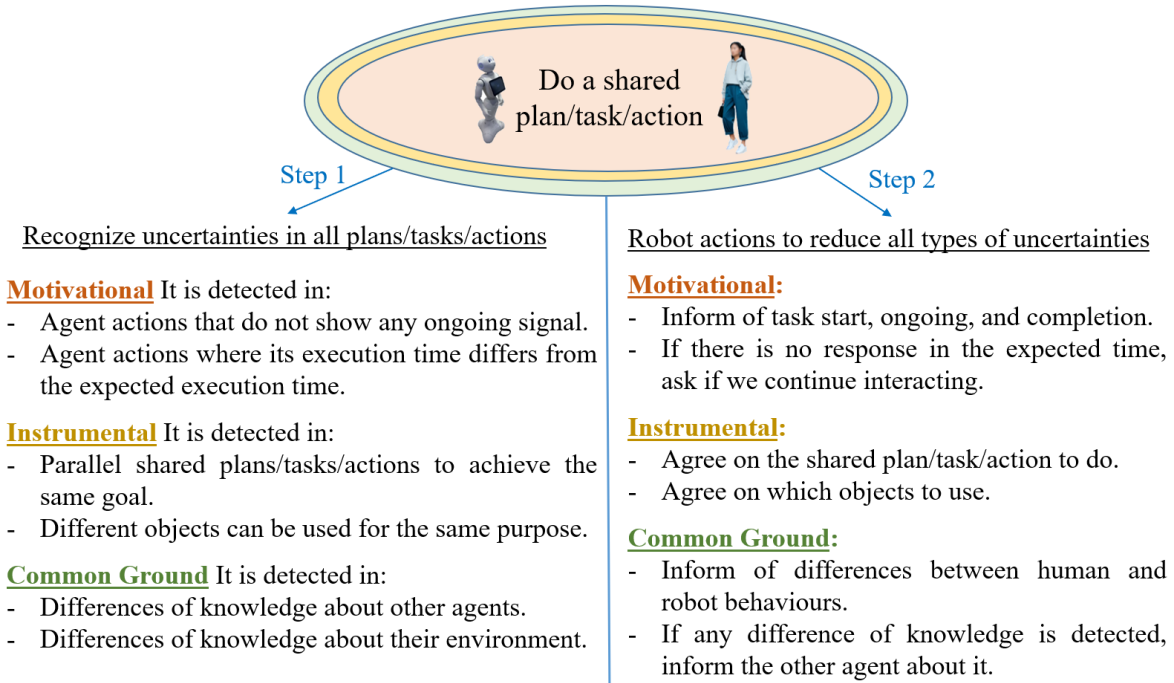


Fig. 1: Overview of the Theoretical Commitment Manager. *Left*: How the robot can recognize the uncertainties that can appear in any plan/task/action. *Right*: How the robot can act to reduce these types of uncertainties.

People may have uncertainties about the objects available in their environment, among other things. The robot can detect this uncertainty when it identifies a discrepancy between the information requested by the person and its database. To address these common ground uncertainties related to the agents' knowledge of their environment, the robot should inform the person about the mismatch in information between their respective databases and work towards resolving it (detailed examples can be found in Sec. IV).

IV. ROBOTIC SYSTEM TO DEAL WITH COMMITMENTS

Once we have developed the general theoretical commitment model for managing and resolving the agents' uncertainties, we have implemented our system that includes the commitment manager (Sec. IV-A) and a proof of concept to demonstrate that the theoretical model can be applied in practice in social robots (Sec. IV-C). These two are included in a concrete framework (Sec. IV-B), which determines the agents' roles, possible shared plans, the interaction location, and collaboration time. System available in LAAS-GitLab¹.

A. System implementation

We have implemented the system of Fig. 2 to enable the robot to perform Human-Customer support tasks within a shop setting. Such a system consists of several blocks:

- The planner and execution software allows the robot to perform all physical tasks and actions, such as navigating around the environment or pointing at products

¹https://gitlab.laas.fr/ai4hri4laas/pepperdemo/Lisp_commitment_manager/. This is the first version with many limitations, please, ask the main author to know if there are better versions in c++ or if there is any problem with the link.

