

Negotiation of Assignment Plans in Human-Robot Team Task Scheduling

Llum Fuster-Palà¹, Marc Dalmaso^{1,2}, Artur Aubach-Altes¹, Silvia Izquierdo-Badiola^{2,3},
Alberto Sanfeliu^{1,2} and Anaís Garrell^{1,2}

Abstract—In recent years, considerable attention has been given to improving human-robot collaboration. Despite advances in robotic capabilities and interaction techniques, achieving a fair distribution of tasks remains challenging due to the dynamic nature of human preferences and situational constraints. This paper presents a novel negotiation framework that enables robots to effectively communicate with humans to facilitate fair and adaptive task allocation. Our approach leverages automated planning techniques with the Planning Domain Definition Language (PDDL), explicitly encoding tasks, constraints, and preferences from both human and robotic perspectives. Task allocation is optimized based on three key criteria: the robot’s effort, the human’s effort, and overall task success. Additionally, we integrate a Natural Language Processing (NLP) model that interprets human preferences and informs the negotiation process, ensuring that the robot generates task proposals aligned with human input. The negotiation follows an alternating-offer protocol, with the robot employing a sigmoid conceder strategy to iteratively refine task allocation, leading to balanced and mutually acceptable plans. To evaluate our approach, we conduct a comprehensive user study with non-trained volunteers interacting with the robot, assessing the effectiveness, fairness, and adaptability of the proposed system in real-world scenarios.

I. INTRODUCTION

Human-robot collaboration has the potential to transform the complexity of task execution across various domains. Beyond mechanical efficiency, effective collaboration requires a mutual understanding of each agent’s preferences, limitations, and objectives. Advances in human-robot interaction (HRI) have facilitated more coordinated methodologies for integrating robots and humans in shared work environments [1], [2].

Despite these advancements, achieving fluid and balanced decision-making between humans and robots remains a significant challenge. Relying solely on explicit human directives restricts robot autonomy, limiting its ability to adapt and leverage its full capabilities in dynamic environments [3]. In addition, requiring humans to dictate task plans unilaterally can impose high cognitive effort and

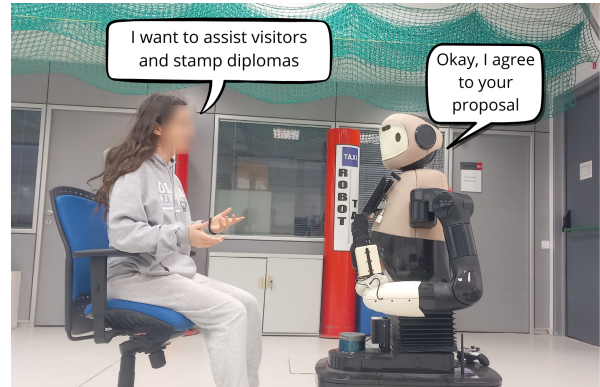


Fig. 1. IVO Robot engaging in a negotiation with a volunteer, demonstrating human-robot interaction in collaborative decision-making.

time demands, often requiring extensive knowledge of the robot’s capabilities [4]. Furthermore, effective collaborative task planning should integrate technical performance and ethical dimensions, aligning the capabilities of the system with broader operational and social considerations. Robotics decisions can reflect surrogate moral reasoning influenced by social norms or interests of third parties, causing additional complexities in shared decision making [4]. Therefore, we argue that collaborative plans should be co-developed through a negotiation process in which all parties contribute to a mutually acceptable task allocation [1], [3].

To address this challenge, we present a negotiation framework for task allocation in human-robot teams. Our shared planning approach enables humans and robots to negotiate task assignments, balancing individual preferences with overall efficiency [5]. The model combines automated planning techniques with the Planning Domain Definition Language (PDDL) to generate optimized plans and schedules [6]. To support interaction, a Large Language Model (LLM) interprets natural language feedback, enabling a negotiation process via an alternating-offers protocol. The framework was validated through a user study where participants negotiated task allocations with the robot.

The remainder of this paper is organized as follows. Section II introduces the problem definition. Section III presents the methodology used. Section IV describes the experimental setup and evaluation methodology, with results reported in Section V. Finally, Section VI discusses the findings and concludes the paper.

