

# Towards a Hybrid LLM/Model-Based Architecture for Robot Coaching: An Instance of Human-Machine Collaboration

Luigi Gargioni<sup>1,2,\*</sup>, Rachid Alami<sup>2</sup> and Daniela Fogli<sup>1</sup>

<sup>1</sup>University of Brescia - Department of Information Engineering, Via Branze 38, Brescia, 25123, Italy

<sup>2</sup>LAAS-CNRS, 7 Av. du Colonel Roche, Toulouse, 31400, France

## Abstract

Human-Robot Collaboration (HRC) presents significant challenges in assessing situations correctly, adapting robotic behavior to human intentions, ensuring explainability, pertinence, and acceptability, and managing uncertainty. Traditional model-based approaches offer reliability but struggle with human unpredictability and approximate humans with specific models that do not consider all the possible situations. At the same time, probabilistic methods like Large Language Models (LLMs) provide adaptability but lack deterministic guarantees. This paper proposes a hybrid architecture that integrates structured techniques with the flexibility of LLMs to enhance robot coaching in dynamic environments. By bridging deterministic and probabilistic techniques, our architecture aims to advance HRC towards safer, more transparent, flexible, and adaptive interactions. The paper provides a detailed description of the framework's specifications; however, it should be noted that it has not yet been fully implemented.

## Keywords

Human-Robot Collaboration, Robot Coaching, Large Language Model, Model-Based, Hybrid Architecture

## 1. Introduction

Human-Robot Collaboration (HRC) is a multidisciplinary research area that studies and designs interactions between humans and robots. This field encompasses principles from artificial intelligence, robotics, cognitive science, psychology, and human factors engineering to create systems that enable natural, effective, and intuitive collaboration between humans and autonomous machines. As robots become increasingly integrated into everyday life, from industrial automation [1, 2] to personal assistance and healthcare [3, 4], ensuring effective collaboration between humans and robots is crucial. However, HRC presents significant challenges beyond conventional automation. It involves robots and humans working together to achieve common goals, requiring advanced reasoning, planning, and adaptability mechanisms to ensure seamless cooperation and effective task completion.

One of the primary challenges in HRC is enabling machines to reason about human beliefs and intentions and adapt their behavior accordingly. Unlike traditional automation, where predefined rules govern robot actions, effective HRC demands that robots infer and respond dynamically to human actions, preferences, and situational changes. Another critical aspect is the ability of robots to plan and coordinate their actions with humans in a way that ensures legibility, explainability, and acceptability.

Several contributions in the literature address the challenges of creating more capable and adaptive robots by exploring control architectures and cognitive-interactive systems[5, 6, 7]. These systems are designed to integrate decisional and functional components into a unified structure that efficiently manages the flow of information. The decisional components typically involve higher-level processes such as situation assessment and planning, which are crucial for the robot's ability to make decisions

---

*AutomationXP25: Hybrid Automation Experiences, April 27, 2025, Yokohama, Japan. In conjunction with ACM CHI'25.*

\*Corresponding author.

\*\*This paper has been partially supported by Agence Nationale de la Recherche (French National Research Agency) for France 2030, reference INNOVCARE ANR-23-PA VH-0005.

✉ luigi.gargioni@unibs.it (L. Gargioni); rachid.alami@laas.fr (R. Alami); daniela.fogli@unibs.it (D. Fogli)

id 0000-0003-4354-916X (L. Gargioni); 0000-0002-9558-8163 (R. Alami); 0000-0003-1479-2240 (D. Fogli)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

based on the presence of the human and the surrounding environment [8, 9]. On the other hand, functional components, including perception and action, allow the robot to interact with the physical world and respond to its surroundings in real-time [10, 11]. The complexity lies in organizing these various components into a coherent architecture that can effectively handle the diverse needs of the system. A limitation of these architectures is that they rely on models of human beliefs, intentions, and preferences to guide interaction, even though human behavior is inherently unpredictable and difficult to model with precision. This fundamental challenge restricts the system’s ability to fully anticipate and adapt to human actions, introducing a layer of uncertainty that remains difficult to overcome.

