Preference and Performance-Based Adaptive Task Planning in Human-Robot Collaboration

by

Ali Noormohammadi Asl

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2024

© Ali Noormohammadi Asl 2024

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Rachid Alami Senior Scientist, LAAS-CNRS, Laboratory for Analysis and Architecture of Systems, The National Center for Scientific Research
Supervisor(s):	Stephen L. Smith Professor, Dept. of Electrical and Computer Engineering, University of Waterloo
	Kerstin Dautenhahn Professor, Dept. of Electrical and Computer Engineering, University of Waterloo
Internal Member:	Nasser Lashgarian Azad Associate Professor, Dept. of Systems Design Engineering, University of Waterloo
Internal-External Member:	Christopher Nielsen Professor, Dept. of Electrical and Computer Engineering, University of Waterloo
	Yue Hu Assistant Professor, Dept. of Mechanical and Mechatronics Engineering, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis encompasses five published manuscripts, all achieved through collaborative efforts with various researchers. Professor Kerstin Dautenhahn and Professor Stephen L. Smith supervised all research activities and reviewed drafts of the manuscripts. The specific contributions to each manuscript are outlined below:

- A. Noormohammadi, A. Dahiya, A. M. Aroyo, S. L. Smith, and K. Dautenhahn, "The effect of robot decision making on human perception of a robot in a collaborative task - a remote study," in Proceedings of the 9th International Conference on Human-Agent Interaction, 2021, p. 423–427.

  I developed the research questions, designed the simulation environment, formulated the planning problem, conducted the user study, analyzed the data, and wrote the initial draft of the paper. My coauthors, Abhinav Dahiya and Alex Aroyo, both affiliated with the Department of Electrical and Computer Engineering at the University of Waterloo, contributed by assisting in the design of the user study and providing feedback throughout the writing process. Abhinav also helped in writing section 3 of this paper. The content of Chapter 3 primarily derives from this paper.

- A. Noormohammadi-Asl, A. Ayub, S. L. Smith, and K. Dautenhahn, "Task selection and planning in human-rob collaborative processes: To be a leader or a follower?" in 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 2022, pp. 1244–1251.

  I took the lead in developing the task planning framework, conducting the corresponding mathematical modeling, designing the simulation environment for method evaluation, analyzing the data, and writing the initial draft of the paper. Ali Ayub, from the Department of Electrical and Computer Engineering at the University of Waterloo, collaborated by providing feedback throughout the design of the collaborative scenario and the writing process. Parts of Chapters 4 and 5 are adapted from this publication.

- A. Noormohammadi-Asl, A. Ayub, S. L. Smith, and K. Dautenhahn, "Adapting to human preferences to lead or follow in human-robot collaboration: A system evaluation," in 2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 2023, pp. 1851–1858.

  As the main author of this work, I implemented the planning framework on a real mobile manipulator robot, programmed the robot for autonomous pick-and-place, conducted the system evaluation, and prepared the first complete draft of the paper. Ali Ayub

contributed by helping to prepare the robot for the study, refining the collaborative scenario, and assisting through the writing process. Parts of Chapter 6 have been drawn from this publication.

- A. Noormohammadi-Asl, S. L. Smith, and K. Dautenhahn, "To lead or to follow? adaptive robot task planning in human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01483 (Under review at a journal)

  The main parts of the work presented in Chapters 4, 6, and 7 are from this publication.

- A. Noormohammadi-Asl, K. Fan, S. L. Smith, and K. Dautenhahn, "Human leading or following preferences: Effects on human perception of the robot and the human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01466 (Under review at a journal)

  For the last two manuscripts, I took the lead in undertaking tasks such as implementing the planning framework on a real mobile manipulator robot, designing the user study, preparing the experimental setup, creating a graphical user interface, recruiting participants, conducting the user study, analyzing the data, and drafting the paper in its entirety. My co-author, Kevin Fan, assisted in running the user study and refining the collaborative scenario. The main parts of the work presented in Chapters 6 and 7 are from this publication.

Additional Publications: I have also made contributions to the publications not included in this thesis

- K. Fan, M. Jouaiti, A. Noormohammadi-As, C. L. Nehaniv and K. Dautenhahn, "A Social Referencing Disambiguation Framework for Domestic Service Robots," 2023 IEEE International Conference on Robotics and Automation (ICRA), London, United Kingdom, 2023, pp. 11979-11985, doi: 10.1109/ICRA48891.2023.10161168.

- K. Fan, M. Jouaiti, A. Noormohammadi-As, C. L. Nehaniv and K. Dautenhahn, "Social Referencing Disambiguation with Robot Mental Imagery for Domestic Service Robots - System Implementation and Validation." (Under review at a journal)

Abstract

This thesis delves into a central challenge in human-robot collaboration (HRC): the adaptive task planning of robots to enhance team performance, fluency, and the human agent's perception of both the robot and the collaboration. This thesis tackles the challenge of proactive task planning and allocation in collaborative scenarios, involving a single human and a single robot working together to accomplish a task. Recognizing the existing gaps in the literature, our focus revolves around balancing human agents' leading/following preferences and their performance, with the aim of fostering collaboration while maintaining a high level of human perception of the robot.

After an in-depth review of related work, we initiate our exploration with an online user study, in a simulation environment using a manipulator robot. This study is designed to evaluate the impact of the robot's planning strategy on participants' perception of the robot and collaboration. This study incorporates three distinct planning strategies: prioritizing the human's objectives, prioritizing the robot's objectives, and achieving a balance between both agents' objectives. The results guide our assessment of how the balancing strategy, in particular, can uphold both team performance and a high level of participants' perception of collaboration and the robot, in comparison to the other strategies. However, a limitation arises as the study employs fixed strategies, randomly assigned to participants, irrespective of their preference and performance.

Building upon the results of the first user study, we address the limitations identified in the initial study by enabling the robot to estimate the human agent's leading/following preference. However, the human agent's preference is not the sole factor influencing the robot's decision-making process; the human agent's performance is also crucial for adjusting the team's overall performance, particularly in cases of the human agent's poor or suboptimal performance. Consequently, the robot also estimates the human agent's performance. Furthermore, the robot needs to be capable of updating the task state based on both agents' actions and mistakes. With an updated understanding of the human agent's performance, leading/following preference, and task state, the robot updates its plan for task allocation and scheduling to minimize collaboration costs.

Next, we evaluate the adaptability of the task planning framework and algorithm in a simulation environment, demonstrating its effectiveness across various human performance and preference scenarios. Yet, recognizing the unique challenges posed by human participants, the complete evaluation of the algorithm's effectiveness requires real-world scenarios, considering uncertainties inherent in human behavior and decision-making.

Subsequently, we tackle the challenges of implementing the task planning framework on a real robot, a mobile manipulator robot, within a carefully designed collaborative scenario.

Providing details on the experimental setup and methodology, a system evaluation study highlights the robot's ability to adapt based on human behavior. Finally, we conduct a user study involving 48 participants, evaluating results from multiple perspectives, including participants' perception of the robot, tasks, and collaboration, participants' actions and performance, and the robot's actions and performance. Results from the study affirm the success of the task planning framework in achieving its objectives: enhancing team fluency by considering the human agents' preferences and performance while maintaining a high level of participants' perception of the robot and the human-robot collaboration.

This thesis also explores participants' leading/following preferences in collaboration, revealing a dominant preference to lead the robot. This finding can assist robotics and autonomous systems designers in considering this factor in their designs. Additionally, we evaluated the influence of participants' leadership and followership styles on their collaboration, warranting further and more in-depth future studies.

In summary, this thesis contributes a proactive task planning framework that takes into account both human leading/following preferences and performance, signifying an advancement in the field of human-robot collaboration. The validation through user studies offers valuable insights, laying the groundwork for future research and applications in the continually evolving domain of human-robot collaboration.

Acknowledgements

I would like to express my sincere appreciation to my supervisors, Professor Kerstin Dautenhahn and Professor Stephen Smith, for their invaluable assistance, encouragement, and guidance during my doctoral studies. Their expertise has significantly influenced the trajectory of my research and greatly contributed to my academic development.

I am also grateful to the members of the examination committee for their time and contribution to evaluating my work. Additionally, I extend my thanks to the members of the Autonomous Systems Laboratory (ASL) and the Social and Intelligent Robotics Research Laboratory (SIRRL) for their constructive feedback.

Furthermore, I am deeply thankful to my friends Alireza Dehghani Sanij, Ali Yaghoubi, Amir Moghadas Jafari, and Armina Soleymani for their support throughout my Ph.D. journey.

Table of Contents

Examining Committee Membership ii

Author's Declaration iii

Statement of Contributions iv

Abstract vi

Acknowledgements viii

List of Figures xiii

List of Tables xviii

1 Introduction 1

1.1 Research Questions 4

1.2 Contributions and Organization 5

2 Related Work 8

2.1 Cobots 8

2.2 Collaborative Scenarios and Systems 8

2.2.1 Types of collaboration with cobots 10

2.2.2 Commercial and industrial applications 11

2.2.3 Research studies . . . 12
2.3 Cobot Programming . . . 14
2.3.1 Trade-offs between human and task objectives . . . 15
2.4 Task Allocation and Adaptation in HRC . . . 16
2.4.1 Task Allocation . . . 16
2.4.2 Adaptation . . . 18

3 Evaluation of Robot Strategies in HRC: A Remote Study 20
3.1 Hypothesis . . . 20
3.2 Experimental Method . . . 21
3.2.1 Materials and setup . . . 21
3.2.2 Experiment task . . . 22
3.2.3 System Modeling and Planning: . . . 23
3.2.4 Procedure and steps . . . 27
3.3 Results . . . 28
3.3.1 Participants' strategy . . . 31
3.4 Summary and Discussion . . . 33

4 Task Allocation & Planning Framework 35
4.1 Problem Statement . . . 35
4.2 Planning Architecture . . . 36
4.3 Planning Strategy . . . 37
4.4 Algorithm . . . 39
4.5 Summary and Discussion . . . 40

5 Adaptive Task Planning: A Simulation Study 42
5.1 A Sample Scenario . . . 42
5.2 Simulation Experiments . . . 45
5.3 Simulation Results . . . 48
5.4 Summary and Discussion . . . 60

6 Experimental Setup and System Evaluation **63**
6.1 Setup & Methodology . . . 63
6.1.1 Setup . . . 64
6.1.2 Tasks . . . 68
6.1.3 Agents' Actions . . . 70
6.1.4 Human-robot communication . . . 70
6.1.5 Robot . . . 72
6.2 Adaptation & Planning . . . 74
6.3 System Evaluation . . . 80
6.3.1 System Evaluation Results: Enacting Different Human Behaviors . 81
6.4 Summary and Discussion . . . 85

7 User Study Procedure & Results **86**
7.1 Related Work & Hypotheses Development . . . 86
7.1.1 Human Perception of the Robot & Collaboration . . . 87
7.1.2 Being in or under control . . . 89
7.1.3 Leadership and Followership Styles . . . 90
7.2 Procedure . . . 92
7.2.1 Recruitment . . . 92
7.2.2 Phase 0 . . . 92
7.2.3 Phase 1 (In-Person) . . . 92
7.2.4 Phase 2 (Online) . . . 94
7.3 Measures . . . 94
7.4 Results & Discussion . . . 98
7.4.1 Participants' Perception of the Tasks, Robot, & Collaboration . . . 98
7.4.2 Participants' Actions and Performance . . . 110
7.4.3 Leadership and Followership style . . . 112
7.4.4 Robot's Actions and Estimation . . . 114
7.4.5 Participant-Specific Analysis . . . 120
7.5 Summary and Discussion . . . 123

8 Conclusion and Future Work **126**

8.1 Summary 126

8.1.1 Takeaways 129

8.2 Contributions to Original Knowledge 130

8.3 Limitations & Future Work 133

References **136**

APPENDICES **149**

A Questionnaires **150**

A.1 Leadership Style 150

A.2 Followership Style 152

B Ethics Approvals **154**

C Flyer **157**

List of Figures

1.1 Spectrum of leading/following roles in human-robot collaboration 2

2.1 Number of publications on the topic of human–robot collaboration (cooperation) from 1996 to 2023; **Search Criteria:** "Human-robot collaboration" OR "Human-robot cooperation" in title and abstract; **Source:** https://app.dimensions.ai 13

3.1 Online experiment environment 22

3.2 State transition 25

3.3 Change in participants' trust before and after the experiment. 30

3.4 Averages of participants' answers to the questions in Table 3.2. 30

3.5 Expected (thick lines) and average of observed number of orange blocks passed to the robot (thin lines) along with their confidence intervals. . . . 32

4.1 Task selection and planning architecture encompassing the state estimator and robot planner. The robot planner takes into account human preference, performance, and task state, followed by the implementation of two planning steps: task selection and task scheduling. 36

5.1 Simulation environment: In a $17 \times 9m$ room, there are four workspaces (W1-W4), each containing five boxes. There are 28 objects of four different colors (7 of each) on the table. Additionally, there is a tray for the human and one for the robot. Each tray has four compartments labeled W1-W4. A video of the simulation is available at https://youtu.be/gEzNoiiGP-w 43

5.2 Humans are shown patterns with different levels of memorization difficulty. 44

5.3 Task graph of the experiment 46
5.4 Temporary task graph of the experiment after task allocation. Blue: Robot's tasks, Orange: Human's tasks, Cyan: Assigning tasks to the human, Red: Correcting human errors, Green: Already assigned tasks, Gray: Finished tasks 47
5.5 Cumulative number of assigned tasks to the human by the robot 50
5.6 Cumulative number of assigned tasks to the human by the robot 51
5.7 The robot's estimations of the human's following preference and error-proneness for a human with a strong preference to follow the robot $P_{follow} = 0.9$ . . . 53
5.8 The robot's estimations of the human's following preference and error-proneness for a human with a moderate preference to follow the robot $P_{follow} = 0.6$. 54
5.9 The robot's estimations of the human's following preference and error-proneness for a human with a slight preference to follow the robot $P_{follow} = 0.3$ . . . 55
5.10 Comparison of the results for the cases with adaptation (proposed algorithm) and without adaptation (ignoring the human's performance and preferences) when the human prefers to lead ($P_{follow} = 0.3$). 57
5.11 probability distribution for the case $P_{follow} = 0.4$, $P_{error} = 0.5$ 59
5.12 The robot's estimates of the human agent's preference and error-proneness, with the set values of $P_{follow} = 0.3$, $P_{error} = 0.4$ 60
5.13 Considering fairness, in terms of travel distance, in task allocation for the case of $P_{follow} = 0.9$, $P_{error} = 0.1$ 61

6.1 The layout of the experimental setup includes two tables in the robot's workspace and two tables in the human agent's workspace. Additionally, there is a shared area (table) where both agents need to arrange blocks. A conveyor belt beside human table 1 allows the robot to return a block to the human agent. 65

6.2 A view of the experiment environment, taken from the location of the camera in Fig. 6.1. The Fetch robot is positioned in its workspace between its two designated tables, as illustrated in Fig. 6.1. The nearest table in the figure, located next to the conveyor belt, belongs to the human and contains orange and green blocks. On this table, there is also a tablet on which the GUI is installed. The human agent's other table is situated at the corner of the room, containing blue and pink blocks. Four workspaces are present on the shared table, along with a light bulb. Safety tapes and cones separate the work areas of the agents. 66

6.3 **a, c, e, g:** Patterns A_1, B_1, C_1, and D_1 represent the set of patterns presented on paper sheets, requiring participants to memorize them within a 45-second timeframe before returning them to the experimenter. **b, d, f, h:** Patterns A_2, B_2, C_2, and D_2 are variations with some partially known spots, serving as cues for participants throughout the collaborative task, and they are permitted to retain these patterns until task completion to aid in recalling the first pattern. 69

6.4 State graph of a single subtask. There are six potential states for each subtask based on the actions detailed in Table 1. Initially, each subtask is in the "Initial state" state. To complete a subtask, it must transition to the "Placed correctly" state and remain in this state. The states "Assigned to robot correctly" and "Assigned to robot incorrectly" occur when the human agent respectively assigns a subtask to the robot with the correct or wrong color. The "Misplaced" state is reached when the human places a wrong color on the shared area for a subtask. The state "Assigned to human" happens when the robot assigns a subtask to the human agent. 71

6.5 A screenshot of the graphical user interface (GUI) that enables participants to convey their actions to the robot and obtain information regarding the robot's decisions and actions (refer to Table 6.2 for a list of actions). 73

6.6 Task graph of the experiment, including twenty subtasks (τ_1 - τ_{20}) and two dummy nodes representing starting (τ_0) and finishing points (τ_{21}). 74

6.7 Temporary task graph of the experiment after task allocation. Blue: Robot's tasks, Orange: Human's tasks, Cyan: Assigning tasks to the human, Red: Correcting human errors, Green: Already assigned tasks, Gray: Finished tasks 77

6.8 Estimating the human's accuracy: Transition probability, T_y 80

6.9 The robot's belief regarding the human's preference and performance (computed as the expectation of the posterior distribution) 82
6.10 Robot's belief regarding the human's performance and preference. The human first had high accuracy but then made more errors 84

7.1 Kelley's five followership styles 96
7.2 Participants' trust in the robot before collaboration and after each task . . 99
7.3 Participants' self-confidence in completing the tasks after seeing the patterns but before starting them based on a) the tasks 100
7.4 Participants' expected helpfulness from the robot 101
7.5 Participants' perception of the robot's traits and working alliance 102
7.6 Participants' perception of the robot's performance and collaboration . . . 102
7.7 Scores of user experience questionnaire 103
7.8 Team fluency and performance 104
7.9 Participants' perceived a) mental, b) physical demand, c) temporal demand, d) performance, e) frustration, and f) effort based on the tasks 106
7.10 Correlation between participants' initial trust and perceived robot traits . 107
7.11 Ranking of task difficulty by participants 108
7.12 Participants' a and b) perceived mental demand and effort, c) self-confidence, and d) relative trust (self-confidence - trust in the robot) based on the patterns (Patterns B, C, and D) 109
7.13 Number of subtasks assigned by participants to themselves in each pattern. 111
7.14 Number of blocks assigned by participants to Fetch in Tasks 1, 2, and 3 . . 112
7.15 Caption 113
7.16 Correlations related to participants' leadership/followership styles 114
7.17 The robot performance in estimating participants' preference by comparing with the participants' actual preference, gathered through interviews. The numbers above each bar show the number of participants falling into that group. 116
7.18 Participants' overall estimated leading/following preference based on the patterns. The robot's estimation of participants' preference to follow was significantly higher in Pattern B compared to Patterns C and D. 117

7.19 Subasks assigned to participants by the robot based on the patterns. The number of assigned subtasks by the robot to participants was significantly higher in Pattern B compared to Patterns C and D. 117

7.20 Color of block that were done by human and the robot as well as were assigned by the robot or human to each other 118

7.21 Correlation between estimated participants' following preference and assigned/done tasks by the robot and participants 119

7.22 The robot's estimates of the participants' preference and performance: a, b) Leading or collaborating-leading preference and high accuracy; c, d) leading preference with occasional robot support; e, f, g) leading preference with occasional errors; h) leading preference with sudden performance drop; i, j) following or collaborating-following preference 123

List of Tables

2.1 A number of cobots deployed in the industry 9
2.2 Traditional industrial robots VS. cobots [1] 9

3.1 Muir's questionnaire for humans' trust in the robot 28
3.2 Questionnaire for human perception of the robot 28
3.3 Average of participants' answers to the Muir's questionnaire about their trust in the robot before and after the experiment. Question 1: predictability, Question 2: dependability, Question 3: reliability over time, Question 4: overall degree of trust 29

5.1 Simulation results: The mean value of result after ten repetitions for each case. 49

6.1 Distribution of blocks with respect to the distance to the shared area . . . 67
6.2 The human and robot's sets of actions 72
6.3 Evaluation results: Four different enacted behaviors (one run per case) . . 81

7.1 Muir's questionnaire - measuring participants' trust in the robot [2] 96
7.2 The questions about the participants' self-confidence and expected helpfulness of the robot. Asked after seeing the patterns but before starting the task 96
7.3 NASA-TLX questionnaire asked after each task 97
7.4 A Questionnaire Assessing Working Alliance, Robot Teammate Traits, Robot's Performance, and Collaboration. Q1-Q7 adopted from [3] 97

7.5 The robot and team's overall performance. Q1-3 are adopted from [3] . . . 97
7.6 Participants' preference based on interviews. 110

Chapter 1

Introduction

Cobots, short for collaborative robots, have signified a transformative leap from traditional industrial robots, working isolated from humans, to robots that can share their workspace with their human coworkers, laying the ground to exploit the synergy of human-robot collaboration (HRC). Traditional industrial robots are known for their speed, power, accuracy, reliability, and ability to perform repetitive tasks for long hours. However, they are limited in terms of flexibility, adaptability, and performance of some tasks that require human dexterity. In contrast, although the cobots are currently slower and less powerful, mainly due to their proximity to humans and safety concerns, they are easy to install and relocate and are productive and cost-effective automation solutions for diverse work environments, even small enterprises [4]. Leveraging these capabilities and establishing a seamless collaboration, however, can be achieved only through human-aware programming of cobots [5], empowering them to learn, adapt, and work robustly.

One of the main challenges in programming cobots is enabling them to adapt to their human teammate, especially their preferences. This topic has been extensively researched in the context of human-robot interaction (HRI) and collaboration, primarily centering on considering human preferences and enhancing human satisfaction and perception of the robot [6–9]. Clearly, a high level of human perception of the robot facilitates effective long-term collaboration between humans and cobots, but a question remains: "Do only the human teammate's satisfaction and perception matter?".

Cobots are typically less expensive compared to traditional industrial robots, yet they still need to demonstrate sufficient productivity to convince business owners to invest in them [5]. In other words, the performance of the human-robot collaboration must be high enough to ensure a return on investment. Hence, team performance, in addition to human

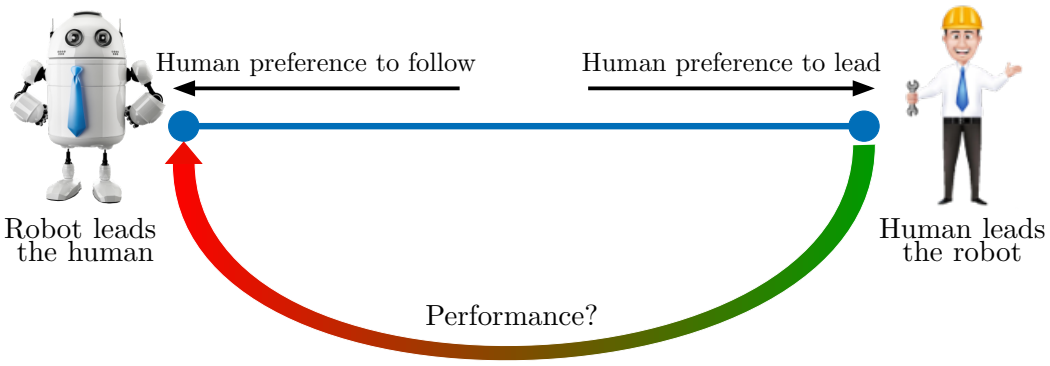


Figure 1.1: Spectrum of leading/following roles in human-robot collaboration

perception of the robot and collaboration, are two important factors to consider in cobot programming.

Both the human and robot contribute to team performance with their complementary skill, but in the context of task planning and scheduling, the robot's high computational and planning abilities allow it to take a more significant and leading role, ensuring good team performance in the short term. However, due to human presence and dynamic environment uncertainties, it is sometimes more efficient for the human interaction partner to plan for the team [5]. However, the objectives of maximizing team performance and human perception may be conflicting. This leads to the following problem, which is the focus of this work: *How do we enable a cobot to adapt to its human coworker's preferences to optimize human perception while keeping the team performance at an acceptable level?*

The fundamental idea to address this problem is to let the robot initiate the collaboration according to its plan but let the human assume the leading role if they prefer, with the robot conforming to this (leading/following) preference in pursuit of long-term performance. This adaptation, in which the human takes the lead, remains valid as long as it does not result in a significant deviation from the optimal plan, or compromise performance beyond acceptable limits.

Thus, I have proposed a framework that equips the robot to estimate and adapt its planning to

1. the human leading/following preferences, and
2. the human performance.

In this framework, in contrast to most common approaches, which strictly focus on one end of the leading/following spectrum, the cobot gradually and continuously adapts to human performance and (leading/following) preferences and can encompass the entire spectrum between acting solely as a leader or solely a follower. Fig. 1.1 encapsulates the main idea of this framework. The framework allows the robot to monitor and adapt to changes in human preferences and performance rather than relying on fixed, long-term preferences and performance. This framework can be applied to different collaborative scenarios; however, my research focuses specifically on task selection/allocation and scheduling in human-robot collaboration.

In a basic task allocation problem, agents need to be assigned tasks with associated payoffs, aiming to optimize overall team payoff. This is a widely explored topic in multi-robot teams [10], as well as HRC, involving a blend of human and robot agents [11–15]. Our work focuses on a single-human, single-robot teams, and what makes it different from conventional task allocation systems is: *the human and cobot's agency in selecting their own tasks and assigning tasks to each other instead of being assigned merely by one of the agents or a third party (e.g., a manager/central controller).*

The provided agency allows the human to demonstrate and implement their leading/following preference in their collaboration with the cobot. Subsequently, the cobot needs to estimate the latent human preference throughout the collaboration continuously. Moreover, based on the proposed framework, the cobot needs to monitor and estimate the human's preference consistently, adapt its planning accordingly, and re-assume the leading role when the human's performance is poor, even if the human agent prefers to lead the team. It is crucial to point out that the cobot's goal is not to estimate the long-term and overall human preference and performance, but it needs to estimate those changes during the collaboration. For example, consider a person with a following preference who becomes fatigued at some point and then prefers to assign more tasks to the robot, or a scenario in which a person with a high level of performance becomes confused and makes a few mistakes.

Nevertheless, knowing the general human preference and level of performance can assist the robot in forming a more accurate initial belief about them. Incorporating the robot's beliefs about human preference and performance necessitates online planning in the belief space. Thus, the robot updates its beliefs about human preference and performance at each decision step. Subsequently, the robot performs two-step task planning at each step: first, task allocation considering human preference and performance, and second, task scheduling.

1.1 Research Questions

In the first research question, we delve into how the human teammate perceives various robot behaviors, spanning from self-serving to purely assistive. Addressing this question provides insights that illuminate the design of a framework for human-robot collaboration. This framework would prioritize adaptability to the human teammate's preferences and performance, rather than rigidly adhering to the human agent's preference and lead.

RQ1. How does the robot's level of cooperation influence the human's perception of the robot (e.g., trust and satisfaction)?

After addressing the first research question, the second research question will focus on developing a task planning framework that enables the robot to adaptively plan based on the task state and the human agent's performance and preference.

RQ2. How can we design a planning architecture that takes into account both human preference and performance and enables the robot to update the task state based on both agents' actions?

After proposing the adaptive task planning framework, we need to evaluate the practicability of the proposed framework in terms of implementation in a real scenario, improving team performance, and maintaining a high level of human users' perception of the robot and collaboration. Hence, we will explore the following research questions:

RQ3. Does the proposed adaptive task planning for the robot improve overall team performance in human-robot collaboration?

RQ4. To what extent can the proposed task planning framework enable the robot to adapt to the human agent's performance and preferences successfully?

RQ5. How does collaboration with the robot influence the human agent's actions and their perception of the robot and collaboration?

RQ6. To what extent do the human agent's perception of the robot and collaboration impact their interaction with the robot?

Additionally, I will explore how the human agent's leadership and followership styles can be related to their collaboration with their robot teammate. Therefore, the last research question of this thesis is:

RQ7. Is there any correlation between the human agent's leadership/followership styles and their collaborative approach with the robot?

These questions shed light on the strengths and limitations of the task planning method and pave the way for future research directions.

1.2 Contributions and Organization

The contributions and organization of this thesis are as follows.

In Chapter 2, we first provide literature on cobots and their applications in industry, and trends in research studies. Then, we focus on the problem of leading/following in human-robot collaboration to see its importance in real applications and how the existing work has addressed this problem. Subsequently, we specifically focus on task allocation in human-robot collaboration and the effect of human preference on it. We also provide a brief literature on adaptation methods used in the area of human-robot interaction and collaboration.

Chapter 3 provides the procedure and results of an online user study addressing the first research question, RQ1. This study involves a scenario with a single human and a robot, requiring the two parties to collaborate to accomplish a given task [16]. We considered three robot strategies: prioritizing the human (Strategy 1), prioritizing the robot (Strategy 2), and balancing both (Strategy 3). The work presented in this chapter is part of the following publication [16]:

- A. Noormohammadi, A. Dahiya, A. M. Aroyo, S. L. Smith, and K. Dautenhahn, "The effect of robot decision making on human perception of a robot in a collaborative task - a remote study," in Proceedings of the 9th International Conference on Human-Agent Interaction, 2021, p. 423–427.

In Chapter 4, building upon the insights gathered from the initial study, addressing RQ2, a task planning framework and algorithm are proposed. These equip the robot, updating the task state based on the agents' actions, to estimate and adapt its planning to both 1) the human leading/following preferences and 2) the human performance. The content discussed in this chapter is part of the following publications [17, 18]:

- A. Noormohammadi-Asl, A. Ayub, S. L. Smith, and K. Dautenhahn, "Task selection and planning in human-robot collaborative processes: To be a leader or a follower?"

in 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 2022, pp. 1244–1251.

- A. Noormohammadi-Asl, S. L. Smith, and K. Dautenhahn, "To lead or to follow? adaptive robot task planning in human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01483

Chapter 5 delves into implementing the proposed task planning framework and algorithm in a collaborative scenario within a simulation environment. This chapter considers a simplified human model and evaluates the algorithm's adaptation to various cases of simulated human leading/following preferences and performance. This chapter will address part of RQ3 and RQ4. The content of this chapter is primarily derived from the following publication [17]:

- A. Noormohammadi-Asl, A. Ayub, S. L. Smith, and K. Dautenhahn, "Task selection and planning in human-robot collaborative processes: To be a leader or a follower?" in 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 2022, pp. 1244–1251.

In Chapter 6, the design and setup of a user study involving a collaborative scenario with a real mobile manipulator robot are explained. This chapter also presents the results of a system evaluation study based on various behaviors enacted by the experimenter before recruiting participants. This chapter continues the addressing of RQ3 and RQ4 that started from the previous chapter. The content of this chapter is part of the following publications [18, 19]:

- A. Noormohammadi-Asl, S. L. Smith, and K. Dautenhahn, "To lead or to follow? adaptive robot task planning in human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01483
- A. Noormohammadi-Asl, A. Ayub, S. L. Smith, and K. Dautenhahn, "Adapting to human preferences to lead or follow in human-robot collaboration: A system evaluation," in 2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 2023, pp. 1851–1858.

Chapter 7 discusses the results of a user study involving 48 participants. This chapter delves into the findings from three main aspects: 1) Participants' perception of the robot, tasks, and collaboration, 2) Participants' actions and decisions, and 3) the robot's task

planning and actions. This not only completes answering of RQ3 and RQ4 but also incorporates RQ5 and RQ6. Additionally, addressing RQ7, this chapter explores the correlation between participants' leadership/followership styles and their collaboration style with the robot. The content of this chapter is derived from the following publications [18, 20]:

- A. Noormohammadi-Asl, S. L. Smith, and K. Dautenhahn, "To lead or to follow? adaptive robot task planning in human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01483
- A. Noormohammadi-Asl, K. Fan, S. L. Smith, and K. Dautenhahn, "Human leading or following preferences: Effects on human perception of the robot and the human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01466

Finally, Chapter 8 provides an overview of the research presented in this thesis and explores possible avenues for future research.

Chapter 2

Related Work

2.1 Cobots

A US patent [21], titled "Cobots", first introduced a type of robot that is "[a]n apparatus and method for direct physical interaction between a person and a general purpose manipulator controlled by a computer." Initially developed as a project by the GM Robotics Center and Professors Edward Colgate and Michael Peshkin in 1994 [22], cobots now have a market that is predicted to grow from $981 million in 2020 to $7,972 million by 2026 [23]. Cobots have sparked a new perspective on automation and have drawn the attention of corporations, such as automotive industries, and even small businesses [5].

Although we are most familiar with cobots' industrial applications, some modern applications utilize them for unconventional purposes such as surgery and healthcare [24–26], service (bartending [27], kitchen assistance [28]), and flight control [29]. Table 2.1 lists the name, manufacturer, and specifications of some cobots now used in the industry. Table 2.2 includes a number of cobot characteristics that differentiate them from traditional industrial robots. These characteristics, and above all safety concerns, have resulted in a growing trend toward the design of human-robot collaborative systems.

2.2 Collaborative Scenarios and Systems

A variety of definitions have been proposed for collaborative interaction [5], some of which differentiate among responsive, cooperative, collaborative, and supportive interaction [30].

Robot	Company, Country	DOF
YuMi	ABB, Sweden-Switzerland	Dual-arm, 7-7
ROBERTA	ABB, Sweden-Switzerland	6
BioRob Arm	Bionic Robotics, Germany	4-5
UR3,5,10	Universal Robotis, Denmark	6
JACO2, MICO2	Kinova, Canada	6
LBR IIWA	KUKA, Germany	7
Baxter	Rethink Robotics, Germany	Dual-arm, 7-7
Sawyer	Rethink Robotics, Germany	7
Nextage	Kawada Industries, Japan	6 for each arm, 2 for neck, 1 for waist
P-Rob 3	F&P Personal Robotics, Switzerland	6
APAS	BOSCH, Germany	6
Speedy 10	MABI, Switzerland	6

Table 2.1: A number of cobots deployed in the industry

Traditional Industrial Robots	Cobots
No safe interaction	Safe interaction with humans
Fixed and difficult installation	Easy to relocate and install
Difficult to change task	Easy to change task
Difficult to teach	Easy to teach
Rare interaction	Frequent interaction
No shared workspace	Shared workspace
Profitable for large and mid-size enterprises	Profitable for large, mid-size, and small enterprises
Small/big and fast	Small and slow

Table 2.2: Traditional industrial robots VS. cobots [1]

However, I ground my research on a broad definition, provided by [31] and the International Federation of Robotics (IFR) [32], which considers any scenario without a fence around a robot as a collaborative scenario.

2.2.1 Types of collaboration with cobots

Based on our definition of collaborative scenarios and the categorization in [31–33], the collaborative scenarios are divided into four groups:

- Coexistence — There is no shared space between the cobot and the human, and the two work independently in their own workspaces. However, the robot is not fenced and works in close proximity to the human.
- Sequential collaboration — Both cobot and human work in the workspace, but there is a time dependency between their actions, and their movements are sequential.
- Cooperation — Both cobot and human simultaneously work in the workspace and on the same piece without any time or task dependency. The cobot, however, is aware of the human's spatial position and plans accordingly to move safely.
- Responsive collaboration — Both robot and human take part in a process on the same work piece, and the cobot responds to the human's actions and movements.

Despite the above-mentioned explicit categorization, the scenarios that I will discuss in the next chapters do not fall cleanly under one of the four categories. In other words, the scenarios as designed are coexistence, sequential collaboration, and responsive collaboration, with some overlap, in the sense that:

1. the cobot(s) and human work independently in their own workspaces;
2. there is time dependency between their actions, and one cannot work in the shared space at the same time as the other;
3. the robot needs to be responsive and consider the human's actions and act accordingly to improve team performance and efficiency.

2.2.2 Commercial and industrial applications

Manufacturing companies, especially small and mid-sized enterprises, are becoming more interested in deploying cobots due to the improved productivity, flexibility, quality, safety, ergonomics, and trust that they provide and their potentially critical role in industry 4.0 and 5.0 [34–37]. Based on the International Federation of Robotics's reports [38], the automotive industry and electrical/electronics industries are the largest adopters of industrial robots by 28% and 24% of total installations, ahead of the metal and machinery (12%), plastics and chemical (5%), and food and beverages (3%) industries. The following provides examples of cobot applications in a number of industries. Herein, I focus on only one of the leading companies in collaborative robotics, Universal Robots, and include a few cases in which this company's products and solutions have been deployed. More cases can be found on the company website.

- **Automotive industry:** The Nissan Motor company has deployed two lines of UR10 robot arms at its Yokohama plant for loosening bolts on cylinder head cam brackets and for installing engine block intake manifolds [39]. Both tasks need to be done in collaboration with real human workers without safety fences. UR10 cobots have also been used in one of the major product lines of the Ford Motor Company in Romania for greasing camshafts, filling engines with oil, and performing quality inspections [40].

- **Electrical/Electronics industry:** The Schneider Electric Bulgaria plant, which focuses on mini circuit breakers (MCB), deploys UR5 cobots for automating the process of picking and placing power breakers and boxes [41]. PT JVC Electronics Indonesia (JEIN), part of the JVCKENWOOD Group, deploys UR3 cobots to perform screwing, soldering and, pick and place tasks [42].

- **Metal and machinery:** The Yokota Corporation, in Japan, is active in designing and manufacturing bearing races; factory automation (FA) equipment; machines for assembly, packing, and inspections; and machine tools. Deploying UR5 cobots has allowed them to reduce labor costs for two workers and improve productivity by 20% [43]. Fluidics Instruments, a Dutch manufacturer of oil burner components, deploys seven UR3 cobots in their assembly line that produces 1,000 oil nozzles per hour [44]. Swiss kitchen manufacturer Franke uses UR5 cobots for gluing blocks in the sink production line [45].

- **Plastics and chemical:** Trelleborg sealing solutions, at its Danish production site, has deployed 42 UR robots to tend mainly CNC machines [46]. Glidewell Laboratories

uses five UR5 robots in the process of picking a crown from dispensers, placing it in the milling lathe, picking it back out, and placing it on a conveyor after the 10 minute milling cycle [47].

- **Food and beverages:** Atria Scandinavia, manufacturers of vegetarian and gourmet products, has integrated two UR5 robots and one UR10 robot to improve and speed up their production [48]. Nordic Sugar uses three UR5 cobots for automating the analysis of raw material samples [49].

- **Medical and Cosmetics:** the Albrecht Jung company uses eight UR cobots for pick-and-place, screw insertion, parts assembly, and packaging [50]. RNB Cosméticos has adopted six Universal Robots UR10 robot arms in its packing and packaging plant to perform end-of-line palletizing tasks [51].

- **Cobots vs Covid-19:** Researchers at Nanyang Technological University (NTU) have introduced the eXtremeDisinfection roBOT (XDBOT) for disinfection and deep cleaning purposes. XDBOT consist of a UR5 cobot on a mobile platform. Lifeline Robotics, in collaboration with robotics researchers from the University of Southern Denmark (SDU), have designed the world's first autonomous throat swabbing robot launched for medical testing. This robot is based on a UR3 cobot arm with a custom 3D-printed end-effector [52].

2.2.3 Research studies

Although most HRC scenarios in industry fall under the categories of "coexistence" and "sequential collaboration" [32], recent research focuses on cooperation and responsive collaboration. These studies mainly revolve around topics that are rarely seen in industrial settings: improving humans' perception of the cobots, learning from the human teammate, modeling the human teammate, providing effective communications, and decision making and planning [5, 53]. Figure 2.1 shows the growing research trend in human-robot collaboration (cooperation) since the first time the term "cobot" appeared in 1996. In addition to the mentioned categories, we can also view these studies from a different perspective that classifies them as:

- Single-Human Single-Robot scenarios — This team composition, involving only one human and one robot, is the most conventional one in human-robot collaboration scenarios [5, 53, 54].