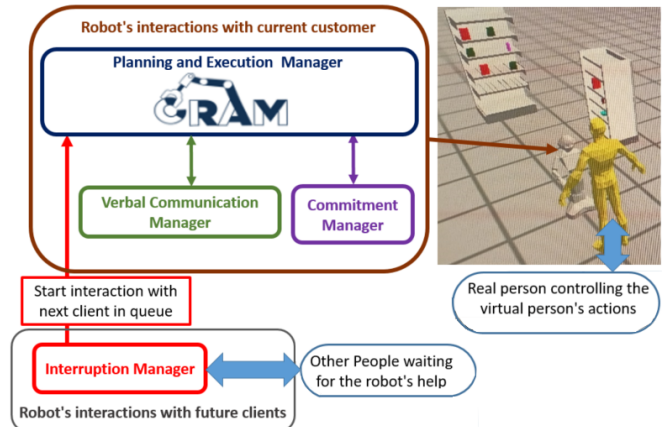


Fig. 2: System Structure that enables the robot to utilize commitments during the employee-client task.

on a shelf. For this purpose, we have used the CRAM² planning software [20].

- The verbal communication manager is responsible for all communications between the human and the robot to agree on what to do and to be informed of the status of the joint tasks, included in images of Sec. IV with the name Interaction Prompt.
- The commitment manager is the core of our system, responsible for handling all the robot actions related to resolve uncertainties of both agents during their collaborative interaction. This allows both agents to remain involved in their shared plan until they finish it. To solve the uncertainties concerning the agents' actions, we extract from each action the concrete types

²https://cram-system.org/tutorials/intermediate/pepper_shopping

of expectations that have important uncertainties to be solved for the good development of their shared plan, see Fig. 3, 4, 5, 6, 7 of Sec. IV.

- The interruption manager is in charge of building a queue of Human-Clients to be served by all employees. In this case, only our simulated robot. Then, this manager allows dealing with multiple client interactions. Then, if a client arrives at the shop when the robot is serving another client (busy), she can wait until it will be her turn to be served by the robot. Its behavior can be seen in Fig. 8 of Sec. IV, labeled as "prompt to serve people interruptions."

Notice that the blocks of verbal communication manager, commitment manager, and interruption manager are entirely developed for the current paper.


B. Framework

The use case involves a client-employee scenario within a simulated food shop. The client, a human using a computer, interacts with the system, while the employee is represented by a robot. In this environment and with these roles, we delimit their interactions to shared plans for client shopping services. It is a turn-take interaction where each agent needs the other agent's action to continue with their actions. If one agent does not continue with their shared plan, the other agent can only cancel their shared plan. However, the commitment manager implementation will be the same in the case of parallel tasks when the agents have uncertainties concerning the other agents' actions. Also, we solely address uncertainties related to normative expectations, as these trigger the robot's corrective actions to ensure both agents complete the joint plan. It is important to note that we simplify the agents' interaction to concentrate on implementing the Commitment Manager.

In this use case, the robot takes on the role of the employee and is responsible for executing the portion of the shared plan that pertains to client service. Specifically, there are two possible shared plans: Greetings and assisting human clients in locating products within the shop. These shared plans involve several sub-tasks, including recognizing the client, greeting the client, retrieving or searching in a database for product's location within the shop, providing verbal instructions or physical guidance to arrive to the product's location, inquiring if the customer requires further assistance, and expressing gratitude and bidding farewell if no additional help is needed.

In our use case, the person does the part of the shared plan expected for customers, where the client can freely select different options related to her role. Concretely, she can greet the robot-employee, decide which product type among all the available ones is looking for, decide if she wants verbal or physical guidance to locate the product, decide whether to search for more products or not, and decide to say goodbye to the robot to close their interaction.

1) Robot starts interaction with Person X



Commitment Manager Get Information:
[Common Ground]

Interaction Prompt:

Hello my name is Pepper.
I'm an employee of this food store
What is your name?


Client 1

=>If client answers go to 2)

=> If client takes more than 30 seconds to answer go to

Commitment Manager Action Execution [Motivational] in 3).

2) Robot asks which help the person needs



Commitment Manager Agree [Instrumental]

Interaction Prompt:

How may I help you **Client 1**?
Please use the number
corresponding to your request.

-----Select a number-----

1. **Help me find a product in the Food Store.** [Go to 3)]
2. I just wanted to say hi. [Go to 7): Finish Interaction]

=>If client takes more than 30 s to answer go to **Commitment Manager Action Execution [Motivational]** in 3).

Fig. 3: Implementation of the steps of greeting and ask about what type of help the person needs.