Traditional control architectures often rely on static models of human intentions and behaviors, which can limit their capacity to adapt to rapidly changing contexts. In contrast, cognitive frameworks incorporating belief management and Theory of Mind (ToM) offer a more flexible approach by allowing robots to model and interpret the mental states of human agents [12, 13]. Indeed, it is important to endow the robot with the ability to permanently estimate the humans beliefs, to reason about them, use them to predict human decisions and actions, and to act accordingly [14, 15, 16].

In addition to control architectures, several planning schemes have been proposed to facilitate the synthesis of action plans for achieving collaborative tasks between humans and machines [17, 18, 19]. These planning systems are designed to allow robots to generate effective action plans while working in tandem with human counterparts. The key challenge in these systems is ensuring that the generated plans account for both the robot’s capabilities and the human’s behaviors and expectations. Most planning systems, including adaptation or learning mechanisms [20], are fundamentally model-based. They rely on pre-defined models of the human, the task at hand, and the interaction between the human and the machine. These models are crucial for predicting the human’s actions, understanding their intentions, and anticipating the evolution of the task. However, due to the inherent unpredictability of human behavior, uncertainty plays a significant role in these systems. A variety of sophisticated methods have been proposed to address such uncertainty. Markov Decision Processes (MDPs) [21] and Partially Observable Markov Decision Processes (POMDPs) [22] are widely used to model uncertainty in observation and estimate the effects of actions. MDPs are helpful when the outcomes of actions are uncertain, and the decision-making process needs to account for immediate and future rewards. POMDPs extend this framework by incorporating scenarios where the robot has incomplete information about the environment or the human’s state, often in human-robot collaborations. Furthermore, epistemic planning models the belief divergence between the robot and the human [23]. By considering the knowledge each agent has about the other’s beliefs and intentions, epistemic planning helps manage the coordination, ensuring more seamless collaboration between the human and the machine. Other contributions are based on non-deterministic planning schemes where the human decisions and actions are dealt with as contingent [24]. Also, ethical planning is an interesting approach since it ensures the production of a plan that satisfies ethical properties [25].

Despite these advancements, existing approaches struggle with the intrinsic difficulty of mind reading and the impossibility of representing human beliefs and decisional processes, making it difficult for robots to interpret complex or nuanced human intentions. The efficacy of these methodologies is contingent upon an accurate model of the environment, the task, and possible ways of accomplishing the task. However, developing a robust model of the human participant remains a significant challenge due to the inherent unpredictability of human behavior.

To address these limitations, integrating Large Language Models (LLMs) offers a promising direction. These models can significantly expand the range of situations that robots can handle, allowing for more flexible specifications concerning explainability and acceptability. By leveraging common sense reasoning and contextual understanding, LLMs can help robots interpret human inputs and the surrounding environment, infer unspoken intentions, and generate adaptive responses that align with human expectations. However, a fully probabilistic approach remains insufficient for ensuring reliability and safety in HRC. It is crucial to design an architecture where the plan has to be validated and key aspects of the system, such as safety constraints and ethical considerations, are explicitly defined and secured. Given the strengths and limitations of traditional architectures and planning approaches and the emerging capabilities and challenges associated with LLMs, this paper explores

a hybrid framework that integrates the structured reliability of deterministic models with the adaptability, common sense, and contextual reasoning offered by LLMs. In this context, *hybrid* refers to the combination of deterministic techniques, which provide formal guarantees and rule-based precision, with non-deterministic, probabilistic methods that enable learning-driven adaptability and robustness in uncertain environments.

In this work, we explore the pertinence of a hybrid architecture based on LLMs and deterministic approaches to enhance task execution, adaptability, and user interaction in HRC while ensuring that the decisions and actions of the machine are safe and pertinent.