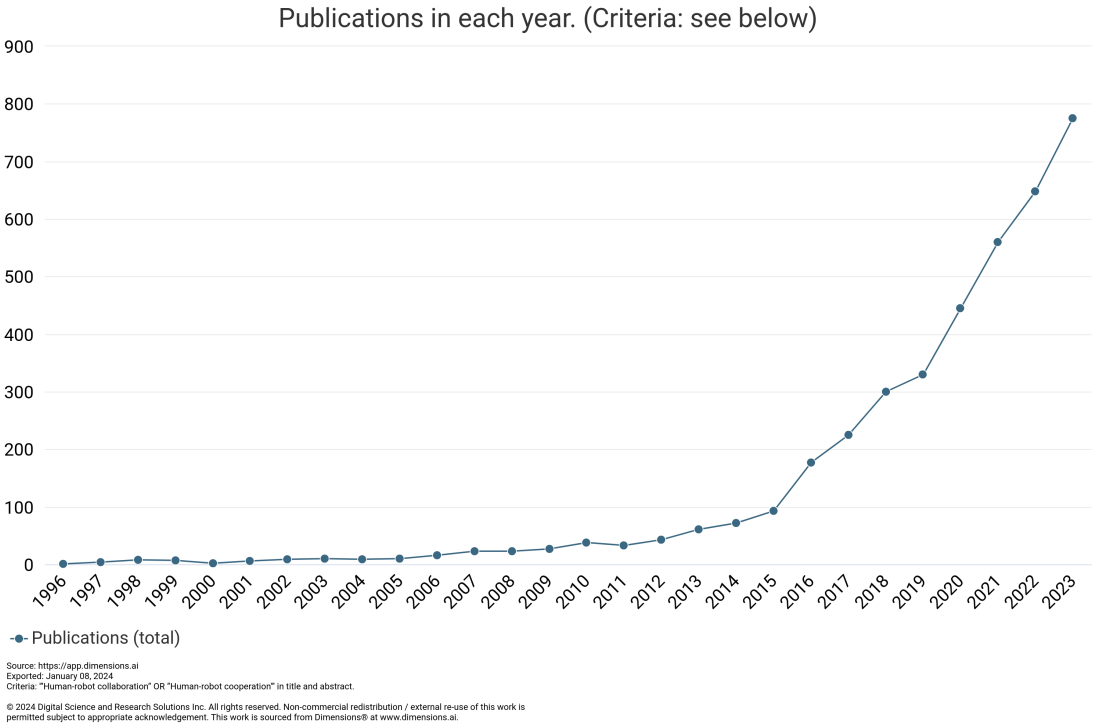


Figure 2.1: Number of publications on the topic of human–robot collaboration (cooperation) from 1996 to 2023; **Search Criteria:** "Human-robot collaboration" OR "Human-robot cooperation" in title and abstract; **Source:** https://app.dimensions.ai

- Single-Human Multi-Robot scenarios — These scenarios comprise one human and at least two robots operating in a fenceless environment. In most scenarios in the literature, such as search and rescue [55–57] and warehouse operation [55, 58], the human plays the role of a supervisor or operator. However, in some other scenarios, there is no supervisory role, and the human and robots collaborate on a joint task (e.g., navigating in an environment [59]) or coexist while working safely on different tasks [60] (e.g., some industrial settings).

- Multi-Human Single-Robot scenarios — In these scenarios, the team consists of one robot and more than one human. These scenarios include applications in which the robot has to serve or interact with a group of human teammates [61, 62] or has to be controlled by them, either because it cannot control itself or because its control is beyond a single human's ability [63, 64]. Again, as an example in industrial settings, we can mention coexistence and sequential collaboration scenarios.

- Multi-Human Multi-Robot scenarios — These scenarios, the most complex ones, involve more than one human and robot. This team composition's possible applications can be seen in search and rescue, exploration, and service tasks [65–67]. Coexistence and sequential collaboration scenarios can also be considered as industrial applications of this type of team composition.

2.3 Cobot Programming

In a scenario where a robot works in isolation or without any direct interaction with humans, the control objectives are based on the robot's performance metrics — time, energy, traversed distance, covered area, etc. However, in HRC, humans may take a collaborative or even leading role, and the control objectives thus need to be modified to meet the humans' comfort needs, both physical and mental. Many studies have focused on the human teammate's physical comfort, such as their ergonomic posture [68] and fatigue levels [69]. Trust [70], intention [71, 72], knowledge of the task [73, 74], and acceptance of cobots are some key topics of interest in HRC specifically, and human-robot interaction (HRI) in general. Programming robots to consider humans' mental states in their decision making can significantly enhance the quality and efficiency of collaboration and humans' perception of robots. For instance, having the human's knowledge of a task helps robots detect where, when, and whom to assist [74], and being able to predict human intentions will help robots act accordingly [75]. As another example, in [76], human trust in the robot is considered in

robot planning to prevent unnecessary human interventions so that the human can focus on their task.

2.3.1 Trade-offs between human and task objectives

Clearly, deploying cobots in industry is economically justifiable only when they can offer long- or short-term benefits and returns on investment [5]. Hence, HRC systems have to not only improve and maintain humans' perception of robots at a reasonably high level, but need to improve task efficiency as well. In [77], a probabilistic approach was adopted to predict human actions in order to optimize a cobot's help and reduce wait times. In [78], workloads per subtask were minimized by distributing assembly sub-tasks between a human and a cobot. In another study [79], an optimal task distribution between humans and robots was proposed to minimize the completion time of a realistic production process.

Most HRC scenarios have considered human-centered designs that optimize human objectives and task objectives that are aligned with the human's goals. However, the problem is more challenging in cases where these two objectives conflict and pull in opposite directions, or where the human's plans and actions are not optimal. To cope with these problems, the robot needs to balance the two objectives and adopt a policy to guide the human to make optimal or compromise policies and decisions. Probabilistic and game theory approaches are commonly used to achieve a mutual benefit, or to guide (persuade) adaptable agents toward the (sub)optimal actions [80–82].

For these cases, little is known about the influence of the robot's strategy on human perception of the robot. Some studies have focused on the effects of a robot's blame attribution in HRC after the collaboration or after doing its actions [83–85]. These studies have considered three different personalities where the robot blames either itself, the human, or the human-robot team. What we are interested in this thesis is the robot's strategy and how it plans and makes decisions. The authors in [86] have investigated how a robot's prosocial and selfish behaviors in a collaborative game can affect the human's perception of competence, responsibility, trust, and preference for future collaborations. However, our study also considers a balanced behavior to compromise between the robot and the human's objectives and priorities.

2.4 Task Allocation and Adaptation in HRC

While human-robot collaboration is a relatively recent topic, discussions regarding human automation have persisted for many years. The emergence of automation systems has gradually supplanted humans in various control systems, ranging from simple to complex ones. Various contributing factors (e.g., human preference and abilities) come into play when determining the level of autonomy within these systems, which can include manual control, with human operators as active controllers, or semi-autonomous and fully autonomous systems, where human operators assume a more supervisory role. In the context of HRI, the results of a user study in [87] show that, in different scenarios, participants mostly chose a semi-autonomous robot, allowing them to be more in control. In [88], the researchers have considered the problem of being in control or under control in a warehouse order-picking scenario. Their study shows a higher productivity when the human leads and a higher accuracy when the human follows the robot.

On the contrary, "6 River Systems," a company actively involved in collaborative robotics and warehouse automation, asserts that associates' self-control over their pace results in reduced speed and efficiency [89]. To address this issue, they have developed a "system-directed picking" approach in which robots direct associates through their tasks. As another industrial example, we can name the robotic picking assistive system in DHL warehouses, where the robot (provided by Locus Robotics) guides the human workers to the pickup location [90]. In these systems, the robot is responsible for planning tasks, while the human handles picking up items. In what follows, we will investigate related research regarding task allocation among agents and the consideration of human preferences in task allocation and planning.

2.4.1 Task Allocation

Task allocation in HRC involves a suitable allocation of tasks to the human and robot agents and finding a proper chronological order for completing tasks based on problem-dependent decision factors and constraints. We can categorize task allocation approaches into two main groups: offline and online.

Offline Task Allocation

In offline task allocation, typically, the goal is to assign tasks based on prior knowledge of the suitability of agents for each task. As a measure of suitability, one can consider

the proportionality of agents' abilities and constraints to the tasks' requirements and constraints. After determining suitability measures, the task allocation can be done by an expert [91, 92], via simulation studies [93, 94], or through mathematical modeling and optimization [14, 95]. In [93], after deciding possible suitable agents for each task, different cases are evaluated by simulation, and the one with the best utility value is selected. Similarly, in [94], multiple criteria are calculated through simulations, and a depth-first search algorithm is employed to find an optimal solution. In [14], the authors designed a human capability-based cost function to minimize human risk factors. Then, they applied their method in a user study, using the A* algorithm for role assignment. In [95], the problem of disassembly sequence planning is formulated as an optimization problem to minimize the disassembly time while considering resource and safety constraints.

Online Task Allocation

Offline task allocation approaches force the human-robot team to adhere to the (optimal) plan obtained offline. However, in many real-world scenarios, uncertainties that might arise due to individual preferences and behaviors, alterations in the workspace, and changes in task requirements question the applicability of these approaches. Online task allocation methods aim to cope with this limitation by endowing the system with online replanning abilities. Typically, these methods fall into two groups. The first group of approaches, similar to the offline methods, finds an optimal task allocation, and the agents must follow the obtained plan. However, these methods are able to reactively replan online when the current task allocation is not valid anymore [12, 13, 96, 97].

The second group of methods mainly relies on the human agents' decisions, and the robot agents play more of a supportive role by adapting to the human agents. These methods do not create a fixed allocation requiring agents to adhere to it. In these methods, the human agents can make their own decisions and manage arising uncertainties, and the robot agents need to adapt their decisions proactively. In [98], the robot infers the human preference by a two-stage clustering approach and provides the parts for the human in an assembly task. Other work proposed a real-time decision-making mechanism for a cobot based on the human's short-term and long-term behaviors [99]. In [100], the robot infers human preferences and chooses one of the three operation modalities: 1- the robot plans, 2- the human plans, and 3- the robot adapts.

Human Preferences in Task Allocation

As discussed in the preceding paragraph, it is crucial for the robot to proactively infer human preferences when taking on a supportive role. However, incorporating human preferences in situations where the robot or a central unit (e.g., a manager, central computer, etc) is actively involved in task allocation and planning is an aspect that is often underestimated. This aspect, central to our research, has been demonstrated as a critical element in enhancing humans' positive perception of the robot [100–102]. In [101], different scenarios were explored in which the manager, the robot, or the participants themselves assign tasks. In [102], human participants' task type preferences are considered while balancing the assigned workloads. In [100], the system can adjust to user preferences, enabling the human partner to issue commands, or take a more passive role and seamlessly transition between different modes as needed. Nonetheless, most of these studies have considered human agents' preferences during offline planning, often by direct inquiry. What is notably missing is an assessment of participants' performance to determine whether their preferences and performance align with the team's overall performance or potentially hinder it.

2.4.2 Adaptation

As noted in the task allocation literature, robot adaptation is an integral ability for a robot to collaborate smoothly and efficiently with humans. This adaptation can be achieved through manually teaching the robot by experts (i.e., how to collaborate effectively while adhering to human preferences) [103, 104], equipping the robot with the ability to learn and adapt itself autonomously, or a combination of both. Manual instruction by experts, while fast and effective, faces practical limitations, including time consumption, challenges in conveying nuanced instructions, and scalability issues. [105].

Learning-based adaptation approaches typically employ supervised or unsupervised learning-based algorithms. The latter involves identifying decision factors (features) influencing human behavior and collecting behavioral data, which is then enriched through expert annotation or participant surveys [106–108]. The former, however, enables machines or robots to learn human preferences by observing their behavior and actions without the explicit need for annotation or data labeling [98, 109–111]. Some research also leverages both learning methods and experts' knowledge [105, 112].

However, these studies often emphasize the robot's adaptation to the human agent. In contrast, mutual adaptation — where both the human and robot adjust based on each

other's actions and feedback — can be essential in human-robot collaboration. The authors of [80] explored mutual adaptation between humans and robots, where the robot initially guides the adaptable human but may unilaterally adapt if the human insists on their poor or inadequate performance or decisions.

Chapter 3

Evaluation of Robot Strategies in HRC: A Remote Study

In this chapter, focusing on our first research question (RQ1), we investigate how a cobot's strategy influences the human perception of the robot and its collaboration. In the experimental scenario, with a single-human single-robot collaborating, the following three different strategies are adopted: the robot either prioritizes the human teammate, prioritizes itself, or uses a balanced strategy (i.e., considers both human and robot objectives). Then, based on the data collected from the questionnaires and participants' actions, we analyze and discuss the influences of each strategy on their perception of the collaboration and the robot. The result of this study will help us determine whether a balanced strategy can maintain both team performance and human perception of the robot at a relatively high level.

3.1 Hypothesis

This study seeks to learn how a robot's decision making in a collaborative task can affect its human teammate's perception of the robot. To do so, we consider three different strategies for the robot:

Strategy 1 (SH) - The robot completely prioritizes the human's task and objective;

Strategy 2 (SR) - The robot completely prioritizes its own task and objective;

Strategy 3 (SHR) - The robot adopts a compromising strategy that balances its and the human's priorities.

The first two strategies are at opposite ends of the planning spectrum, where we expect to see the highest and the lowest levels of satisfaction and trust. What we are trying to understand through this experiment is to what extent the third strategy can maintain the human's trust and satisfaction while optimizing the robot's (or team's) objectives.

Based on these three strategies, we developed the following hypothesis:

Hypothesis 1: The human's

a. trust in the robot,

b. perception of the robot's performance,

c. perception of the robot's collaboration, and

d. willingness to collaborate with the robot,

will all be the highest when the robot uses strategy 1 (SH), second highest when using strategy 3 (SHR), and lowest when using strategy 2 (SR).

3.2 Experimental Method

For this experiment, we have designed an online interactive simulation environment in which the human and the robot need to organize a certain number of objects.

3.2.1 Materials and setup

We chose Webots, a free, open-source, and multi-platform robot simulator for the simulation environment. We used the UR5 robot, a well-known 6-degree-of-freedom cobot manufactured by Universal Robots. UR5 is a safe, agile, easy-to-program platform used extensively in various applications.

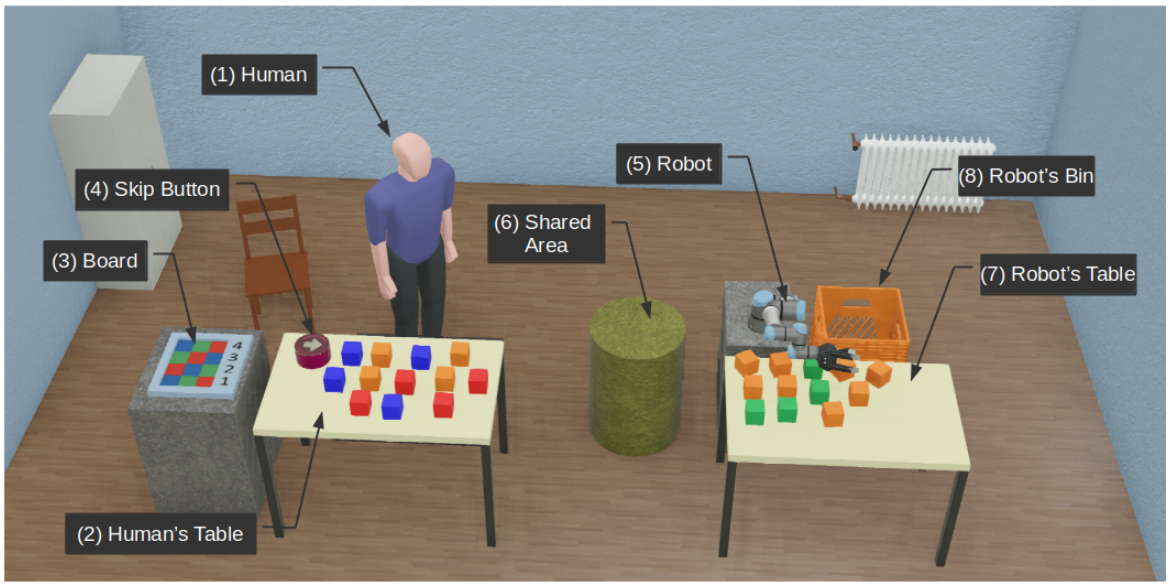


Figure 3.1: Online experiment environment

3.2.2 Experiment task

The human participants collaborate with a robot to organize certain objects. As shown in Figure 3.1, the human (1) stands beside their table (2). Their task is to put the corresponding colored blocks on the squares marked on the board (3) beside them in the shortest time. Participants have to fill each row before moving to the next one, starting with the first one and ending at the fourth. The colors of the blocks and the squares have to be matched. For example, a blue block has to be placed on a blue square. The robot (5) also needs to put orange objects in the bin. This is a turn-taking activity, meaning that the robot takes the first turn, and when it is done, the human chooses their action; then, the robot makes its next move, and so on.

There are four red, four blue, and four orange blocks on the human's table (2). Participants need blue, red, and green blocks to complete the board (3), but they only have access to blue and red ones on their table. Thus, the robot has to place the green objects on the shared space (6), the green table beside their table, so that participants can use them to complete the board. Participants also need to place the orange blocks from their table on the shared area (6) so that the robot can take them and put them in its bin (8). The shared space has a capacity of two objects. There is also a skip button (4) on the participants' table, which they can press if want to skip their turn.

The participants' goal is to fill the board in the shortest time, while the robot's goal is to put all the orange blocks in the bin. The task finishes when all objects are removed from the tables and the shared space. Based on the explained scenario, there is a set of possible actions for both the human and the robot, from which they can choose a feasible action in each turn.

Human's possible actions A^h:

1. Pick up a blue block from the table and place it on the board.
2. Pick up a red block from the table and place it on the board.
3. Pick up a green block from the shared space and place it on the board.
4. Pick up an orange block from the table and place it on the shared space.
5. Skip their turn.

Robot's possible actions A^r:

1. Pick up an orange block from the table and place it in the bin.
2. Pick up an orange block from the shared space and place it in the bin.
3. Pick up a green block from the table and place it on the shared space.
4. Skip its turn.

3.2.3 System Modeling and Planning:

System Modeling:

The complete state of the system at any time step k is defined by three components: the state of the board $\mathbf{x}_k^b$, the state of the shared space $\mathbf{x}_k^s$, and the state of the orange objects on the tables $\mathbf{x}_k^o$.

State of the board: To represent the state of the board, it is enough to know which row a participant is currently working on, $row_k \in \{1, 2, 3, 4\}$, and whether the blue ($b_k \in \{0, 1\}$), red ($r_k, \in \{0, 1\}$), and green ($g_k \in \{0, 1\}$) squares in this row are filled (a value of 1) or

empty (a value of 0). Thus, $\mathbf{x}_k^b = [row_k, b_k, r_k, g_k]^T$ is the state vector of the board. The simulation environment is designed in such a way that when participants click on the blocks they intend to move, the blocks are automatically moved to their designated place. This is done to make any mismatched placement (e.g., placing the red block on the blue square of the board) impossible.

Orange objects: As the experiment starts, there are respectively, four and eight orange objects on the human's and the robot's tables. The vector $\mathbf{x}_k^o = [o_k^{rr}, o_k^{rh}]^T$ is the state of the orange objects on both tables, and o_k^{rr} and o_k^{rh} are the numbers of orange objects on the robot's and the human's tables, respectively.

Shared space: The robot (human) can place green (orange) objects on the shared space or pick up orange (green) ones from it. As mentioned, the shared space has the capacity of two blocks. $\mathbf{x}_k^s = \left[d_k^r, d_k^h\right]^T$ represents the state vector of the shared space, in which $d_k^r, d_k^h \in \{0, 1, 2\}$ denote the number of the robot's and the human's blocks on the shared space. Considering the shared space capacity, $d_k^r + d_k^h \leq 2$.

Thus, the state of the system $\mathbf{x}_k \in X$ can be represented by the augmented state

$$\mathbf{x}_k = \begin{bmatrix} \mathbf{x}_k^b \\ \mathbf{x}_k^o \\ \mathbf{x}_k^s \end{bmatrix} = \left[row_k, b_k, r_k, g_k, o_k^{rr}, o_k^{rh}, d_k^r, d_k^h\right]^T. \tag{3.1}$$

The experiment starts from the initial state, $\mathbf{x}_0 = \left[1, 0, 0, 0, 8, 4, 0, 0\right]^T$, where all objects are on the tables, and ends when the system falls into the goal state (absorbing state) $\mathbf{x}^g = [4, 1, 1, 1, 0, 0, 0, 0]^T$. In the goal state, the human's board is completed, and there is no object on the tables and the shared space.

For modeling this experiment, we consider the following deterministic model

$$\mathbf{x}_{k+1} = f\left(\mathbf{x}_k, \mathbf{u}_k\right), \quad k = 0, 1, 2, \ldots, \tag{3.2}$$

where $\mathbf{u}_k \in U$ represents the vectors of the system input at time step k. Let A^h and A^r be the sets of all possible actions for the human and the robot in the experiment. We denote the sets of feasible actions at stage k with $A_k^{fh} \subset A^h$ and $A_k^{fr} \subset A^r$. The input vector, $\mathbf{u}_k = \left[a_k^r, a_k^h\right]^T$, is comprised of the human's action, $a_k^h \in A_k^{fh}$ and the robot's action, $a_k^r \in A_k^{fr}$. U is the set of all possible system input vectors $\mathbf{u}_k$

We assume at any time step k, first, the robot takes its action a_k^r, changing the system state to an intermediate state $\mathbf{x}'_k$. The human observes the robot's action and then decides their action a_k^h, transitioning the system to the next state $\mathbf{x}_{k+1}$. Thereafter, the robot takes

its action for time step $k+1$ and so on. This turn-taking between the robot and the human continues until the system reaches the goal state. Figure 3.2 shows how the system evolves as the robot and the human perform their actions. For example, $\mathbf{x}_n = [2, 1, 1, 0, 6, 3, 1, 0]^T$ means that the human is currently working on the second row of the board and has already placed blue and red blocks on the second row. He has also placed an orange block on the shared space. The robot has also removed two orange objects from its table. Assuming $\mathbf{u}_k = [3, 3]^T$, the robot first picks up a green block from its table and places it on the shared space, $\mathbf{x}'_n = [2, 1, 1, 0, 6, 3, 1, 1]^T$. Then, the human picks up the green block from the shared space and places it on the board, $\mathbf{x}_{n+1} = [3, 0, 0, 0, 6, 3, 1, 0]^T$.

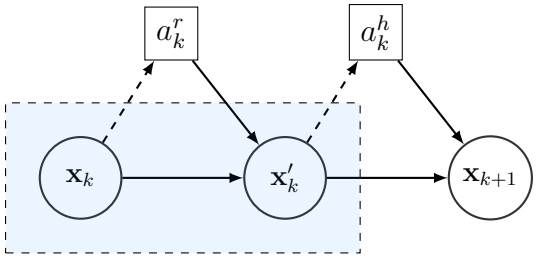


Figure 3.2: State transition

Robot planning:

After each robot action, the robot receives a real-valued reward $r(\mathbf{x}_k, \mathbf{u}_k)$. In this experiment, the human policy is not known in advance, and no online or offline estimation is made during the robot planning. Instead, it is assumed that the human takes the optimal action at every step, i.e., the one that helps maximize the accumulated reward. Define the optimal policy $\pi^*: X \to U$ that maximizes the discounted sum of rewards as,

$$\pi^*(\mathbf{x}) = \arg\max_{\pi \in \Pi} \left[\sum_{k=0}^{\infty} \gamma^k r\left(\mathbf{x_k}, \pi(\mathbf{x}_k)\right) \,\middle|\, \mathbf{x}_0 = \mathbf{x} \right]. \tag{3.3}$$

where Π is the set of all admissible policies $\pi : X \to U$. During the experiment, at each time step, the robot chooses its action according to the policy π^*, regardless of the human's actions.

The reward is only the function of the current state of the experiment and the robot's action, and is a weighted combination of two scalar functions $\text{robot\_pref}(\mathbf{x_k}, a_k^r)$ and $\text{human\_pref}(\mathbf{x_k}, a_k^r)$. These two functions reflect the extent to which the action a_k^r aims to satisfy the robot's

and the human's preferences:

$$r\left(\mathbf{x}_k, \mathbf{u}_k\right) = r\left(\mathbf{x}_k, a_k^r\right) = (1-q) \times \text{robot\_pref}(\mathbf{x_k}, a_k^r) + q \times \text{human\_pref}(\mathbf{x_k}, a_k^r), \tag{3.4}$$

where $q \in [0, 1]$ is a constant value indicating the importance of the human preferences. The three strategies can be implemented by simply setting the value of q. For strategy 1 (pure-assistive), we set $q = 1$; for strategy 2 (self-serving), we set $q = 0$; and for strategy 3 (balanced), we set $q = 0.5$. The functions $\text{robot\_pref}(\mathbf{x_k}, a_k^r)$ and $\text{human\_pref}(\mathbf{x_k}, a_k^r)$ are defined as below:

$$\text{human\_pref}(\mathbf{x_k}, a_k^r) = \begin{cases} 3, & \text{if } G = 0,\ d^h + d^r < 2,\ a^r = 3 \\ 2, & \text{if } G = 0,\ d^r = 2,\ a^r = 2 \\ 1, & \text{if } G = 1,\ a^r = 4 \\ 0, & \text{otherwise,} \end{cases} \tag{3.5}$$

and

$$\text{robot\_pref}(\mathbf{x_k}, a_k^r) = \begin{cases} 3, & \text{if } G = 0,\ a^r = 1,\ d^r = 0 \\ 2, & \text{if } G = 0,\ a^r = 1,\ d^r > 0 \\ 3, & \text{if } G = 0,\ a^r = 2,\ d^r > 0 \\ 1, & \text{if } G = 1,\ a^r = 4 \\ 0, & \text{otherwise,} \end{cases} \tag{3.6}$$

where G denotes whether the system is in goal state ($G = 1$) or not ($G = 0$). The function human_pref encourages the robot to give green objects (if any are available) to the human when the shared space is empty. Furthermore, when the shared area fills up with orange objects, the robot is rewarded by removing one orange object and making an empty space. The function robot_pref encourages the robot to first pick up orange objects from the shared area (if any are available) and then pick up orange objects from its table. In both functions, when the system falls into the goal state, a positive reward is generated to make the robot skip its turn.

This optimization problem can be solved offline by methods like Dijkstra's, A*, and policy iteration algorithms. In the real experiment, the robot determines and executes its action through the obtained policy, π^*, and then waits for the human's action.

3.2.4 Procedure and steps

Participants:

We recruited 50 participants, including 48 graduate and 2 undergraduate students, by using email advertisements. The cohort for each condition consisted of:

- Condition 1: 17 participants (9 men and 8 women) with an average age of 28.71 $\pm$ 9.59 years
- Condition 2: 17 participants (10 men and 7 women) with an average age of 28.07 $\pm$ 4.18 years
- Condition 3: 16 participants (9 men and 7 women) with an average age of 27.56 $\pm$ 4.76 years

The participants were randomly assigned to one of the conditions in a way that has a balance across the conditions. After being recruited and giving their consent, participants needed to take part in three phases of the experiment.

The study received ethics approval from The study received ethics approval from the University of Waterloo Research Ethics Board.

Phase 1 - Pre-experiment Questionnaire: Participants answered four questions regarding their age, gender, education, and field of study. Then, they read a page describing the experiment and required instructions and watched a short video showing a part of a sample collaboration to familiarize themselves with the setting. Afterward, they were asked to answer four questions about their experience in building, controlling, and interacting with robots, to reveal their level of familiarity with robots. Finally, we used Muir's questionnaire [2] with a seven-point Likert scale to evaluate participants' trust in the robot. Muir's questionnaire has been widely used in the literature as a measure of trust in machines. This questionnaire includes the six dimensions of competence, predictability, dependability, responsibility, reliability, and faith, which are essential for measuring trust. Table 3.1 shows a set of four questions that we considered most relevant to our study and that we used in our experiment.

Phase 2 - Main Experiment: In this phase, participants joined the scheduled Zoom meeting and were given remote control over the researcher's computer. After participants learned how to work with the simulation environment, one of the three mentioned between-participant conditions (planning strategies) was randomly assigned to each one. After the

Questions
Q1 - To what extent can the robot's behavior be predicted from moment to moment? (predictability)
Q2 - To what extent can you count on the robot to do its job? (dependability)
Q3 - What degree of confidence do you have that the robot will be able to cope with similar situations in the future? (reliability over time)
Q4 - Overall, how much do you trust the robot? (overall degree of trust)

Table 3.1: Muir's questionnaire for humans' trust in the robot

Questions
Q5 - The robot is able to manipulate objects successfully and without any failure.
Q6 - The robot works efficiently in the sense of minimizing collaboration time.
Q7 - The robot tends to collaborate.
Q8 - I will be happy to collaborate again with this robot in similar future tasks.

Table 3.2: Questionnaire for human perception of the robot

experiment, in order to collect subjective information from participants about the strategies they adopted during the experiment, we verbally asked them the following question to explain their strategy: *Could you please explain the strategy that you used to accomplish the task?*. The experimenter wrote down the answers.

Phase 3 - Post-experiment Questionnaire: Participants were asked to complete a questionnaire as the last step of the study. The four questions of this questionnaire are exactly the same as in the prestudy to compare participants' views before and after the experiment. The other questions, as shown in Table 3.2, were in another seven-point Likert scale (Agree – Disagree) to assess their perception of the robot.

3.3 Results

Hypothesis 1a: We asked the same four questions about participants' trust in the robot before and after the experiment. Table 3.3 shows the average of their answers for each question. To show how the participants' trust changed after the experiment, we took the average of the four questions, both in the pre- and post-study questionnaires, and calculated their differences (see Figure 3.3). Based on Figure 3.3, the first and the third conditions show that participants' trust in the robot increases in most cases. Since the

	Question 1		Question 2		Question 3		Question 4		Q1-4 Average	
	Before	After	Before	After	Before	After	Before	After	Before	After
Strategy 1 (SH)	5.059	6	5.765	6.529	6.059	6.529	5.706	6.353	5.647	6.353
Strategy 2 (SR)	5.235	4.882	6.0	5.882	5.706	5.882	5.647	5.647	5.647	5.573
Strategy 3 (SHR)	4.875	6.313	5.25	6.25	5.063	5.625	4.938	5.688	5.031	5.969

Table 3.3: Average of participants' answers to the Muir's questionnaire about their trust in the robot before and after the experiment.
Question 1: predictability, Question 2: dependability, Question 3: reliability over time, Question 4: overall degree of trust

questionnaire data does not show a normal distribution, we used a Kruskal Wallis test to evaluate the statistical difference between the three different strategies. The analysis, $\chi^2(2, N = 50) = 12.56, p < .002$, reveals at least one significant difference between the strategies. Next, Dunn's test, a non-parametric pairwise multiple comparisons procedure, was used as a Post-hoc test. It indicates a significantly more positive effect in strategy 1-SH ($p < .01$) and strategy 3-SHR ($p < .003$) on the participants' trust in the robot compared to strategy 2-SR. However, no significant difference was found between strategies 1 (SH) and 3 (SHR) ($p = .63$).

Hypothesis 1b: In this experiment design, the robot does not fail and can always successfully manipulate objects. Hence, as expected, participants, in all three conditions, gave a high score to question 5 of Table 3.2 (See Figure 3.4, Q5). However, their answers to question 6, asking about robot performance in minimizing collaboration time, differ under the various strategies (Figure 3.4, Q6) and differ significantly($\chi^2(2, N = 50) = 15.561, p < .001$). Participants gave significantly higher scores to the robot with strategy 1 ($p < .001$) and 3 ($p = .0122$) than the robot with strategy 2. Again, strategies 1 and 3 do not differ significantly ($p = .308$).

Hypothesis 1c: To assess whether the robot's strategy can affect the participants' perception of the robot's collaboration, we analyzed their responses to questions 7 in Table 3.2 (See Figure 3.4, Q7). The results showed a significant effect of the robot's strategy ($\chi^2(2, N = 50) = 12.317, p < .002$). The participants found that the robot with strategies 1-SH ($p = .009$) and 3-SHR ($p = .004$) was more collaborative than the robot with strategy 2-SR. However, like the previous hypotheses, no significant differences were found between strategies 1 and 3 ($p = .7008$).

Hypothesis 1d: Based on the participants' scores to question 8, in Table 3.2, the robot strategy can affect participants' interest in collaborating with the robot ($\chi^2(2, N = 50) =$

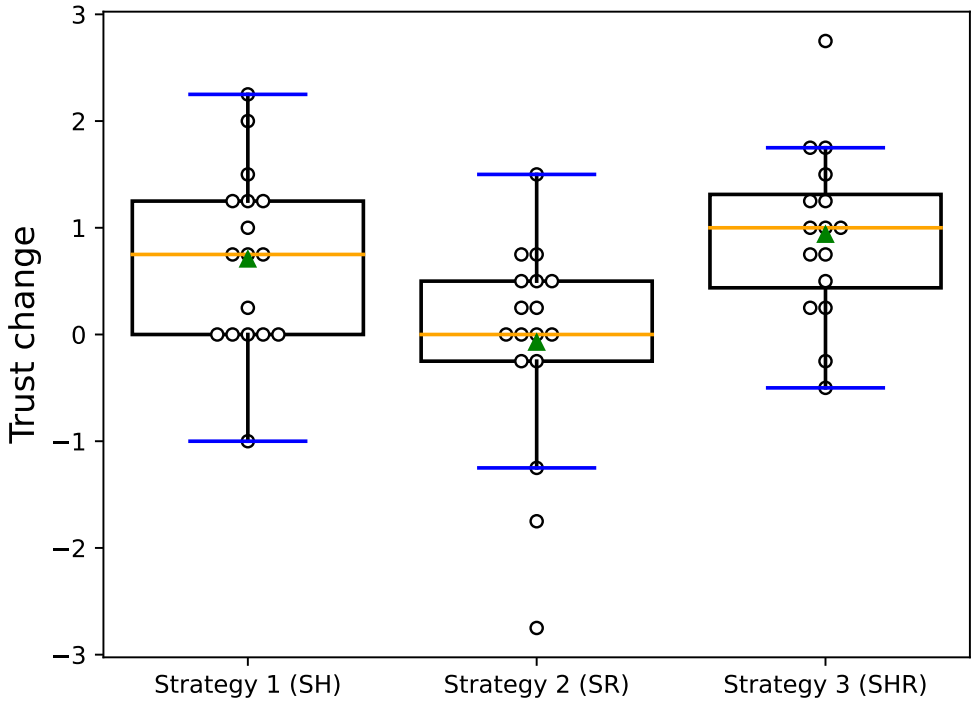


Figure 3.3: Change in participants' trust before and after the experiment.

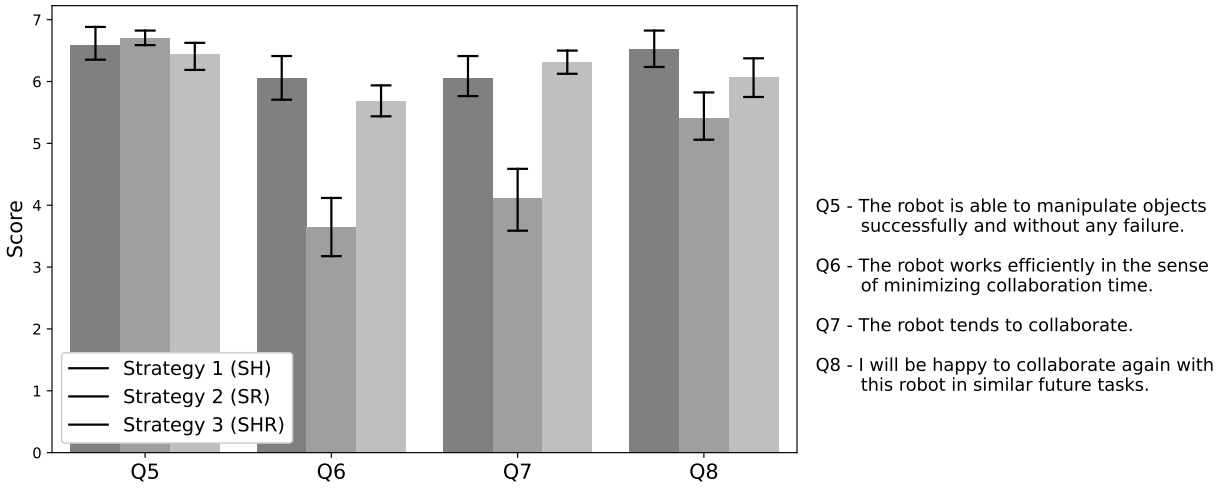


Figure 3.4: Averages of participants' answers to the questions in Table 3.2.

$8.168, p = .009$). The post-hoc test indicates that the participants showed more-positive inclination toward collaborating with the robot with strategy 1 than the robot with strategy 2-SR ($p = .013$). However, there is no difference between strategies 1-SH and 3-SHR ($p = .3$) and strategies 2-SR and 3-SHR ($p = .03$). Figure 3.4, Q8, shows the average of participants' scores to this question.

3.3.1 Participants' strategy

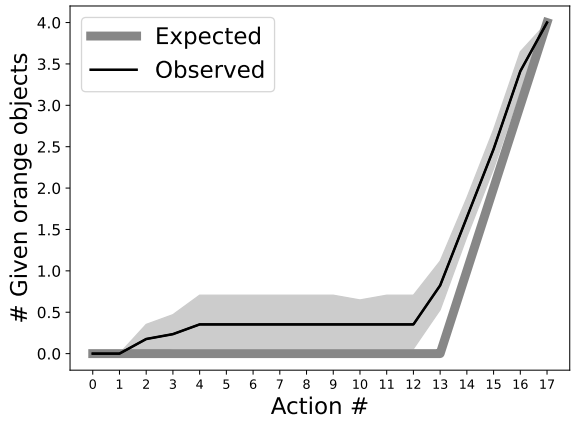
After the experiment, participants were asked to explain the strategy they adopted to accomplish their objective (i.e., finishing their board in the shortest time) and the task (i.e., removing all objects from the tables and the shared space). We also recorded participants' and the robot's actions and analyzed their collaboration with the robot.

Robot strategy 1 (SH): In this case, based on the robot's recorded actions, the robot provides participants with green blocks regardless of their needs. Assuming a self-serving human teammate, we expect they would give the robot the orange objects after finishing their objective. The thick line in Figure 3.5a shows the number of orange objects a self-serving human will give the robot. The thin line shows the average number of orange objects participants gave the robot. Six out of 17 participants followed the self-serving plan, and the others gave at least one orange object before finishing their board. Participants' explanation about their strategy revealed two main reasons for doing so: collaborating with the robot to finish the task and preventing the robot from being idle in cases where the robot had no object to move.