This work was partially supported by JST Moonshot R, D grant number JPMJMS2011 and LENA project (PID2022-142039NA-I00), funded by MCIN/AEI/10.13039/501100011033

¹ Universitat Politècnica de Catalunya (UPC), Jordi Girona 31, Barcelona, 08034, Spain {llum.fuster, marc.dalmaso, artur.aubach, anais.garrell}@upc.edu

² Institut de Robòtica i Informàtica Industrial (CSIC-UPC), Llorens i Artigas 4-6, Barcelona, 08028, Spain {mdalmaso, sizquierdo, agarrell}@iri.upc.edu

³ Eurecat, Centre Tecnològic de Catalunya, Robotics and Automation Unit, Barcelona, 08290, Spain silvia.izquierdo@eurecat.org

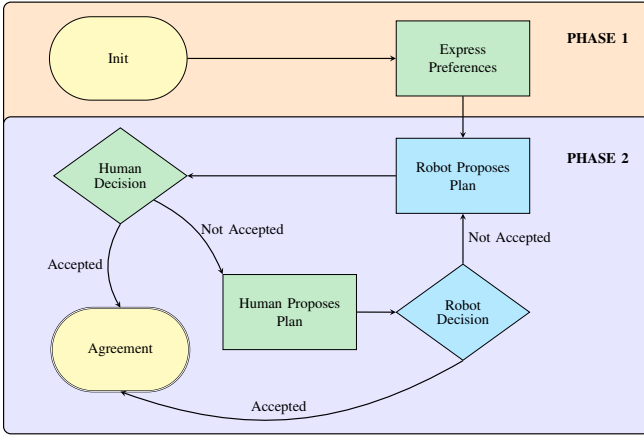


Fig. 2. *Interaction flow*. It shows the protocol between the human (green) and the robot (blue), structured in two phases: **Phase 1** (Preference Elicitation) begins with user initialization and expressing preferences, and **Phase 2** (Alternating Offers Protocol) ends when an *Agreement* is reached.

II. PROBLEM DEFINITION

In this work, we address the challenge of establishing a collaborative negotiation process between two agents – a human (h) and a robot (r) – with the aim of achieving an optimized allocation of tasks. The set of agents is denoted as $A = \{r, h\}$ and the task set is defined as:

$$\mathbf{T} = \{T_1, T_2, \dots, T_n\} \quad (1)$$

encompassing all pending tasks. These tasks have to be distributed among the two agents, thus generating an assignment plan:

$$\pi : \mathbf{T} \rightarrow A \quad \text{with} \quad \pi(T_j) \in A \quad \forall T_j \in \mathbf{T}. \quad (2)$$

A plan π represents a specific allocation of tasks to the agents, meaning that for every task T_j , the function π jointly assigns an agent $a \in A$ responsible for executing it. Formally, the negotiation process entails generating a series of such plans. This sequence is denoted by:

$$\Pi = \{\pi_1, \pi_2, \dots, \pi_m\}, \quad (3)$$

The negotiation process continues until one of the agents accepts the plan proposed by the other agent. Formally, for each proposal, π_i , there exists a feedback $fb_i \in \{accept, reject\}$ linked to that proposal.

A. Application Domain

This study focuses on a reception environment scenario where a human and a robot collaborate to perform management and logistics tasks, such as document archiving, material preparation, mail handling, and visitor reception. In our domain, we define a set of 16 possible tasks: $\mathbf{T} = \{T_1, T_2, \dots, T_{16}\}$.

The team, composed of agents r and h , has the flexibility to dynamically distribute tasks. The methodology remains applicable as long as the agents' preferences and limitations are either known or can be inferred through interaction, and the tasks are explicitly defined in a structured format.

III. METHODOLOGY

In order to validate our scheduling plan negotiation framework for human–robot teams, we adopted an iterative, modular approach that integrates speech processing, natural language understanding, and automated planning within a unified system. The overall system is illustrated in Fig. 3.

A. Communication interface

As illustrated in Fig. 3, the negotiation process is conducted entirely through spoken interaction, ensuring a natural and intuitive communication flow between the user and the robot. The system captures the user's voice input s_h and processes it through the Speech-to-Text (STT) module, which converts the spoken words into textual data τ_h . Then τ_h is processed by the negotiation system, that generates text responses τ_r . These are, in turn, processed by the Text-to-Speech (TTS) module, a synthesizer of responses into spoken language, generating the speech message s_r . Future work could unify these STT and TTS modules using a single multi-modal LLM.

B. Agent & Plan Modelling

Each agent $a \in A$ is assumed to have a set of preferences ρ_a , which are formally defined as follows:

$$\rho_a(T_j) = \begin{cases} +1, & \text{if } T_j \text{ favored} \\ 0, & \text{if } T_j \text{ neutral} \\ -1, & \text{if } T_j \text{ disfavored} \end{cases} \quad (4)$$

Building on that, each plan is characterized using a vector of rewards that quantifies a plan from three perspectives: human preferences, robot preferences, and a common objective. The reward vector is defined in detail in Section III-D.

C. Planning Process

The shared planning process is composed of two phases: the preference elicitation and the negotiation phase.

1) *Preference elicitation - phase 1*: As illustrated in Fig. 2, the first phase focuses on extracting human preferences, which serve as a foundation for the initial task allocation.