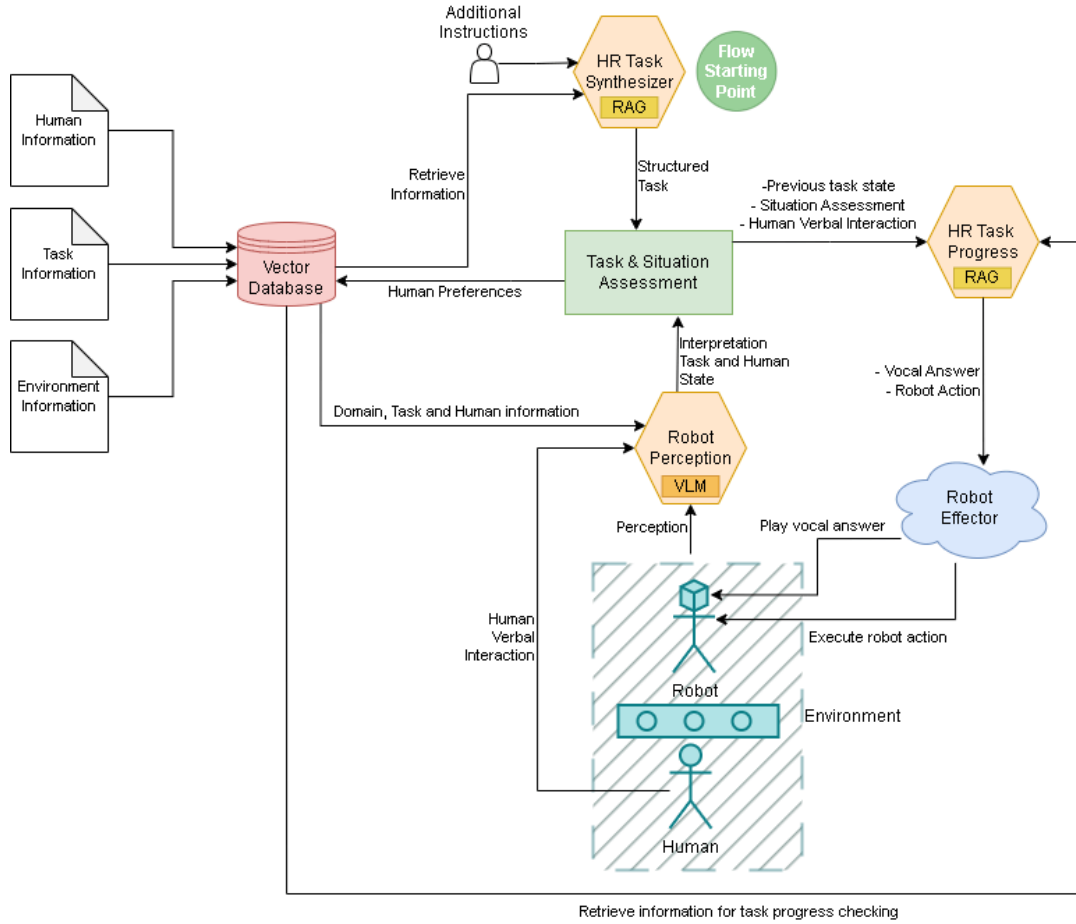
## 2. A Hybrid Approach

HRC presents various challenges, ranging from deterministic task execution to highly dynamic and ambiguous human behaviors. In this context, the goal is defined and shared with the human (i.e., Human-Robot Joint Action [26]), ensuring that both the robot and the human share a common understanding of the desired outcome and can contribute to reaching it. This shared goal guides the robot’s actions and decisions, aligning its behavior with human expectations. Furthermore, a task refers to a specific activity that the robot and the human have to perform to achieve a specific goal, which can vary in complexity from simple, predefined operations to more adaptive and interactive behaviors requiring real-time situation assessment. The decision process is essential because, even if the task is well specified, there are different ways to reach the objective. However, certain assumptions are present in this context. The human and the robot are inherently different, with the robot’s role assisting the human to accomplish the task effectively (i.e., *robot coaching*). Furthermore, human is regarded as an unpredictable entity, yet it is presumed that they will collaborate with the robot and will not deceive it with malicious intent. Both the robot and the human participate in the task, employing multimodal verbal interaction to support the successful achievement of the objective. A hybrid approach can integrate these deterministic and probabilistic methodologies, ensuring robustness in execution while maintaining adaptability in complex and evolving situations. The architecture in Figure 1 is designed to follow the concepts described in the previous section.