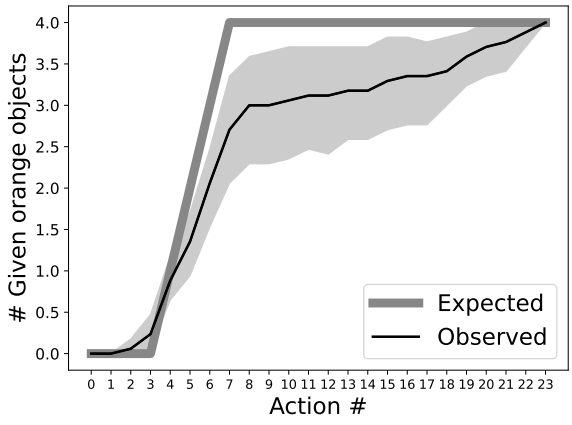
Robot strategy 2 (SR): In this case, the robot optimizes its own objective. The recorded robot's actions show that it does not give green objects at first and continues sorting its own objects. Assuming a purely assistive human teammate who does not press the skip button when they can provide the robot with orange objects, the thick line in Figure 3.5b shows the number of orange objects they will give to the robot. The thin line shows the average number of orange objects participants gave to the robot. Six out of 17 participants followed the purely assistive behavior, and the others pressed the skip button at least once despite having the option to give the robot a block. According to participants' answers to our question about their strategy, two of them were annoyed by the robot's behavior, and others wanted to check whether the robot gave them a green object when they pressed the skip button, withholding orange objects.

Robot strategy 3 (SHR): In strategy 3, the robot starts with its own blocks and tries to give green blocks whenever the human needs them. Similar to strategy 1, the thick line in Figure 3.5c shows the expected number of orange objects that a self-serving human will

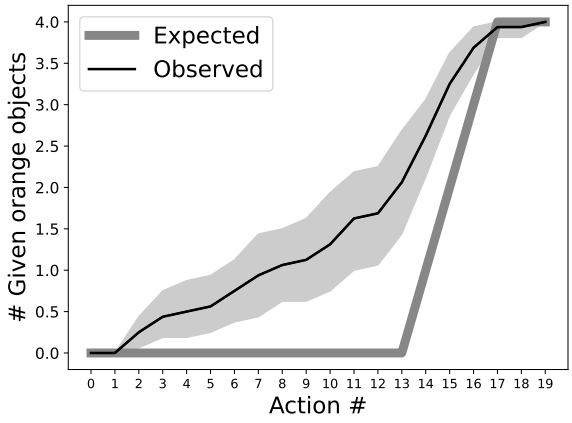
give to the robot. The thin line also shows the number of orange objects participants gave the robot. Three out of 16 participants followed the self-serving behavior, and the others gave at least one orange block to the robot, although they had the option to place their own blocks on the board. Their explanations for doing so were as follows: (1) collaborating with the robot to finish the task, (2) preventing the robot from being idle in cases where the robot had no object to move, (3) testing whether the robot gave them a green object when they gave the robot an orange one.



(a) Strategy 1 (SH)



(b) Strategy 2 (SR)



(c) Strategy 3 (SHR)

Figure 3.5: Expected (thick lines) and average of observed number of orange blocks passed to the robot (thin lines) along with their confidence intervals.

Prior experience with robot: We asked four questions to learn participants' prior experience with robots. We also examined how participants' prior experience with robots affected their initial trust in the robot, but the result showed no significant effects.

3.4 Summary and Discussion

In the experiment design, we have considered a single-human, single-robot scenario that required the two parties to collaborate to accomplish a given task. The team's task was to remove all blocks from the tables and the shared area. However, aligned with this task, we assigned the human an objective: to finish their board in the shortest time. Then, we developed three different offline strategies: 1- prioritizing the human (Strategy 1-SH), 2- prioritizing the robot (Strategy 2-SR), and 3- balancing between the human and the robot (Strategy 3-SHR). The main goal of this study was to evaluate how each strategy could affect participants' perception of the robot and collaboration, as well as how these strategies can influence the way they collaborate with the robot. This study addresses our first research question (RQ1).

We hypothesized that as the robot gave more priority to the human, the human teammate's perception of the robot, in terms of a) trust, b) performance, c) collaboration, and d) willingness to collaborate, would improve. The results partially support Hypothesis 1a concerning the human's trust in the robot. Compared to strategy 2, strategies 1 and 3 showed significant improvement in the human participants' trust in the robot, as measured before and after the experiment. However, this difference was insignificant between strategy 1 (SH) and strategy 3 (SHR). Regarding Hypothesis 1b and 1d, we noticed a significant difference between strategies 1 and 2 and strategies 2 and 3 in participants' perception of the robot's performance and collaboration. Again, we did not observe a significant difference between strategies 1 and 3. Regarding participants' willingness to collaborate with the robot (hypothesis 1d), there is a significant difference only between strategies 1 and 2. In addition, based on participants' actions, we analyzed how they collaborated with the robot, finding that participants gave more blocks to the robot when it used strategy 3, which we tentatively suggest could be an indication of better human-robot teamwork (although this needs to back up in future studies). Furthermore, participants tended to give fewer blocks when the robot used strategy 2. Considering all the above, our findings suggest that, in a collaborating scenario, developing a plan and strategy for the robot that can balance the human teammate's objective (or policy) and the robot's (or team's) objective can benefit the team without negatively affecting humans' perception of the robot. Those results will inform our future work towards improving human-cobot collaboration.

The current work has several limitations. The turn-taking scenario employed might not be applicable in many real industrial settings, and for future work, we will need to consider more realistic scenarios. Furthermore, we did not assess participants' attitudes toward the collaboration and the robot before and during the experiment, which was not feasible due to the remote nature of the experiment (required due to COVID-19).

In this chapter, we addressed the first research question (RQ1), investigating how the robot's planning, prioritization of the human, prioritization of itself, and balancing strategy can influence human perception of the robot. The findings of this study are important from this perspective that confirms the idea that a balancing strategy designed efficiently can improve the team performance while maintaining the human perception of the robot and collaboration at a high level and lead to a smooth and efficient collaboration, exploiting both agents' abilities. Building upon insights gained from this study, the next chapter, Chapter 4, introduces a robot task planning framework and algorithm. This framework takes into account the human agent's performance and leading/following preference, estimating and updating them throughout their collaboration. It further updates the task state dynamically and strategically plans actions aimed at enhancing team performance.

Chapter 4

Task Allocation & Planning Framework

In this chapter, we begin by defining the task allocation problem. Subsequently, we delve into the proposed task planning architecture and strategy. Finally, we discuss the task planning algorithm, which takes into account both the human agent's preference and performance. Reiterating RQ2, this chapter will address the question: "How can we design a planning architecture that takes into account both human preference and performance and enables the robot to update the task state based on both agents' actions?"

4.1 Problem Statement

We consider a collaborative task, τ, comprising n subtasks, and two agents: a human and a robot. These agents must cooperate to complete a set of precedence-constrained subtasks, denoted as $\tau = \{\tau_1, \tau_2, \ldots, \tau_n\}$. Each subtask τ_i has associated completion times, represented as t_i^h for the human and t_i^r for the robot. However, owing to uncertainties stemming from agents and the environment, the task completion time may deviate from the initially specified duration.

In each decision-making step, an agent has the agency to allocate a set of feasible subtasks to the other agent, as well as to assign a subtask to itself and execute it. What sets this problem apart from conventional task allocation and scheduling problems is that the agents here have the autonomy to choose their actions and subtask assignments rather than being assigned specific tasks with predetermined instructions regarding what and when they should execute them. Throughout this collaboration, the robot needs to:

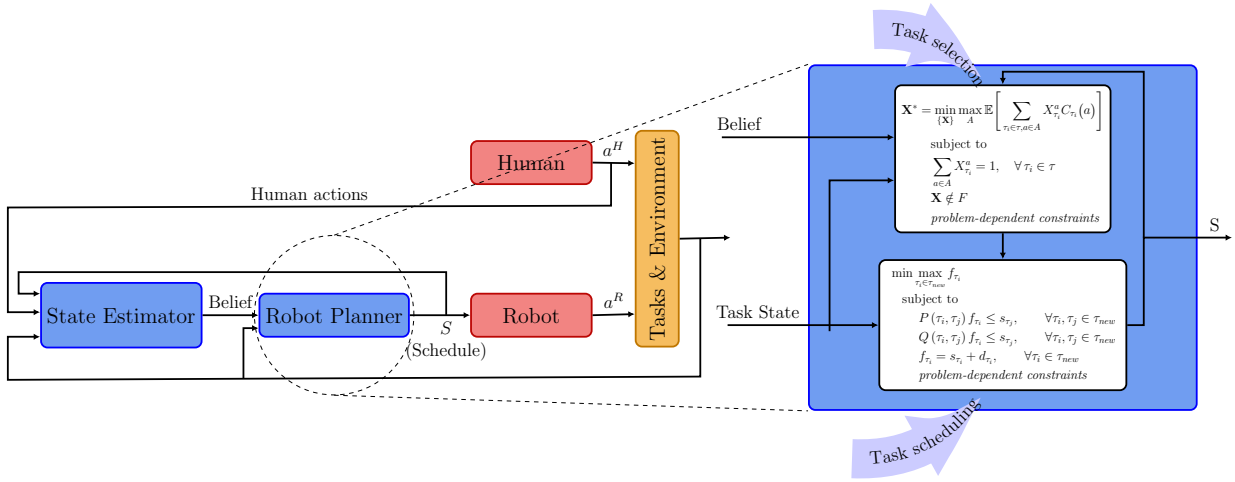


Figure 4.1: Task selection and planning architecture encompassing the state estimator and robot planner. The robot planner takes into account human preference, performance, and task state, followed by the implementation of two planning steps: task selection and task scheduling.

- estimate the human agent's leading/following preference,
- monitor the impact of the human agent's actions on the overall team performance continuously,
- minimize the collaboration cost(e.g., completion time) while adapting to the human agent's preference and performance,
- detect and address human errors, if applicable.

4.2 Planning Architecture

The planning architecture is illustrated in Fig. 4.1.

Tasks & Environment: the overall system or task the human and robot collaborate on.

Human/Robot: These two blocks represent the input provided by the human and robot and applied to the system. It's important to note that this is an asynchronous

decision-making process, where the human and robot agents act and make decisions independently and at different times.

State estimator: During the collaborative process, the robot evaluates the human's actions and infers their inclination towards leading or following. Furthermore, it is responsible for monitoring the human's performance and assessing their level of performance. These states, however, cannot be measured directly, and the robot needs to infer them through the history of the interaction. To do so, the state observer takes the history of the human's actions, the robot's beliefs and schedule, and the tasks' states.

Robot Planner: The robot planner is responsible for providing the robot with a schedule based on the tasks and environment's states and the output of the state estimator block, belief about the human agent's preference and performance. The robot planner consists of two phases: task selection and task scheduling. In each decision step, when necessary, the robot performs task selection and subsequently performs task scheduling to determine its following action.

4.3 Planning Strategy

At each decision step, the robot planner must determine a one-to-one subtask assignment for the agents and establish a task execution schedule to minimize the collaboration cost, injecting both the human agent's preference and performance. Task allocation and scheduling problems can usually be modeled as mixed linear integer programs (MILP). However, the complexity of MILP-based solutions makes them computationally intractable. In addition, involving the robot's belief about the human agent's preference and performance adds more complexity to the problem due to the dynamic and unpredictable nature of human behavior and intentions. These challenges, in concert, make formulating and solving the problem as a single optimization problem increasingly demanding and arduous. Decomposing task allocation and scheduling is a promising and commonly used approach to deal with this complexity [102]. Here, we also take advantage of this idea and split the problem into two subproblems: task allocation and task scheduling.

In the first step, considering the agents' set $A = \{human, robot\}$, the task $\tau = \{\tau_1, \tau_2, \dots, \tau_m\}$, and the subtasks' associated assignment costs to the human and robot $C_{\tau_i}(a), a \in A$, the robot first seeks an optimal task allocation. m is the number of subtasks that are not completed yet or need to be fixed. Subsequently, if necessary, a new set of subtasks $\tau_{new} = \tau \cup \tau_{allocate}$ is generated, including actions needed to allocate subtasks to the human $\tau_{allocate}$. For example, in a sorting task, the robot may place a box on the human agent's side

to indicate that the sorting of this box has been assigned to them. This additional subtask, which requires a certain amount of time, is needed as part of the task allocation process. The task scheduler utilizes the derived optimal task allocation and τ_{new} to determine an optimal task schedule. If the solution achieved in the task allocation phase proves infeasible during the task scheduling phase, the first step must be repeated to obtain a revised allocation.

The task allocation optimization problem can be formulated as (4.1), minimizing the maximum cost of task assignments, between the human and robot,

$$\mathbf{X}^* = \min_{\{\mathbf{X}\}} \max_{A} \mathbb{E}\left[\sum_{\tau_i \in \tau, a \in A} X^a_{\tau_i} C_{\tau_i}(a) \right] \tag{4.1}$$

subject to

$$\sum_{a \in A} X^a_{\tau_i} = 1, \quad \forall\, \tau_i \in \tau \tag{4.2}$$

$$\mathbf{X} \notin F \tag{4.3}$$

$$\textit{problem-dependent constraints.} \tag{4.4}$$

In Eq. (4.1), $\mathbf{X} = \left\{ X^a_{\tau_i} \mid \tau_i \in \tau, a \in A \right\}$, where $X^a_{\tau_i} \in \left\{0, 1\right\}$ is a binary decision variable that indicates whether task τ_i is allocated to agent $a \in A$ ($X^a_{\tau_i} = 1$) or not ($X^a_{\tau_i} = 0$). Function $C_{\tau_i}(a)$ is the cost incurred by assigning task τ_i to the agent $a \in A = \left\{ human, robot \right\}$ taking into account human agent's performance, and preference to follow the robot. Function C_{τ_i} imposes a higher cost for assigning the task to the human who prefers to lead, and a lower cost for allocating the task to a human with low performance, since having the robot allocate tasks limits the human's control and mitigates the potential for human errors. Equation (4.2) specifies each task must be assigned to exactly one agent, the human or the robot. In Constraint 4.3, F represents a set of task allocation solutions that do not result in a feasible task scheduling solution. Consequently, this constraint prohibits the selection of task allocation solutions from this set. Eq. (4.4) indicates that additional problem-dependent constraints can also be added to (4.1). After finding a solution for the task allocation problem, we generate τ_{new} with known required time, d_{τ_i}, to finish each task τ_i and updated task-precedence constraints.

In the subsequent phase, the robot determines an optimal schedule specifying both the tasks to be performed and their corresponding start and finish times. We introduce decision variables s_{τ_i}, representing the start times of subtask $\tau_i \in \tau_{new}$. f_{τ_i} denotes the finish time for subtask τ_i. To account for task precedence, we utilize a binary function $P(\tau_i, \tau_j)$, which takes the value of 1 when τ_i must be completed before τ_j. Additionally,

we employ a binary decision variable $Q(\tau_i, \tau_j)$, where $Q(\tau_i, \tau_j) = 1$ indicates that both τ_i and τ_j are assigned to the same agent, and τ_i precedes τ_j.

Considering that we solely focus on collaboration time, the task scheduling problem can be formulated as minimizing the overall processing time:

$$\min \max_{\tau_i \in \tau_{new}} f_{\tau_i} \tag{4.5}$$

subject to

$$P(\tau_i, \tau_j) f_{\tau_i} \leq s_{\tau_j}, \qquad \forall \tau_i, \tau_j \in \tau_{new} \tag{4.6}$$

$$Q(\tau_i, \tau_j) f_{\tau_i} \leq s_{\tau_j}, \qquad \forall \tau_i, \tau_j \in \tau_{new} \tag{4.7}$$

$$f_{\tau_i} = s_{\tau_i} + d_{\tau_i}, \qquad \forall \tau_i \in \tau_{new} \tag{4.8}$$

$$\textit{problem-dependent constraints.} \tag{4.9}$$

Inequality (4.6) guarantees the precedence constraints. Inequality (4.7) maintains the constraint that each agent can handle only one subtask at a time. In (4.8), f_{τ_i} is determined by the time needed by the assigned agent, d_{τ_i}, to complete subtask τ_i. It's important to note that the constraints in this optimization problem are specific to the problem being addressed, and as such, the mentioned constraints can be adjusted, and new ones can be introduced (4.9).

Depending on the constraints in optimization problems (4.1) and (4.5), they can be formulated as mixed-integer linear or nonlinear programs. Nevertheless, in the collaborative scenario explained in the next chapters, we will formulate the problem as a mixed-integer linear program and subsequently use the Gurobi optimizer as the solver.

4.4 Algorithm

The proposed task selection and planning procedure is explained in Algorithm 1.

As indicated in Lines 1, the robot initializes its belief about the human agent's performance and following preference. Throughout the collaboration, until completing all subtasks, the robot monitors and records the human agent's actions and errors (Lines 3-4). The robot also has to update the task and precedence constraints based on the finished subtasks and the subtasks need to be fixed (Line 5). Then, based on the planning structure shown in Fig. 4.1, the robot updates its belief about the human agent's following preference and performance (Line 6). Next, if the situation requires a new plan, as discussed before, the robot first solves for an optimal task allocation (Line 9) and then creates τ_{new} (Line

Algorithm 1: Task selection and planning

input : Precedence-constrained tasks, τ

1 *bel* $\leftarrow$ initial beliefs about human agent's following preference and performance
2 **while** *Tasks are not finished* **do**
3 Monitor the human's actions
4 Detect the human's errors
5 Update task, τ
6 Update *bel*
7 **if** *new schedule is needed* **then**
8 **while** *schedule, $S^\star$, is not found* **do**
9 $X^\star \leftarrow$ `TaskSelection`(τ, *bel*)
10 $\tau_{new} \leftarrow$ `CreatNewTask`(τ, $X^\star$)
11 $S^\star \leftarrow$ `TaskSchedule`(τ_{new}, $X^\star$)
12 $a^R \leftarrow$ `GetAction`($S^\star$)
13 `ApplyAction`(a^R)

10). Next, the robot solves for an optimal schedule (Line 11). If the robot fails to find an optimal schedule with the current task allocation, it proceeds to generate a new task allocation. Subsequently, the robot performs its action, a^R based on the obtained schedule (Lines 12-13). This continues until the team finishes all subtasks.

This chapter has outlined the general framework, the problem formulation, and the algorithm. The implementation of them in a collaborative scenario is mainly problem-dependent. However, in Chapter 6, I will discuss in more detail how to implement and derive each part of the algorithm and the problem formulation for a designed collaborative scenario. Section 6.1 will discuss the collaborative task, its corresponding task graph, and both agents' actions. Then, in Section 6.2, the task allocation and scheduling formulation and their corresponding linearized optimization problems will be elaborated. This section also discusses the details of estimating and updating the robot's belief on human preference and performance.

4.5 Summary and Discussion

This chapter addresses RQ2 by introducing the robot planning framework and algorithm designed for task allocation and scheduling in human-robot collaboration. These tools em-

power the robot to estimate the human agent's performance and leading/following preference, dynamically update tasks, and proactively plan actions to enhance team performance. Although this chapter outlines the general framework and algorithm for a wide range of collaborative scenarios, Chapter 5 and, more meticulously, Chapter 6 focus on applying the task planning algorithm and estimation methods to specific collaborative scenarios. However, the versatility of the proposed framework is not confined to these scenarios alone, as it can be adapted to a wide range of contexts and applications. The next chapter starts with translating this theoretical framework into a practical application by implementing it in a collaborative scenario within a simulation environment.

Chapter 5

Adaptive Task Planning: A Simulation Study

In this chapter, we implement the proposed framework from Chapter 4 in a collaborative scenario within a simulation environment. Within this simulation, we consider a simplified human decision-making model, taking into account two factors: the human agent's leading/following preference and error-proneness. Subsequently, we apply the proposed planning algorithm to examine its performance for nine different combinations of human error-proneness levels and following preferences. Through this simulation study, we commence the exploration of RQ3 and RQ4, an exploration that will be further developed in the following chapters (Chapters 6 and 7). This simulation study facilitates an environment for investigating diverse scenarios, allowing for a comprehensive evaluation of the adaptive planner's efficiency in handling varying conditions.

5.1 A Sample Scenario

Here, I provide a possible real-world collaborative scenario of the explained problem to help readers better understand the objectives of this study. In this scenario (Fig. 5.1), the human and robot have to collaborate to fill a certain number of boxes in each workspace (W1-W4) with the correct objects.

As shown in Fig. 5.1, there are four workspaces, W1-W4, with five empty numbered boxes on each. We assume at the beginning of the experiment, a pattern (e.g., Fig. 5.2a or 5.2b) is displayed to the human-robot team depicting how the boxes must be filled.

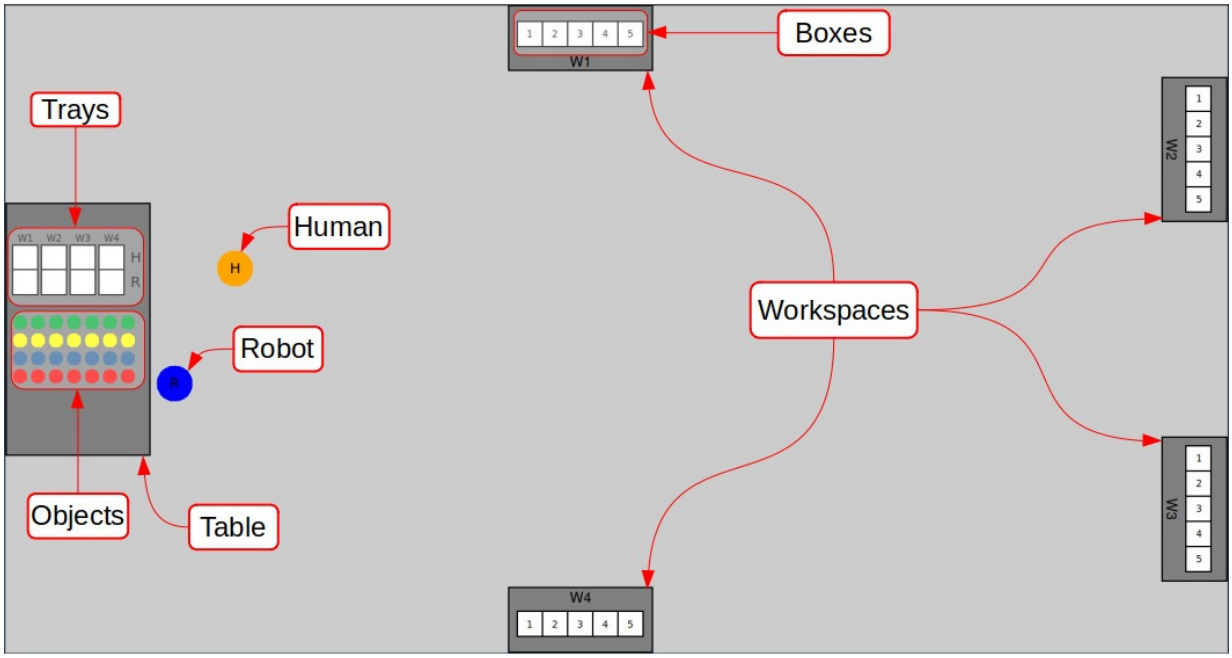


Figure 5.1: Simulation environment: In a $17 \times 9m$ room, there are four workspaces (W1-W4), each containing five boxes. There are 28 objects of four different colors (7 of each) on the table. Additionally, there is a tray for the human and one for the robot. Each tray has four compartments labeled W1-W4. A video of the simulation is available at https://youtu.be/gEzNoiiGP-w

.



(a) Pattern 1: An easy one to memorize (Difficulty: Easy)



(b) Pattern 2 (Difficulty: Medium)

Figure 5.2: Humans are shown patterns with different levels of memorization difficulty.

The pattern is shown for a short amount of time, and the human needs to memorize it. Depending on the pattern's difficulty level, the human may make errors or need the robot's help to complete the tasks. The boxes on each workspace are numbered, and objects must be placed in the correct order. For instance, the agents cannot place an object into box 3 before filling boxes 1 and 2.

On the large table, there are two trays, which are empty at the beginning of the experiment. Each tray has four compartments labeled W1-W4. The upper one in Fig. 5.1, labeled *'H'*, is where the robot can assign objects to the human. Similarly, the human can assign objects to the robot and put them on the tray, labeled *'R'*. For example, when the robot places an object in the human's tray labeled 'W1', the next object that the human has to put in workspace 1 (W1) is the one that is in the tray. In addition, when one of the compartments in a tray is full, the other agent is not allowed to put another object in the same-labeled compartment of the other tray. For instance, when the robot has already placed an object in one of the human tray's compartments labeled 'W2', the human cannot place any object in the compartment with the same label, 'W2', on the robot's tray.

We assume that both the robot and the human decide when they approach the table. In addition, a safety zone is imposed around the table, which an agent cannot enter when the other agent is inside (to account for, e.g., safety considerations). Upon approaching the table, the human has to take one or more feasible actions from the following set of actions:

1. picking up an object from the table or the human's tray and placing it in the

workspace,

2. placing an object in the robot's tray,

3. picking up an object from the robot's tray and placing it in the workspace.

In each step, the human needs to take one of Actions 1 or 3 to exit the table area so that the robot can enter. The human can select Action 3 only when Actions 1 and 2 are not feasible.

The robot is also responsible for a set of actions similar to those of the human. The robot also needs to correct possible human errors. Thus, the robot's set of actions are as follows:

1. picking up an object from the table or the robot's tray and placing it in the workspace,

2. placing an object in the human's tray,

3. bringing back a misplaced object from the workspace,

4. bringing back a misplaced object from the robot's tray,

5. picking up an object from the human's tray and placing in the workspace.

Similar to the human, the robot also needs to take one of Actions 1, 3, or 5 to exit the table area. Action 5 can be chosen only when Actions 1-4 are not feasible. The collaboration objective is accomplished when all objects are put in the correct boxes.

5.2 Simulation Experiments

Collaboration task: The tasks and precedence constraints are encoded by a directed acyclic graph, called a task graph, where vertices are tasks and edges represent precedence constraints between them. The initial task graph of the experiment, with two dummy nodes as starting (τ_0) and finishing (τ_{21}) points, is depicted in Fig. 5.3. In this experiment, we assume that the human's speed does not change during the experiment and is higher than that of the robot. Thus, the nominal required processing times for tasks are fixed. However, in a future real-world scenario, this assumption can be easily relaxed if the robot can measure the human's speed and update processing times accordingly.

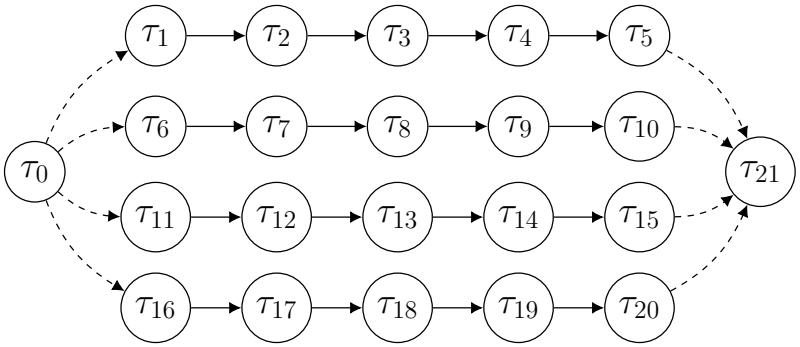


Figure 5.3: Task graph of the experiment

Human following preference: In this study, with a single scalar random variable α_f, we capture the human preference for following the robot. In addition, we assume a tuple of possible discrete values for α_f, given by $V = \big(0.0, 0.2, 0.4, 0.6, 0.8, 1\big)$. The values $\alpha_f = 0$ and $\alpha_f = 1$ respectively mean that the human very likely prefers to lead and follow the robot. At the beginning of the experiment, the robot assumes that the human will follow it and sets its initial belief about that as follows:

$$P[\alpha_f = v_i] = b(i; n = 5, p = 0.8),\ V = (v_0, \ldots, v_5),$$

where $b(i; n, p)$ is a binomial distribution.

Human performance: In this experiment, we only consider the probability that the human may make a mistake and put wrong objects in the workspaces and the robot's tray. A single scalar random variable p_e, with a tuple of discrete values $W = (0, 0.1, 0.2, \ldots, 1)$, is adopted to capture the probability of human error. Humans with $\alpha_e = 1$ make too many errors, and in contrast, $\alpha_e = 0$ indicates they are accurate. The robot first assumes that the human is almost accurate and, during the collaboration, updates its initial belief, which is a binomial distribution

$$P[p_e = w_i] = b(i; n = 10, p = 0.2),\ W = (w_0, \ldots, w_{10}).$$

Task allocation: The task allocation problem is modeled as (4.1), where

$$C_{\tau_i}(a) = \begin{cases} t_i^h p_f + c_f(1 - p_f), & \text{if } a = \textit{Human} \\ t_i^r + p_e c_e + c_h x_{\tau_i}^{human}, & \text{if } a = \textit{Robot}. \end{cases} \tag{5.1}$$

In Eq. 5.1, c_f is the penalty value for assigning the task to the human who prefers to lead, and c_e is the penalty imposed for not assigning tasks to the human making mistakes. By

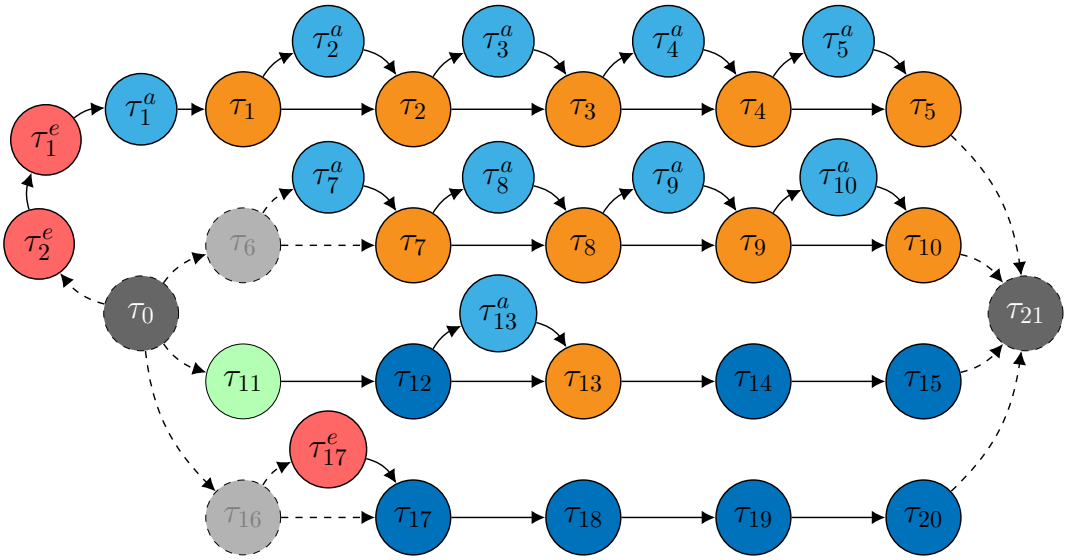


Figure 5.4: Temporary task graph of the experiment after task allocation.
Blue: Robot's tasks, Orange: Human's tasks, Cyan: Assigning tasks to the human, Red: Correcting human errors, Green: Already assigned tasks, Gray: Finished tasks

assigning a task to the human, the robot informs the human about the next object that has to be put in a workspace. This restricts the human's agency and prevents the human from making mistakes. In addition, penalty c_h is incurred for assigning the robot a task whose corresponding object has already been placed on the human's tray ($x_{\tau_i}^{human} = 1$). Fig. 5.4 shows an example of task allocation and the temporary task graph. In this case, the human has made a mistake in tasks 1, 2, and 17. Thus, the robot needs to fix them by performing tasks τ_1^e, τ_2^e, and τ_{17}^e. Tasks 6 and 16 are finished, and the robot has already placed an object in the human's tray for task 11. The robot has assigned tasks 1, 2, 3, 4, 5, 7, 8, 9, 10, and 13 to the human and needs to put the tasks' corresponding objects on the human's tray (e.g., τ_7^a). The robot has also selected tasks 12, 14, 15, 17, 18, 19, and 20 to do by itself.

Solving optimization problems: The task allocation and task scheduling problems have been rewritten as a mixed integer linear program. Both problems are considered as NP-hard optimization problems. We implemented the simulation experiments on a computer with an Intel Core i7-8700 CPU 3.20GHz × 6-cores CPU and 16 GB RAM, on Ubuntu 20.04. We utilized the Gurobi mathematical optimization solver and implemented a 4-second time constraint to prompt the termination of the search and provide the current solution if it continues to seek additional solutions. Additionally, a relative gap tolerance of

0.004 was set as another termination criterion. We also employed a warm-start – providing the solver with a (partially) valid initial solution – that uses the current step's solution in the subsequent step. Consequently, for the initial step, the problem can be solved offline with a reduced relative gap tolerance and an extended time limit. To assess the solver's performance, we conducted the simulation five times to determine the frequency with which the solver reached the time limit. The findings indicate that during the task allocation phase, the solver consistently avoided reaching the time limit and terminated due to the second stopping criterion. During the task scheduling phase, the solver terminated only twice due to the time limit, primarily occurring at the initial stages of the task. This is mainly because as the team progressed and completed more subtasks, the problem size reduced, facilitating resolution by the optimizer.

Updating α_e and α_f: We used the method proposed in [80] and created some models to update the robot's belief about the human's performance level and preference to follow the robot. Details of these models and the method will be explained in Chapter 6.

Robot action: The robot chooses its action based on the created plan. Potential field-based motion planning is employed for the robot to avoid colliding with the human.

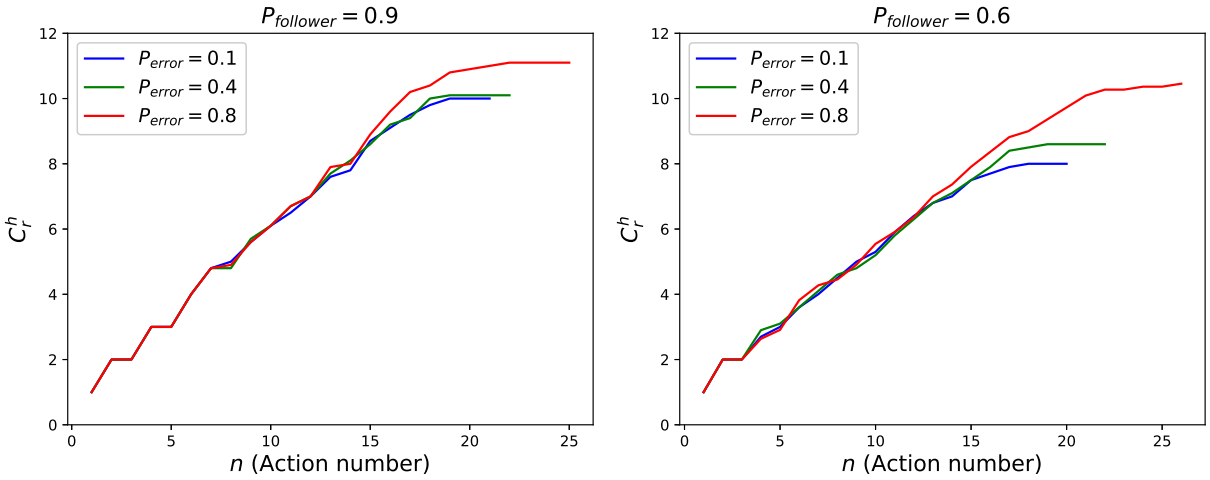
5.3 Simulation Results

This study conducted simulation experiments to evaluate how the proposed algorithm deals with different combinations of human preferences and accuracy levels. Nevertheless, a real-world human-subject study is necessary to assess the algorithm's efficiency fully, and this will be the focus of the next chapters. For this simulation, we created a simplified human model that selects actions from the action set discussed in Section 5.1, based on two different parameters: the preference to follow the robot P_{follow} and the error-proneness P_{error}. For the simulation, we consider three different cases of preferences to follow: strong ($P_{follow} = 0.9$), moderate ($P_{follow} = 0.6$), and slight ($P_{follow} = 0.3$). We also take into account three different levels of error-proneness: low ($P_{error} = 0.1$), moderate ($P_{error} = 0.4$), and high ($P_{error} = 0.8$).

Table 5.1 shows the results of ten simulations for each combination of the human's preferences and accuracy levels. To assess how considering the human's preference and accuracy level affects collaboration performance, we repeated the simulation for three different cases in which the robot takes only the completion time into account and omits the human's error-proneness and preference for leading (i.e., $\alpha_e = 0, \alpha_f = 1$).

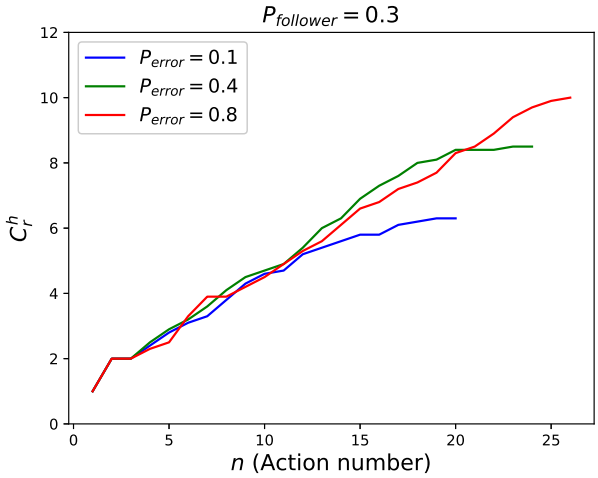
Table 5.1: Simulation results: The mean value of result after ten repetitions for each case.