The preference extraction process is handled by the LLM Preferences module (Fig. 3). This module analyzes natural language input from the human user τ_h , extracting a structured set of preferences ρ_h . Specifically, it identifies tasks that are favoured, neutral, or disfavoured, enabling an adaptive task planning approach.

Mathematically, the preference extraction process can be formulated as:

$$\rho_h = PE(\tau_h) \quad (5)$$

where PE represents the preference extraction function, which processes the human input τ_h and generates a structured preference set ρ_h .

In contrast, the robot also has predefined preferences, but unlike the human, it does not explicitly express them. The robot's preferences ρ_r consist of four disfavoured tasks, while all other tasks are considered neutral. These predefined

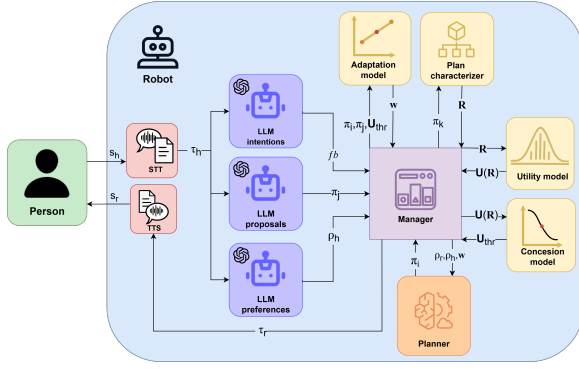


Fig. 3. Diagram illustrating the flow from the human’s spoken input (s.h), processed by STT and LLM modules, to a central Manager that coordinates the negotiation loop before returning to the user.

constraints are directly incorporated into the task allocation process without requiring an extraction phase.

This phase is crucial in tailoring the task allocation process to human preferences, ensuring a more efficient and user-centric collaboration in subsequent phases.

2) *Negotiation Protocol - phase 2*: As depicted in Fig. 2, the second phase of the interaction flow consists of an Alternating Offers Protocol where the robot and the human iteratively negotiate task allocation until an agreement is reached.

The process begins with the robot making the first proposal, suggesting a full allocation of tasks π_1 based on the extracted preferences from Phase 1. The human then responds by either accepting the proposal or counter-proposing a new allocation π_2 . This iterative exchange continues, with the robot adjusting its proposals accordingly, until one of the agents, denoted as a , accepts the proposed allocation—meaning that the feedback value fb associated to τ_a is *accept*.

To support this interaction, two additional LLM modules are used. First, the LLM Feedback Interpreter analyzes the user’s natural language input τ_h to determine whether the plan is accepted. This decision is formalized as:

$$fb = HI(\tau_h) \quad (6)$$

where HI denotes the human intention function, returning a binary feedback fb .

If the response is not an acceptance, a second LLM module, the LLM Plan Parser, extracts the user’s proposed assignment of tasks. This module interprets the subset of tasks the human is willing to execute and derives a complete plan π_h :

$$\pi_h = PH(\tau_h) \quad (7)$$

where PH is the function that processes the human’s textual input τ_h and outputs a structured plan π_h , implicitly assigning the remaining tasks to the robot.

It is relevant to mention that the plans are described differently by each agent: the robot explains all assignments, while the human only specifies their own. This design choice aims to reduce the human’s verbal effort during the negotiation.

D. Plan Generation and Evaluation

Fig. 3 provides an overview of the modules involved in creating and assessing a plan. In essence, the *Planner* generates candidate allocations for a set of tasks (defined in a structured format with their durations and constraints), while the *Plan Characterizer* extracts relevant metrics from each plan to obtain a reward vector \mathbf{R} . Subsequently, a *Utility Model* computes a scalar value representing how desirable the plan is for the robot. Finally, the *Concession Model* determines a threshold for an acceptable plan, and the *Adaptation Model* adjusts the weights $\hat{\mathbf{w}}$ for the Planner to seek a new solution, adapting to a new human proposal.

1) *Objective Function*: For the creation of new plans, the robot considers a cost function that integrates the three evaluation components. Let the weight vector be defined as:

$$\hat{\mathbf{w}} = (w_r, w_h, w_c), \quad \text{such that } \|\hat{\mathbf{w}}\| = 1 \quad (8)$$

where each weight w_i (for $i \in \{r, h, c\}$) reflects the relative importance of the robot, human and common objective, respectively. The common objective, in terms of costs, is directly related to the total task execution time, following a 1-to-1 correspondence. In our experiments (Sec. IV), these weights are independent variables that act as multipliers in the cost function.