It is organized into interconnected modules that work collaboratively to ensure seamless interaction, reliable task execution, and robust adaptability to evolving scenarios. It is interesting to point out in which modules a deterministic approach is used (i.e., green rectangle and blue cloud) and in which an LLM is used (i.e., yellow hexagons) and to explain the reasons for this in each case.

The subsequent section delineates the fundamental modules of the architecture, with each module playing a pivotal role in the processing and execution of designated tasks. These modules transform unstructured information into actionable outputs, ensuring seamless interaction between the user, the environment, and the robotic system.

- **Vector Database:** The vector database is the central knowledge repository, storing embeddings derived from unstructured information (e.g., information about the human, the task, and the environment). It can also be enriched by knowledge gained from interactions (e.g., human preferences emerged from the interaction). The database stores relevant embeddings to provide information for defining the task and supporting the task progress checking.
- **Human-Robot Task Synthesizer:** This module leverages an LLM and Retrieval-Augmented Generation (RAG) to translate unstructured information from the Vector Database into structured tasks. This step is critical to move from unstructured information, such as natural language, to a structured task plan (i.e., JSON format) that can be used programmatically, either by deterministic or probabilistic methods. This is also the starting point of the architecture workflow.
- **Task and Situation Assessment:** The Task and Situation Assessment module is the high-level control unit, parsing the structured task, interpreting the task and human state, and human verbal interaction. Thanks to a rule-based approach, it manages the workflow and the execution of the other modules. For example, it is responsible for receiving information from the environment and the human and passing it to the Human-Robot Task Progress module to check whether the status



**Figure 1:** The proposed Hybrid Architecture

of the task has been updated or to check whether the task has been completed and update the vector database with new knowledge (e.g., human preferences and behavior) from the interaction with the human and the progress of the task. This module updates the state of the world, the task, and the beliefs of the human and their preferences, behavior, and goals.

- **Human-Robot Task Progress:** This module leverages an LLM to monitor task execution and ensure that the task's progression aligns with predefined objectives. This module aims to update the plan status according to the information received and what to do next. The flexibility of the architecture stems from the absence of a predefined algorithm governing task evolution, thereby avoiding excessive rigidity and enhancing adaptability to dynamic conditions. The Human-Robot Task Progress is also responsible for checking the validity of the plan through *Progress Checking* function. This function employs RAG to retrieve complementary information or verify pre-declared ones from the Vector Database, thereby ensuring the validity and compliance of certain Human-Robot Task Progress plan properties. This iterative loop ensures task reliability and completion. The Progress Checking function is very important to provide robustness to errors that may appear from the Human-Robot Task Progress module.
- **Robot Perception:** The Robot Perception module exploits an LLM, specifically a Visual Language Model (VLM), to interpret images, retrieve situational data from the environment, including task and human state, and provide continuous feedback to the situation assessment process. This model is characterized by its ability to draw on general knowledge and apply common sense, enabling it to consistently provide helpful information that facilitates the progression of the task.
- **Robot Effector:** The Robot Effector module executes actions and communicates with the human through vocal responses and physical actions. This module ensures real-world applicability and contextual responsiveness. As far as the execution of voice commands for the user is concerned,

the robot’s speakers will simply be used. Then, when it comes to performing actions in the environment, the robot will have to rely on an integrated motion planner.

The system initialization begins with collecting unstructured information, which is reported in natural language, converted into embeddings, and added to the vector database. At the beginning of the interaction workflow, the Human-Robot Task Synthesizer extracts and processes the relevant information to create a structured task represented in a JSON structure. Thanks to the other modules, the Task and Situation Assessment module integrates this task with situational data, previous task states, and interaction history to develop an updated task plan. As the task progresses, the system monitors its execution through the Human-Robot Task Progress module and validates outcomes using the Progress Checking function. The robot executes the corresponding actions if task outputs meet predefined success criteria. If discrepancies arise, the updated task returns in the Human-Robot Task Progress, creating a feedback loop for iterative refinement. Finally, the Robot Perception module continuously interprets task and human states, ensuring the system remains adaptable and contextually aware. Combined with the Robot Effector module, this enables the robot to interact dynamically with its environment while maintaining safety and user-centric functionality.