	Human param.		Wrong actions		Assigned tasks		Total time (s)		Travel distance (m)		Total actions		Idle time (s)
	P_{follow}	P_{error}	n_h^{wrong}	t_r^{wrong}	n_h^{assign}	n_r^{assign}	t_h^T	t_r^R	d_h^T	d_r^R	n_h^T	n_r^T	t_r^{idle}
Considering human preference and accuracy	0.9	0.1	0.1	1.0	0.1	10.0	133.2	141.3	22310.9	17681.5	11.8	19.3	11.9
	0.9	0.4	1.1	15.5	0.1	10.1	142.8	155.4	24107.8	19415.4	12.7	20.4	11.5
	0.9	0.8	1.9	29.4	0	11.1	154.2	169.8	25938.8	21181.7	13.6	21.9	13.2
	0.6	0.1	0.5	7.7	1.7	8	140.8	146.9	23378.9	18151.1	13.5	17.7	14.0
	0.6	0.4	2.4	23.6	1.6	8.6	146.2	156.7	24771.0	19576.5	14.9	19.7	11.5
	0.6	0.8	4.3	46.9	1.1	10.3	166.3	184.5	28762.6	23921.1	15.8	23.9	20.5
	0.3	0.1	0.9	9.7	4.7	6.3	142.1	149.2	22687.8	17803.4	15.7	17.2	19.5
	0.3	0.4	3.1	22.1	5.5	8.5	152.6	160.5	25351.0	19824.6	18	21.5	15.2
	0.3	0.8	8.5	46.3	6.6	10.1	161.7	172.1	26837.0	20673.6	21.4	27.2	16.1
$\alpha_e = 0$ $\alpha_f = 1$	0.3	0.1	1.4	9.4	6.3	9.3	134.4	140.8	22664.6	17978.0	19.5	18.5	11.0
	0.3	0.4	5	23.2	7.4	10.4	145.6	149.3	24793.9	19073.3	21.3	23.6	10.5
	0.3	0.8	11.4	76.8	7.5	11.7	191.8	200.1	32841.5	25888.7	25.8	30.8	12.9



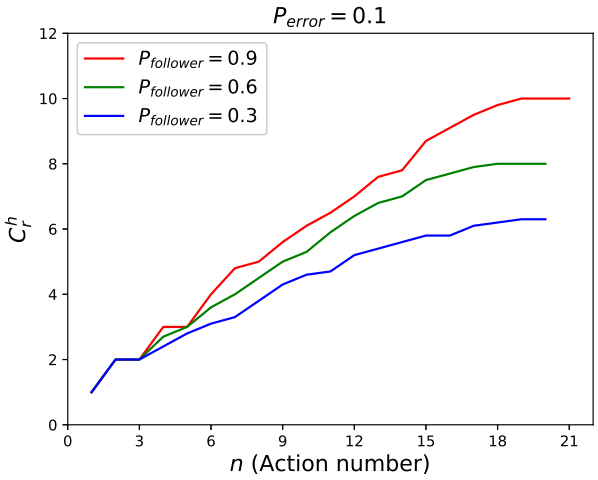
(a) Strong preference to follow the robot

(b) Moderate preference to follow the robot

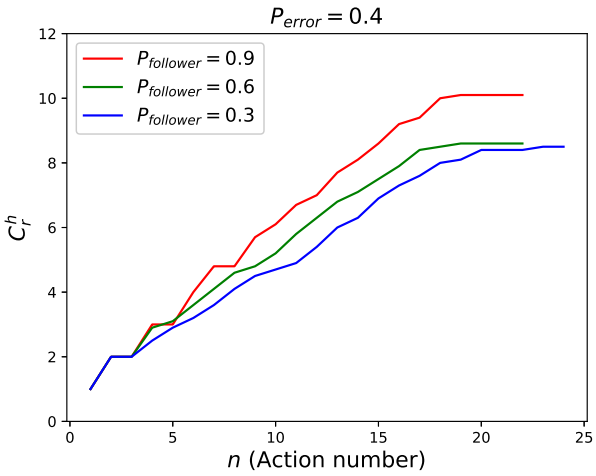


(c) Slight preference to follow the robot

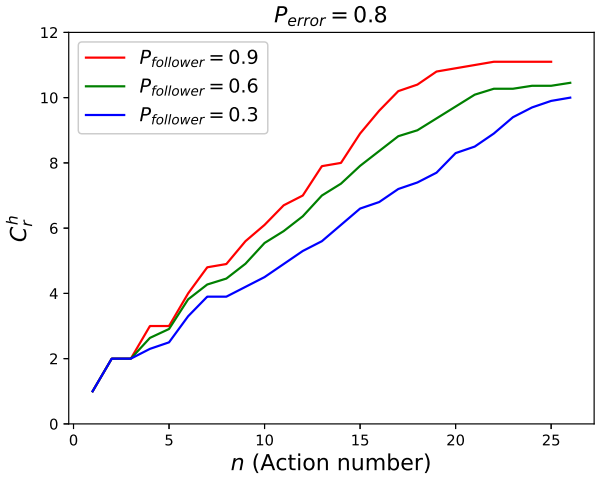
Figure 5.5: Cumulative number of assigned tasks to the human by the robot



(a) Strong preference to follow the robot



(b) Moderate preference to follow the robot



(c) Slight preference to follow the robot

Figure 5.6: Cumulative number of assigned tasks to the human by the robot

Strong preference to follow the robot ($\mathbf{P}_{follow} = \mathbf{0.9}$)

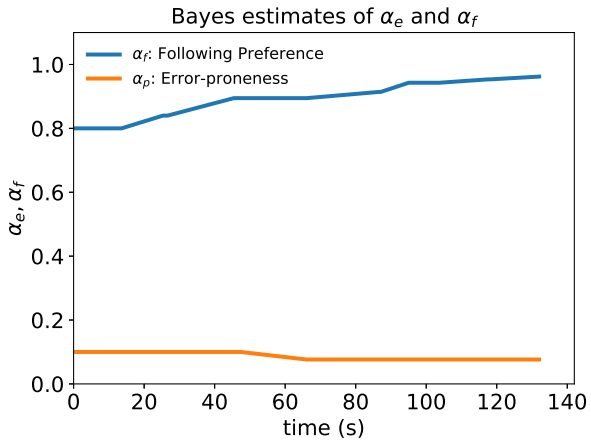
When the human prefers to follow, as the number of assigned tasks indicates n_r^{assign}, the robot takes the leading role and guides the team. In this case, even when the human's accuracy would otherwise be low, the human makes only a few mistakes (n_h^{wrong}) because of following the robot. As P_{error} increases from 0.1 to 0.8, a slight increase can be seen in the number of wrong actions the human takes (n_h^{wrong}), the time that the robot needs to fix them (t_r^{wrong}), the number of tasks that the robot assigns to the human (n_r^{assign}), as well as the time (t_h^T and t_r^T), travel distance (d_h^T and d_r^T), and the total number of actions the human and robot needed to finish the tasks (n_h^T and n_r^T). Fig. 5.5a shows the cumulative number of assigned tasks by the robot to the human. In Fig. 5.7, the robot's estimates of the human agent's error-proneness and following preference are shown for a single run across three different values of the human agent's accuracy levels. The robot's estimates of human error-proneness are almost the same since the human agent follows the robot and makes only a few, if any, mistakes.

Moderate preference to follow the robot ($\mathbf{P}_{follow} = \mathbf{0.6}$)

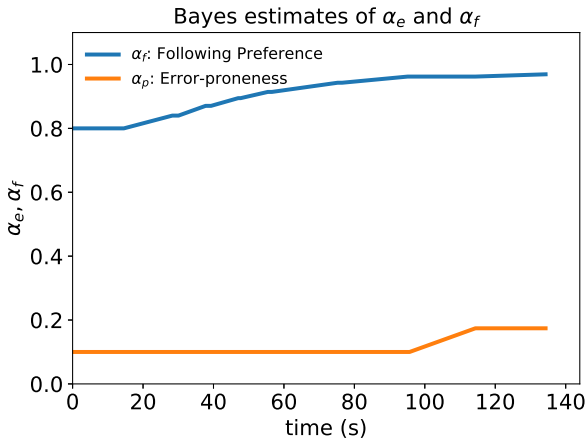
In this case, the robot has a moderate preference for following. The robot estimates the human's preference and assigns fewer tasks to the human compared to $P_{follow} = 0.9$. However, when the human's errors increase in number, the robot assigns more tasks. Since the human does not entirely follow the robot, when P_{error} increases, n_h^{wrong} increases. With the increase of P_{error}, also n_r^{assign}, t_h^T and t_r^T, d_h^T, d_r^T, n_h^T, and n_r^T grow. The cumulative number of assigned tasks by the robot is depicted in Fig. 5.5b. Fig. 5.8 illustrates the robot's estimates of the human agent's error-proneness and following preference for a single run across three different values of the human agent's accuracy levels.

Slight preference to follow the robot ($\mathbf{P}_{follow} = \mathbf{0.3}$)

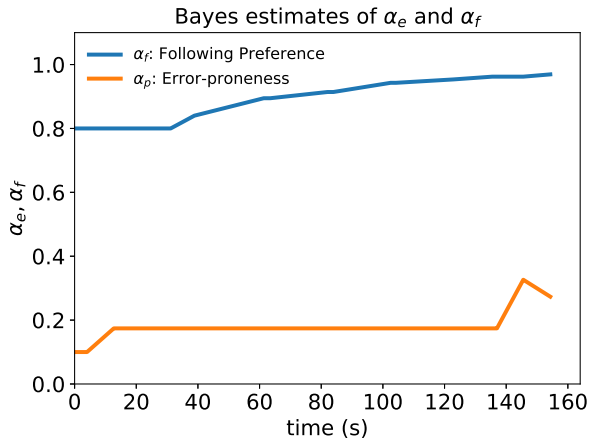
When the human prefers to lead, n_r^{assign} shows that the robot can adapt itself to the human preference and assign fewer tasks. In this case, when the human is inaccurate, the robot assigns significantly more tasks to retake the leading role and improve the team's performance. This case also shows an increase in n_r^{assign}, t_h^T and t_r^T, d_h^T, d_r^T, n_h^T, and n_r^T with the rise of P_{error}. Fig. 5.5c shows the cumulative number of assigned tasks by the robot to the human. Fig. 5.9 shows the robot's estimates of the human agent's error-proneness and following preference.



(a) $P_{error} = 0.1$

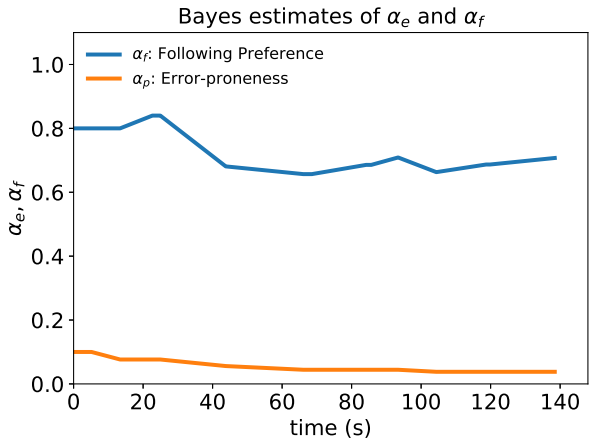


(b) $P_{error} = 0.4$

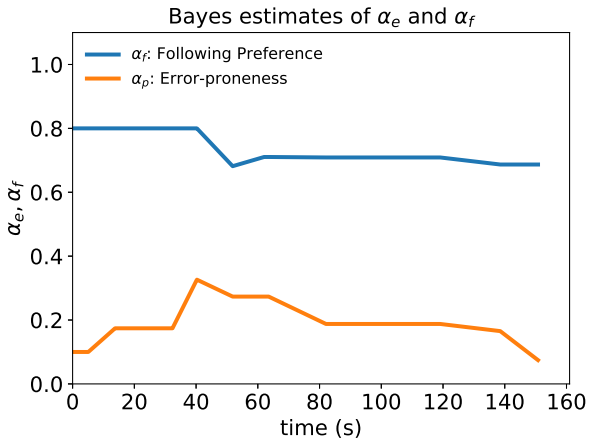


(c) $P_{error} = 0.8$

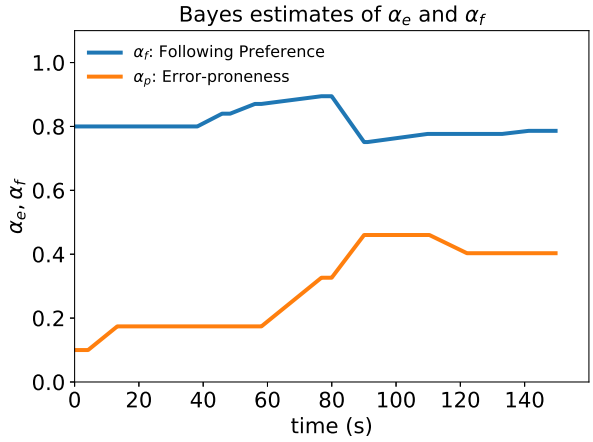
Figure 5.7: The robot's estimations of the human's following preference and error-proneness for a human with a strong preference to follow the robot $P_{follow} = 0.9$



(a) $P_{error} = 0.1$

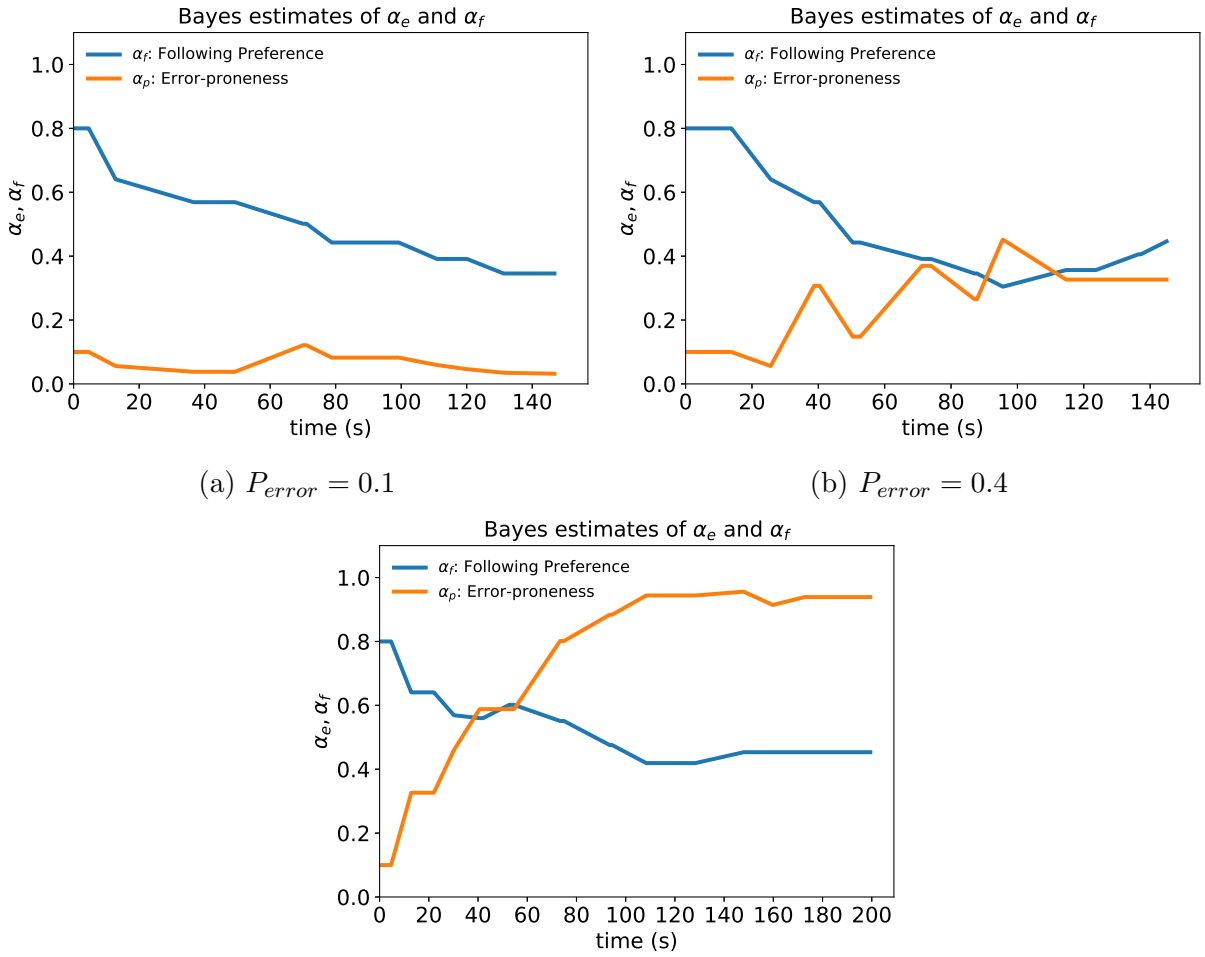


(b) $P_{error} = 0.4$



(c) $P_{error} = 0.8$

Figure 5.8: The robot's estimations of the human's following preference and error-proneness for a human with a moderate preference to follow the robot $P_{follow} = 0.6$



(a) $P_{error} = 0.1$

(b) $P_{error} = 0.4$

(c) $P_{error} = 0.8$

Figure 5.9: The robot's estimations of the human's following preference and error-proneness for a human with a slight preference to follow the robot $P_{follow} = 0.3$

Assigned tasks by the robot

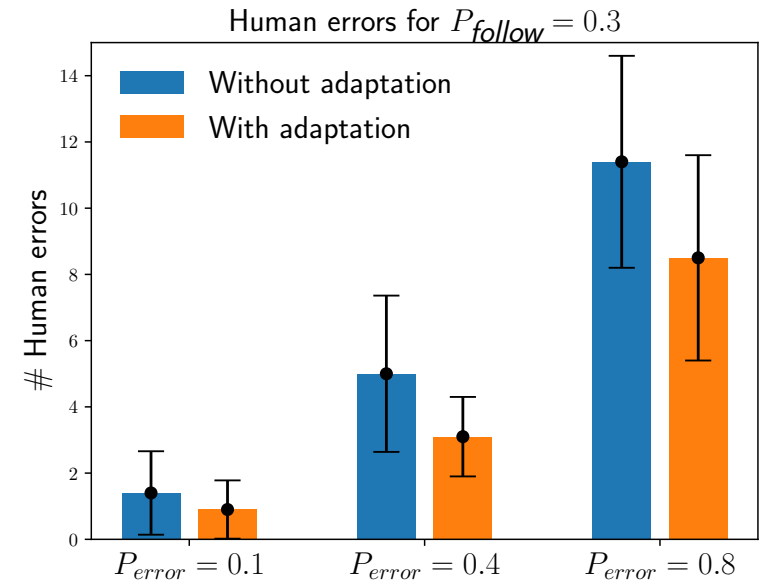
As discussed, Fig. 5.5 displays the cumulative number of assigned tasks based on the human agent's following preference parameter values. In Fig. 5.6, we have rearranged the cumulative number of assigned tasks based on the human agent's error-proneness parameter values. As shown in Fig. 5.6, the number of tasks assigned by the robot is higher when the human prefers to follow the robot. This difference decreases as the human's error-proneness increases because the robot needs to assume the leading role and assign tasks, regardless of the human's following preference.

Travel distance - Total actions

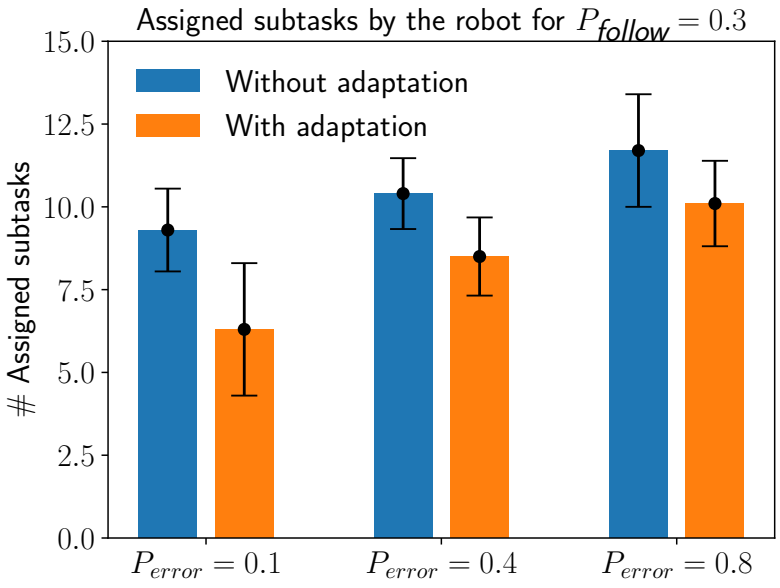
This simulation assumes that the robot is slower than the human. Thus, to reduce collaboration time, the robot completes the tasks related to workspaces 1 and 4 (W1 and W4) and leaves the tasks of workspaces 2 and 3 to the human, as we assume that the human is faster at performing them. Therefore, when the human follows the robot, even though their total working times are the same, the human covers a greater travel distance than the robot. Additionally, the human completes fewer tasks than the robot.

No adaptation ($\alpha_{\mathbf{e}} = \mathbf{0},\ \alpha_{\mathbf{f}} = \mathbf{1}$)

Here we consider a scenario in which the robot does not adapt its planning based on the human's preferences and performance. The robot's goal is only to minimize the task completion time, and the robot perceives the human agent as one who completely follows the robot's lead and makes no errors. The value of n_r^{assign} shows that when the robot ignores the human's preference for leading, it assigns many tasks to the human, regardless of the human's accuracy level. Ignoring human preferences can result in human distrust and dissatisfaction in real applications. More interestingly, even though the robot assigns even more tasks in the same cases where it considers the human's preference and accuracy, the number of wrong actions by the human increases. These errors occur because the robot does not assign the right actions at the right time, allowing the human agent to make more mistakes. Fig. 5.10 depicts the number of the human's wrong actions and the tasks assigned by the robot. The results indicate that the adaptive planner improves team performance compared to the scenario where the robot plans for the team without considering the human agent's preferences and performance.



(a) The human's errors



(b) Assigned tasks by the robot

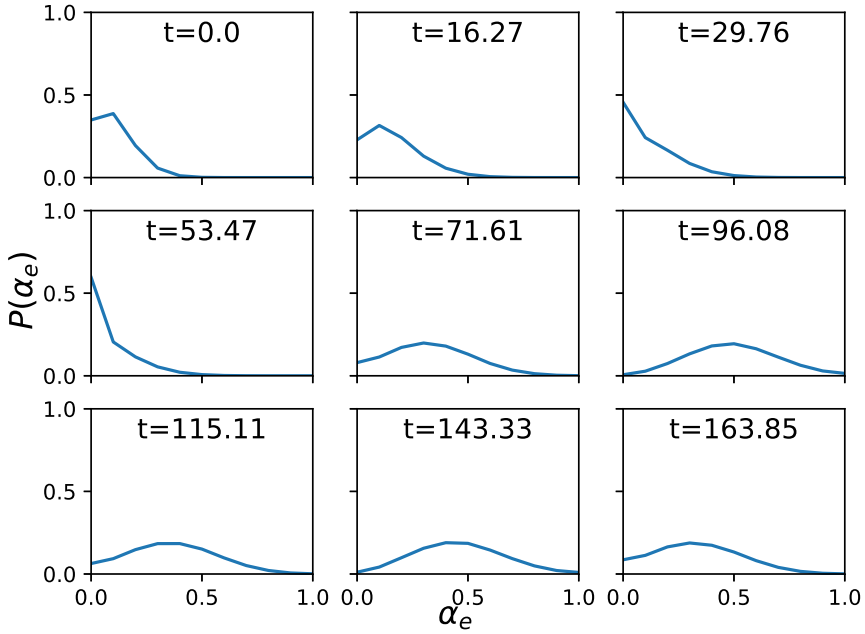
Figure 5.10: Comparison of the results for the cases with adaptation (proposed algorithm) and without adaptation (ignoring the human's performance and preferences) when the human prefers to lead ($P_{follow} = 0.3$).

Probability distributions

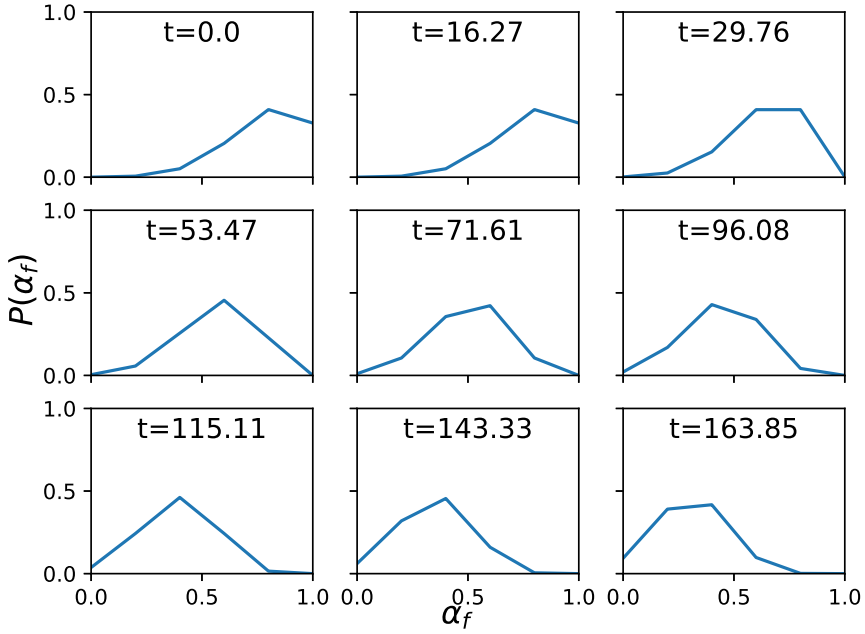
For brevity, we show only the changes of the probability distributions over the human's error-proneness (Fig. 5.11a) and following preference (Fig. 5.11b) for the case $P_{follow} = 0.3$, $P_{error} = 0.4$. As they show, the robot can successfully estimate the human's preference and accuracy and get close to the values set for the human model. Fig. 5.12 also shows the robot's estimates of the human following preference and error-proneness.

Fairness in task allocation

We also considered a scenario in which the robot takes fairness in travel distance into account. A constraint was introduced to the task allocation optimization problem to limit the disparity in travel distances between the two agents. In the simulation study, we configured the following preference and error-proneness parameters: $P_{follow} = 0.9$ and $P_{error} = 0.1$, and ran it five times. Fig. 5.13a illustrates the travel distances of the agents, which are close, as expected. However, Fig. 5.13b demonstrates that the robot completed a higher number of tasks than the human since it executed tasks in closer proximity and assigned more distant tasks to the human.



(a) Robot belief about the human's accuracy



(b) Robot belief about the human's preference to follow

Figure 5.11: probability distribution for the case $P_{follow} = 0.4$, $P_{error} = 0.5$

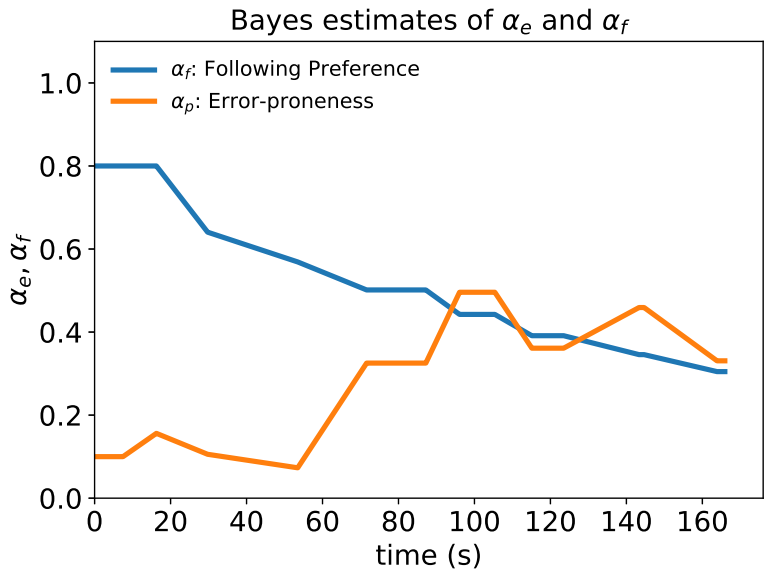
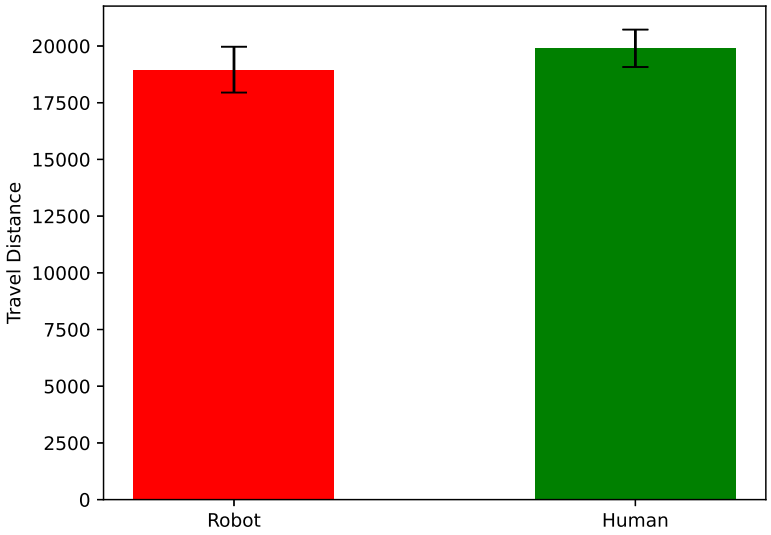


Figure 5.12: The robot's estimates of the human agent's preference and error-proneness, with the set values of $P_{follow} = 0.3$, $P_{error} = 0.4$

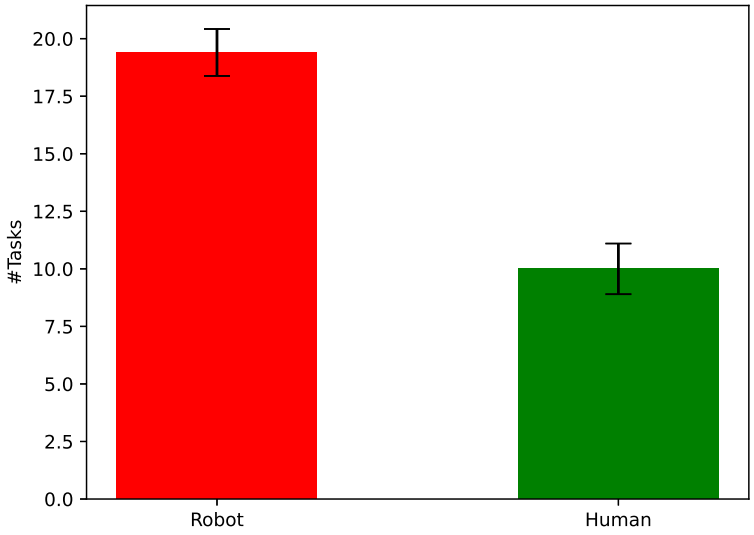
5.4 Summary and Discussion

In this chapter, to test the efficiency of the framework and the algorithm proposed in Chapter 4 for the task selection and planning problem, a collaborative scenario was designed wherein the human and the robot can select tasks and assign them to one another. Then, the collaborative task was presented in a form following the problem formulation in Chapter 4. A state estimator was also designed to enable the robot to estimate the human agent's preference and performance. The proposed algorithm was employed and evaluated in a simulated collaborative scenario.

The goal of this study was to evaluate how this planning framework could enable the robot to adapt to the human's performance level and preference by updating its plan and its belief about the human agent online. To do so, we considered a simple human model and applied the algorithm for nine different combinations of human accuracy and preference to lead. This simulation study provided the opportunity to explore various cases and assess the adaptive planner's ability to cope with them. The results have shown that the proposed algorithm enables the robot to adjust its leading role based on the human's performance and leading/following preference. In addition, comparing our algorithm with the case that the robot only considers completion time shows a significant improvement in the team's



(a) Travel distance



(b) Number of tasks completed by the agents

Figure 5.13: Considering fairness, in terms of travel distance, in task allocation for the case of $P_{follow} = 0.9$, $P_{error} = 0.1$

performance.

Although this chapter begins to explore the answers to RQ3 and RQ4, the study is limited to a simulation environment with a simplified human model. The subsequent chapter shifts its focus to the design and system evaluation of a collaborative scenario in a real environment to assess the practicality and performance of the task planning framework and algorithm.

Chapter 6

Experimental Setup and System Evaluation

This section provides a comprehensive explanation of the setup and collaborative scenario. It also details the implementation of the task planning framework tailored for the designed collaborative scenario. Finally, the section elaborates on the results of a system evaluation, in which the experimenter enacts some possible collaboration styles of the human agent. This chapter continues the exploration of RQ3 and RQ4, which was initiated in Chapter 5 but is now implemented on a real mobile manipulator robot.

This real implementation offers some advantages over the simulation study we conducted in Chapter 5. Firstly, real robots offer a more authentic and immersive interaction experience, facilitating the observation of human responses and behaviors in real-world contexts. Their social presence can evoke emotional and social responses crucial for studying contributing factors in human-robot collaboration, including trust, acceptance, and user experience. Secondly, implementing on a real robot provides the opportunity to better capture the complexities and nuances of human-robot collaboration, environmental factors, and human preference and performance. Thirdly, employing real robots enables us to validate our findings in real-world settings and better assess the practical feasibility of the adaptive planner.

6.1 Setup & Methodology

The design of the user study scenario was shaped by three pivotal factors:

1. **Collaboration:** Focusing on the planning abilities of both the human and robot, leveraging the cobot's advanced memory capacity and the human's quicker pace,

2. **Leading/following preference:** Providing the human with the agency to adjust their role as either a leader or follower,

3. **Performance:** Introducing a task that imposes cognitive load and penalizes errors.

6.1.1 Setup

In this specific scenario, the human-cobot team collaborates to select colored blocks from tables, adhering to a specified pattern, and arrange them within a shared area with a predetermined number of spots. The designed collaborative scenario closely resembles the kitting task. The kitting task involves gathering a specific set of components for a defined purpose, which is then directed to workstations for the assembly of intermediate or end products [113,114]. Additional details about the setup are visualized in Fig.6.1 and Fig.6.2. The camera's placement in Fig.6.1 corresponds to the viewpoint from which the image in Fig.6.2 was captured. Additionally, Fig. 6.1 illustrates the location of the experimenter's table, indicating the position of both the computer and the experimenter themselves.

- **Shared area:** The shared area consists of four workspaces denoted as W_1 through W_4, each containing five numbered spots. As per the predetermined pattern, participants are required to place a colored block in each spot, adhering to the numerical sequence, noted on each spot, from one to five.

- **Patterns:** The pattern depicted in Fig. 6.3a is an example of the pattern the human-robot team must follow when filling the spots. Participants are initially presented with a pattern (Patterns A_1, B_1, C_1, and D_1) on a sheet of paper, and they are given 45 seconds for memorization. Subsequently, they return the pattern, and the experimenter provides them with a partially known version of the same pattern (Patterns A_2, B_2, C_2, and D_2), again on a sheet of paper, in which certain spots contain two colors, one of which is correct. This second pattern acts as a cue to recall the initial pattern, and participants are allowed to retain it until the end of the task. Patterns A_2, B_2, C_2, and D_2 feature 9, 12, 6, and 9 partially known spots, respectively, introducing varying levels of difficulty and cognitive load. We asked a small number of our lab members to test these patterns and rank their difficulty levels to ensure that they present varying levels of difficulty in terms of memorization and recall. The

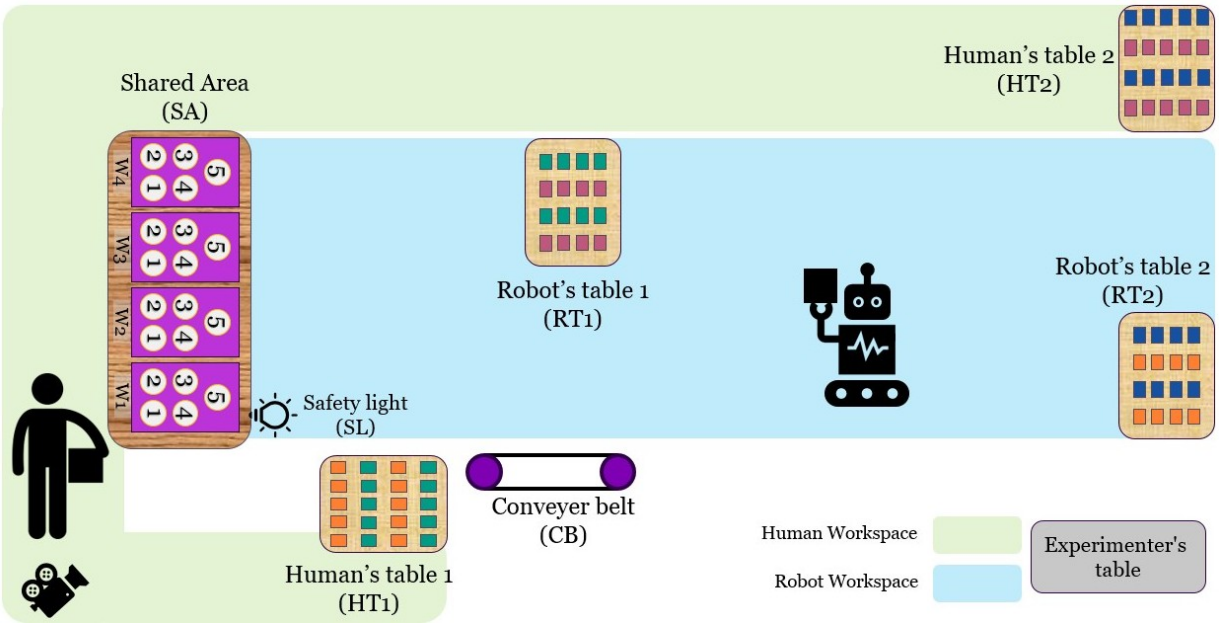


Figure 6.1: The layout of the experimental setup includes two tables in the robot's workspace and two tables in the human agent's workspace. Additionally, there is a shared area (table) where both agents need to arrange blocks. A conveyor belt beside human table 1 allows the robot to return a block to the human agent.



Figure 6.2: A view of the experiment environment, taken from the location of the camera in Fig. 6.1. The Fetch robot is positioned in its workspace between its two designated tables, as illustrated in Fig. 6.1. The nearest table in the figure, located next to the conveyor belt, belongs to the human and contains orange and green blocks. On this table, there is also a tablet on which the GUI is installed. The human agent's other table is situated at the corner of the room, containing blue and pink blocks. Four workspaces are present on the shared table, along with a light bulb. Safety tapes and cones separate the work areas of the agents.

Table 6.1: Distribution of blocks with respect to the distance to the shared area

Color	**Human**	**Robot**
Green	Close	Close
Pink	Close	Far
Orange	Far	Close
Blue	Far	Far

decision to incorporate a pattern memorization task was made to introduce a cognitive challenge while adhering to the limitations of a relatively short collaborative scenario. Conducting an extensive experiment that might place excessive mental and physical burdens on participants was considered impractical and challenging to obtain ethics approval. Therefore, concise tasks with durations of around 12-20 minutes were chosen to ensure participants' sustained mental engagement.

- **Tables:** The human and robot have separate work areas, with two tables in each. In the human work area, one table is close to the shared area and contains blocks of two colors: green and orange. The other table is far from the shared area (i.e., requires traveling a longer distance) and contains blocks of two other colors: blue and pink. Likewise, the robot's work area has two tables, one distant from the shared area, with blue and orange blocks, and one close to the shared area, including green and pink blocks. Table. 6.1 summarizes the distribution of the blocks. Participants were also told to pick up only one block at a time, as the cobot's gripper only has the capacity to pick up a single block. This table and block distribution help evaluate the robot and the human's decision-making and task planning (i.e., selecting and assigning tasks) regarding their preference, travel distance, completion time, and proximity to the optimal plan.

- **Blocks:** Four unique block colors are employed: green, blue, pink, and orange. The human's tables are equipped with ten blocks of each color, whereas the robot's tables contain eight blocks of each color. This ensures a total of eighteen blocks for each color, exceeding the necessary five blocks per pattern to account for possible errors. A larger quantity of blocks is placed on the human agent's table, taking into account their swifter working pace. Each block has an ArUco marker affixed to it, aiding the robot in locating and handling the blocks within the room.

- **Conveyor Belt:** I have built a conveyor belt that has been integrated to allow the robot to return any misplaced blocks to the human. The robot places a block

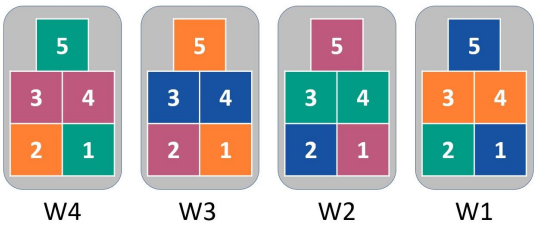
at one end of the conveyor belt, which then transports the block to the human's side. This design decision accomplishes two objectives: it separates the workspaces of the human and the robot, mimicking a real industrial environment, while also enhancing safety by reducing the need for the robot and human to interact closely. This conveyor

- **Light Bulb:** On the shared table, there is a red light bulb that lights up when the robot approaches to place or retrieve a block from the shared table. Participants are informed that this light serves as a signal to avoid approaching the table due to safety precautions. However, they can continue planning, navigating within their designated work area, and collecting blocks from other tables. In addition to safety concerns, the light bulb also helps regulate the participants' pace, prompting them to wait for the robot and ensuring that the collaboration is not perceived as a competitive race.
- **Safety Measures:** Alongside the safety light bulb, we have set up safety tape and floor cones to create a well-defined physical boundary that separates the work areas of both agents. These measures are complemented by continuous vigilance from the experimenter, who is prepared to intervene as needed. Intervention options include manually controlling the robot using a joystick or promptly activating the emergency safety stop button. During the sessions, there were instances where manual control of the robot was necessary, particularly when it came too close to the tables.