2) *Plan Generation via PDDL*: We formalize the task-allocation problem in a PDDL problem, associating each task to durations $d_a(T)$ and costs $c_a(T)$ for each agent a . The costs are determined as follows:

$$c_a(T) = d_a(T) \cdot f_a(T) \quad (9)$$

where $f_a(T)$ is the cost adjustment factor, defined as:

$$f_a(T) = 1 - \rho_a(T) \cdot \frac{k}{N_a} \quad (10)$$

where k is a parameter that controls the intensity of the adjustment and N_a represents the total number of tasks with preferences (either positive or negative) for agent a :

$$N_a = \sum_{T_j \in \mathbf{T}} |\rho_a(T_j)| \quad (11)$$

A weighted sum of these costs (robot effort, human effort, and total execution time) is minimized to produce a plan π . This plan assigns tasks to each agent and schedules their execution to respect temporal constraints.

3) *Plan Characterization*: For a given plan π (either generated by the Planner or received from the human), we extract three key measures: the robot total effort $C_r(\pi)$, the human total effort $C_h(\pi)$ and the plan duration cost $C_c(\pi)$. Each one of these measures is normalized to using the single agent allocation plans ($\pi(T_i) = a \forall i$) and inverted to obtain a reward vector

$$\mathbf{R}(\pi) = (R_r(\pi), R_h(\pi), R_c(\pi)). \quad (12)$$

where higher values in each component indicate lower cost.




Fig. 4. Rewards of the plans π_1 (red), π_2 (blue), π_3 (green) in the 3D space and weight vector used to weight the utility function to obtain π_3

4) *Plan Utility*: A utility function U then maps the reward vector to a scalar value. In our approach, we define $U(\cdot)$ as a smoothed multidimensional Gaussian, that peaks at the robot’s ideal reward vector. Formally,

$$U(\mathbf{R}) = \frac{\text{Gauss}(\mathbf{R})}{\text{Gauss}(\bar{\boldsymbol{\mu}})}, \quad (13)$$

where $\text{Gauss}(\cdot)$ denotes a multivariate normal density centered at $\bar{\boldsymbol{\mu}}$ and the denominator ensures $U(\bar{\boldsymbol{\mu}}) = 1$. This $\bar{\boldsymbol{\mu}}$ is set as $\mathbf{R}(\pi_1)$ meaning that robot’s utility will be 1 for its first proposal, considering it as its “ideal” plan.

5) *Acceptance and Concession*: To decide if a human-proposed plan π is acceptable at a given time t (in seconds), the robot compares $U(\mathbf{R}(\pi))$ against a sigmoidal threshold U_{thr} :

$$U_{thr} = U_{max} + \frac{U_{min} - U_{max}}{1 + e^{-k(t-t_{ip})}} \quad (14)$$

where U_{min} is the initial value of the sigmoid function, defined by previous experiments, U_{max} is the final value, defined by expert knowledge, k controls the steepness of the sigmoid curve, and t_{ip} determines its inflection point. The plan is accepted if

$$U(\mathbf{R}(\pi)) > U_{thr}. \quad (15)$$

which means whether the plan’s utility reaches the utility threshold U_{thr} . Otherwise, the robot adapts its priorities to move toward a plan that meets or slightly differs the same utility threshold U_{thr} . This adaptation is achieved by gradient descent on the reward space:

$$\mathbf{R}_{new} = \text{GradientDescent}(\mathbf{R}(\pi_i), \mathbf{R}(\pi_{i-1}), U_{thr}). \quad (16)$$

After identifying the target rewards \mathbf{R}_{new} , the cost normalization process is inverted to obtain a corresponding weight vector \mathbf{w} , which will guide the planner’s minimization function. The goal is to steer the search toward a plan whose rewards (and resulting utility) approximate the desired target—although such a plan may not exist exactly within the

reachable state space, the new weights bias the search toward a close alternative. Fig. 4 shows an example where $R(\pi_1)$, $R(\pi_2)$, and $R(\pi_3)$ are plotted in a 3D space, along with the weight vector used to approximate the third plan.

This iterative cycle continues until either the human’s proposal achieves sufficient utility or the system converges on a mutually acceptable plan, ensuring that both efficiency (in terms of time and effort) and the robot’s internal criteria are adequately balanced.

E. Implementation Details

The complete system follows a modular architecture, integrating existing components and custom-developed modules as illustrated in Fig. 3. Speech interaction is facilitated by STT and TTS components based on Vosk and gTTS, respectively, adapted from Hriscu et al.[7]. Natural language understanding is performed using a Large Language Model (LLM), specifically OpenAI’s GPT-4o-mini. In this context, preference extraction is achieved through an adapted module from Izquierdo-Badiola et al.[8], which employs few-shot learning techniques. Other GPT-based modules—including proposal interpretation and feedback identification—leverage the structured outputs functionality to ensure consistent and interpretable responses from the LLM [9]. Plan generation use a tailored PDDL 2.1 formulation, solved with the OPTIC temporal planner [10]. Finally, all modules are coordinated via a centralized Manager, implemented as a ROS node.