C. Proof of concept: Help multiple clients to find products

In our proof of concept, we have already implemented robotic actions to solve three types of uncertainties, specifically targeting the uncertainties that need to be resolved for the successful completion of the shared plan. Our shared plan is divided into X shared actions. First, the Robot and person greet and recognize each other (step 1 in Fig. 3-left). Second, they agree on a shared plan (step 2 in Fig. 3-right). Third, they agree on the product to search for (step 3 in Fig. 4). Fourth, the Robot should search for the product location (step 4 in Fig. 5). Five, they should agree on the way to perform the guidance task (step 5 in Fig. 6). Six, they should return to agree if they make another shared plan of the possible ones (step 6 and Fig. 7-left). Seven, they close their interaction if the client does not want to perform more shared plans, and the Robot continues to serve the next client in the shop's queue if needed (step 7 and Fig. 7-right). This last action is related to our Interruption Manager (Fig. 8) that allows to serve customers who arrive while the Robot is busy. If we did not have this manager, we would lose interactions with those other clients because they could not wait to be served by the Robot if they wanted it. In all interaction steps, the client is allowed to finish her shared plan with the Robot. In the figures of this example, we have three types of actions for the commitments. First, the actions that only get information or inform the other agent are included in the interaction prompt to allow a more fluent agent interaction. Second, when the agents must

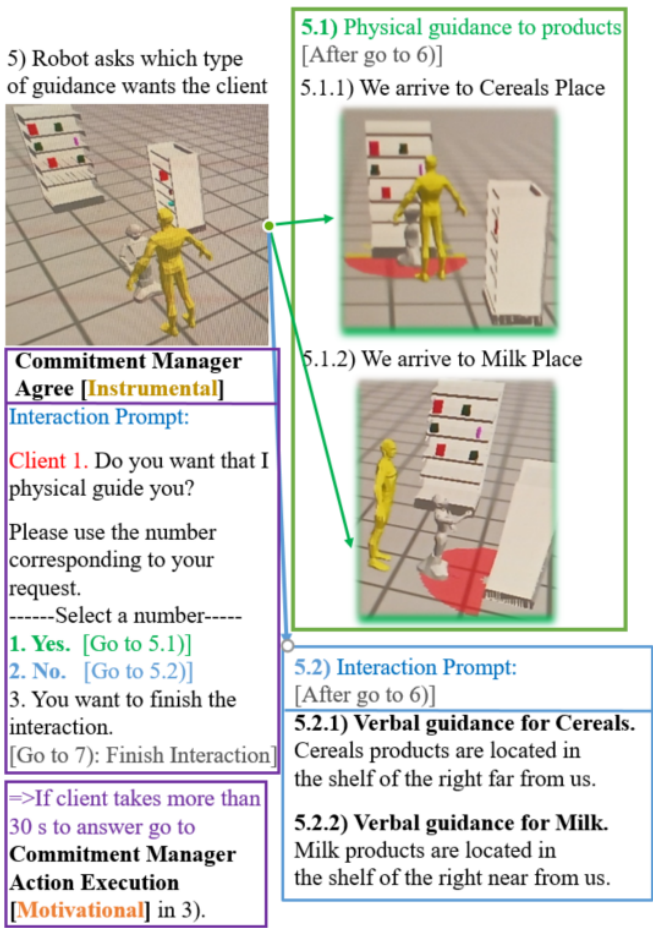


Fig. 6: Implementation of the step for the guide to the product physical location in shop.

assigned to the robot. Then, the duration required for the robot to execute a task is not commonly known by people. To address this, the robot needs to inform the person about the average execution time for actions that may take longer than an instant. In our example, this is relevant for the action of the robot searching for an object's position in its knowledge database, See the Interaction prompt in Fig. 5. This initial notification about the robot's execution time is purely informative and is included in the interaction prompt to ensure smoother agent interaction without requiring a response from the person. However, if the robot action takes longer than the initially informed time but is progressing as expected, the robot should provide a response to address any uncertainty the person may have about the task's performance. This response would reassure the person that the task is progressing well and only requires additional time. In this case, the robot may need a response from the person to agree to continue waiting or to conclude their shared plan. This response is executed by the Commitment Manager, as shown in the center of Fig. 5). Notice that to be able to do these types of actions, the robot should know the average time of execution for all the actions it can perform. Finally, as part of the interaction prompt for this common ground uncertainty, the robot should explicitly acknowledge its comprehension of the search item and initiate the search

- 6) Return to initial position. Robot asks if client needs more help.
- 7) **Finish Interaction with Client 1.** Starts interaction with **Client 2** in queue.