6.1.2 Tasks

Each participant is asked to complete four tasks:

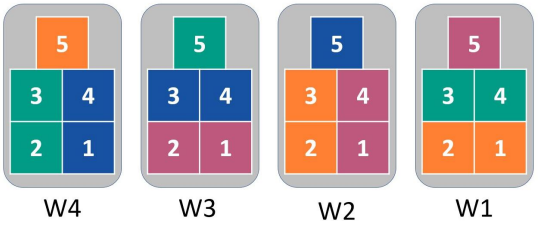
- **Task 0:** In this initial task, participants engage in solo work without the robot's involvement. For all participants, we provide pattern A and follow a consistent procedure. They are tasked with memorizing pattern A_1, printed on a sheet of paper, within a 45-second timeframe and returning it. Subsequently, we furnish them with pattern A_2, which they can retain until the task's conclusion. It is essential to note that errors made by participants in this task do not incur penalties; rather, this task primarily serves to acquaint participants with the setup and offer them practice in block placement. This practice encompasses aspects such as adhering to the correct block order, picking up a single block at a time, and ensuring the markers on the blocks face the robot. While focusing on task practice, we also assess participants' accuracy, self-confidence, and their perception of workload.



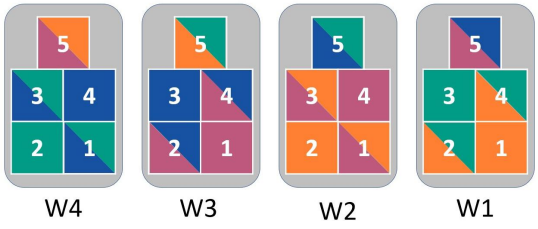
(a) Pattern A_1



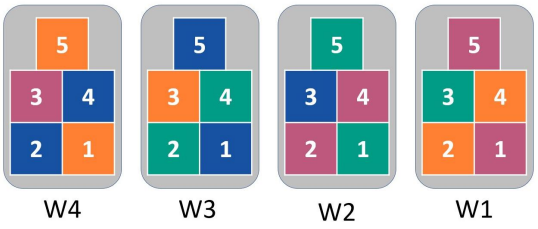
(b) Pattern A_2



(c) Pattern B_1



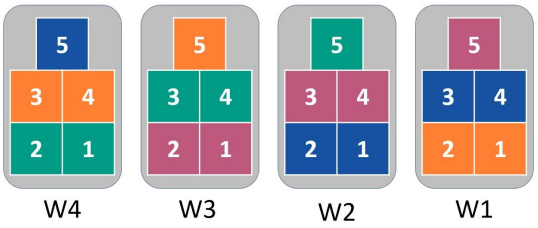
(d) Pattern B_2



(e) Pattern C_1



(f) Pattern C_2



(g) Pattern D_1



(h) Pattern D_2

Figure 6.3: **a, c, e, g:** Patterns A_1, B_1, C_1, and D_1 represent the set of patterns presented on paper sheets, requiring participants to memorize them within a 45-second timeframe before returning them to the experimenter. **b, d, f, h:** Patterns A_2, B_2, C_2, and D_2 are variations with some partially known spots, serving as cues for participants throughout the collaborative task, and they are permitted to retain these patterns until task completion to aid in recalling the first pattern.

- **Tasks 1, 2, 3:** The robot collaborates with the human in these tasks. We consider six distinct permutations (modes) that determine the order of presenting patterns B, C, and D (e.g., B, C, D, C, B, D). Participants are randomly allocated to one of these six modes (permutations) to ensure an equal distribution among them. Following a consistent procedure, participants are tasked with memorizing the initial pattern within a 45-second timeframe and subsequently returning it. Then, they are provided with the second pattern. To emulate real-world collaborative scenarios in which errors carry consequences, participants are informed that, upon declaring the task as completed, for each block incorrectly placed on the table, a deduction of $1 will be applied to their total remuneration. This method of conveying disinformation, classified as a form of deception, received approval from the University of Waterloo Human Research Ethics Board.

6.1.3 Agents' Actions

We considered a set of six different actions for each agent, the human and the robot. These actions, listed in Table 6.2, are the same for both agents, providing them with a similar level of agency. Note, the feasibility of actions depends on the state of the task, and at each decision step, some of them may not be feasible. For example, when no task has been allocated to the robot, action H4, performing a task assigned by the human, is not applicable. Additionally, to provide the robot with greater autonomy in adapting its leadership role or reassuming it when human performance is poor, we intentionally designed the rejecting option (Action H6) to be less straightforward for the human (will be explained later). Action H6 involves the consecutive performance of H4 and H3 without actually performing them physically. Further elaboration on this will be provided in the subsequent paragraph. Fig. 6.4 displays the state graph of a single subtask, considering the possible actions of both human and robot agents as outlined in Table 6.2.

6.1.4 Human-robot communication

Both agents are required to uphold communication to keep each other informed about their forthcoming actions, as specified in Table 6.2. This communication is facilitated through a purpose-built graphical user interface (GUI) installed on a tablet. Participants have the flexibility to position the tablet on a table within the room or hold it according to their preference. This GUI serves as a medium for participants to assign tasks to the robot and convey information about their intended actions. Similarly, the robot can employ

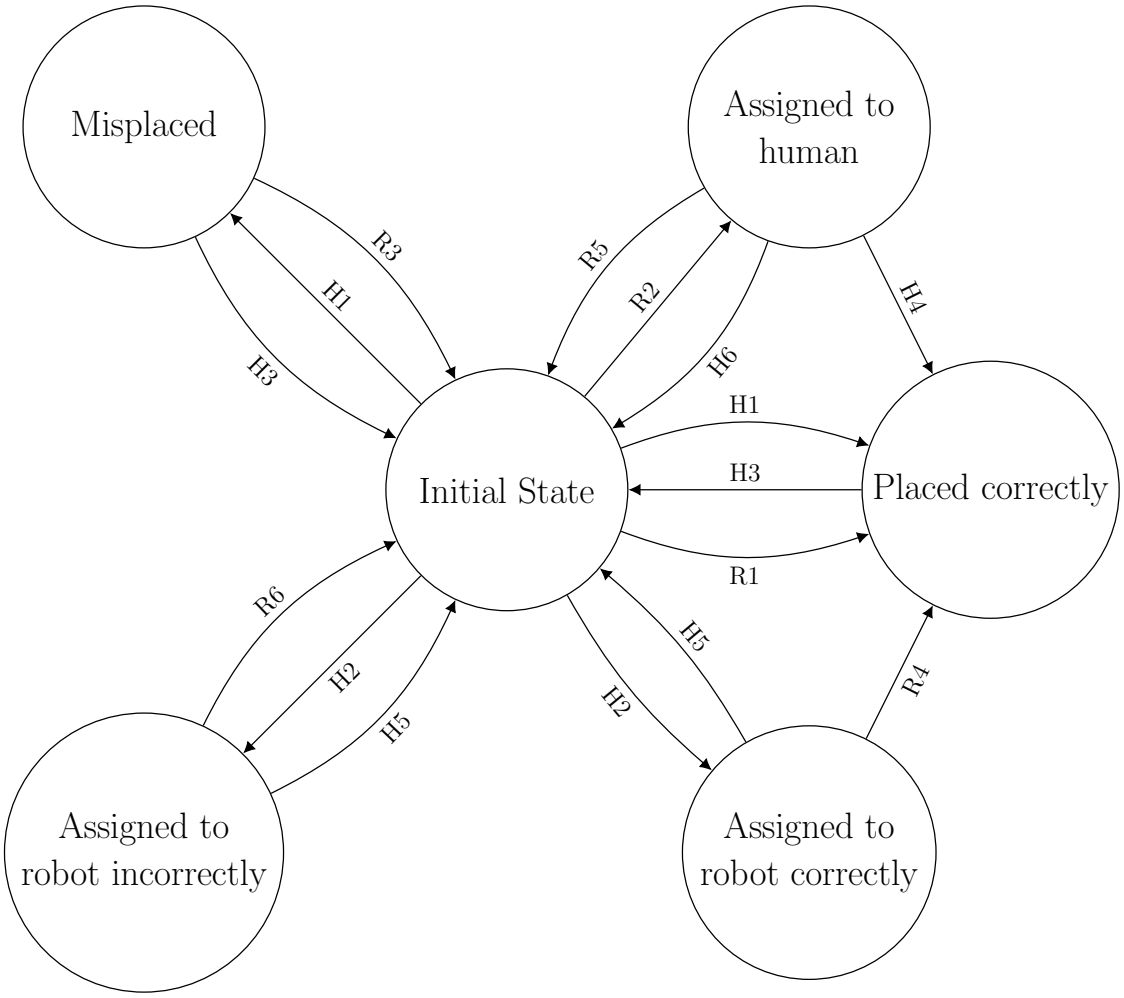


Figure 6.4: State graph of a single subtask. There are six potential states for each subtask based on the actions detailed in Table 1. Initially, each subtask is in the "Initial state" state. To complete a subtask, it must transition to the "Placed correctly" state and remain in this state. The states "Assigned to robot correctly" and "Assigned to robot incorrectly" occur when the human agent respectively assigns a subtask to the robot with the correct or wrong color. The "Misplaced" state is reached when the human places a wrong color on the shared area for a subtask. The state "Assigned to human" happens when the robot assigns a subtask to the human agent.

Table 6.2: The human and robot's sets of actions

Agent	Actions
Human	H1- Selecting a task for themselves H2- Assigning a task to the robot H3- Returning a block from the shared workspace H4- Performing a task assigned by the robot H5- Canceling a task assigned to the robot H6- Rejecting a task assigned by the robot
Robot	R1- Selecting a task for itself R2- Assigning a task to the human R3- Returning a *wrong* block from the shared workspace R4- Performing a *correct* task assigned by the human R5- Canceling a task assigned to the human R6- Rejecting a task assigned by the human

the GUI to assign tasks to the participant and communicate its own planned actions. An image of the GUI is provided in Fig. 6.5. Notably, the GUI enforces constraints to prevent the human agent from selecting actions that would violate precedence constraints or interfere with tasks already initiated by the robot. Participants received instructions on operating the GUI and had the opportunity to practice using it before commencing the tasks. Additionally, participants were instructed to scan the marker on the blocks before placing them on the shared area, as this action enables the robot to identify the block's ID for potential future retrieval. The GUI streamlines this process by automatically activating the tablet's camera, allowing participants to scan the marker. To maintain brevity, the details of the GUI's design and implementation are omitted due to space limitations.

Remark. In case participants wish to decline a task assigned by the robot, the sequence of actions involves initially performing the assigned task (H4) and subsequently engaging in the action for returning (H3). Both of these actions are executed through the GUI, without any requirement for physical execution.

6.1.5 Robot

For our experiments, we utilized the Fetch mobile manipulator robot [115]. The robot was configured to perform pick-and-place operations autonomously. While creating a fully

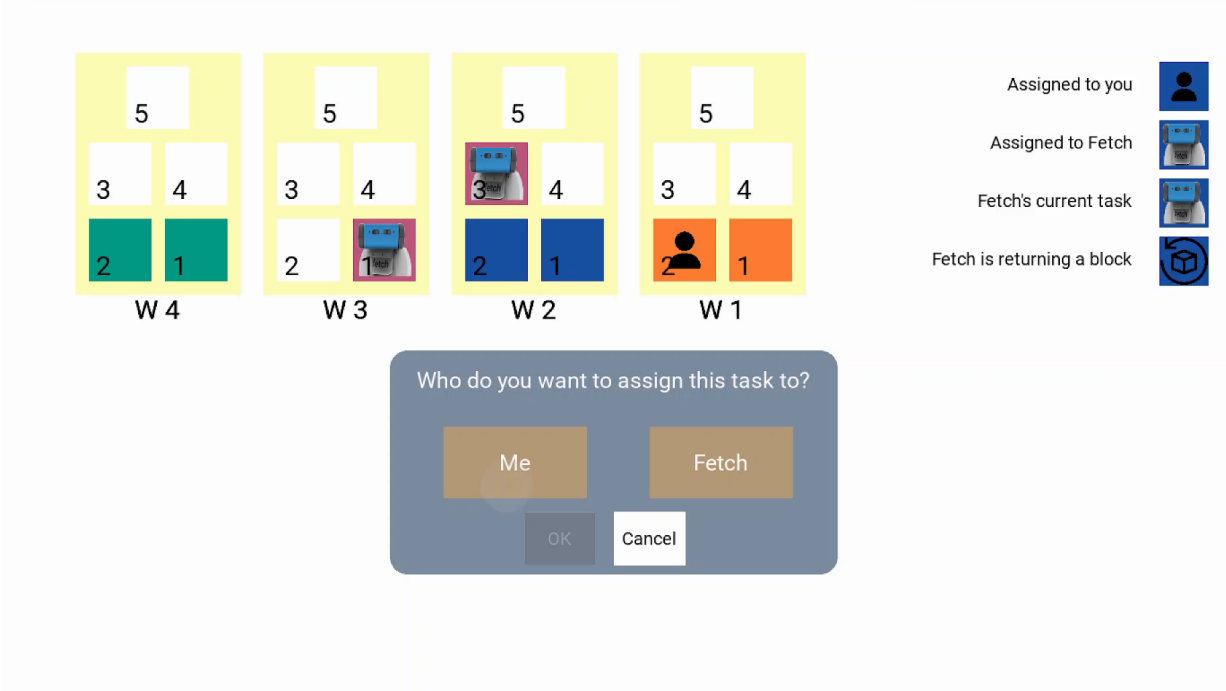


Figure 6.5: A screenshot of the graphical user interface (GUI) that enables participants to convey their actions to the robot and obtain information regarding the robot's decisions and actions (refer to Table 6.2 for a list of actions).

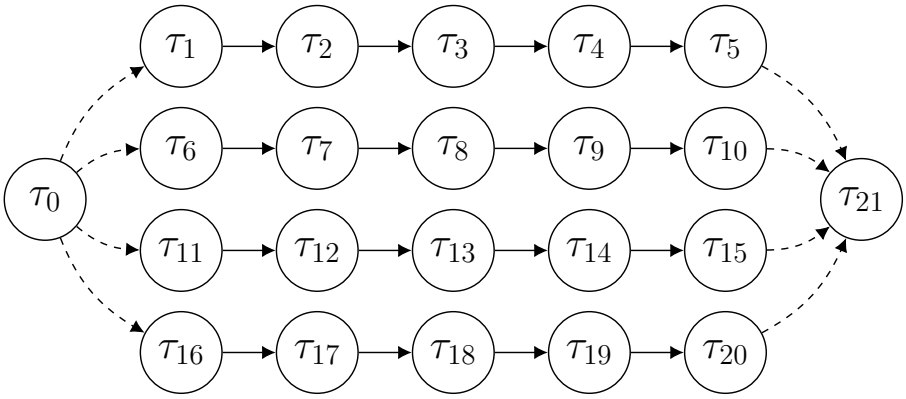


Figure 6.6: Task graph of the experiment, including twenty subtasks (τ_1 - τ_{20}) and two dummy nodes representing starting (τ_0) and finishing points (τ_{21}).

autonomous pick-and-place system demanded a considerable amount of time, it was crucial to ensure that participants experienced a work environment resembling real industrial settings. To maintain brevity, we skip the explanation of the autonomous pick-and-place implementation on the robot.

6.2 Adaptation & Planning

Here, we briefly explain the components and some details of implementing the task planning architecture for the scenario designed.

Collaborative task: The tasks and their associated precedence constraints are conveyed using a directed acyclic graph termed the task graph. This graph depicts tasks as vertices, while edges signify their precedence constraints. Fig. 6.6 provides a visual representation of the initial task graph for the experiment, including two dummy nodes designating the starting (T_0) and finishing (T_{21}) points.

In this experiment, we assume the same and time-invariant speed for all participants, the average human walking speed (1.3 m/s). In addition, we tested the required time for the robot to complete the pick-and-place scenario from different tables and used the average time for the near and far tables. Thus, the nominal required processing times for tasks are fixed for both the human and the robot.

Following preference and Performance: The robot uses the single scalar random variables α_f and α_e, respectively, to capture the human agent's following preference and

performance (i.e., error-proneness or inaccuracy). For both, we consider a tuple of possible discrete values, $W = (w_0, \ldots, w_{10}) = (0, 0.1, 0.2, \ldots, 1)$ with steps of 0.1. In other words, there exists $i \in \{0, \ldots, 10\}$ such that $w_i = \alpha_f$ (the same for α_e). For α_f, a value closer to zero indicates a leading preference, while values closer to one suggest a following preference. Additionally, values of α_e closer to one represent higher error-proneness (low accuracy), while values closer to zero indicate lower error-proneness (higher accuracy).

At the beginning of the session, as a default, the robot assumes that the human agent prefers to follow it and has high accuracy, so the robot sets its initial belief about the human agent as follows:

$$P[\alpha_f = w_i] = b(i; n = 10, p = 0.7),\ i \in \{0, \ldots, 10\},$$

and similarly,

$$P[\alpha_e = w_i] = b(i; n = 10, p = 0.1),\ i \in \{0, \ldots, 10\},$$

where $b(i; n, p)$ is a binomial distribution.

Task allocation: The problem of task allocation is formulated using (4.1), where

$$C_{\tau_i}(a) = \begin{cases} t_i^h \alpha_f + c_f(1 - \alpha_f) + c_v x_{\tau_i}^{robot} & a = Human \\ t_i^r + \alpha_e c_e + c_v x_{\tau_i}^{human} & a = Robot \end{cases}. \tag{6.1}$$

The parameter c_f represents the penalty incurred when assigning the subtask to a human agent who prefers to lead, while c_e denotes the penalty imposed for not allocating subtasks to the human agent who makes mistakes. Assigning subtasks to the human agent allows the robot to inform them about the next blocks to be placed on the shared table, minimizing the risk of wrong decisions by the human agent. Additionally, a penalty cost c_v is applied when the robot (human) is assigned a subtask that has already been allocated to the human (robot) agent, $x_{\tau_i}^{human} = 1$ ($x_{\tau_i}^{robot} = 1$). This penalty prevents frequent and substantial changes in task allocation. In addition, as a problem-dependent constraint, we need to add a constraint ensuring the allocation of at least one subtask from τ to the robot at each decision step. This constraint ensures the robot will start doing another subtask or fixing the human agent's errors.

The derivation of optimization problem 4.1 for the task within the designed scenario is equivalent to the optimization problem represented by 6.2, which is formulated as a Mixed-Integer Linear Program (MILP). In this context, we introduce decision variables $q_i \in Q$ for $i \in I = \{1, \ldots, n_t\}$ to determine the assignment of agents to subtasks, where n_t represents the total number of remaining tasks. Specifically, when $q_i = 1$, it signifies

that subtask i is assigned to the human agent, while $q_i = 0$ indicates the allocation of the subtask to the robot.

$$Q^* = \arg\min_{\{q_i\}} z \tag{6.2}$$

subject to

$$z - \sum_{i \in I} q_i \Big(t_i^h p_f + c_f (1 - p_f) + c_h x_{\tau_i}^{robot} \Big) \geq 0 \tag{6.3}$$

$$z - \sum_{i \in I} (1 - q_i) \Big(t_i^r + p_e c_e + c_r x_{\tau_i}^{human} \Big) \geq 0 \tag{6.4}$$

$$\sum_{i \in U} q_i \leq |U| - 1, \tag{6.5}$$

where U is the set of indexes of all feasible subtasks that the robot can perform immediately by placing the blocks on the shared area. We ignore details of obtaining U. This inequality ensures that the robot will start placing a new block after placing the previous one.

Fig. 6.7 illustrates an example of the updated task graph following task allocation. Subtasks 6 and 16 (τ_6 and τ_{16}) are shaded in gray to indicate their completion. However, for subtasks 1 and 17, the human agent mistakenly placed the wrong colored blocks on the shared area. Consequently, the robot must rectify these errors by executing subtasks τ_1^e and τ_{17}^e. Subtask 11 (τ_{11}) has already been assigned to the human. Additionally, the robot has allocated subtasks 1, 2, 4, 5, 7, 9, 10, 13, 18, and 20 to the human. Considering precedence constraints, the robot can initiate subtask τ_7^a to allocate and instruct the human agent through the GUI to perform subtask 7 (τ_7). The robot has also allocated subtasks 3, 8, 12, 14, 15, 17, and 19 to itself.

Task scheduling: After obtaining the optimal task allocation, the robot needs to find the optimal schedule by solving the optimization problem in 4.1. After allocating tasks to both the robot, denoted as τ_{robot}, and the human, denoted as τ_{human}, such that $\tau_{new} = \tau_{robot} \cup \tau_{human}$, the robot proceeds to solve the optimization problem 6.6 in order to find an optimal task schedule. Within this problem, the decision variables $S = \{s_{\tau_i}\}$ dictate the start times of the subtasks. Binary decision variables $O = \{o_{i,j}\}$ are employed to determine if subtask τ_i comes before or after τ_j. The set V comprises indices for all feasible subtasks in τ_{robot} that the robot can immediately perform by placing the blocks on the shared area. Moreover, binary decision variables $B = \{b_{\tau_i}\}$ are used to indicate whether $\tau_i \in V$ begins at $s_{\tau_i} = 0$.

$$\min_{\{S,O,B\}} z \tag{6.6}$$

subject to

$$P\left(\tau_i, \tau_j\right).f_{\tau_i} \leq s_{\tau_j}, \qquad \forall \tau_i, \tau_j \in \tau_{new} \tag{6.7}$$

$$s_{\tau_i} - f_{\tau_j} + M(1 - o_{i,j}) \geq 0 \qquad \forall \tau_i, \tau_j \in \tau_{human} \tag{6.8}$$

$$s_{\tau_j} - f_{\tau_i} + Mo_{i,j} \geq 0 \qquad \forall \tau_i, \tau_j \in \tau_{human} \tag{6.9}$$

$$s_{\tau_i} - f_{\tau_j} + M(1 - o_{i,j}) \geq 0 \qquad \forall \tau_i, \tau_j \in \tau_{robot} \tag{6.10}$$

$$s_{\tau_j} - f_{\tau_i} + Mo_{i,j} \geq 0 \qquad \forall \tau_i, \tau_j \in \tau_{robot} \tag{6.11}$$

$$s_{\tau_i} - Mb_{\tau_i} \leq 0 \qquad \forall i \in V \tag{6.12}$$

$$\sum_{i \in V} b_{\tau_i} \leq |V| - 1 \tag{6.13}$$

$$z - f_{\tau_i} \geq 0 \qquad \forall \tau_i \in \tau_{new} \tag{6.14}$$

where M is a large positive constant. In this problem, we assume that rejecting the human agent's assignments and allocating tasks to the human takes zero seconds.

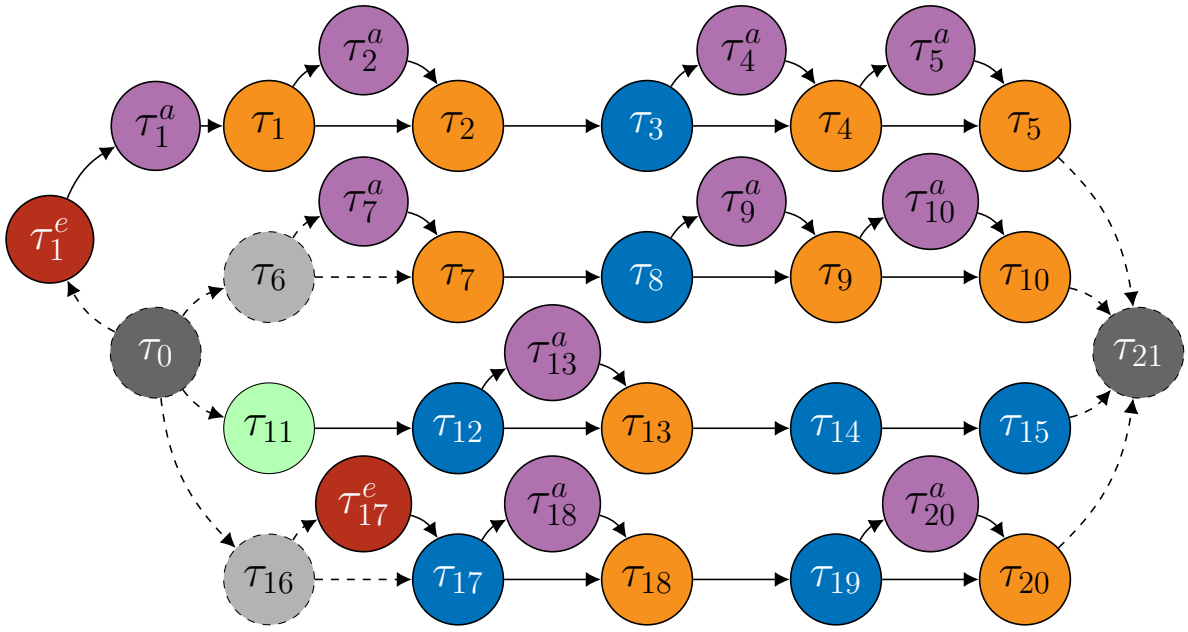


Figure 6.7: Temporary task graph of the experiment after task allocation.
Blue: Robot's tasks, Orange: Human's tasks, Cyan: Assigning tasks to the human, Red: Correcting human errors, Green: Already assigned tasks, Gray: Finished tasks

Solving optimization problems: The task allocation and task scheduling problems, including their constraints, have been reformulated into mixed-integer linear programs, which are recognized as NP-hard optimization problems. Our simulation experiments were conducted on a computer running Ubuntu 18.04, equipped with an Intel Core i7-11700 CPU with 8 cores operating at 2.5GHz and 16 GB of RAM. We utilized the Gurobi mathematical optimization solver, imposing a time limit (4 seconds) to terminate the solver if it continues searching for additional solutions. We also set a relative gap tolerance of 0.004 as an additional stopping criterion. We adopted a warm-start approach, providing the solver with a partially valid initial solution derived from the previous step's solution. Consequently, for the initial step, we could solve the problem offline.

Updating α_f and α_e: In updating the robot's estimate of the human agent's preference to follow and their performance, the robot needs to detect changes and adapt its planning accordingly. That is, the estimation method has to be sensitive enough to consider these changes. However, it also must not be oversensitive to affect the planning abruptly. For example, a single mistake in the human's decision making will not be reacted upon precipitously. Taking advantage of the "bounded memory adaptation model" approach, proposed in [80], we use a history of k-step in the past to estimate the human agent's preference and performance. A small value of k makes the estimation sensitive to changes, and a large value leads to measuring the overall leading/following preference and performance. In this work, we chose $k = 3$.

To update the belief, we employ the belief update method employed by [80] to estimate human adaptability. According to this method, the system needs to be considered as a factorization of observable (X) and unobservable (Y) state variables of the system $S : X \times Y$. Subsequently, belief update can be computed as:

$$\begin{aligned} b'(y') =& \eta Z\big(x', y', a^R, o\big) \sum_{y \in Y} T_x\Big(x, y, a^R, a^H, x'\Big) \\ & T_y\big(x, y, a^R, a^H, x', y'\big) \pi^H\big(x, y, a^H\big) b\big(y\big), \end{aligned} \tag{6.15}$$

where T_x and T_y are the transition functions, z is the observation function, and π^H is the human action model (policy).

Following preference: Here to update α_f, we use Eq. 6.15, where $y = \alpha_f$ and b represents belief on α_f. We consider $Z = 1$ and $T_x = 1$. We also consider $T_y\big(x, y, a^R, a^H, x', y'\big) = \mathbb{I}(y = y')$, where $\mathbb{I}$ is an indicator function. This is based on the reasonable assumption that the human agent's preference changes infrequently and is usually fixed. The human agent's strategy is determined based on the analysis of the preceding three steps $(k = 3)$ in their actions. The actions taken into consideration by the robot for updating

its policy regarding human following preferences (P_f) include the actions of assigning a subtask to the robot (F_1), performing a subtask assigned by the robot (F_2), or refraining from performing a subtask assigned by the robot (F_3). By denoting the occurrences of F_1, F_2, and F_3 as f_1, f_2, and f_3, respectively, within the sequence of human actions spanning a history of three steps, the resulting human policy can be defined as follows:

$$\pi_f^H(y, a^H) = \begin{cases} \frac{\alpha f_1 + f_2}{\alpha f_1 + f_2 + f_3} y & a^H \in F_1 \cup F_2 \\ \frac{f_3}{\alpha f_1 + f_2 + f_3} y & a^H \in F_3 \end{cases}, \tag{6.16}$$

where $\alpha > 1$ is a parameter that weighs cases where the human assigns a subtask to the robot more heavily. Therefore, based on Eq. 6.15, the update of belief on α_f can be written as:

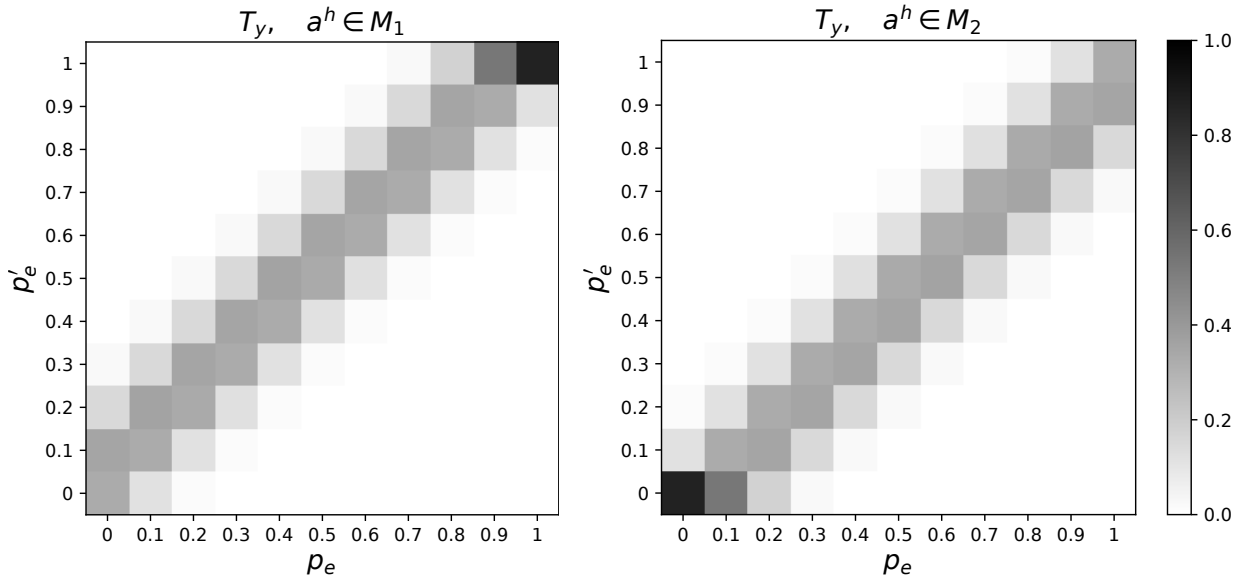
$$b'(y') = \eta \sum_{y \in Y} \mathbb{I}(y' = y) \pi_f^H(y, a^H) b\big(y\big), \tag{6.17}$$

Human error: To update α_e, we utilize Eq. 6.15, where $y = \alpha_e$, and b denotes the belief on α_e. We consider $Z = 1$ and $T_x = 1$. Modeling human error, specifically the human memory model in this particular scenario, presents significant challenges and falls outside the primary focus of this thesis. Nevertheless, we adopt a simplified model for T_y and π^H as necessary components for estimating p_e. We define two functions, $g_l(y)$ and $g_u(y)$, which respectively identify the closest value less than and greater than y within the set Y. Consequently, we formulate T_y as follows:

$$T_y = \begin{cases} p\big(y' \leq Z < g_u(y')\big), \\ \qquad Z \sim \mathcal{SN}\big(g_u(y),\, \sigma^2, \beta_1\big) & \text{if } a^H \in M_1 \\ p\big(g_l(y') \leq Z < y'\big), \\ \qquad Z \sim \mathcal{SN}\big(g_l(y),\, \sigma^2, \beta_2\big) & \text{if } a^H \in M_2 \end{cases}, \tag{6.18}$$

where β is the skewness factor of the skew-normal distribution function $\mathcal{SN}(g_l(y),\, \sigma^2, \beta)$. The erroneous actions (M_1) and correct ones (M_2), if they are not assigned to the human by the robot, are taken into account for updating α_p. The transition probability T_y heatmap is illustrated in Fig. 6.8. Considering m_1 and m_2 as the frequency counts of M_1, M_2, in the preceding three steps ($k = 3$), the human error model is as follows:

$$\pi_e^H(y, a^H) = \begin{cases} \frac{m_2}{m_1 + m_2} y & a^H \in M_2 \\ \frac{m_1}{m_1 + m_2} y & a^H \in M_1 \end{cases}. \tag{6.19}$$



(a) Wrong action (b) Correct action

Figure 6.8: Estimating the human's accuracy: Transition probability, T_y

Therefore, based on Eq. 6.15, the update of belief on α_e can be written as:

$$b'(y') = \eta \sum_{y \in Y} T_y\big(y, a^H, y'\big)\pi_e^H(y, a^H)b\big(y\big). \tag{6.20}$$

6.3 System Evaluation

Concerning the explained user study, we performed a system evaluation study prior to the user study experience to assess the algorithm performance, identify issues, optimize the system, improve the user experience, and validate the system. This can ultimately lead to more accurate and reliable results in our user study.

To do so, we mainly focused on the in-person phase, where participants need to collaborate with the robot, and enacted four possible different behaviors that participants may have during the experiment. These behaviors were intentionally enacted by the experimenter and included:

1. a preference for following with low accuracy,

Table 6.3: Evaluation results: Four different enacted behaviors (one run per case)

Human		Wrong Actions	Assigned Tasks		Travel Distance (m)		Total Tasks		Total Time (min)
Preference	Accuracy	n_h^{wrong}	n_h^{assign}	n_r^{assign}	d_h^T	d_r^R	n_h^T	n_r^T	t_{total}
Follow	High	0	0	9	286	36	13	7	7.65
Follow	Low	2	0	11	300	48	14	10	10.1
Lead	Low	8	7	12	300	64	14	10	13.1
Lead	High	0	10	0	180	76	11	9	11.1

2. a preference for following with high accuracy,
3. a preference for leading with high accuracy,
4. a preference for leading with low accuracy.

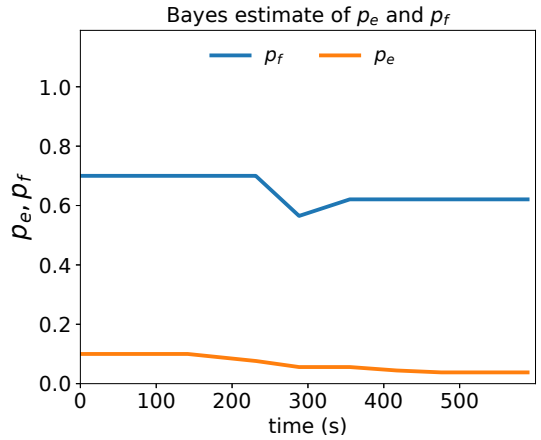
We also took into account the scenario that the participant preferred to be the leader and performed well initially, but then forgot the pattern and made errors while still wanting to lead the robot. By considering this case, we aimed to assess the robot's ability to adapt quickly to new conditions. Furthermore, this scenario represents a possible situation where participants remember the pattern better in the beginning but gradually forget it over time.

6.3.1 System Evaluation Results: Enacting Different Human Behaviors

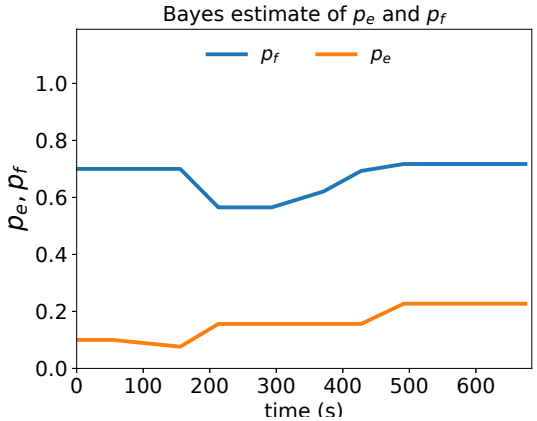
In this part, we evaluated the effectiveness of the proposed framework by conducting experiments that enacted four distinct types of human behavior.

Results are shown in Table 6.3, where n_h^{wrong}, n_h^{assign}, and n_r^{assign} are the number of the human's wrong actions, the number of actions assigned to the robot by the human, and the number of actions assigned to the human by the robot. In addition, we measured the total number of actions done by the human and robot, denoted respectively by n_h^T and n_r^T. The travel distance by the robot and human were also measured approximately by measuring the distance between the picking and placing tables. The last column of the table, t_{tot}, shows the task completion time.

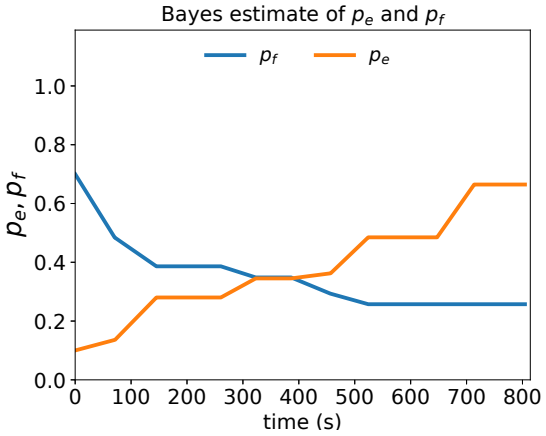
Based on Table 6.3, as expected, the team has the best performance in terms of collaboration time when the human preferred to follow the robot and had high accuracy. In this



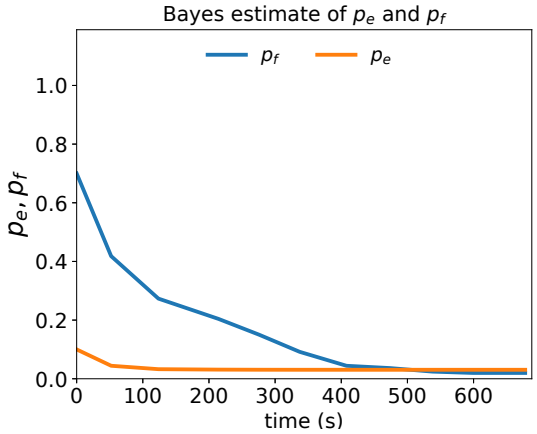
(a) Following preference with high accuracy



(b) Following preference with low accuracy



(c) Leading preference with low accuracy



(d) Leading preference with high accuracy

Figure 6.9: The robot's belief regarding the human's preference and performance (computed as the expectation of the posterior distribution)

case, the human did not have any mistakes and followed the optimal plan generated by the robot. The robot assigned more tasks related to blue objects (which were far from both the human and the robot) to the human agent since they could complete those tasks more quickly. The robot also gave the human agent more tasks involving orange objects, which were closer to the human and farther from the robot. Meanwhile, the robot handled the tasks with pink objects because they were nearer to the robot and distant from the human agent. As a result, the human agent traveled a greater distance than the robot, since they completed more tasks that were farther away.