IV. EXPERIMENTAL SETTING

Through these experiments, we assess the negotiation system in terms of robot behavior, human response, and overall interaction quality.

A. Scenario

The experimental setup was conducted at the Mobile Robotics Laboratory of the Institute of Robotics and Industrial Informatics (IRI), located in the Faculty of Mathematics and Statistics (FME). This facility serves as a controlled experimental space for mobile robotics research, enabling practical exploration and development of collaborative human-robot tasks (see Fig. 1) .

The robot used in this study is IVO (see Fig. 1), a social robot specifically designed for human interaction and collaboration [11]. IVO has an adjustable height (1.20 m–1.55 m), two 7-DOF robotic arms with wrist force/torque sensors, and a five-finger robotic hand alongside a gripper. Its Communication capabilities include an integrated speaker and microphone for verbal interaction, complemented by a touchscreen interface for visual communication.

B. Participants

The experimental study included the participation of 28 individuals, each of whom engaged in one negotiation with the robot. Their age ranged from 19 to 80 years old (mean 30.5, std 16.22). Their gender distribution was balanced, with 46,4% male and 54,6% female representation, minimizing

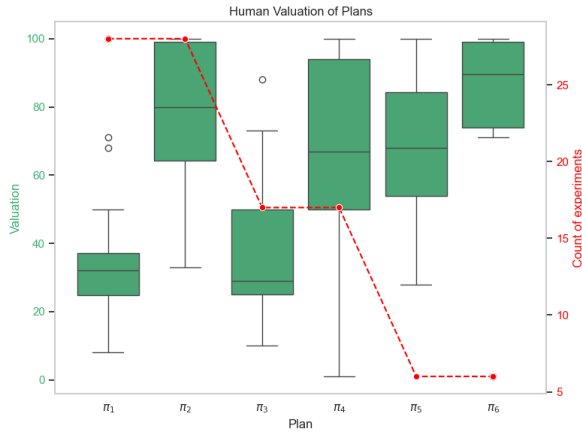


Fig. 5. Human Evaluation of Π , represented with boxplots (green, left axis). The red line (right axis) shows the number of experiments that reached the corresponding iteration for each plan

potential biases in the analysis of results. Additionally, the participants’ level of knowledge in robotics (median 2.0, std 1.12) has been assessed using a scale from 1 (no knowledge) to 5 (expert), as summarized in Table II.

C. Experimental Groups

We defined four experimental groups (G1–G4), each combining two independent variables: the weight vector \hat{w} and the inflection point influencing the acceptance threshold. Table II outlines the parameters and demographic distributions.

D. Data Collection

After each negotiation session, all participants completed a comprehensive post-experiment questionnaire. It included NASA -TLX [12], [13] to assess perceived workload, along with subsets adapted from Goodspeed [14], ROCI II [15], Rahim (JGP) [16], and Barbuto Scholl [17] with the aim of evaluating human HRI and negotiation perception.

V. RESULTS

This section summarizes the user study results (Table I). We begin by examining participants’ evaluations of the different plan proposals, extracted from the questionnaire, Fig. 5. Notably, the robot’s initial proposal π_1 received the lowest median evaluation. However, subsequent robot proposals (π_3 and π_5) show higher evaluations, suggesting that the robot gradually proposes better plans and adapts more effectively to the human’s feedback. In contrast, the human’s proposals start with π_2 at the highest median evaluation, followed by π_4 , which reflects some degree of concession but it is not seen on π_6 . This rise corresponds to the right y-axis, indicating more iterations were needed for participants to reach π_6 , suggesting that firm stances led to greater robot concessions.

Similarly, Fig. 6 compares the final negotiated plan (π_a) with unilateral decision-making by the robot (π_1) and the first human proposal (π_2). Under the unilateral scenario (π_1), the human’s utility is low, indicating dissatisfaction. When the human presents their own plan (π_2), their utility improves

significantly, while the robot’s utility decreases only slightly. Finally, negotiation (π_a) achieves a more balanced outcome for both.

To analyze the questionnaire responses and identify patterns in behavior and perception of the robot, ANOVA tests were applied, Table I. The NASA-TLX item Q1.3, which measures mental demand, showed a positive and statistically significant correlation with the robotics expert profile, suggesting that as participants’ robotics knowledge increases, they tend to perceive the task as more mentally demanding. In a similar vein, Rahim(JGP)’s Q4.1, which assesses whether the robot withholds critical information needed for a collaborative agreement, also exhibited a significant positive correlation with robotics knowledge. This indicates that participants with more expertise are more sensitive to potential biases or specific attitudes displayed by the robot.