To better illustrate this workflow, consider an assistive robot in a healthcare setting as a scenario where a robot must assist a patient in following a prescribed therapy. We conducted preliminary testing on selected steps of the proposed scenario using a separate single LLM (*Llama 3.3 70B*) that was not yet integrated into the system architecture. This was done to assess the quality of potential outputs and interactions. Some examples of generated outputs and interactions are provided in the scenario below. The process unfolds as follows:

1. The daily therapy of the patient is defined (in natural language), and a caregiver (e.g., a doctor) provides additional instructions in natural language on how to follow it, such as *"It is morning before breakfast, human and robot are present in the room. Assist and ensure that the patient takes their daily medication as specified in the prescription."* Optionally, the doctor could add some additional specifications for the day. For example: *"Drink more water and avoid sugar today," "Skip Paracetamol today,"* or *"Do physical exercise today."*
2. These inputs are processed into structured information using embeddings stored in the vector database, where they are combined with existing information about the patient and the context, which is also stored in the vector database.
3. The Human-Robot Task Synthesizer module exploits relevant details (e.g., medication type, timing constraints, assumptions notes, etc.) and formulates a structured task (i.e., JSON structure).
4. The robot, acting as both a supportive and authoritative assistant, monitors task execution through the Task and Situation Assessment and Human-Robot Task Progress modules. It dynamically adapts to the human’s latitude and emotional state, guiding them progressively toward task completion. If the patient correctly takes the medication, the system validates the action and updates the task status. If a discrepancy is detected, such as the patient refusing or missing a dose, the robot assesses their commitment level and responds accordingly. It may initiate a gentle reminder (e.g., *"Taking it now will help you stay on track with your treatment, and it will make you feel better. Please take this pill."*) or escalate to a more authoritative intervention, such as alerting a caregiver (e.g., *"The patient refuses to take their medication at the scheduled time. Please check it with them."*). This adaptive strategy, informed by past interactions, ensures effective assistance while respecting the patient’s autonomy.
5. The Robot Perception module continuously interprets task and human states, ensuring real-time adaptability. For instance, it reports whether the patient has taken medicine (e.g., *"The patient took the red pill."*) or whether the patient is sleeping and can’t reply or do anything.
6. The Robot Effector module executes the corresponding physical actions, such as referring the medication in an accessible manner or guiding the patient through the assumption process (e.g., *"This medicine should be taken with water just before lunch."*).

This scenario highlights how deterministic rule-based mechanisms enforce critical constraints (e.g., medication timing) while LLMs enhance interaction quality, contextual understanding, and user adaptation. By structuring the system into modules, the architecture can balance safety, adaptability, and user-centric interaction, aligning with the overarching goal of advancing HRC research.

### 3. Discussion and Conclusions

The proposed hybrid architecture addresses key challenges in HRC by integrating deterministic methods with probabilistic approaches. This balance ensures the safety and reliability of task execution and the flexibility and adaptability needed to operate in dynamic human-centered environments.

The rationale behind this architecture is the following:

- Some aspects of the task specification and decisions for its achievement can be (and are often) formalized (e.g., a doctor’s written prescription) and efficiently handled with model-based algorithms.
- The actions and decisions of the patient are contingent on the machine and highly unpredictable.
- The patient’s beliefs and potential behavior are not precisely known; information about them can only be obtained through observation and verbal interaction in natural language.
- Interaction often requires non-structured or non-predefined verbal communication.
- Other aspects of the task and its human performance can only be specified in natural language.
- Additionally, the specification of the machine’s desired behavior and the criteria defining an acceptable machine response are also verbally defined. They can only be refined incrementally through verbal interaction and observation of human behavior.