In the second scenario, where the human had low accuracy and had the potential to make mistakes, the team's performance was still somewhat similar to the previous scenario since the human followed the robot and didn't make any mistakes. In other words, despite the human's potential to make errors, following the robot's plan prevented the human from making any mistakes. However, a few mistakes by the human agent caused the robot to assign more tasks to it, and the collaboration time increased as the robot needed to fix the human's errors.

In the scenario where the human preferred to lead but had low accuracy, more tasks were assigned to the robot, and most of them were wrong. Additionally, the human agent placed some wrong objects on the shared workspace. The robot was able to recognize the human's low accuracy and gradually assumed more responsibility for leading and assigning tasks. However, the human's preference to lead, combined with their low accuracy, resulted in a longer completion time and greater travel distance compared to the other scenarios. In the case where the human preferred to lead and had high accuracy, the robot allowed the human to take the leading role. However, the human agent assigned more tasks related to blue objects (which were far from both the human and the robot) to the robot. As a result, the robot traveled a greater distance, while the human agent traveled less, leading to a longer overall completion time.

Fig. 6.9 shows the changes in the robot's belief regarding the human agent's leading preference and accuracy in each scenario. In Figures 6.9a and 6.9b, the robot estimated that the human preferred to follow. In the second scenario, it was assumed that the human had low accuracy and forgot the pattern. However, in both the first and second scenarios, the human agent followed the robot without making significant mistakes, leading the robot to consider them as nearly accurate. As depicted in Fig. 6.9a and 6.9b, the robot perceived the human as an accurate agent, despite its accuracy estimation being slightly lower in the second scenario compared to the first. In Figures 6.9c and 6.9d, the robot accurately estimated the human agent's both leading preference and accuracy, allowing it to plan accordingly.

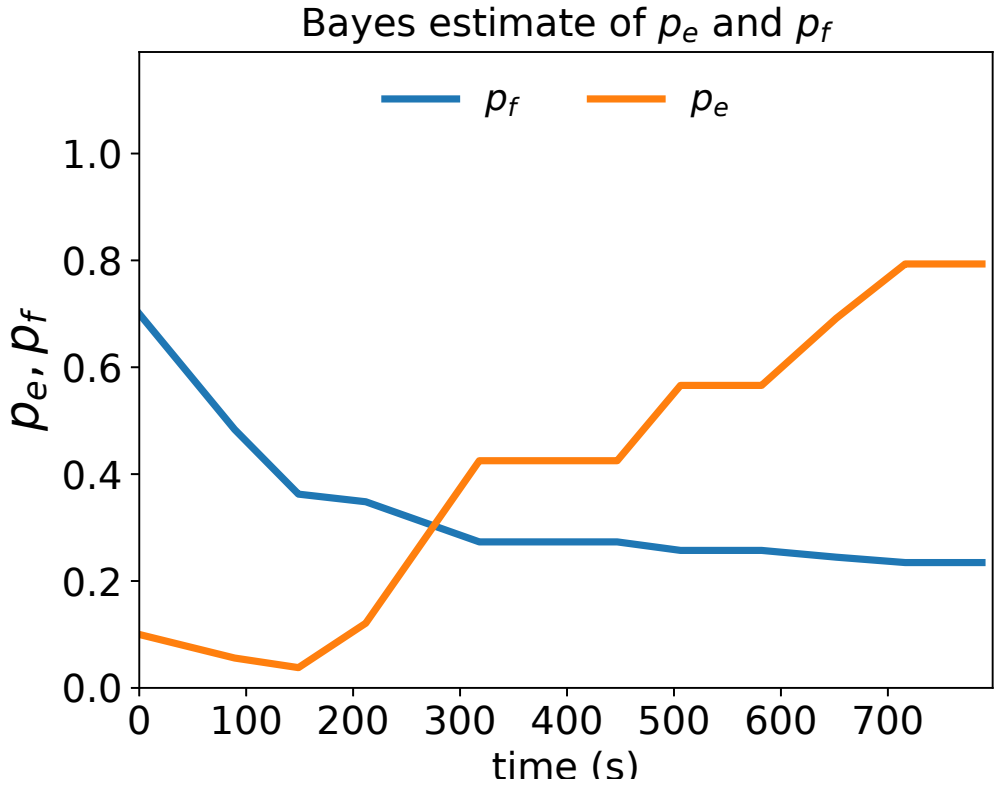


Figure 6.10: Robot's belief regarding the human's performance and preference. The human first had high accuracy but then made more errors

In addition to the four different cases previously discussed, we also examined a scenario where the human initially performed well but later began to make errors. The video of this scenario is available[1]. Fig. 6.10 illustrates the evolution of the robot's belief regarding the human's performance and preference throughout the experiment. The robot's belief about the human's preference decreased consistently since the human wanted to be the leader throughout the experiment. However, the robot's belief about the human's performance initially decreased, indicating that the human was accurate, but later increased as the human made more errors.

[1] https://youtu.be/gbyDUV0ZLfI

6.4 Summary and Discussion

This chapter focuses on implementing the proposed task planning framework for human-robot collaboration in a real scenario to further delve into the answers to RQ3 and RQ4. The proposed planning framework offers an approach that considers human preference and performance, and our previous study showed the effectiveness of the proposed framework in a simulation environment. However, to assess its efficiency, it was necessary to test it in a real environment with possible uncertainties. To do so, a collaborative scenario was designed and represented as the problem formulation in Chapter 4. Similar to the state estimator used in Chapter 5, a state estimator was applied for estimating the human agent's preference and performance. Additionally, the robot was equipped with an online task state updater, based on the agents' actions and errors. We used the Fetch robot for this collaboration requiring the autonomous pick-and-place ability.

The successful implementation of the algorithm with the Fetch robot in a real-world scenario showed that our planning framework has enabled the robot to adapt itself based on different human preferences and performances (enacted by the experimenter). In this study, we applied the framework to another collaborative scenario showing the potential of the proposed framework to be applied in various settings and human-robot collaborative scenarios. In this chapter, we only considered four different variations of human behavior, including different combinations of leading/following preference and high/low accuracy, emulated by the experimenter. However, in the next chapter, we will report the results of the user study and evaluate the planning framework with participants working in a team with the Fetch robot. Moreover, we will investigate the impact of different factors, such as the human's leadership/ followership styles or task difficulty, on collaboration and the human perception of the robot.

Chapter 7

User Study Procedure & Results

In the previous chapter, after providing details of the setup, the collaborative scenario, and the derivation of the task planning framework, we performed a system evaluation by enacting different human behaviors by the experimenter. However, this chapter goes beyond by delving into the results of a comprehensive user study involving 48 participants, aiming to address multiple research questions in this thesis. By analyzing the team's performance and the robot's task planning and actions, we complete the exploration started and continued through Chapters 5 and 6. Subsequently, through an examination of participants' perceptions of the robot, tasks, and collaboration, as well as their actions and decisions, we address RQ5 and RQ6. Finally, this chapter explores RQ7 by investigating any possible correlation between the human agents' leadership/followership styles and their collaborative approach with the robot.

7.1 Related Work & Hypotheses Development

In this section, we initially conduct a review of pertinent literature focusing on the human perception of the robot and collaboration in terms of trust, self-confidence, reliance, workload, and helpfulness, and develop our hypotheses based on the insights gathered from the existing literature. We also explore the human perception of control and the various leadership and followership styles. Finally, we briefly review research concerning task allocation and adaptation to human preferences.

7.1.1 Human Perception of the Robot & Collaboration

Trust, self-confidence & reliance

Trust is a complex and multifaceted concept in human-human, human-technology, and, more specifically, in human-robot interactions (HRI). In the context of human-robot trust, contributing factors to human trust in robots can be categorized into three groups [116]:

- robot characteristics and design elements (e.g., reliability)
- human characteristics (e.g., self-confidence)
- environmental-based factors (e.g., task type)

Typically, in most human-robot collaboration scenarios, it is assumed that humans and robots share common goals, and the human agent is aware of the collaborative context of the interaction, which prevents them from considering deception. Sanders et al. mainly focussed on the first two factors mentioned above and investigated how robot performance can affect human trust [116].

Trust is a topic of interest in HRI and HRC and has been studied extensively. Some studies focus on how to model and measure trust and its dynamics and, in some cases, use it to adjust the robot's behavior or improve human distrust or overtrust in the robot [76,107,117]. Some other studies, including this work, use objective or subjective measures (e.g., questionnaires) to evaluate human trust in robots before or after interaction with a robot. In our study, using Muir's questionnaire [2], we are interested in how participants' trust in the robot evolves during the sessions over time. We established the following hypothesis:

Hypothesis 1. *Participants' trust in the cobot, measured before and during the collaborative tasks, will improve during the session while engaged in different tasks.*

Furthermore, it is well-established that human reliance on an autonomous system is intertwined with their trust in the system and their self-confidence (i.e., their trust in their capability to accomplish the task). Studying and modeling this relationship has been the focus of many studies. The "confidence vs. trust" hypothesis is a well-accepted concept and serves as the foundation for many proposed models [118]. It posits that a person's reliance on an automation system is influenced by their relative trust, which refers to the difference between their trust in automation and their self-confidence. That is, if an individual's trust

in the system exceeds their self-confidence, they will opt to rely on the system; otherwise, they will rely on their own decisions and perform the task on their own. However, the findings of some studies contradict this hypothesis [119]. Authors in [120] have proposed a computational model that outperforms the "confidence vs. trust" hypothesis and signifies nuances between trust and self-confidence in predicting human reliance on automation. It is also noteworthy that in a collaborative scenario, the goal is to exploit human and robot abilities, not bypassing one by employing the other. This point may affect the full applicability of the "confidence vs. trust" hypothesis in human-robot collaboration.

We state the associated hypotheses as follows:

Hypothesis 2. *Participants' self-confidence will improve during the collaborative tasks over time as they collaborate more with the robot.*

Hypothesis 3. *For demanding tasks, participants will trust the robot more than their own ability and will rely more on the robot.*

Perceived workload

There are three fundamental measures for evaluating workload: physiological (e.g., heart rate), performance-based (e.g., time), and perceptual (subjective) [121]. In this chapter, we consider performance-based and subjective measures. For the subjective measure, we use NASA-TLX, which defines the workload experience in terms of "the sources of loading imposed by different tasks" [122]. We seek to study whether collaborating with the robot could positively influence participants' perceived workload, and derive the following hypothesis.

Hypothesis 4. *Participants will experience less workload when collaborating with the cobot rather than working alone, in terms of the dimensions in the NASA task load index, including: (a) mental demand (b) physical demand (c) performance (d) effort (e) frustration. However, the temporal demand will be similar, as there was no time limit for completing tasks.*

Helpfulness

In [123] helpfulness is defined as "with honest intentions, trying to play a positive role with the task at hand." In this study, we aim to investigate how humans' expected helpfulness from the robot, which refers to their anticipation of the robot's ability to provide effective assistance and enhance their overall collaboration experience, changes during collaboration.

Hypothesis 5. *Participants' expected helpfulness of the cobot will increase over time as they collaborate more with the cobot.*

7.1.2 Being in or under control

While the concept of human-robot collaboration is a relatively recent subject, conversations surrounding human automation have endured for many years. The rise of automation systems has gradually replaced humans in various control systems, spanning from simple to complex setups. Different factors, such as human preferences and abilities, contribute to determining the autonomy levels within these systems. These can range from manual control, where human operators actively control, to semi-autonomous and fully autonomous systems, where human operators take on more supervisory roles.

The perception of control by humans is a crucial area of research, spanning psychology, social science, and human-computer (robot) interaction. In the late 1990s, researchers like Ben Shneiderman and Pattie Maes discussed human-computer interface design from distinct perspectives [124]. Shneiderman emphasized user-controlled direct manipulation, highlighting human control, predictability, and responsibility. In contrast, Maes advocated for long-lived, proactive, adaptive software agents aligned with individual user habits and preferences. As of 2023, traditional input methods like mice and keyboards remain predominant in computer interactions, suggesting a prevailing inclination toward control and responsibility when dealing with technology. This underscores the significance of design decisions and user anticipations when incorporating intelligent technologies into human-computer interactions, comprehending users' behaviors and control perceptions across diverse levels of agency and automation [125, 126]. When extending this concept to human-robot interaction, a preference is likely for humans to maintain control over robots.

The perception of robot control by humans can be influenced by factors such as human characteristics (e.g., locus of control), trust in robots, self-confidence, task difficulty, task criticality, human abilities, and robot abilities. In [88], the issue of being in control or under control was explored in a warehouse order-picking scenario, demonstrating higher productivity when the human leads and higher accuracy when the human follows the robot. In contrast, "6 River Systems," a company actively engaged in collaborative robotics and warehouse automation, contends that allowing associates to have self-control over their pace leads to decreased speed and efficiency [89]. In response to this challenge, they introduced a "system-directed picking" approach, wherein robots guide associates through their tasks. Another instance in the industrial domain is the robotic picking assistive system implemented in DHL warehouses, where the robot (supplied by Locus Robotics)

directs human workers to the designated pickup location [90]. In these setups, the robot is tasked with planning the tasks, while humans are responsible for picking up items.

The findings from a user study in [87], in the context of household robots, indicate that participants felt a greater sense of control when the robot operated in semi-autonomous mode. They observed a preference among participants for reduced autonomy in critical tasks (e.g., scheduling a doctor's appointment). However, when it came to a robot transporting biscuits from the kitchen to the living room, contrary to their stated preference, participants leaned towards the autonomous robot condition. The authors attributed this shift to the perceived advantages of higher physical efficiency and comfort. It is worth noting that a more comprehensive analysis could be conducted if the researchers had also inquired about participants' trust in the robot's ability to move autonomously. Furthermore, in [127], the authors explore the connection between automation accuracy and task controllability, emphasizing its impact on user satisfaction. Their study underscores a strong preference for manual control over automation.

In our work, based on the collaborative and cognitively critical nature of the scenario designed (due to a penalty for mistakes), we state the next hypothesis of this chapter as follows:

Hypothesis 6. *Overall, participants will prefer retaining control over the robot and being a leader.*

7.1.3 Leadership and Followership Styles

In this study, we are also interested in how participants' general leadership and followership styles can affect their collaboration with the robot.

Leadership Styles [128]: Leadership style pertains to the behaviors leaders demonstrate, encompassing their actions and interactions with followers. It characterizes what leaders do and how they behave in diverse circumstances. Leadership style is intricately linked to an individual's personal leadership philosophy. Commonly observed leadership styles include [128]: 1. authoritarian 2. democratic 3. laissez-faire

- **authoritarian** – involves leaders perceiving followers as requiring guidance, exerting control over their actions, emphasizing their authority, and directing communication toward themselves rather than fostering interaction among group members.
- **democratic** – involves leaders treating followers as capable collaborators, guiding rather than controlling, and fostering open, inclusive, and supportive communication. They prioritize helping followers achieve personal goals, provide guidance and

suggestions without giving orders, and deliver objective praise and criticism in evaluations.

- **laissez-faire** – is a "hands-off, let it ride" approach where leaders exert minimal influence and allow followers significant freedom. They do not control or guide followers' activities and may serve as temporary or interim leaders with limited authority or interest in shaping outcomes.

It is worth emphasizing that these leadership styles are not separate categories; they can overlap and coexist. Leaders may adopt different styles depending on the situation. For example, a leader might exhibit an authoritarian approach in certain contexts while adopting a democratic style in others, showcasing a dynamic and multifaceted approach to leadership. Northouse proposed the leadership style questionnaire in [128] to identify participants' tendency to the leadership styles.

Followership Styles [129]: Kelley (1992) identified five followership styles based on two key dimensions: engagement and critical thinking. Engagement ranges from passive (waiting for direction) to active (taking the initiative to participate actively in tasks). Critical thinking spans from dependent uncritical thinking (accepting information without evaluation or questioning) to independent critical thinking (evaluating and analyzing information to identify consequences and opportunities). These five follower styles are:

- **Exemplary**: These followers show high levels of active engagement and independent critical thinking, thinking for themselves, challenging leaders constructively, proactively supporting organizational goals, and assuming extra responsibilities.
- **Conformist:** Conformist followers are actively engaged but tend to be dependent uncritical thinkers, enthusiastically following leader directions without questioning.
- **Passive:** Passive followers have low levels of both engagement and critical thinking, needing constant direction and acting only after explicit instructions.
- **Alienated:** Alienated followers are independent critical thinkers but lack engagement. They are skeptical, may oppose management, and see themselves as mavericks.
- **Pragmatist:** Pragmatists have a moderate level of both engagement and critical thinking. They are uncommitted, prefer maintaining the status quo, and act cautiously, often waiting for crises before taking action.

Kelley also designed the questionnaire in [129] for determining followership styles.

7.2 Procedure

7.2.1 Recruitment

We started distributing flyers and recruiting participants after securing ethics approval from the University of Waterloo Human Research Ethics Board (see Appendix B). In the flyer (see Appendix C), the four different colors used in the study were shown, and distinguishing them was mentioned as an eligibility condition to participate, to ensure the exclusion of individuals with color blindness. This study involves three phases for each participant.

7.2.2 Phase 0

The first phase of our study commenced with prospective participants responding to our recruitment flyers and establishing contact with the experimenter. Following this, we sent the participants a consent form via email, along with a questionnaire aimed at gauging their background and experience in the domains of robotics and artificial intelligence.

7.2.3 Phase 1 (In-Person)

In Phase 1, we scheduled a 90-minute in-person session with each participant, where we conducted the following procedure:

Step 1: We began by welcoming participants and providing an overview of the setup, showing them their designated work area. We used presentation slides to elaborate on the specifics and informed participants that:

- The robot's decision-making may involve occasional errors (deception).
- Upon confirming the task's completion, There is a penalty of $1 for the team per misplaced block (deception).

Step 2: *Task 0* - To begin, we gave participants Pattern A_1, as a sheet of paper, and granted them 45 seconds to memorize it. We then collected Pattern A_1 and provided participants with Pattern A_2, a partially known version of Pattern A_1, and indicated that they could keep it until they finished the task. Participants were prompted to assess their self-confidence in successfully completing the task and subsequently initiated the retrieval

and placement of the blocks. The experimenter was available to provide guidance and address any deviations from the instructions or questions. It was emphasized that the task had no time constraint, and the completion time wouldn't be evaluated. After successfully placing all the blocks and declaring task completion, the experimenter conducted an inspection of the blocks based on the provided pattern, counting any misplaced blocks. Importantly, participants were assured that no penalties would be applied for errors in this task. Following this, participants were requested to complete a questionnaire regarding their perceived task load.

Step 3: They watched a video of the pick-and-place of some blocks by the Fetch robot[1], and then answered the questionnaire about their trust in the robot.

Step 4: Participants familiarized themselves with the GUI and practiced its operation.

Step 5: *Task 1* - Depending on the assigned mode (a permutation of patterns B, C, D), participants were provided with the corresponding pattern (e.g., B_1) and instructed to memorize it within a 45-second timeframe. Subsequently, they returned the pattern and received the second pattern, a partially known version of the first pattern (e.g., B_2). Before commencing the task, participants were asked to respond to two questions about their self-confidence and the expected helpfulness of the robot. Following this, they initiated the task, collaborating with the Fetch robot. Upon task completion, participants completed three questionnaires assessing their perceived task load, trust, and perception of the robot.

Remark. Participants initiate the task, and the robot awaits their instructions. They can assign tasks to the robot, and the robot starts working when participants allocate a task to themselves. This approach allows the robot to establish an initial understanding of participants' preferred leading/following roles in the collaboration.

Step 6: *Task 2* - This task followed the same procedure as *Task 1*.

Step 7: *Task 3* - This task followed the same procedure as *Task 1*.

Step 8: Lastly, participants were asked to complete two questionnaires. The first questionnaire aimed to evaluate their performance as a team with the robot, while the second one explored their collaborative experience and utilized the shortened version of the User Experience Questionnaire (UEQ) [130, 131]. Additionally, participants were asked to rank the difficulty of the tasks (Tasks 0-3) and respond to an open-ended question: *"Which abilities would you improve or add to Fetch if you were to use it in a manufacturing setting?"*.

[1] https://youtu.be/ahZDo0_iyjg

7.2.4 Phase 2 (Online)

In the online phase, we presented each participant with a video recording of their interaction with the robot, with a specific focus on Pattern B. These videos included synchronized content from two different camera angles, a screen recording of the graphical user interface (GUI), and the display of Patterns B_1 and B_2. One of the participant's videos is accessible online[2]. During the online interview, we played the video and encouraged participants to discuss their strategies, plans, and preferences, and how these evolved throughout their collaboration. We also requested them to complete two questionnaires regarding their leadership and followership styles. In line with the approved ethics application, we also informed participants about the existence of *"Deception"* elements in the study and sought their consent to use their data. To express our gratitude for their participation, participants received a $30 gift card.

7.3 Measures

As part of our analysis aimed at assessing participants' perception of the robot, tasks, and collaboration, we utilized a range of subjective measures, primarily derived from the items in the questionnaires they completed at various phases of the study, as well as insights from interviews. Additionally, we considered objective measures such as the number of tasks allocated to the robot by the human.

To assess participants' trust in the robot, we employed Muir's questionnaire [2] (refer to Table 7.1) using a ten-point Likert scale. Muir's questionnaire has been widely used in the literature as a measure of trust in machines, in terms of predictability, dependability, reliability, and overall trust. We administered this questionnaire at multiple points in the study, including after they watched a video of the robot (to gauge initial trust) and following each collaboration with the robot in Tasks 1, 2, and 3. In addition, participants provided responses to two questions on a 21-point Likert scale, as shown in Table 7.2. These questions pertained to their self-confidence in completing each task (Tasks 0-3) and their expected helpfulness of the robot (Tasks 1-3). Participants answered these questions after reviewing the patterns but prior to commencing the task of fetching and placing the blocks. High and low values on the Likert scale represented high and low self-confidence, as well as the expected helpfulness of the robot.

Following the completion of each task, participants evaluated their workload, considering factors such as perceived mental demand, physical demand, temporal demand, per-

[2] https://youtu.be/X6Rj0zwQhz8

formance, effort, and frustration. They did so by responding to the NASA-TLX questionnaire, as shown in Table 7.3, employing a scale ranging from 0 to 100. We employed the NASA TLX questionnaire because it is widely recognized across diverse domains as an established tool for evaluating task demands. Additionally, participants completed a questionnaire designed to assess aspects of their working alliance, the traits of the robot teammate, the robot's performance, and the collaboration experience. This questionnaire aimed to further evaluate participants' perception of the robot and the collaboration. The first seven questions in this questionnaire, relevant to our study, were adapted from [3]. This questionnaire, as displayed in Table 7.4, used a 7-point Likert scale, presented in an agree-disagree format.

At the end of the last task (Task 4), participants were asked to complete the user experience questionnaire (UEQ). The User Experience Questionnaire (UEQ) was chosen for its demonstrated effectiveness in capturing diverse aspects of user experience, providing valuable insights into participants' interactions with the system. The UEQ encompasses six paired aspects, each evaluated on a 7-point semantic differential scale: obstructive/supportive, complicated/easy, inefficient/efficient, clear/confusing, boring/exciting, not interesting/interesting, conventional/inventive, and usual/leading edge. Participants also answered a 7-point Likert scale questionnaire about team fluency (i.e., efficient, synchronized collaboration with seamless coordination in shared activities [3]) and their willingness to collaborate again with fetch.

During the online phase of the experiment, we conducted interviews with participants and categorized them into four distinct groups based on their preference for leading or following in collaboration with the robot. These categories included "lead," "collaborative-lead," "collaborative-follow," and "follow." Additionally, we administered two questionnaires to assess their leadership and followership styles. The leadership style questionnaire utilized a 5-point Likert scale [128] (see Appendix A.1). Scores for each leadership style (authoritarian, democratic, laissez-faire) reflect the degree of inclination toward that style, with potential scores ranging from very low (6-10) to very high (26-30), as well as intermediate levels such as low (11-15), moderate (16-20), and high (21-25). To determine participants' followership style, we employed a seven-point Likert scale questionnaire from [129], which determines the score for independent thinking and active engagement (see Appendix A.2). Based on participants' responses, they received scores for independent thinking and active engagement, with scores ranging from 0 to 60. Subsequently, as illustrated in Fig. 7.1, they were categorized into one of five distinct followership styles.

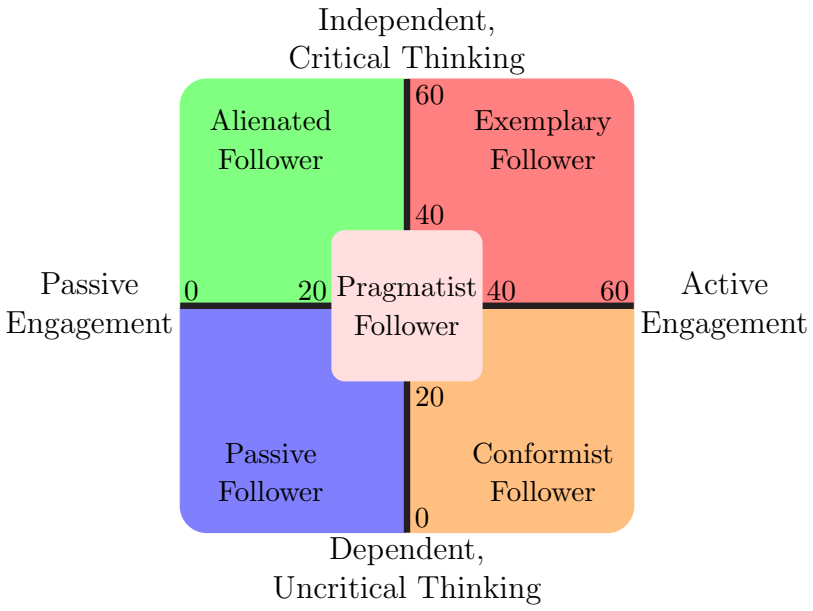


Figure 7.1: Kelley's five followership styles

Table 7.1: Muir's questionnaire - measuring participants' trust in the robot [2]

1- To what extent can the robot's behavior be predicted from moment to moment? (predictability)
2- To what extent can you count on the robot to do its job? (dependability)
3- What degree of confidence do you have that the robot will be able to cope with similar situations in the future? (reliability over time)
4- Overall, how much do you trust the robot? (overall degree of trust)

Table 7.2: The questions about the participants' self-confidence and expected helpfulness of the robot. Asked after seeing the patterns but before starting the task

Q1 (Self-confidence): To what extent are you confident about your ability to successfully accomplish the task?
Q2 (Expected Helpfulness): To what extent do you think Fetch can help you improve your performance?

Table 7.3: NASA-TLX questionnaire asked after each task

(Mental demand) How mentally demanding was the task?
(Physical demand) How physically demanding was the task?
(Temporal demand) How hurried or rushed was the pace of the task?
(Performance) How successful were you in accomplishing what you were asked to do?
(Effort) How hard did you have to work to accomplish your level of performance?
(Frustration) How insecure, discouraged, irritated, stressed, and annoyed were you?

Table 7.4: A Questionnaire Assessing Working Alliance, Robot Teammate Traits, Robot's Performance, and Collaboration. Q1-Q7 adopted from [3]

Q1: Fetch was intelligent
Q2: Fetch was committed to the task
Q3: Fetch perceived accurately what my goals are
Q4: Fetch did not understand how I wanted to do the task.
Q5: Fetch and I worked towards mutually agreed-upon goals.
Q6: Fetch and I respected each other.
Q7: I feel that Fetch appreciates me.
Q8: Fetch and I collaborated well together.
Q9: Fetch worsened the team's performance in terms of completion time.
Q10: Fetch improved the team's performance in terms of accuracy.
Q11: The robot's decisions were reliable.
Q12: The robot failed to pick up and place objects reliably.

Table 7.5: The robot and team's overall performance. Q1-3 are adopted from [3]

1- Our team improved over time.
2- Our team's fluency improved over time.
3- The robot's performance improved over time.
4- I will be happy to collaborate again with this robot in similar future tasks

7.4 Results & Discussion

We recruited 58 participants from our university. However, we had to omit data from 10 participants due to reasons such as a bug in the robot's program[3] and the robot's malfunction during the sessions (e.g., router disconnection). Consequently, our data analysis is centered on the remaining 48 participants, comprising 22 females, 24 males, and 2 individuals self-identifying as 'others,' with an average age of 24.02 $\pm$ 3.93. The majority were University students (44), while 3 were postdoctoral or visiting researchers, and 1 was a staff member. Initially, we delve into participants' perception of the tasks, the robot, and collaboration. Subsequently, we shift our focus to the analysis of participants' actions and performance.

7.4.1 Participants' Perception of the Tasks, Robot, & Collaboration

This part mainly focuses on subjective measurements through questionnaires. Throughout this section, the Kruskal–Wallis H test, a nonparametric statistical test, is applied to ascertain if there are statistically significant variations among two or more groups. If a noteworthy overall difference is detected among multiple groups, the Dunn test is subsequently employed as a post hoc analysis to pinpoint specific group differences. In addition, whenever we state that the analysis is based on the tasks, it implies an examination following the chronological sequence, starting from Task 0 and progressing through Task 3.

Trust

We assessed participants' trust in the robot at four stages: before collaborating with the robot, after watching a video of the robot (initial trust), and after each collaboration for Tasks 1, 2, and 3. We found a statistically significant difference among these stages ($H(3) = 13.85$, $p = .003$). Fig. 7.2 illustrates participants' trust in the robot, representing the average of their responses to questions in Table 7.1. The figure reveals a notable difference between participants' initial trust and trust levels after Tasks 1, 2, and 3. Furthermore, it

[3]A bug occurred in the code due to a discrepancy between Python 2 and Python 3 when dividing integers. In Python 2, division of two integers produced an integer result, while in Python 3, it resulted in a floating-point number.

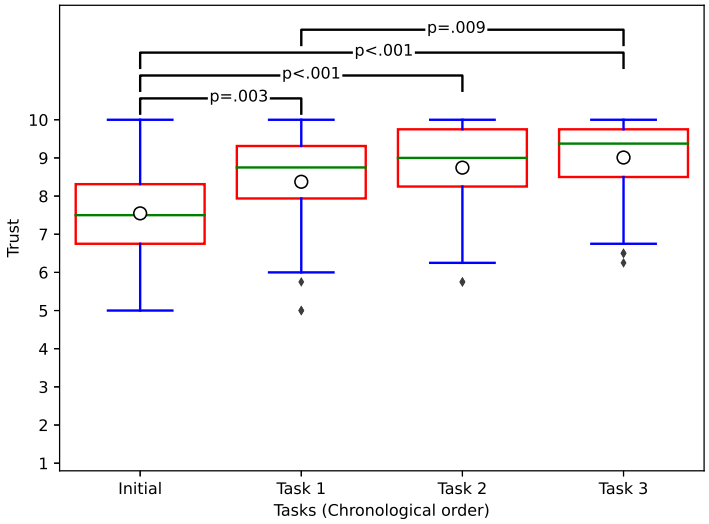


Figure 7.2: Participants' trust in the robot before collaboration and after each task

demonstrates a gradual increase in participants' trust over the sessions, with a significant difference observed between Task 1 and Task 3.

Self-confidence

Participants were also asked about their self-confidence in accomplishing the tasks before collaborating with the robot, but after seeing the pattern they needed to work on. In Task 0, participants were told they had to work alone, without the robot. To answer this question in Task 1, participants knew that the robot would work alongside them, but they had not yet worked with the robot, and it was their first experience. Fig. 7.3 shows the results for this question, revealing a significant difference in self-confidence across different tasks ($H(3) = 13.85$, $p = .003$). As expected, despite an increase in the average self-confidence of participants in Task 0 and Task 1, there is no significant difference between them, as both tasks involved participants working without prior experience with the Fetch robot. However, the results indicate a significant difference between Task 0 and Task 2, Task 0 and Task 3, as well as Task 1 and Task 3.

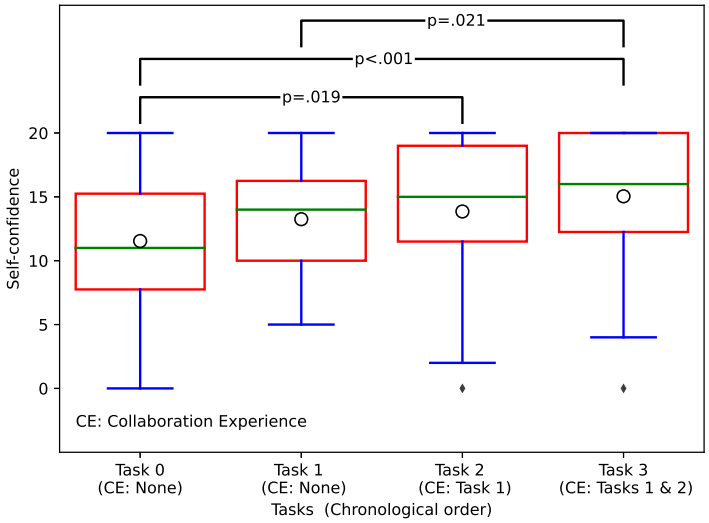


Figure 7.3: Participants' self-confidence in completing the tasks after seeing the patterns but before starting them based on a) the tasks

Helpfulness

Participants' self-confidence improvement can be attributed to becoming more familiar with the task and finding the robot more helpful over time. However, to demonstrate that this improvement is not solely due to the former, in Tasks 1, 2, and 3, in addition to inquiring about their self-confidence, we also asked about their expectations regarding the robot's helpfulness in completing the task. This revealed a significant difference among these responses ($H(2) = 19.76$, $p \ll .001$). As depicted in Fig. 7.4, participants' expected helpfulness of the robot increased over time, with a significant difference observed between their expectations in Tasks 1 and 2, as well as Tasks 2 and 3. It's worth noting that when participants answered this question in Task 1, they had not yet collaborated with the robot. Thus, participants perceived the robot as more helpful after the initial collaboration. However, there was no significant difference in participants' expected helpfulness of the robot across different patterns.

Relative Trust

We also evaluated participants' relative trust, i.e., the difference between participants' trust in the robot and their self-confidence. The result shows no significant difference in

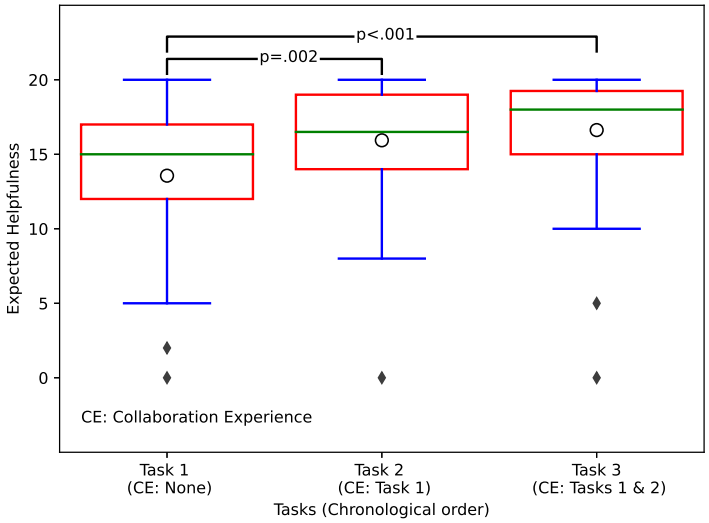


Figure 7.4: Participants' expected helpfulness from the robot

participants' relative trust between Tasks 1, 2, and 3.

Robot's Performance and Traits

After each task (Tasks 1-3), participants answered the questions in Table 7.4. In Fig. 7.5, participants' scores indicate a positive working alliance and favorable robot teammate traits. We also observe a significant improvement in how participants perceive the robot's intelligence between Tasks 1 and 2, as well as between Tasks 1 and 3. We asked participants to rate their collaboration with the robot and its performance in terms of reliability, accuracy, and operating speed. Based on Fig. 7.6, participants found the robot's decisions accurate and reliable, leading to a positive collaboration experience. It's worth noting that Fetch was noticeably slower than the human agent in pick-and-place tasks, which is reflected in some participants' scores. However, overall, the average of participants' scores is low (reverse scale), indicating they believed Fetch did not significantly extend task completion time. This perception aligns with participants' understanding of the robot's role in the collaboration, as one participant expressed during the interview: "Mentally, it can compensate for me, but physically, I can compensate for it (Fetch)."

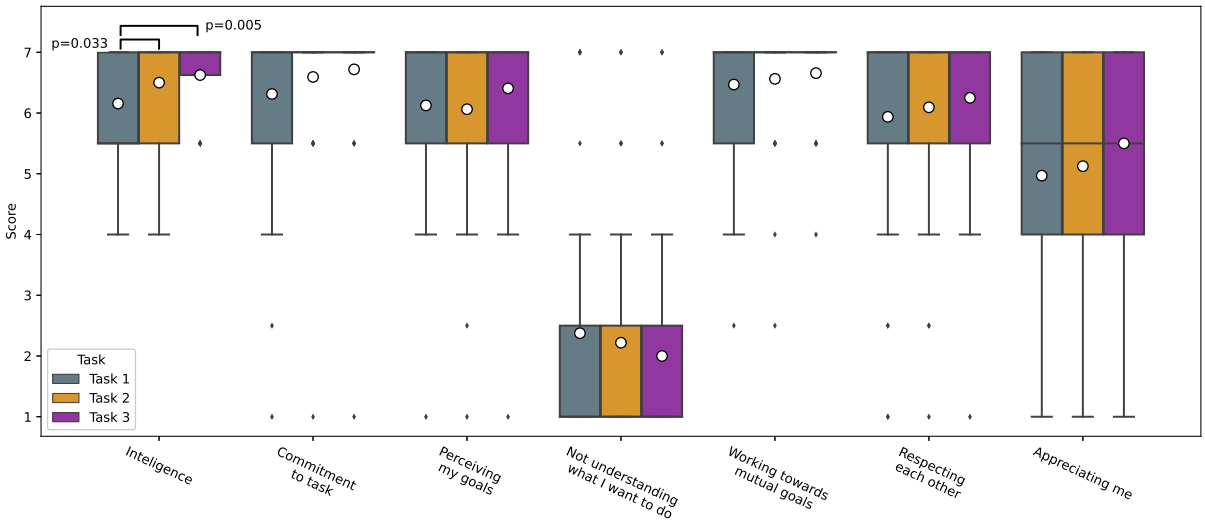


Figure 7.5: Participants' perception of the robot's traits and working alliance

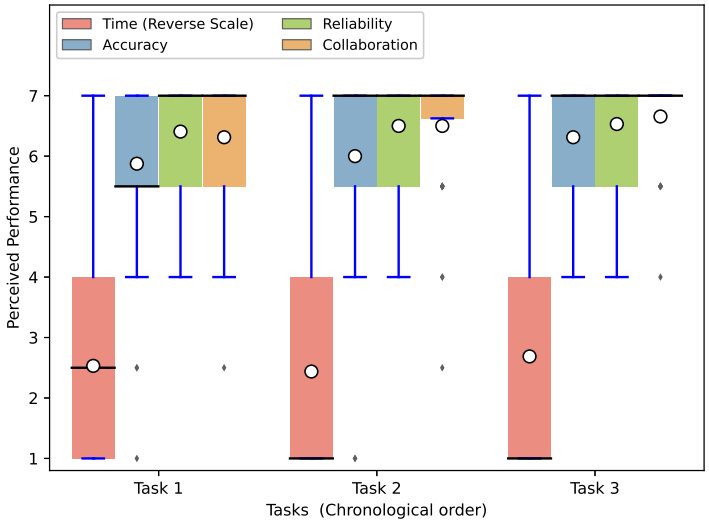


Figure 7.6: Participants' perception of the robot's performance and collaboration

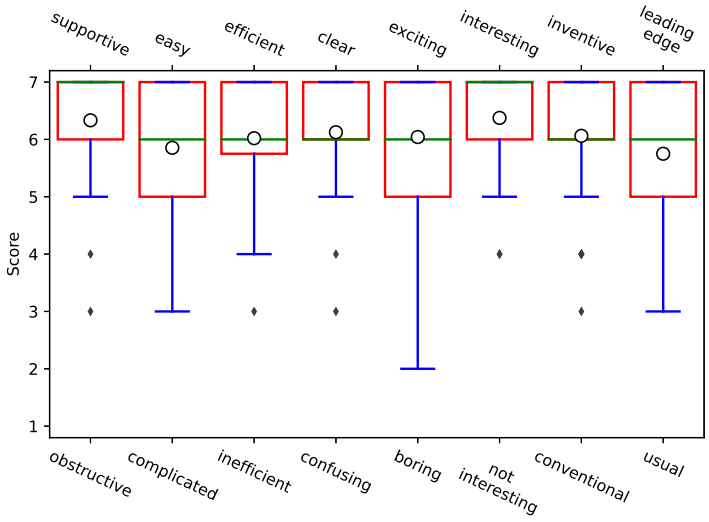


Figure 7.7: Scores of user experience questionnaire

User Experience & Team Fluency

At the end of the in-person phase of the experiment, participants answered the user-experience questionnaire. Fig. 7.7 shows participants' scores on this questionnaire, indicating a positive user experience. They also answered four questions in Table 7.5, and Fig. 7.8 shows their scores for these four questions. Based on it, participants perceived an improvement in their team performance and fluency and expressed a willingness to collaborate again with Fetch. Despite a high average, participants' scores regarding the improvement in the robot's performance cover the whole range of the scale. However, this is reasonable as the robot had the same plan for all three tasks.