Regarding total task time, Goodspeed items Q2.2 and Q2.3 showed significant negative correlations, indicating that shorter negotiations are linked to perceptions of the robot as more knowledgeable and competent. In the ROCI II subset, Q3.1 (robot accommodating user preferences) and Q3.2 (robot yielding to human proposals) also showed significant negative correlations with total time, suggesting that quicker negotiations enhance the robot’s perceived flexibility. Notably, Q3.2 had a particularly strong effect. In contrast, Rahim(JGP) item Q4.4 (perceived internal disagreement) and Q4.6 (divergence in proposal evaluation) showed significant positive correlations with time, implying that longer negotiations may increase perceptions of conflict in task allocation.

Although Q3.7, which gauges the inclination to seek a middle-ground agreement, did not reach a high level of statistical significance with respect to robotics knowledge, there was a noticeable trend: more expert participants tended to show less willingness to negotiate compromises. Overall, these findings suggest that both the negotiation duration and the participants’ level of robotics expertise are critical factors in shaping perceptions of the robot’s performance. Participants with higher robotics knowledge appear to adopt a more assertive, challenger role, while shorter negotiation interactions tend to foster more favorable impressions of the robot’s intelligence and adaptability.

In Fig. 7, a progressive increase is observed in the user’s ability to detect the robot’s priorities as the accepted plan index increases, i.e., the more iterations carried out before concluding the negotiation. For π_2 , the precision and recall values are low; with π_4 , both metrics improve; and, finally, with π_6 , the highest precision and an intermediate recall are achieved. This indicates that, as the human interacts more with the robot, it becomes easier to correctly identify its priorities.

As shown in Fig. 8, we find no significant correlation between preference elicitation understanding errors and the resulting utility for the agents, suggesting that the negotiation process provides an implicit alignment mechanism robust to modelling errors.

To investigate this further, a Pearson correlation analysis

TABLE I
QUESTIONNAIRE COMPILATION

Questionnaire	mean	std	Total time		Robotics knowl.	
			corr.	p-value	corr.	p-value
NASA-TLX [12], [13] (1 - 100)						
<i>Q1.3</i> MENTAL DEMAND	37.43	21.90	0.188	0.339	0.388	0.041
Goodspeed Subset [14] (1 → 5)						
<i>Q2.1</i> No emotion → 5 Very emotional	1.857	1.044	0.023	0.909	-0.276	0.155
<i>Q2.2</i> Ignorant → Knowledgeable	3.679	0.723	-0.416	0.028	-0.062	0.754
<i>Q2.3</i> Incompetent → Competent	3.750	0.844	-0.435	0.021	-0.137	0.487
ROCI II Subset [15] (Strongly disagree (1) - Strongly agree (7))						
<i>Q3.1</i> The robot usually accommodates my preferences.	5.321	1.492	-0.511	0.005	0.008	0.968
<i>Q3.2</i> The robot has given in to my proposals.	4.929	1.904	-0.761	3e-06	0.159	0.419
<i>Q3.7</i> I try to find a middle course to reach an agreement.	4.464	1.753	0.020	0.918	-0.319	0.098
Rahim(JGP) Subset [16] (Strongly disagree (1) - Strongly agree (7))						
<i>Q4.1</i> The robot withholds information necessary for the attainment of a collaborative agreement.	3.429	1.731	0.027	0.891	0.481	0.010
<i>Q4.3</i> There has been cooperation between the robot and me.	4.679	1.278	-0.098	0.620	-0.216	0.269
<i>Q4.4</i> In our group, we have lots of bickering over who should do what job.	3.357	1.909	0.714	2e-05	0.074	0.707
<i>Q4.5</i> There is "good" match between the teams's (human and robot) preferences and the tasks	5.179	1.416	-0.311	0.108	0.108	0.583
<i>Q4.6</i> There is a difference of opinion in the evaluation of the proposals between the robot and me.	4.036	1.401	0.456	0.015	-0.025	0.898

TABLE II
FINDINGS IN GROUP STUDY

Group	Age		Gender		Independent Variables				Goodspeed Subset [14]			Rahim(JGP) Subset Adaptation [16]				
	mean	std	M	F	w_r	w_h	w_c	T_{ip}	<i>Q2.1</i>	<i>Q2.2</i>	<i>Q2.3</i>	<i>Q4.1</i>	<i>Q4.3</i>	<i>Q4.4</i>	<i>Q4.5</i>	<i>Q4.6</i>
G1	28.7	18.7	43%	57%	0.55	0.60	0.20	228	1.714	4.000	4.286	3.000	5.429	2.143	5.857	3.571
G2	36.1	16.5	43%	57%	0.55	0.60	0.20	312	1.143	3.714	3.857	4.000	4.143	4.714	5.143	5.143
G3	24.3	4.4	71%	29%	0.40	0.80	0.70	228	2.333	3.000	3.000	4.667	5.167	4.500	3.667	4.333
G4	33.0	23.2	29%	71%	0.40	0.80	0.70	312	2.667	3.833	3.667	2.000	4.333	2.500	5.500	3.667
Anova test p-value									0.034	0.079	0.055	0.039	0.065	0.013	0.023	0.083