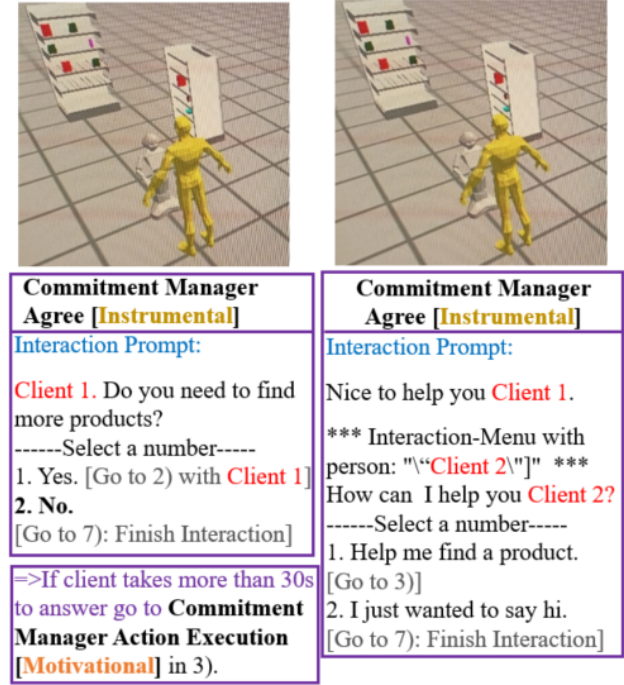


Fig. 7: Implementation of the step to find more products, finish their interaction and starts an interaction with the next client.

Prompt to serve people Interruptions or Multiple People interactions:
*** Menu to serve Interruptions***

Sorry, the robot is busy now.
Do you want to wait for the robot's help?

If yes, please write your name to include you in the waiting queue.

Client 2

Fig. 8: The image shows the implementation of the prompt to have a queue of clients waiting until the robot can help them. When the robot finishes the interaction with the current person, it looks at this queue. Then, if clients are waiting for its help, it starts the interaction with the next client in the queue until no clients are waiting for the robot's help, in Fig. 7.

process. This will alleviate any person's uncertainty about the robot's task status, who may not be able to perceive the robot's progress due to the non-physical nature of the task.

In addition, clients can have uncertainties about the shop, for example, which products can be found there. Therefore, if the person asks for a product that is not inside the robot's product-database, it should do an action to correct this person's uncertainty (see the bottom Common Ground Commitment Manager of Fig. 5). In our example, if the person requests the product car wheel, that is not available in our food shop, the robot's actions involve informing the person about the non-existence of this product and allowing the person to ask for another product if needed.

V. DISCUSSION

The general theory of this method, Sec. III, can be used in Human-Robot Shared Plans that can be performed in other type of services (hotels, airports, malls), or in the medical sector (hospitals, elderly, and children care), in the education sector (teachers' assistants or children's tutors), and industrial or agricultural settings.

A. Limitations

The method is general and the only limitation, maybe, are missing of some subcategories inside each uncertainty type. The limitations of our implementation are due to that it is only a proof of concept and the first implemented version: We only want to demonstrate the theoretical effectiveness of this approach in addressing uncertainties for robots. Then, it is not a fluent real-life HRI interaction, also it uses several text-prompts and not speech that will be more natural. Additionally, the method is embedded within the CRAM planner, which is implemented using Lisp, a programming language that is not widely used by researchers and the authors are not expert on it (due to the AI4HRI project). Consequently, the method is currently applicable only to the specific use case of the shop, also includes messages of CRAM planner in the information inside the interaction-prompt that it will be better to remove. Then, in the future, the method will be coded in a ROS-node using C++ to be able to obtain a general and fluent HRI, Sec. V-B.

B. Potential for Generalization

This method has significant potential for generalization. By separating the commitment manager from the planning process, it can be applied to any system that involves similar shared plans. The method will subscribe to the input information and create general behaviors to address uncertainties, outlined in Sec. III. It also will allow for customization through scripting, enabling adaptation to specific interaction scenarios that require tailored speech or actions for users who may not be experts in robotics. For example, the motivational uncertainty can refer to actions that the robot/human does not detect the start of the action, said in sec III. To solve human's uncertainties, the robot should always advertise verbally when it starts actions. To solve robot's uncertainties, it should recognize the start of a human action. If it can not recognize it. It should ask the human to obtain a confirmation. For further details about inputs and output actions for other uncertainties, we refer the reader to Sec. III.

VI. CONCLUSIONS AND FUTURE WORK

We have presented a theoretical model and a proof of concept system that applies commitment theories to address uncertainties experienced by both agents in a joint action. Our robot performs actions to effectively resolve uncertainties related to motivation, instrumentation, common ground. Using our proof of concept, we have demonstrated that this theoretical model can be used in a robotic system to solve the agents' uncertainties. Additionally, we have implemented an Interruption Manager to handle interruptions from other

shop clients, enabling multiple interactions with the same robot in our system.