Perceived Workload

After each task (Tasks 0-3), participants answered the NASA-TLX questionnaire. Hereinafter, we discuss the results for each of the six measures in this questionnaire.

- **Mental Demand:** The mental demand for each task started from the point when participants were given the first pattern (i.e., Patterns A_1, B_1, C_1, and D_1) to memorize until the end of the task where they needed to recall the pattern with the help of the second given pattern (i.e., Patterns A_2, B_2, C_2, and D_2). We analyzed participants'

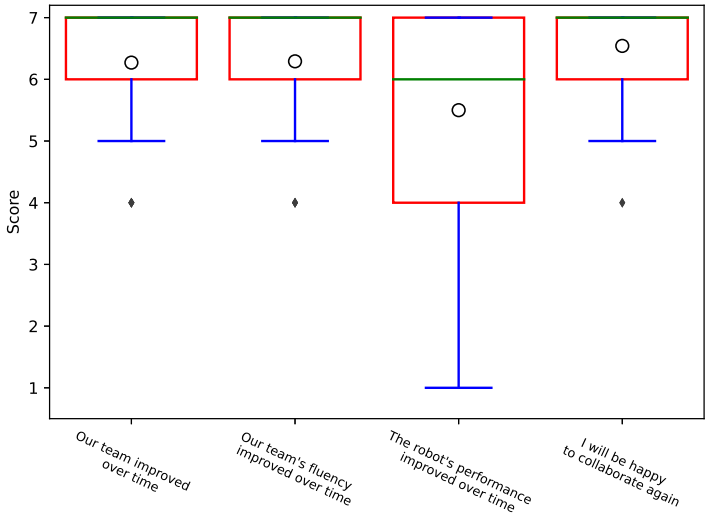


Figure 7.8: Team fluency and performance

scores for this question based on the tasks and the patterns. Fig. 7.9a shows the result for Task 0 to Task 3 (in chronological order), with a significant difference between them ($H(3) = 18.05$, $p < .001$). The results show a significant required mental demand for Task 0 (working alone) compared to Tasks 1, 2, and 3. This indicates that collaboration with the robot could decrease participants' perceived mental demands, supporting part of Hypothesis 4.

- **Physical Demand:** Fig. 7.9b illustrates the results for physical demand scores based on the chronological order of the tasks. Despite the higher mean value of physical demand for Task 0, and in contrast to Hypothesis 4, we observe no significant difference between Task 0 and Tasks 1, 2, and 3 ($H(3) = 3.35$, $p = .34$). The reason can be that participants placed most of the blocks on the table, as the robot was slower than them and, therefore, did not experience a higher physical demand compared to Task 0.

- **Temporal Demand:** Participants were told several times during the experiment that there was no time limit, and we did not measure their time. However, as shown in Fig. 7.9c, there is a significant difference between them ($H(3) = 9.58$, $p = .023$), and participants perceived a higher temporal demand in Task 0 compared to Tasks 2 and 3, rejecting part of Hypothesis 4.

- **Performance:** In Tasks 1, 2, and 3, as the robot could fix participants' errors, all

participants completed the tasks successfully. Conversely, in Task 0, participants made some mistakes and had no teammate to fix the errors. Therefore, as shown in Fig. 7.9d, participants significantly perceived themselves as more successful in completing Tasks 1, 2, and 3 compared to Task 0 ($H(3) = 52.75$, $p \ll .001$). This supports part of Hypothesis 4.

- **Frustration:** The results show that there is a significant difference in perceived frustration in the tasks ($H(3) = 19.25$, $p < .001$), specifically in Task 0 in comparison with Tasks 1, 2, and 3 (Fig. 7.9e. However, there is no significant difference in perceived frustration between Tasks 1-3 and Patterns B, C, and D ($H(3) = 0.27$, $p = .87$). This result also confirms part of Hypothesis 4, which mentions that participants experienced less frustration when collaborating with the robot.

- **Effort:** The results show a significant difference in perceived effort based on the tasks ($H(3) = 15.74$, $p = .001$) and patterns ($H(3) = 6.89$, $p = .03$). We can see in Fig. 7.9f that participants perceived more required effort to accomplish Task 0 in comparison with Tasks 1, 2, and 3. This confirms part of Hypothesis 6, which states collaboration can decrease perceived effort.

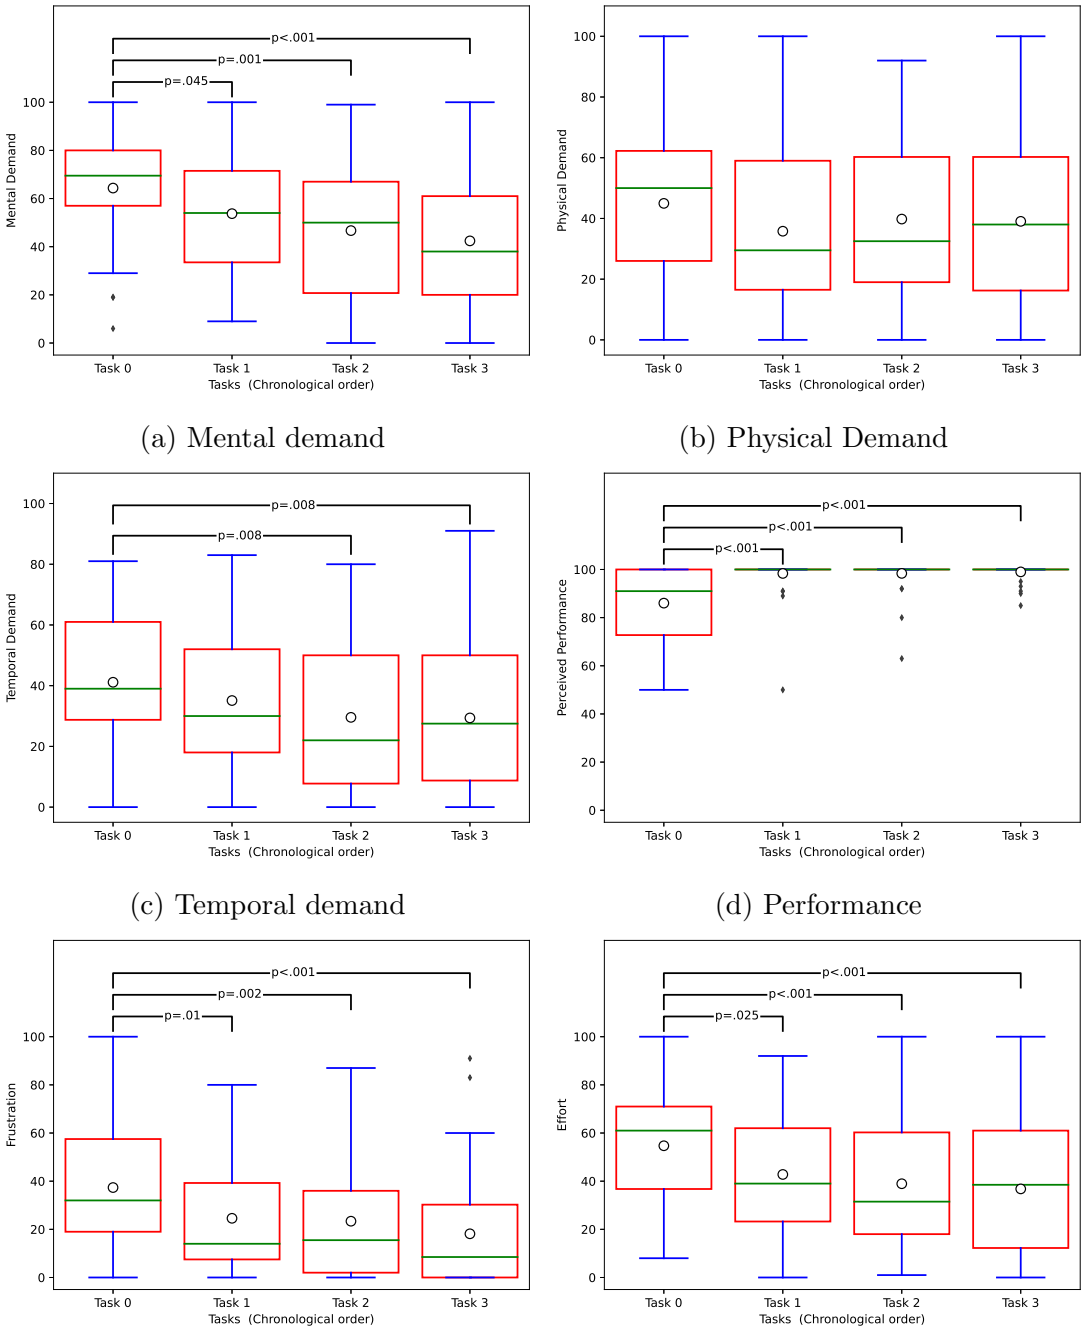
Participants' Initial Trust vs. Perception of the Robot

We assessed the correlation between participants' initial trust and their perceptions of the robot's traits, helpfulness, and collaboration (Fig. 7.10). The findings indicate weak correlations between participants' initial trust and their expected helpfulness of the robot, perceived robot's intelligence, understanding, and respect for them. Additionally, we observed moderate correlations between their initial trust and perceived robot's commitment to the task, understanding of goals, and collaboration.

Task Difficulty

Here, we investigate how task difficulty influenced participants' perception of collaboration, performance, and decision-making.

- **Participants Ranking:** At the end of the last task, we asked participants to rank the tasks based on their difficulty level. Fig. 7.11 shows their rankings based on the patterns used in the tasks. Pattern A was used only in Task 0 when they worked without the robot. Typically, the perceived difficulty of the task depends on the challenge of



(a) Mental demand

(b) Physical Demand

(c) Temporal demand

(d) Performance

(e) Frustration

(f)

Figure 7.9: Participants' perceived a) mental, b) physical demand, c) temporal demand, d) performance, e) frustration, and f) effort based on the tasks

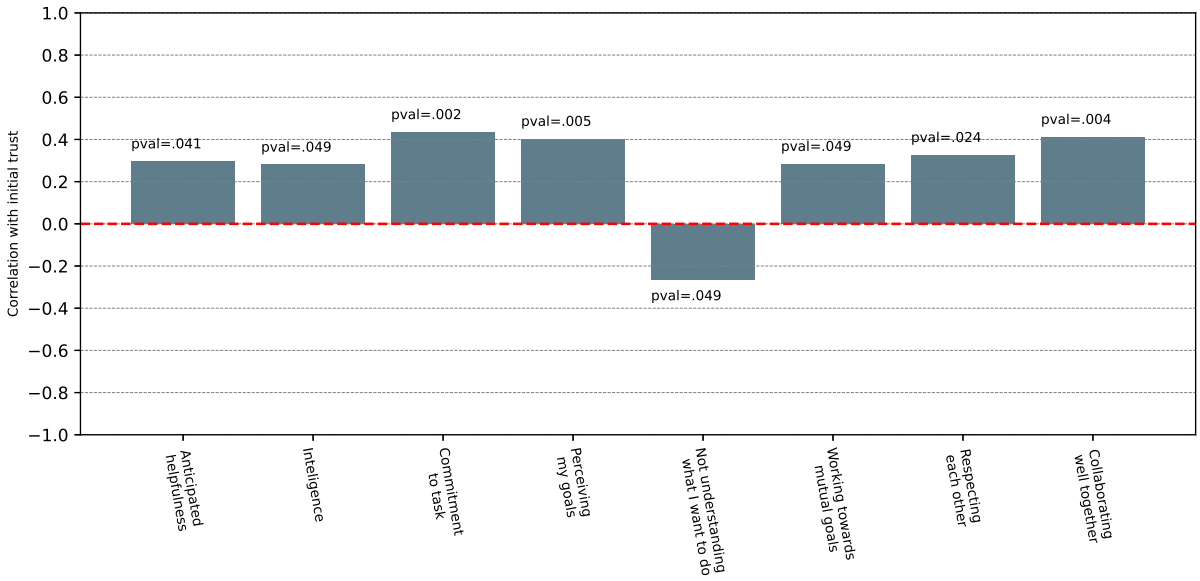


Figure 7.10: Correlation between participants' initial trust and perceived robot traits

memorizing the first pattern within 45 seconds and then recalling it while arranging the blocks, with the second given pattern as a hint. The second patterns, Patterns A_2, B_2, C_2, and D_2, had 9, 12, 6, and 9 partially known spots, respectively. However, the number of these unknown spots is not the sole factor in determining their difficulty level; the complexity of memorizing them also plays a crucial role. This is why most participants considered Pattern D_2 with 9 unknown spots simpler than Pattern C_2 with 6 unknown spots because it was more challenging for them to memorize Pattern C_1 compared to Pattern D_1. Participants also found Pattern B to be one of the most challenging patterns, as expected.

- **Self-confidence:** We have also analyzed participants' self-confidence based on the patterns, revealing a significant difference ($H(3) = 24.42$, $p \ll .001$). Pattern A was consistently used for Task 0 when participants worked alone. Fig. 7.12c illustrates that participants' self-confidence in accomplishing Pattern 0 (Task 0) is significantly lower than that for Patterns 2 and 3 (No significant difference observed between Patterns 0 and 1). Furthermore, we observed that participants exhibited significantly higher self-confidence in performing Pattern 3 compared to Patterns 1 and 2 (No significant difference was noted between Patterns 1 and 2).

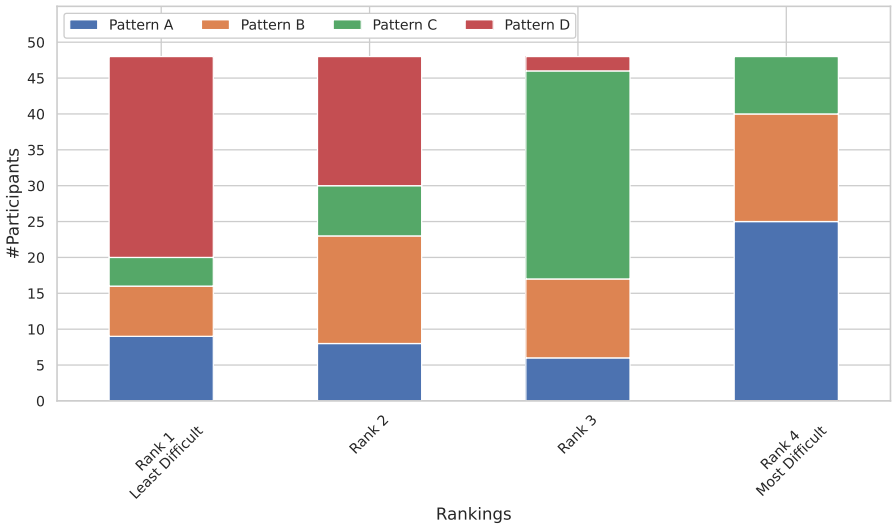


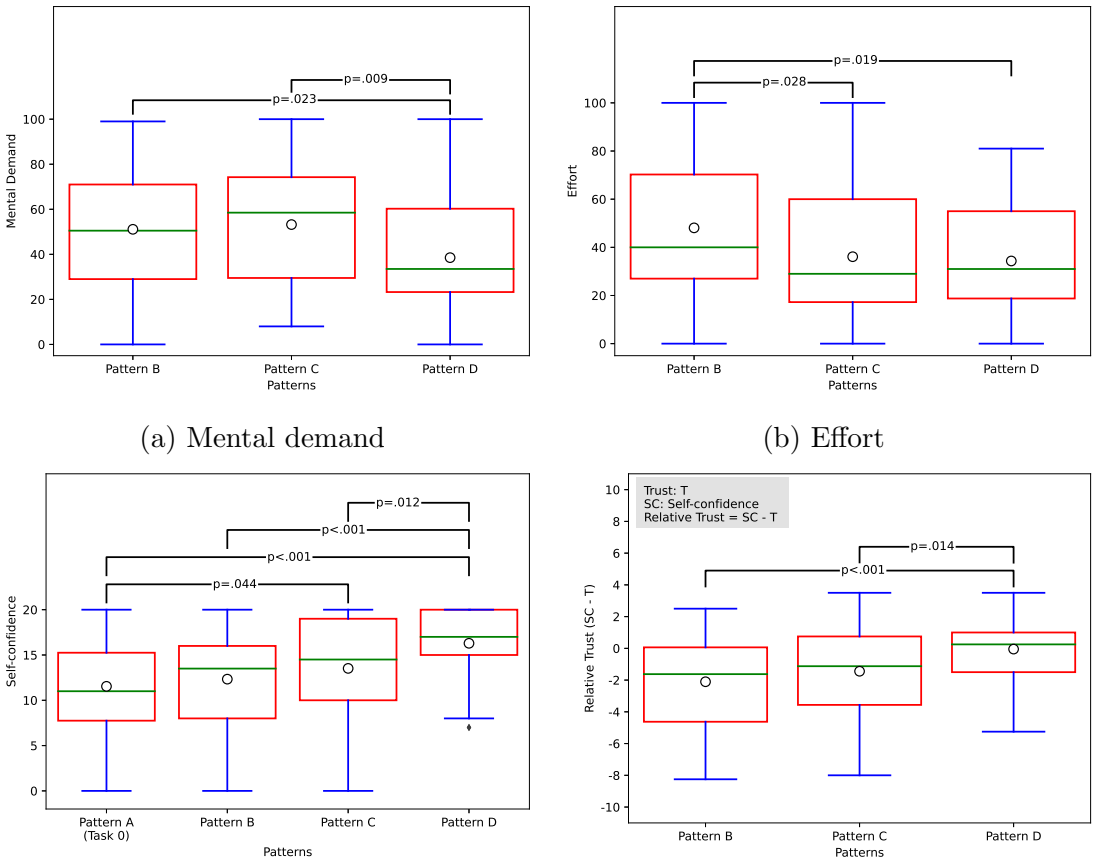
Figure 7.11: Ranking of task difficulty by participants

- **Relative Trust:** A significant difference was found when evaluating participants' relative trust for the patterns ($H(2) = 12.9$, $p = .001$). Fig. 7.12d shows participants' relative trust based on the patterns, where we can observe a significant difference between Patterns B and D as well as Patterns C and D.

- **Perceived mental demand:** In Fig. 7.12a, as Pattern A was used only for Task 0, we only considered Patterns B, C, and D, revealing a significant difference ($H(2) = 8.05$, $p = .019$). The results show that participants perceived Pattern D as less mentally demanding than Patterns B and C.

- **Perceived effort:** Additionally, in Fig. 7.12b, participants perceived significantly higher effort for completing Pattern B compared to Patterns C and D.

Discussion: The subjective measurements support Hypothesis 1, suggesting that participants' trust in the robot grows over time during the sessions. Additionally, according to the results, participants' self-confidence improved during the sessions as they collaborated more with the robot. This supports Hypothesis 2. Participants' answers also indicate that they expected more help from the robot during the sessions, supporting Hypothesis 5. The results also partially support Hypothesis 4. As discussed, participants perceived less workload in terms of mental demand, performance, effort, and frustration, supporting Hypothesis 4. Contrary to Hypothesis 4, participants also perceived less temporal demand



(a) Mental demand

(b) Effort

(c) Self-confidence

(d) Relative trust

Figure 7.12: Participants' a and b) perceived mental demand and effort, c) self-confidence, and d) relative trust (self-confidence - trust in the robot) based on the patterns (Patterns B, C, and D)

Table 7.6: Participants' preference based on interviews.

Preference	Lead	Collaborative-lead	Collaborative-follow	Follow	Other
# Participants	17	20	4	3	4

when collaborating with the robot. However, participants' perceived physical demand did not significantly change when they completed the task in collaboration with the robot. *To summarize, the most important conclusion drawn from this study is that the collaboration could enhance participants' perception of the robot and collaboration and decrease their perceived workload.*

We also evaluated the participants' self-confidence, relative trust, and workload (mental demand and effort) based on the patterns. The results show that, as planned, Pattern B was more challenging for the participants than Patterns C and D. This also supports Hypothesis 3, which states that participants will trust the robot more than their own abilities for challenging tasks.

7.4.2 Participants' Actions and Performance

Overall preference: In the online phase, we interviewed participants and asked about their strategies and preferences for collaborating with the robot. Probing their inclination towards leading or following, we inquired whether they preferred assigning tasks to the robot or being assigned tasks. Through detailed discussions, we discerned distinct preferences, categorizing participants into four groups: lead, collaborative-lead, collaborative-follow, and follow. The groups "lead" and "follow" represent participants with completely leading or following preferences, and the groups "collaborative-lead" and "collaborative-follow" represent participants who preferred to collaborate while having a slight preference to lead or follow, respectively. Four participants fell under neither of these four groups. One of them preferred to neither lead nor follow, but do his own tasks, and he did not assign any subtask to Fetch and did all subtasks assigned by Fetch. Three other participants quickly began doing subtasks because they didn't want to forget the pattern. They mostly chose subtasks for themselves but were also willing to do subtasks assigned by Fetch. Table 7.6 shows the number of participants for each group. This indicates that 37 out of 48 participants had more of a leading preference, supporting Hypothesis 6.

Task difficulty: At the end of Task 0, we counted the number of misplaced blocks on the shared table. We also recorded the number of wrong actions by participants, including

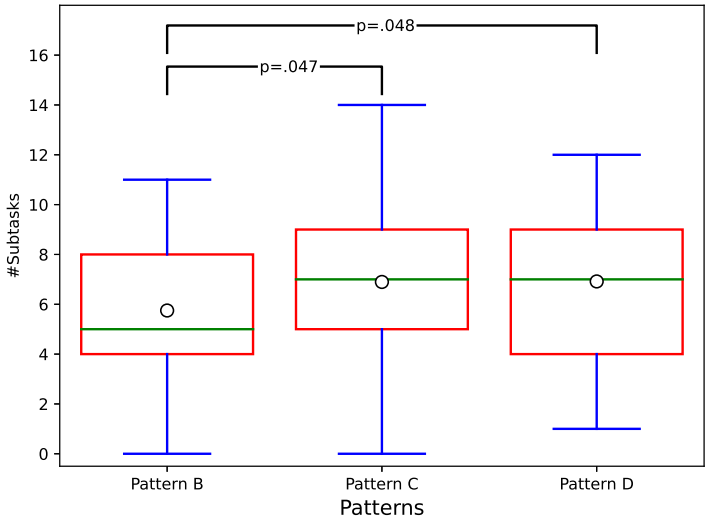


Figure 7.13: Number of subtasks assigned by participants to themselves in each pattern.

placing or assigning wrong blocks and returning a correct block from the shared table during the experiment. Regardless of the number of errors, 28, 12, 9, and 3 participants made at least a mistake in respectively Patterns A, B, C, and D. Pattern A was used only in Task 0, with the most participants who made mistakes.

Task assignment: We counted the number of subtasks that participants assigned to themselves based on the patterns. As depicted in Fig. 7.13, participants selected significantly fewer tasks for themselves in Pattern B compared to Patterns C and D. This indicates that they allowed the robot to assign more tasks to them. This finding aligns with their perception of task difficulty discussed earlier.

Additionally, we measured the number of subtasks that participants assigned to the robot, revealing a significant difference between the tasks ($H(2) = 11.56$, $p = .003$). Fig. 7.14 shows that participants preferred to assign fewer blocks to the robot, and there is a significant difference between Tasks 1 and 3. This can indicate participants' adaptation to the robot's slower speed in performing the pick-and-place task and having a better understanding of the robot's abilities.

Discussion: The results of the interviews show that most of the participants preferred to take on the leading role and have more control over the robot, which aligns with Hypothesis 6. This finding is insightful for designing future human-robot (or human-machine in general) collaborative systems. In addition, according to the interviews, participants

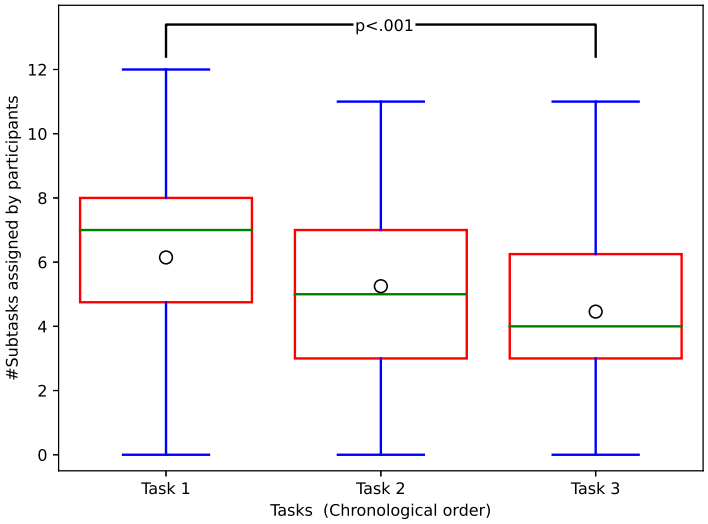


Figure 7.14: Number of blocks assigned by participants to Fetch in Tasks 1, 2, and 3

found the robot to be slower than themselves and preferred to handle more blocks. However, this preference might change if they were faced with a longer task requiring more physical effort. Some participants also pointed out that if the task were longer or the blocks were further apart, they would choose to assign tasks to the robot, even at the cost of completion time, to reduce their physical effort. Furthermore, the higher number of participants who made at least one mistake in Task 0 (Pattern A) compared to the other tasks indicates that the robot could assist them in improving their performance and making fewer errors. Except for Pattern A, which was only used in Task 0, participants made the most mistakes in Pattern B, a fact reflected in the subjective analysis of their perception of the task difficulty. This led them to rely more on the robot's decisions and help.

7.4.3 Leadership and Followership style

The analysis of followership styles reveals that the majority of participants (37 out of 48) are classified as "exemplary" followers, and ten of them fall into the "pragmatist" category. Only one participant exhibited traits of an "alienated" follower, and we did not encounter any instances of passive or conformist followers among the participants. This observation can likely be attributed to their shared academic background, which places a strong emphasis on critical thinking and active engagement.

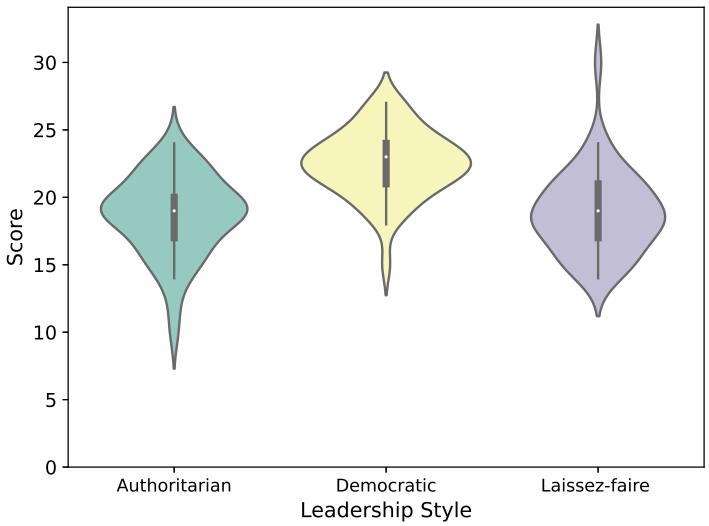


Figure 7.15: Caption

It is worth highlighting that all four participants listed in Table 7.6, who neither assumed leadership nor followed the robot, are categorized as "pragmatist" followers. However, due to the limited number of these observations, no definitive conclusion can be drawn, and more investigation is needed.

The distribution of leadership styles is depicted in Fig. 7.15. The majority of participants exhibited a predominant "democratic leadership" style. Among the four participants with no clear inclination towards leading or following, three displayed a dominant "laissez-faire leadership" style. However, we cannot make conclusive statements based on this limited sample size. The sole participants characterized by an alienated followership style also demonstrated a notably dominant "laissez-faire leadership" style and exhibited a leading-collaborative approach during the collaborative interactions.

We conducted an analysis of the correlation between engagement, critical thinking, and leadership styles in relation to various participants' factors. Intriguingly, our findings revealed a positive (negative) correlation between the number of tasks assigned by the robot to participants (the number of self-assigned tasks by participants) and their authoritarian leadership style. In other words, participants with authoritarian leadership tendencies were more inclined to accept additional subtasks assigned by the robot while selecting fewer tasks for themselves. Additionally, a moderate positive correlation was identified between participants' initial trust in the robot and their authoritarian leadership style.

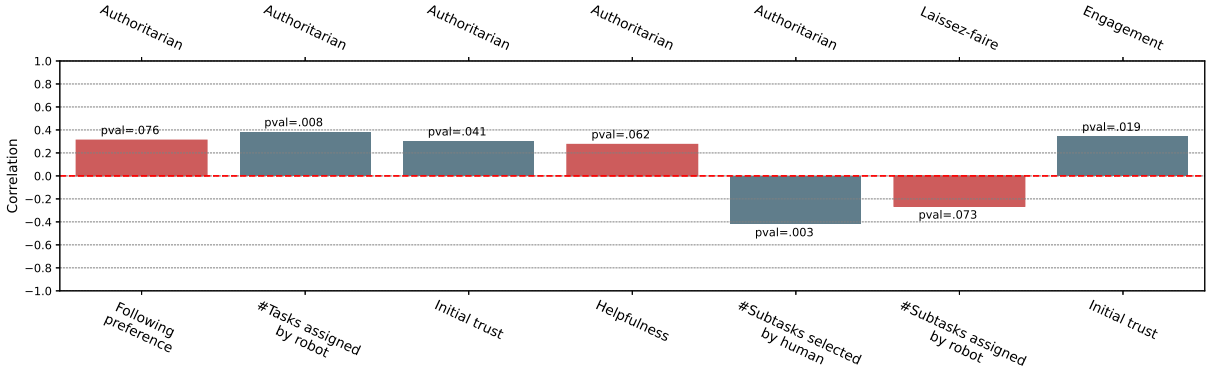


Figure 7.16: Correlations related to participants' leadership/followership styles

A positive, albeit nonsignificant, correlation was observed between these participants' following preference and their perceived helpfulness of the robot in relation to their authoritarian leadership style. However, this connection warrants further investigation in future research. We also encountered a nonsignificant negative correlation between participants' laissez-faire leadership style and the number of tasks assigned by the robot. Moreover, there was a positive correlation between participants' initial trust in the robot and their level of engagement. These correlations are shown in Fig. 7.16.

7.4.4 Robot's Actions and Estimation

Highlights from Subjective & Objective Analysis

Here, we summarize the findings from analyzing 1- participants' perception of the tasks, the robot, and collaboration and 2- participants' preference and performance.

1. Subjective assessments reveal an improvement in participants' perception of the robot and collaboration following their interaction, along with a reduction in perceived workload.

2. Both subjective and objective assessments demonstrate that the robot effectively assisted participants in enhancing their performance and reducing errors.

3. The interview results show that most participants preferred to take on the leading role and have more control over the robot. Based on participants' preferences, we

categorized them into four groups (from highest leading preference to highest following preference): 1- lead (17 participants), 2- collaborative-lead (20 participants), 3- collaborative-follow (4 participants), and 4- follow (3 participants).

4. According to the interviews, participants found the robot to be slower than themselves and preferred to handle more blocks.

5. The results indicate that participants, in general, found Pattern B more difficult compared to Patterns C and D. In addition, there was the highest number of participants who made at least one mistake in Pattern B. This can be attributed to the fact that Pattern B was both challenging to memorize and had the most unknown spots, leading participants to rely on the robot.

We briefly highlighted participants' preferences and performance based on the interviews and recorded data from the user study. However, we must also explore how the robot could adapt to participants' preferences and performance.

Participants Preference

The robot updated its estimation of the human preference based on the 3-step history of the human's actions. The robot needs to consider the human agent's preference changes and adapt accordingly. However, we needed to create a measure to evaluate the robot's performance in estimating the human overall preference. To do so, first, we normalize the completion time and then fit a polynomial (e.g., degree 4) on the estimated values, $f(t), t \in [0, 1]$. Then, we calculate the area under the curve in a certain range, as a measure of participants' overall preference, $op = \int_{t_0}^{1} f(t)$, where we considered $t_0 = 0.2$.

To evaluate the robot's ability to estimate participants' preference, for each of them, we measured the average of the overall preference (op) in Tasks 1, 2, and 3. Combining them with the information gathered in interviews (i.e., participants' actual preferences) leads to Fig 7.17, which shows that the robot effectively estimated participants' actual preferences.

In addition, we analyzed participants' estimated overall preference for each task and noticed no significant difference among them ($H(2) = 2.29$, $p = 0.32$). This is justifiable as the majority of participants preferred to lead the robot.

Task Difficulty - Following Preference

As stated in the highlights of the participants' subjective and objective measurements, participants found Pattern B more challenging than Patterns C and D. Analyzing partici-

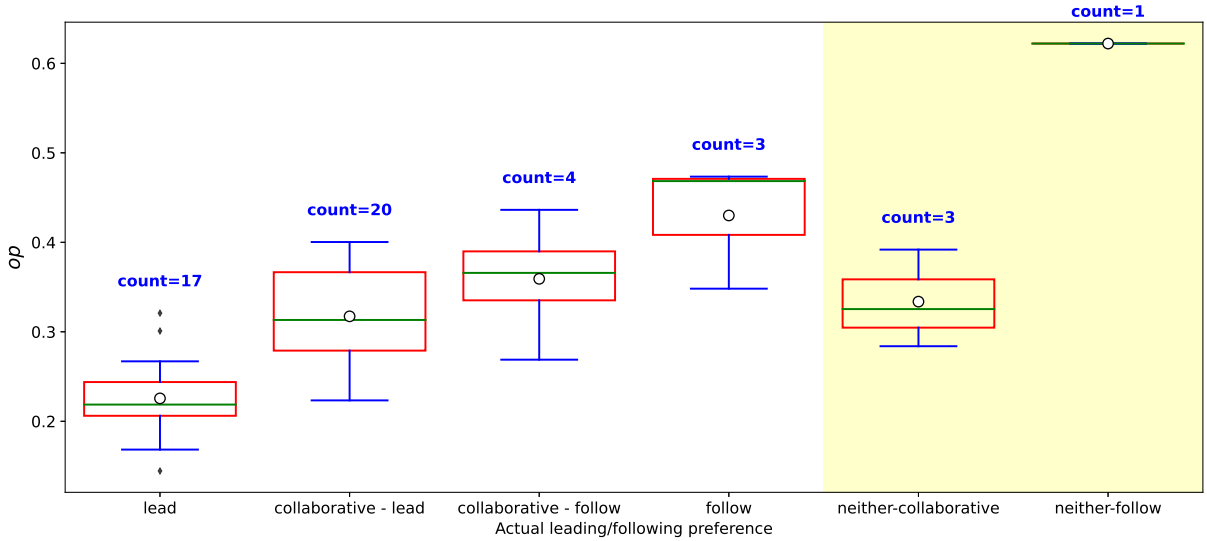


Figure 7.17: The robot performance in estimating participants' preference by comparing with the participants' actual preference, gathered through interviews. The numbers above each bar show the number of participants falling into that group.

pants' estimated overall preference based on the patterns, as illustrated in Fig. 7.18, shows that there is a significant difference in participants' following preference based on the patterns ($H(2) = 6.5$, $p = 0.039$), indicating a significantly higher preference for following in Pattern B compared to Patterns C and D.

Based on the robot's task planning algorithm, the greater the human agent's preference for following or the occurrence of errors, the more tasks are assigned to them by the robot. Fig. 7.19 shows that the robot significantly assigned more subtasks to the human in Pattern B compared to Patterns C and D ($H(2) = 6.32$, $p = 0.042$). Likewise, this is justifiable as Pattern B was more difficult than the two others, and participants made more mistakes or preferred to follow the robot, and the robot could accordingly adapt its planning.

Task Assignment & Distribution

Referring back to the distribution of the blocks (Table 6.1) and the lower speed of the robot compared to the human, the optimal task allocation will be closer to assigning pink blocks to the robot and blue and orange blocks to the human. Fig. 7.20 shows the colors of blocks completed by participants and those assigned by the robot. This distribution, resulting

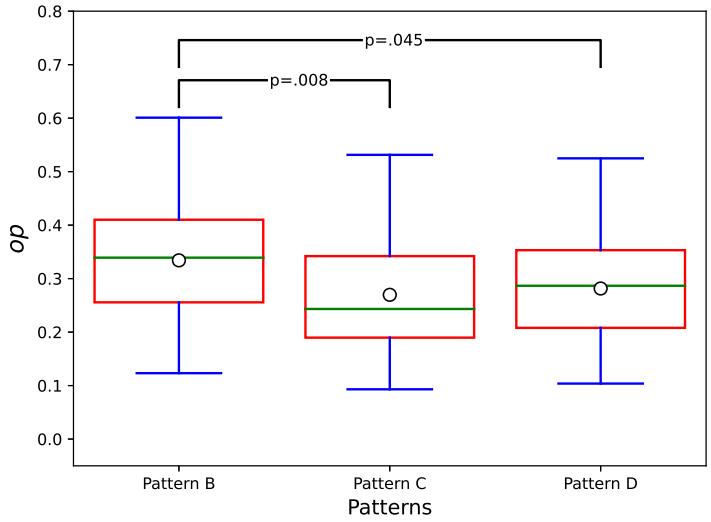


Figure 7.18: Participants' overall estimated leading/following preference based on the patterns. The robot's estimation of participants' preference to follow was significantly higher in Pattern B compared to Patterns C and D.