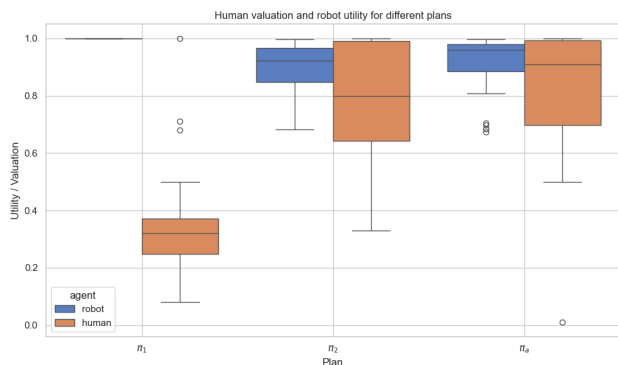


Fig. 6. Comparison of robot’s utility (blue) and human evaluation (orange) for three key plans: the robot’s initial proposal (π_1), the human’s first counter-proposal (π_2) and the final agreement (π_a).

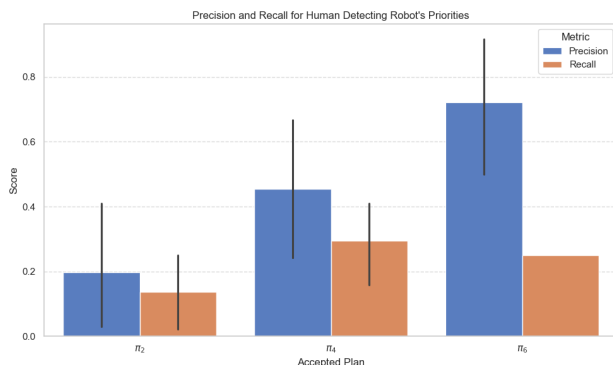


Fig. 7. Human precision and recall in identifying the robot’s priorities, by the agreement iteration (π_2 , π_4 , or π_6). Increasing scores demonstrate improved mental model alignment as the negotiation extends.

examined the relationship between the robot’s understanding of human preferences (*LLM_preferences_metric*) and its evaluation of its own plans. When the robot poorly understood human preferences, it tended to rate its plans more negatively. While this correlation was not significant for π_1 ($r = -0.2035$, $p = 0.2990$), significant negative correlations

were found for π_3 and π_5 ($r = -0.6517$, $p = 0.0046$ and $r = -0.8451$, $p = 0.0342$).

Although these results might appear to conflict with Fig. 8, they are complementary: early misunderstandings may lower the perceived quality of individual robot plans, yet the negotiation still leads to agreements with no significant loss

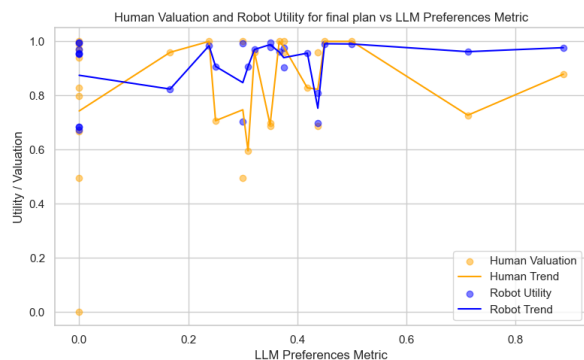


Fig. 8. Human evaluation and Robot utility for the final plan accepted vs preferences detection metric (mean for correctness and completeness for both favor and disfavor)

in utility or perceived evaluation by human. This is reinforced by the weak, non-significant correlation for human’s final plan rating ($r = 0.1852$, $p = 0.3454$).

A. Group Study

Table II summarizes the ANOVA results assessing group differences in participants’ perceptions based on the Goodspeed and Rahim (JGP) questionnaire subsets. The Tukey HSD test highlights significant pairwise differences (in bold).

In the Goodspeed subset, Q2.1 shows a significant difference in perceived robot emotionality. Q2.2 and Q2.3, with p -values near 0.05, suggest a potential trend in perceived robot intelligence, particularly between G1 and G3.

In the Rahim (JGP) subset, which evaluates human-robot interaction dynamics, Q4.1 reveals differences between G3 and G4, possibly influenced by gender imbalance. Q4.4 shows a significant difference between G1 and G2 in perceived conflict, potentially linked to higher T_{ip} in G2. Q4.6, though not significant, suggests similar trends.