Based on this, we have structured a *hybrid* architecture that identifies several decisional and functional processes involved in task performance. It determines when and how model-based algorithms and representations are better suited and when relying on LLM abilities to assess, decide, or predict is more pertinent. A key aspect of the presented system lies in its modular structure, which allows specialized modules to handle different aspects of HRC. Deterministic modules, such as the Task and Situation Assessment module, provide a robust foundation to ensure compliance and predictable behaviors. Simultaneously, integrating advanced LLMs enhances the robot’s ability to interpret human behavior and environmental context.

Despite its advantages, the proposed architecture faces certain limitations. While improving adaptability, the reliance on LLMs introduces challenges related to computational resource demands and potential biases and errors in model outputs. Ensuring real-time performance, error robustness, and addressing ethical considerations, such as fairness and transparency in decision-making, will require further optimization and rigorous testing.

Future developments will focus on further specifying the duties of the different LLMs and determining whether additional models are needed to divide each task step better. Another area of future work is optimizing the Progress Checking function and knowledge management. Enhancing these two aspects will strengthen the approach’s fault-tolerant adaptability and scalability. Finally, to ascertain the viability of the proposed architecture, a real-world scenario must be conducted and subsequently evaluated by users. To establish a reference point, a comparative study will be carried out with systems based on planning and strictly deterministic methods. Ablation experiments will also be performed to assess the importance of each module in the architecture.

As robots become integral to human-centered environments, advancing HRC systems with hybrid approaches will be crucial to ensure seamless and effective integration. The presented approach aims to contribute to the broader goal of creating intelligent, adaptable, and user-centric robotic systems, paving the way for safer and more efficient human-robot collaborations.



## Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT and Grammarly in order to: grammar and spelling check, paraphrase, and reword. After using these services, the authors reviewed and edited the content as needed, thus, they take full responsibility for the publication's content.

## References

- [1] V. Villani, F. Pini, F. Leali, C. Secchi, Survey on human–robot collaboration in industrial settings: Safety, intuitive interfaces and applications, *Mechatronics* 55 (2018) 248–266. doi:10.1016/j.mechatronics.2018.02.009.
- [2] J. E. Michaelis, A. Siebert-Evenstone, D. W. Shaffer, B. Mutlu, Collaborative or simply uncaged? understanding human-cobot interactions in automation, in: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 1–12. doi:10.1145/3313831.3376547.
- [3] W. Weiss, F. Fuhrmann, H. Zeiner, R. Unterberger, Towards an architecture for collaborative human robot interaction in physiotherapeutic applications, in: *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 2017, pp. 319–320.
- [4] L. Gargioni, D. Fogli, P. Baroni, Preparation of personalized medicines through collaborative robots: A hybrid approach to the end-user development of robot programs, *ACM Journal on Responsible Computing* (2025).
- [5] J. G. Trafton, L. M. Hiatt, A. M. Harrison, F. P. Tamborello, S. S. Khemlani, A. C. Schultz, Act-r/e: An embodied cognitive architecture for human-robot interaction, *Journal of Human-Robot Interaction* 2 (2013) 30–55.
- [6] P. E. Baxter, J. de Greeff, T. Belpaeme, Cognitive architecture for human–robot interaction: Towards behavioural alignment, *Biologically Inspired Cognitive Architectures* 6 (2013) 30–39.
- [7] A. Umbrico, R. De Benedictis, F. Fracasso, A. Cesta, A. Orlandini, G. Cortellessa, A mind-inspired architecture for adaptive hri, *International Journal of Social Robotics* 15 (2023) 371–391.
- [8] S. Lemaignan, M. Warnier, E. A. Sisbot, A. Clodic, R. Alami, Artificial cognition for social human–robot interaction: An implementation, *Artificial Intelligence* 247 (2017) 45–69.
- [9] K. Darvish, E. Simetti, F. Mastrogiovanni, G. Casalino, A hierarchical architecture for human–robot cooperation processes, *IEEE Transactions on Robotics* 37 (2020) 567–586.
- [10] J. Páez, E. González, Human-robot scaffolding: An architecture to foster problem-solving skills, *ACM Transactions on Human-Robot Interaction (THRI)* 11 (2022) 1–17.
- [11] P. Foggia, A. Greco, A. Roberto, A. Saggese, M. Vento, A social robot architecture for personalized real-time human-robot interaction, *IEEE Internet of Things Journal* (2023).
- [12] J. G. Trafton, N. L. Cassimatis, M. D. Bugajska, D. P. Brock, F. E. Mintz, A. C. Schultz, Enabling effective human-robot interaction using perspective-taking in robots, *Systems, Man and Cybernetics* 35 (2005) 460–470.
- [13] L. M. Hiatt, A. M. Harrison, J. G. Trafton, Accommodating human variability in human-robot teams through theory of mind, in: *Int. Joint Conf. on Artificial Intelligence*, volume 22, 2011, p. 2066.
- [14] S. Devin, R. Alami, An Implemented Theory of Mind to Improve Human-Robot Shared Plans Execution, in: *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, The Eleventh ACM/IEEE International Conference on Human Robot Interaction, Christchurch, New Zealand, 2016, pp. 319–326. doi:10.1109/HRI.2016.7451768.
- [15] M. Romeo, P. E. McKenna, D. A. Robb, G. Rajendran, B. Nasset, A. Cangelosi, H. Hastie, Exploring theory of mind for human-robot collaboration, in: *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, IEEE, 2022, pp. 461–468.
- [16] A. Favier, S. Shekhar, R. Alami, Models and Algorithms for Human-Aware Task Planning with Integrated Theory of Mind, in: *IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2023. doi:10.1109/RO-MAN57019.2023.10309437.