Our future work will include enlarging the theoretical commitment manager and its implementation with new ways to solve agents' uncertainties, as well as translating it to C++ and ROS to be able to implement a more realistic HRI.

REFERENCES

- [1] C. L. Sidner, C. Lee, C. D. Kidd, N. Lesh, and C. Rich, "Explorations in engagement for humans and robots," *Artificial Intelligence*, vol. 166, no. 1-2, pp. 140–164, 2005.
- [2] C.-M. Huang and B. Mutlu, "The repertoire of robot behavior: Enabling robots to achieve interaction goals through social behavior," *Journal of Human-Robot Interaction*, vol. 2, no. 2, pp. 80–102, 2013.
- [3] A. Garrell, M. Villamizar, F. Moreno-Noguer, and A. Sanfeliu, "Teaching robot's proactive behavior using human assistance," *International Journal of Social Robotics*, vol. 9, no. 2, pp. 231–249, 2017.
- [4] M. E. Bratman, *Shared agency: A planning theory of acting together*. Oxford University Press, 2013.
- [5] P. R. Cohen and H. J. Levesque, "Teamwork," *Nous*, vol. 25, no. 4, pp. 487–512, 1991.
- [6] C. Breazeal, G. Hoffman, and A. Lockerd, "Teaching and working with robots as a collaboration," in *AAMAS*, vol. 4, 2004, pp. 1030–1037.
- [7] C. Breazeal, J. Gray, G. Hoffman, and M. Berlin, "Social robots: Beyond tools to partners," in *RO-MAN 2004. 13th IEEE International Workshop on Robot and Human Interactive Communication (IEEE Catalog No. 04TH8759)*. IEEE, 2004, pp. 551–556.
- [8] K. Strabala, M. K. Lee, A. Dragan, J. Forlizzi, S. S. Srinivasa, M. Cakmak, and V. Micelli, "Toward seamless human-robot handovers," *Journal of Human-Robot Interaction*, vol. 2, no. 1, pp. 112–132, 2013.
- [9] H. Clark, "Social actions, social commitments in (eds) nj enfield, stephen c," *Levinson, . Roots of human sociality: culture, cognition and interaction*. New York, NY: Berg, 2006.
- [10] T. Mioch, M. M. Peeters, and M. A. Neerincx, "Improving adaptive human-robot cooperation through work agreements," in *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2018, pp. 1105–1110.
- [11] M. P. Singh, "An ontology for commitments in multiagent systems," *Artificial intelligence and law*, vol. 7, no. 1, pp. 97–113, 1999.
- [12] A. Vignolo, A. Sciutti, F. Rea, and J. Michael, "Spatiotemporal coordination supports a sense of commitment in human-robot interaction," in *International Conference on Social Robotics*. Springer, 2019, pp. 34–43.
- [13] M. L. Chang and A. Thomaz, "Valuable robotic teammates: Algorithms that reason about the multiple dimensions of human-robot teamwork," in *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 2021, pp. 580–582.
- [14] V. F. Castro and E. Pacherie, "The credibility problem in human robot interaction," in *27th Conference of the European Society for Philosophy and Psychology (ESPP 2019)*, 2019.
- [15] J. Michael and E. Pacherie, "On commitments and other uncertainty reduction tools in joint action," *Journal of Social Ontology*, vol. 1, no. 1, pp. 89–120, 2015.
- [16] V. F. Castro, A. Clodic, R. Alami, and E. Pacherie, "Commitments in human-robot interaction," *arXiv preprint arXiv:1909.06561*, 2019.
- [17] V. Fernández Castro and E. Pacherie, "Joint actions, commitments and the need to belong," *Synthese*, vol. 198, no. 8, pp. 7597–7626, Jan. 2020. [Online]. Available: <https://doi.org/10.1007/s11229-020-02535-0>
- [18] B. Kathleen, F.-C. Víctor, M. Amandine, C. Aurélie, P. Elisabeth, G. Michèle, A. Rachid, and C. Hélène, "Addressing joint action challenges in hri: Insights from psychology and philosophy," *Acta Psychologica*, vol. 222, p. 103476, 2022.
- [19] J. Michael and A. Salice, "The sense of commitment in human-robot interaction," *International Journal of Social Robotics*, vol. 9, no. 5, pp. 755–763, Sept. 2016. [Online]. Available: <https://doi.org/10.1007/s12369-016-0376-5>
- [20] M. Beetz, L. Mösenlechner, and M. Tenorth, "Cram—a cognitive robot abstract machine for everyday manipulation in human environments," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 1012–1017.