Perceptions of task fit (Q4.5) differ between G1 and G3, possibly due to weight vectors less favorable to the human; this warrants further investigation, especially considering gender effects. Finally, Q4.3 approaches significance, hinting at cooperation discrepancies between G1 and G2.

VI. CONCLUSIONS

This work presents a negotiation framework for task allocation in human-robot teams, integrating automated planning (PDDL) with natural language processing. Using an alternating offers protocol and a sigmoidal concession strategy, the system iterates until both parties reach mutually acceptable plans. Tests in a reception setting demonstrated improved efficiency, perceived fairness, and adaptability compared to unilateral decision-making.

Experiments with non-expert participants show that iterative negotiation enhances preference alignment and user satisfaction. While results are promising, the negotiation process can be further refined for more complex and dynamic scenarios. Future research should explore advanced preference models, real-time adaptation, and multi-agent collaboration, as well as assess the long-term effects of negotiation on team performance.

Ultimately, combining formal planning with LLM-driven negotiation enables robotic systems that are not only technically proficient but also better aligned with human expectations and collaborative norms.

REFERENCES

- [1] S. Nikolaidis and J. Shah, “Human-robot cross-training: Computational formulation, modeling and evaluation of a team training strategy,” in *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. ACM, 2013, pp. 33–40.
- [2] A. Ali, Y. Wang, and P. Doshi, “Using trust for heterogeneous human-robot team task allocation,” *Autonomous Agents and Multi-Agent Systems*, vol. 35, pp. 1–18, 2021.
- [3] E. Sebastiani, A. Cesta, and G. Cortellessa, “Dealing with on-line human-robot negotiations in hierarchical task planning,” in *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*. AAAI Press, 2017.
- [4] J. van Diggelen, M. Johnson, J. Bradshaw, et al., “Moral decision making in human-agent teams: Human control and the role of explanations,” *Frontiers in Robotics and AI*, vol. 8, pp. 289–294, 2021.
- [5] S. Rothfuss, J. Milliken, M. Brown, and J. Lee, “A negotiation-theoretic framework for control authority transfer in mixed-initiative robotic systems,” in *IEEE Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2022.
- [6] M. Helmert, “Concise finite-domain representations for pddl planning tasks,” *Artificial Intelligence*, vol. 173, no. 5, pp. 503–535, 2009, advances in Automated Plan Generation. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0004370208001926>
- [7] L. Hriscu, “Large language models to enhance the interaction between humans and robots,” Master’s thesis, Universitat Politècnica de Catalunya, 2024.
- [8] S. Izquierdo-Badiola, G. Canal, C. Rizzo, and G. Alenyà, “Plancolabl: Leveraging large language models for adaptive plan generation in human-robot collaboration,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [9] M. Pokrass, C. Colby, M. Guan, T. Sanders, and B. Zhang, “Introducing structured outputs in the api,” 9 2024. [Online]. Available: <https://openai.com/index/introducing-structured-outputs-in-the-api/>
- [10] J. Benton, A. Coles, and A. Coles, “Temporal planning with preferences and time-dependent continuous costs,” *ICAPS 2012 - Proceedings of the 22nd International Conference on Automated Planning and Scheduling*, vol. 22, 05 2012.
- [11] J. Laplaza, N. Rodríguez, J. E. Domínguez-Vidal, F. Herrero, S. Hernández, A. López, A. Sanfeliu, and A. Garrell, “Ivo robot: A new social robot for human-robot collaboration,” in *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2022, pp. 860–864.
- [12] S. G. Hart, “Nasa-task load index (nasa-tlx); 20 years later,” in *Proceedings of the human factors and ergonomics society annual meeting*, vol. 50, no. 9. Sage publications Sage CA: Los Angeles, CA, 2006, pp. 904–908.
- [13] I. de Arquer and C. Nogareda, “Ntp 544: Estimación de la carga mental de trabajo: el método nasa tlx,” *Madrid: Instituto Nacional de Higiene y Seguridad en el Trabajo, Ministerio de Trabajo y Seguridad Social del Gobierno de Espana*, 2000.
- [14] C. Bartneck, “Godspeed questionnaire series: Translations and usage,” in *International handbook of behavioral health assessment*. Springer, 2023, pp. 1–35.
- [15] M. A. Rahim, “Rahim organizational conflict inventory–ii,” *Journal of Applied Psychology*, 1983.
- [16] —, “Measurement of organizational conflict,” *The Journal of General Psychology*, vol. 109, no. 2, pp. 189–199, 1983.
- [17] J. E. Barbuto Jr and R. W. Scholl, “Motivation sources inventory: Development and validation of new scales to measure an integrative taxonomy of motivation,” *Psychological reports*, vol. 82, no. 3, pp. 1011–1022, 1998.