- [17] G. Hoffman, C. Breazeal, Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team, in: *Proceedings of the ACM/IEEE international conference on Human-robot interaction*, 2007, pp. 1–8.
- [18] F. Pecora, M. Cirillo, F. Dell’Osa, J. Ullberg, A. Saffiotti, A constraint-based approach for proactive, context-aware human support, *Journal of Ambient Intelligence and Smart Environments* 4 (2012) 347–367.
- [19] M. C. Gombolay, R. A. Gutierrez, S. G. Clarke, G. F. Sturla, J. A. Shah, Decision-making authority, team efficiency and human worker satisfaction in mixed human–robot teams, *Autonomous Robots* 39 (2015) 293–312.
- [20] K. Ramachandruni, C. Kent, S. Chernova, Uhtp: A user-aware hierarchical task planning framework for communication-free, mutually-adaptive human-robot collaboration, *ACM Transactions on Human-Robot Interaction* 13 (2024) 1–27.
- [21] B. Hayes, J. A. Shah, Improving robot controller transparency through autonomous policy explanation, in: *Proceedings of the 2017 ACM/IEEE international conference on human-robot interaction*, 2017, pp. 303–312.
- [22] V. V. Unhelkar, S. Li, J. A. Shah, Decision-making for bidirectional communication in sequential human-robot collaborative tasks, in: *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 2020, pp. 329–341.
- [23] G. Buisan, A. Favier, A. Mayima, R. Alami, Hatp/ehda: A robot task planner anticipating and eliciting human decisions and actions, in: *2022 International Conference on Robotics and Automation (ICRA)*, IEEE, 2022, pp. 2818–2824.
- [24] A. Favier, R. Alami, A model of concurrent and compliant human-robot joint action to plan and supervise collaborative robot actions, in: *Advances in Cognitive Systems (ACS)*, 2024, pp. 1–16.
- [25] T. Parker, U. Grandi, E. Lorini, A. Clodic, R. Alami, Ethical planning with multiple temporal values, in: *Social Robots in Social Institutions*, IOS Press, 2023, pp. 435–444.
- [26] A. Clodic, E. Pacherie, R. Alami, R. Chatila, Key elements for human-robot joint action, *Sociality and normativity for robots: philosophical inquiries into human-robot interactions* (2017) 159–177.