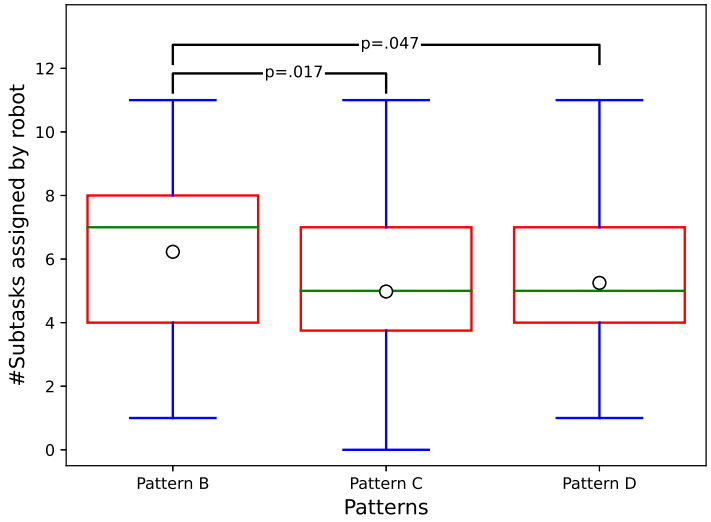


Figure 7.19: Subasks assigned to participants by the robot based on the patterns. The number of assigned subtasks by the robot to participants was significantly higher in Pattern B compared to Patterns C and D.

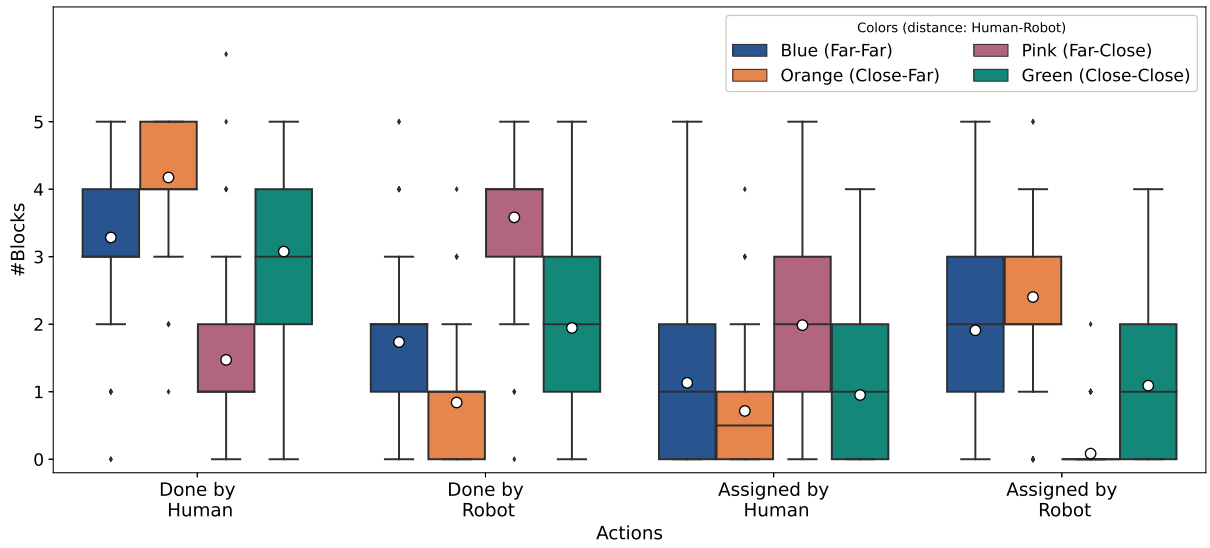


Figure 7.20: Color of block that were done by human and the robot as well as were assigned by the robot or human to each other

from the interplay between the robot and participants, is close to optimal. As expected, most of the orange blocks were completed by the human agent due to participants' expected rational decision-making and the robot's assignments. Similarly, pink tasks were completed by the robot. Regarding the blue subtasks, on average, participants completed most of them. However, the distribution of assigned and completed blue subtasks by the robot and human ranges from 0 to 5. This variation is due to some participants, as indicated by interviews, who preferred to assign more blue tasks to the robot, reducing their physical effort at the expense of longer collaboration time.

Following preference vs. Task distribution: Based on the designed algorithm, we expect the robot to assign more subtasks to a human with the following preference, which is particularly important for blue subtasks. Additionally, we are interested in examining how the robot's estimation of participants' following preference relates to the number of subtasks participants assign to it. As shown in Fig. 7.21, there is respectively a strong positive correlation and a strong negative correlation between the number of subtasks the robot assigned to participants and the number of subtasks that the human assigned to the robot with the estimated participants' following preference. This aligns with the designed algorithm.

Furthermore, we can observe a moderate negative correlation between the number of

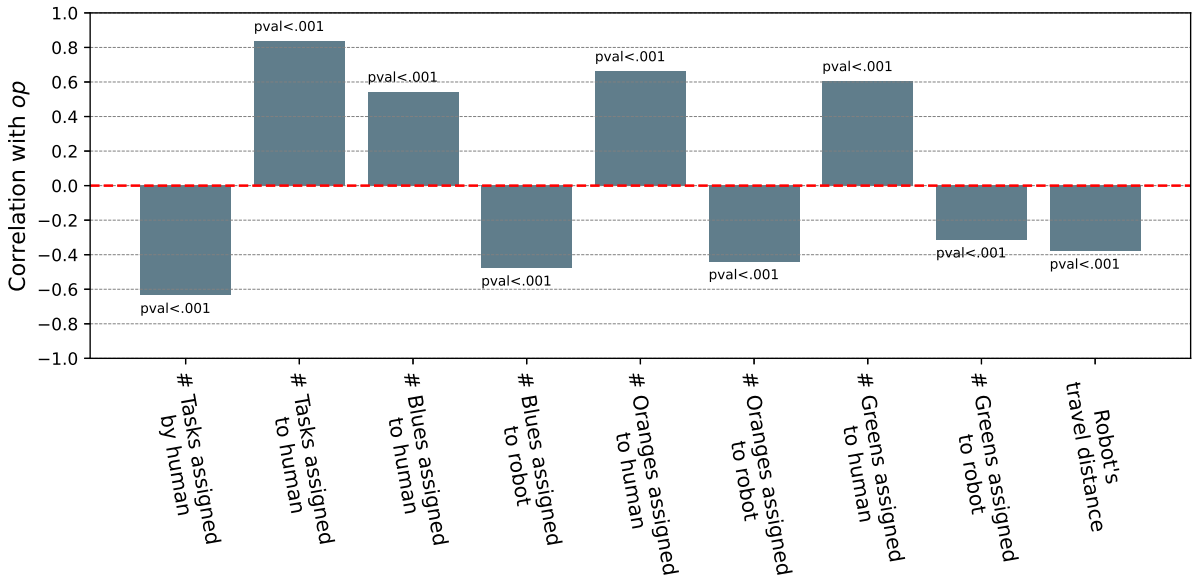


Figure 7.21: Correlation between estimated participants' following preference and assigned/done tasks by the robot and participants

assigned blue and orange tasks to the robot by humans and their following preference. Conversely, a moderate to strong positive correlation exists between the orange, green, and blue blocks allocated to the human agent and the estimated following preference. A weak negative correlation exists between the number of assigned green tasks and participants' following preference. As expected, the correlation between the approximate robot's travel distance and the estimated following preference shows a moderate negative correlation, as when the robot has a more leading role, it assigns tasks that are farther from itself to the human agent.

Discussion: The results indicate that the robot could successfully identify participants' following preferences and adapt its planning accordingly. We also observed that the robot's estimation of participants' following preference was higher for Pattern B than for Patterns C and D, resulting in more tasks being assigned to participants. This aligns with the previous finding that Pattern B was more challenging than the other two patterns, and participants required more support from the robot. All of these, in concert, showed that the robot was more of a leader in challenging tasks. This could result from either the participants' choice to follow the robot or the significant number of mistakes they made. Additionally, the robot could guide participants toward optimal task allocation to minimize

collaboration time, taking into account the block's location and the fact that it is slower in comparison to the human agent.

7.4.5 Participant-Specific Analysis

We discussed the overall participants' preferences and performance. The main goal of the proposed framework, however, is to track changes in participants' performance and preferences. To evaluate the framework's online adaptation and planning ability, we discuss the results for some specific participants.

Leading or collaborating-leading Preference with a high accuracy

These participants preferred to lead the team and had a high accuracy. Thus, the robot estimated their leading preference and gave the leading role to them. Fig. 7.22a and 7.22b show the robot estimate of two participants' error-proneness and following preference, with respectively leading and collaborating-leading preferences.

Leading preference with occasional robot support

One of the participants preferred to lead the robot and minimize her physical effort by assigning most of the tasks to the robot. She also approved this in her interview and mentioned that she followed the same strategy in all three tasks. However, for Pattern B, she forgot the part of the pattern and thus let the robot assign her some tasks. She was also unsure about the last row of Patterns C. In Fig. 7.22c, the robot's estimation of her preference and performance is shown. Fig. 7.22d shows the robot estimations for another participant with a leading preference who was unsure about pattern B.

Leading preference with occasional errors

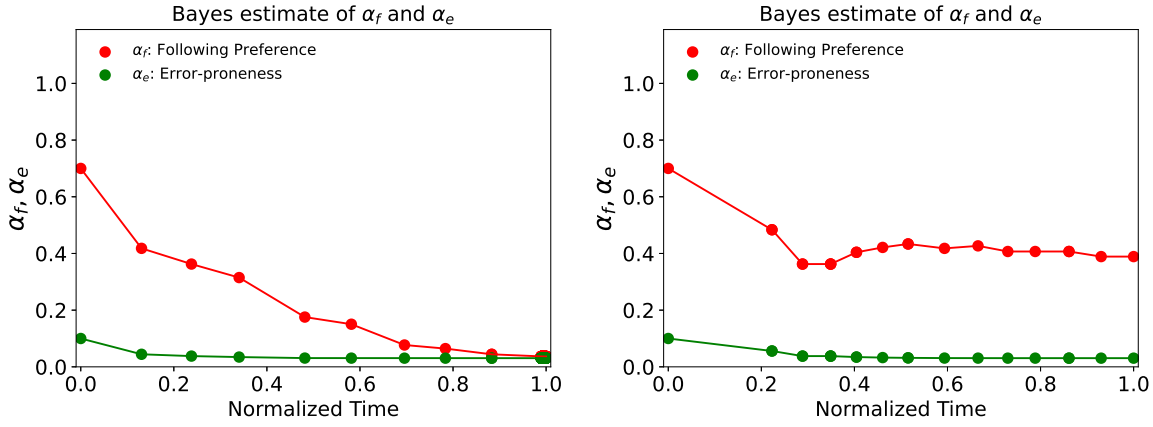
In this case, some participants took the leading role, and, in total, they could recall the pattern, but they made a few mistakes at some points. The robot detected the errors and updated its belief about their performance with usually a slight and temporary rise of α_e. Fig. 7.22e, 7.22f, and 7.22g show the robot's estimate of three participants' preference and performance who had occasional errors. In Fig.7.22e, the participant made some consecutive mistakes at a point, but she improved her performance with the help of the robot.

Leading preference with Sudden Performance Drop

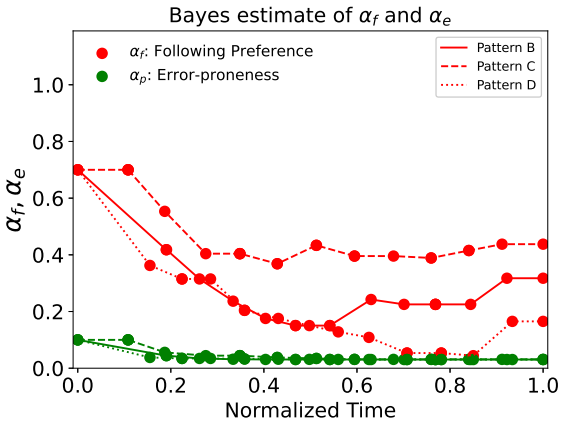
This participant led the robot in Pattern B, and her performance was good until almost the end of the task. However, she started making mistakes at the end as she had forgotten and was confused about the last row of the pattern. She initially insisted on her wrong decisions, although the robot rejected her assignment and fixed her mistakes. Subsequently, the robot updated its belief about her performance (Fig. 7.22h). While considering her as a person with a leading preference, it reassumed the leading role and assigned more tasks to her, guiding her to the correct pattern.

Following or collaborating-following Preference

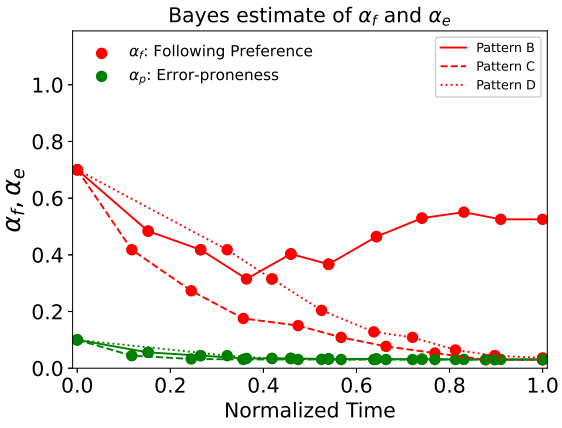
Fig. 7.22i and 7.22j show the robot estimate of two participants' error-proneness and following preference, who preferred to follow the robot. Therefore, the robot took the leading role and assigned subtasks to them.



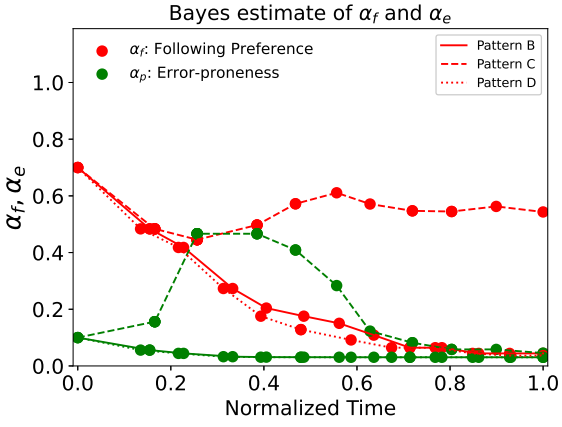
(a) (b)


Bayes estimate of α_f and α_e
α_f: Following Preference
α_p: Error-proneness
Pattern B
Pattern C
Pattern D
α_f, α_e
Normalized Time

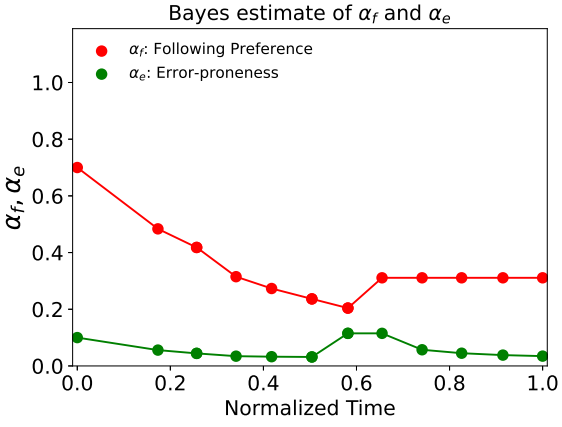
(c)


Bayes estimate of α_f and α_e
α_f: Following Preference
α_p: Error-proneness
Pattern B
Pattern C
Pattern D
α_f, α_e
Normalized Time

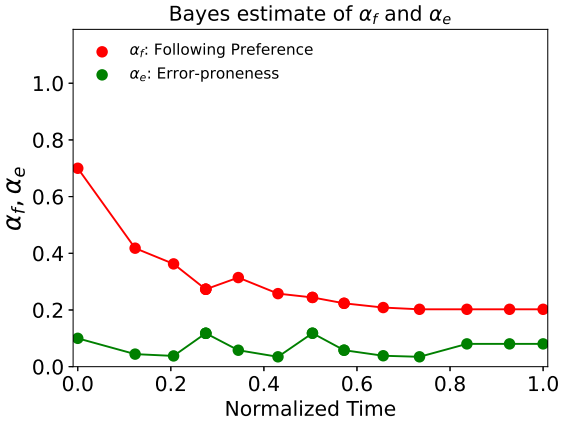
(d)


Bayes estimate of α_f and α_e
α_f: Following Preference
α_p: Error-proneness
Pattern B
Pattern C
Pattern D
α_f, α_e
Normalized Time

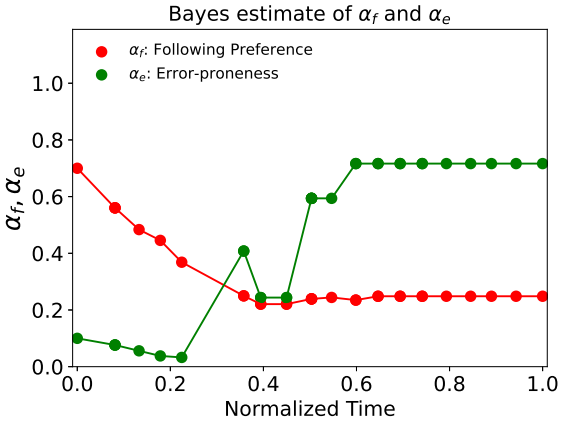
(e)


Bayes estimate of α_f and α_e
α_f: Following Preference
α_e: Error-proneness
α_f, α_e
Normalized Time

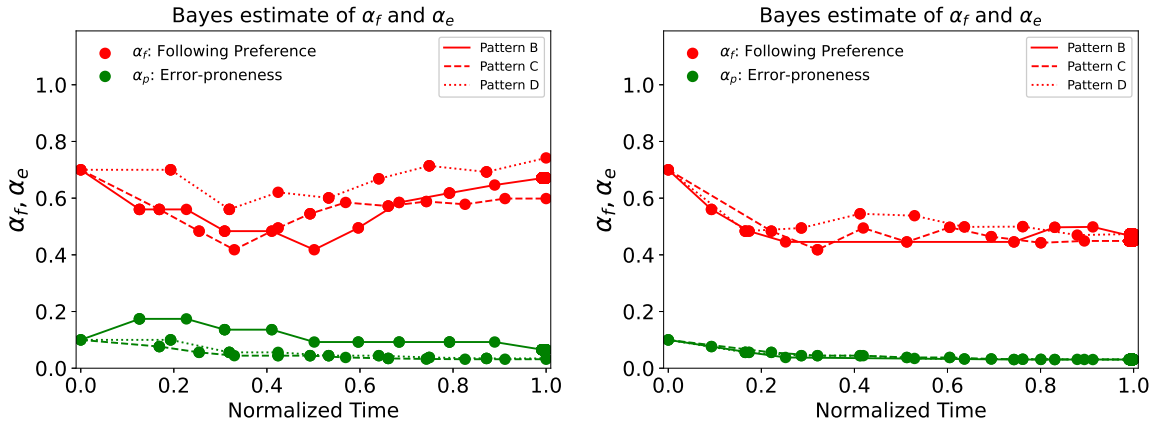
(f)


Bayes estimate of α_f and α_e
α_f: Following Preference
α_e: Error-proneness
α_f, α_e
Normalized Time

(g)


Bayes estimate of α_f and α_e
α_f: Following Preference
α_e: Error-proneness
α_f, α_e
Normalized Time

(h)



(i) (j)

Figure 7.22: The robot's estimates of the participants' preference and performance: a, b) Leading or collaborating-leading preference and high accuracy; c, d) leading preference with occasional robot support; e, f, g) leading preference with occasional errors; h) leading preference with sudden performance drop; i, j) following or collaborating-following preference

7.5 Summary and Discussion

In this chapter, we focused on addressing RQ5, RQ6, and RQ7. We first discussed and developed the study hypotheses. Then, we delved into the user study procedure and the measures considered for evaluating and analyzing the results of participants' actions and their perception of the robot, task, and collaboration. Finally, we explored the ability of the adaptive task planning framework of the robot to infer participants' preferences and performance and planning based on them.

Trust is an integral part of this collaboration. We assessed how the robot's adaptive and proactive planning can improve participants' trust in the robot. The subjective measurement of participants' trust in the robot confirmed our Hypothesis 1, showing trust enhancement as they collaborated more with the robot. Additionally, we hypothesized that participants' self-confidence would improve over collaboration with the robot, which was confirmed by analyzing their self-reported self-confidence (Hypothesis 2). Furthermore, we observed an increase in the expected helpfulness of the robot by participants, supporting Hypothesis 5. According to participants' answers to the questionnaires, they found the robot to be a teammate with positive traits and a strong working alliance, with

whom they were willing to collaborate again in the future.

In support of Hypothesis 3, the results indicated that for more difficult tasks, participants trusted the robot more than their own abilities, which led them to take relatively more following roles. Our proposed task planning method properly inferred this need and provided more help to participants by taking on more leading roles. The robot could also identify when participants struggled to remember the correct patterns and made errors, and accordingly, it fixed their errors and provided more help. This ability of the robot to adapt its task planning led us to hypothesize that participants would perceive, overall, a lower workload (Hypothesis 4). The results partially supported our hypothesis, and we observed an enhancement in participants' perceived mental demand, performance, effort, frustration, and unexpectedly temporal demand. However, contrary to our hypothesis, participants did not perceive a lower physical demand, likely due to the designed scenario focusing more on mental ability than physical abilities.

Based on interviews with participants, we categorized them into four groups: "lead", "collaborative-lead", "collaborative-follow", and "follow", with the majority falling into the first two categories. This supports Hypothesis 6 that participants, while having a high level of trust in the robot, would prefer to take on more of a leading role and have more control over the collaboration. This finding can guide the design of collaborative scenarios and collaborative robots. Furthermore, we compared this result, showing actual participants' leading/following preferences, with the robot's estimation of their preference. This analysis showed that the robot successfully inferred their preference in most cases.

Additionally, the results showed that participants selected fewer tasks for themselves and were allocated more subtasks in tasks that were more difficult since they preferred to follow the robot or made many errors, causing the robot to take back the leading role. In other words, participants trusted the robot more than their own abilities, which led them to take relatively more following roles. Our proposed task planning method properly inferred this need and provided more help to participants by taking on more leading roles. The robot could also identify when participants struggled to remember the correct patterns and made errors, and accordingly, it fixed their errors and provided more help.

We also inquired about participants' leadership and followership styles using questionnaires, addressing RQ7. Our observations revealed that individuals with an authoritarian leadership style exhibited greater trust in the robot. They also allowed the robot to assign more subtasks to them. Additionally, we identified a positive correlation between their initial trust and the level of engagement. Nevertheless, a more comprehensive investigation involving a diverse participant group is warranted to delve deeper into this subject.

We also analyzed the distribution of subtasks between the robot and the human agent,

showing that, overall, the interplay of participants and robot agents led to near-optimal task allocation. The results also indicated a moderate to strong correlation between estimated participants' preferences and measures of task distribution. Additionally, the results showed that participants were allocated more subtasks in tasks that were more difficult since they preferred to follow the robot or made many errors, causing the robot to take back the leading role.

Chapter 8

Conclusion and Future Work

8.1 Summary

After reviewing the general problem, defining the research questions, and clarifying the contributions, the thesis proceeded to Chapter 2, where cobots and their applications in industry were explored alongside a review of the current research landscape. The chapter then delved into a review of related work on task planning and allocation in human-robot collaboration, focusing on the integration of human preferences. Subsequently, it explored ongoing work in adaptation within Human-Robot Interaction (HRI) and Human-Robot Collaboration (HRC).

Chapter 3 presented the outcomes of an online study in which participants engaged with a robot in a simulated environment, employing various strategies. The findings revealed no significant difference in participants' perception of the robot when it prioritized the human agent's objectives or both agents' objectives (team performance). In this chapter, we mainly focused on the first research question of this thesis:

RQ 1. How does the robot's level of cooperation influence the human's perception of the robot (e.g., trust and satisfaction)?

The primary finding from investigating this research question highlights that effective robot planning aimed at enhancing team performance can indeed elevate overall team performance while also upholding the human agent's perception of the robot and collaboration. The insights gained from the results, presented in Chapter 3, confirmed the feasibility of

the performance and preference-based planning and laid the foundation for Chapter 4, where we introduced a task planning framework and algorithm. This framework enables the robot to adapt its plans according to the human agent's performance and preference to lead or follow. This chapter addressed the thesis's second research question, stating:

RQ 2. How can we design a planning architecture that takes into account both human preference and performance and enables the robot to update the task state based on both agents' actions?

The task planning architecture and algorithm outlined in Chapter 4 serve as a foundational framework applicable to a wide array of task and collaborative scenarios. However, customization is necessary to tailor the algorithm to specific tasks and collaborative contexts, ensuring optimal performance and adaptability. Thus, to test the practicality of the proposed adaptive task planning method, in Chapter 5, we moved from theory to practice, implementing the task planning framework in a simulated environment with a simplified human decision-making model.

The developed simulation environment serves as a flexible platform for investigating various scenarios, evaluating the performance of task planning algorithms, and refining them iteratively. Although the simulation employs a simplified human model, its utility lies in its ability to simulate scenarios that may be infrequent or challenging to replicate in user studies. This approach facilitates comprehensive exploration and analysis, contributing to the refinement and enhancement of task planning strategies in human-robot collaboration. This chapter helped us see how well the proposed task-planning framework enabled the robot to adapt to different cases of the human agent's performance and leading/following preference. Additionally, a comparison study between the adaptive and non-adaptive planners revealed significant enhancements in task planning and overall team performance with the adaptive planner. This chapter, along with Chapters 6 and 7, focused on research questions RQ3 and RQ4, reiterated as below:

RQ 3. Does the proposed adaptive task planning for the robot improve overall team performance in human-robot collaboration?

RQ 4. To what extent can the proposed task planning framework enable the robot to adapt to the human agent's performance and preferences successfully?

Chapter 6 introduced the challenges associated with implementing the task planning framework on a real robot within a newly designed collaborative scenario. This new collaborative scenario also shows that the general task planning architecture and algorithm

can be applied to various scenarios by applying required customizations. In this chapter, we initiated the discussion with an overview of the experimental setup and methodology. Subsequently, we delved into the details of implementing and integrating the proposed framework into the designed collaborative scenario. This new collaborative scenario also demonstrates that the general task planning architecture and algorithm can be applied to various scenarios with the necessary customizations. Finally, to assess the robot's ability to adapt its planning based on the human agent's performance and leading/following preference, a system evaluation was conducted, involving the experimenter enacting potential behaviors representative of human agents. This chapter also centered on RQ3 and RQ4 by exploring the results obtained from implementing the proposed framework on a real mobile manipulator robot.

Chapter 7 shifted the focus to the procedure and results of the user study involving 48 participants. After introducing the objective and subjective measures, we evaluated the results from three different viewpoints: 1- participants' perception of the robot, tasks, and collaboration, 2- participants' actions and performance, and 3- the robot's actions and performance. The results demonstrated that the main goals of the proposed task planning framework were achieved: adapting the robot's planning to the human agent's performance and leading/following preference to improve team performance while maintaining their perception of the robot and collaboration at a high level.

In terms of participants' perception of the robot, tasks, and collaboration, the results indicated a high level of trust in the robot, expected helpfulness, and self-confidence among participants. They also reported a positive collaboration experience, viewing the robot as a high-performing teammate with positive traits and a strong working alliance, whom they would be willing to collaborate with again. Additionally, the results showed a significant decrease in perceived workload when participants collaborated with the robot compared to working alone. Furthermore, correlations were found between participants' initial trust in the robot and their perceptions of its traits.

With regard to participants' actions and performance, the results indicate that collaboration improved participants' performance compared to when they worked alone. Furthermore, their perception of task difficulty suggests that participants found the task more challenging when working alone without robot collaboration. Additionally, participants' actions, such as assigning fewer tasks to the robot as they collaborated more with it, suggest that they perceived the robot as slower but with a better ability to memorize patterns. This observation indicates adaptation in planning for their collaboration. Another significant observation from the interviews was that the majority of participants preferred to be in control and lead the team. This finding underscores the importance of considering human preferences and their desire for control in the design of future human-robot collaborative

scenarios. These addressed the research question RQ5 and RQ6:

RQ 5. How does collaboration with the robot influence the human agent's actions and their perception of the robot and collaboration?

RQ 6. To what extent do the human agent's perception of the robot and collaboration impact their interaction with the robot?

Ultimately, this chapter explored potential correlations between the leadership or followership styles of participants and their collaborative interactions with the robot, to address RQ7:

RQ 7. Is there any correlation between the human agent's leadership/followership styles and their collaborative approach with the robot?

However, deeper studies involving a broader range of participants are necessary to obtain more conclusive and generalizable results. These studies could include participants from diverse demographic backgrounds, varying levels of expertise and professional experience to ensure the findings are robust and applicable across a wider spectrum.

In terms of the robot's actions and performance, the comparison between participants' actual preferences obtained during the interviews and the robot's estimations confirmed the robot's ability to accurately estimate participants' preferences. Moreover, the robot successfully identified when participants preferred to follow its lead for tasks that posed greater challenges to them. As a result, the robot provided more assistance and assumed a more leading role in these situations. Additionally, analysis of task distributions revealed that participants and the robot, working together as a team, were able to devise plans that closely resembled the optimal offline plan, involving task allocation and scheduling in advance. This finding demonstrates the effectiveness of their collaborative efforts. These findings based on the user study data offer deeper insights into addressing RQ3 and RQ4.

8.1.1 Takeaways

The task planning framework and algorithm described in the thesis provide a foundational architecture that, with customization to suit specific contexts and tasks, can be applied to various tasks and collaborative scenarios. We tested it on the designed collaborative scenarios in a simulation environment and an actual robot in a user study. Although the results discussed are based on these particular studies, we can draw some general takeaway messages from them:

- Efficiently designed adaptive planning, achieved by customizing the proposed task planning framework, can significantly enhance participants' perception of the robot and collaboration (e.g., satisfaction and trust in the robot).
- Additionally, adaptive planning leads to improvements in team performance and reduces participants' perceived task load. This finding underscores the importance of adaptive strategies in enhancing overall collaboration efficiency.
- In the design of future collaborative scenarios, it is crucial to consider that humans typically prefer to maintain a degree of control when collaborating with machines or robots. Understanding and accommodating this preference can enhance the acceptance and effectiveness of human-robot collaborative systems.
- Human agents' preferences and performance can vary based on factors such as task difficulty. The adaptive planner demonstrates the ability to infer these changes and adjust the robot's plan accordingly, optimizing collaboration outcomes.
- Acknowledging the inevitability of human error in collaboration, it is essential to integrate error detection mechanisms into robot planning. The proposed adaptive task planning approach effectively detects errors and adjusts the robot's plan to maintain high team performance levels, even in the face of human errors.

By incorporating these insights into the design and implementation of collaborative systems, we can pave the way for more effective and fluent human-robot interactions and foster greater collaboration efficiency.

8.2 Contributions to Original Knowledge

In this thesis, our primary focus has been the investigation of whether and how proactive task planning and allocation can enhance the efficiency of human-robot collaboration. A noticeable gap in the existing literature lies in the omission of either human agents' leading/following preferences or the human agent's performance. Therefore, our central emphasis in this thesis has been on addressing this gap by achieving a balance between the human agent's preference and performance, all while sustaining a high level of collaboration and human perception of the robot. To initiate this exploration, we conducted an online user study, evaluating three distinct robot planning strategies within a collaborative scenario. The outcomes of this study, presented in a publication ([16]), contribute to a

deeper understanding of how robot decision-making influences human perception of the robot and collaboration in collaborative tasks:

- A. Noormohammadi, A. Dahiya, A. M. Aroyo, S. L. Smith, and K. Dautenhahn, "The effect of robot decision making on human perception of a robot in a collaborative task - a remote study," in Proceedings of the 9th International Conference on Human-Agent Interaction, 2021, p. 423–427.

Next, building upon insights gathered from the initial study, we developed a framework enabling the robot to adapt concurrently to both the human's leading/following preferences and their performance. The key feature of this framework lies in its gradual and continuous adaptation to human performance and preferences, avoiding extremes where the human is limited to either leading or following the robot. Following the establishment of this general planning framework, our focus shifted to the specific problem of task selection and planning. Unlike conventional task allocation approaches, our approach enables both the human and robot to collaboratively choose and assign tasks. A simulated collaborative scenario was created, employing a simplified human model, and the planning algorithm was applied to various combinations of human accuracy and preference to lead. The results demonstrated that our proposed algorithm empowered the robot to adjust its leading role based on the human's accuracy and preference. Furthermore, comparing our algorithm with a scenario where the robot only considered completion time revealed a significant enhancement in the team's performance. These findings are detailed in Chapter 4 and Chapter 5 and have been published as a conference paper:

- A. Noormohammadi-Asl, A. Ayub, S. L. Smith, and K. Dautenhahn, "Task selection and planning in human-rob collaborative processes: To be a leader or a follower?" in 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 2022, pp. 1244–1251.

Despite the satisfactory outcomes of the proposed framework in this simulation study, the effectiveness of a human-related framework can only be thoroughly confirmed when applied to studies with real human participants, accounting for all possible uncertainties that might arise due to individual preferences and behaviors. Thus, I began implementing it on a real robot, envisioning a scenario in which the human and robot collaborate to accomplish a mutual task. The experiment involved a human-robot collaboration scenario where the team needed to work together to achieve a collaborative task. The designed scenario offers flexibility, allowing the human to choose how to interact and collaborate

with the robot, and allows the robot to adapt to the human agent's leading/following preference and performance, taking the leading role when necessary.

However, before recruiting participants, given the complexity of the task, the framework, and the autonomous pick-and-place task, it was imperative to conduct a system evaluation to ascertain the proper and robust functioning of the robot and the system. Therefore, as the experimenter, I enacted various collaborative behaviors that participants might exhibit to evaluate the effectiveness of the proposed framework. The results showed that the proposed framework enabled the robot to adapt itself based on different human preferences and performances. The results of this study were published as a conference paper:

- A. Noormohammadi-Asl, A. Ayub, S. L. Smith, and K. Dautenhahn, "Adapting to human preferences to lead or follow in human-robot collaboration: A system evaluation," in 2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), 2023, pp. 1851–1858.

In the final phase, I recruited participants and conducted a user study with the fully functioning system. The results revealed that participants' perception of the robot was notably elevated, encompassing aspects such as trust, robot intelligence, helpfulness, and effective collaboration. Additionally, the findings indicated that participants' trust in the robot and their perception of the robot's intelligence improved during the sessions over time.

Furthermore, the outcomes demonstrate the robot's capability to adapt its planning and task selection in accordance with human preference and performance. It effectively guided participants toward optimal strategies and plans. Comparing estimated participant preferences and their actual preferences and plans, gathered from interviews, revealed the robot's success in estimating their preferences. This result, in conjunction with the aforementioned findings, validates the successful realization of the proposed framework's original objectives.

- Adapting to human preference and performance
- Maintaining the team performance at a high level
- Establishing a heightened level of human perception of the robot and collaboration

Additionally, we analyzed the correlation between participants' initial trust in the robot and their task performance. Lastly, we evaluated potential correlations between their

leadership/followership style and the connection between their initial trust in the robot and their collaborative interactions with the robot. The details of the user study, results, and their analysis were submitted as two journal papers, which are currently under review:

- A. Noormohammadi-Asl, S. L. Smith, and K. Dautenhahn, "To lead or to follow? adaptive robot task planning in human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01483 (Under review at a journal)
- A. Noormohammadi-Asl, K. Fan, S. L. Smith, and K. Dautenhahn, "Human leading or following preferences: Effects on human perception of the robot and the human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01466 (Under review at a journal)

8.3 Limitations & Future Work

This work has certain limitations in terms of the study design and methods, which, in turn, provide valuable avenues for future research and improvements.

Participant demographics: The majority of our participants were young adults recruited from the University of Waterloo campus, while our ultimate target audience consists of working adults in settings such as manufacturing and warehouses. These two groups may have considerably different expectations and perceptions of a robot teammate. Recruiting actual working adults in those settings could help us create more practical collaborative settings and robots. Additionally, studying human-human collaboration could provide valuable insights into designing collaborative systems for human-robot teams.

Advancing toward real-world collaboration and applications: Additionally, despite our attempts to simulate a working environment, e.g. by having an autonomous robot performing pick-and-place tasks and a setup including a conveyor belt, safety equipment, and a graphical user interface, future work could further improve the scenarios to make them as close as possible to a realistic environment, e.g. a warehouse automation or assembly environment. Moreover, incorporating abilities such as direct communication, human-intention estimation, and human physical state measurement (e.g., speed) could enhance team coordination and effectiveness in dealing with uncertainties and handling challenging tasks.

Dynamic manufacturing challenges: Another prevalent issue in manufacturing settings is the occurrence of sudden changes or unpredictable events that only humans can effectively manage. Moreover, tasks in real scenarios are not always predetermined,

emphasizing the necessity for a high level of adaptability, collaboration, and decision-making skills from the human-robot team, particularly given that the difficulty of certain tasks may differ between humans and robots. Incorporating these cases could further enhance the study's alignment with real-life situations. This necessitates the robot to promptly recognize the human agent's response to such situations and adjust its planning and actions accordingly.

Conflict resolution: One intriguing and important avenue for future research involves expanding the framework to address scenarios where conflicts arise between agents, such as instances where both the human and the robot believe their decisions are correct. In our current study, we operated under the assumption that the robot's decisions were always accurate. However, in real-world settings, the robot may make mistakes, and it may not be aware of them, potentially stemming from issues like perception inaccuracies. A potential solution to navigate such situations could involve empowering the robot to consider a degree of uncertainty in its decisions, which may be contingent on factors such as the (sub)task and the human agent's performance.

Performance metric misinterpretation: In this study, we evaluated the human agent's performance by measuring their accuracy in selecting block colors, a criterion that participants easily understood. While this straightforward measure served well in our controlled study environment, real-world scenarios may necessitate the consideration of alternative metrics, such as completion time or travel distance. The incorporation of these metrics could introduce complexity, given that they may not be as intuitively understandable to human agents. This lack of clarity might potentially lead to conflicts between collaborating agents, as the human agent's interpretation of team performance could differ from that of the robot.

Multi-human multi-robot team: The task planning framework presented in this thesis exclusively addresses single-human – single-robot scenarios. However, real-world scenarios often involve team structures like single-human – multi-robot, multi-human – single-robot, or multi-human – multi-robot configurations. Therefore, an intriguing avenue for future research could involve extending the proposed framework to accommodate scenarios where either or both the human and robot teams comprise more than one agent.

Improving scalability in task allocation and scheduling algorithm: The proposed task planning algorithm divides the problem into two subproblems: 1) task allocation, considering the human agent's preference and performance, and 2) task scheduling. To facilitate finding the optimal solution, a warm start approach was also implemented in the task planning algorithm. However, a potential future research direction could involve devising dedicated solutions for each of these two optimization problems, specifically

addressing scenarios with a high number of subtasks.

Leadership/followership styles: While participants' leadership and followership styles were analyzed in the user study, the results were limited to a specific demographic (i.e., students at the University of Waterloo). Some correlations between certain factors were observed, but due to the limited observations, drawing conclusive findings was unfeasible. This underscores the need for a future study that conducts a more comprehensive investigation into how human leadership and followership styles influence collaboration with a robot.

References

[1] A. M. Djuric, R. Urbanic, and J. Rickli, "A framework for collaborative robot (CoBot) integration in advanced manufacturing systems," *SAE International Journal of Materials and Manufacturing*, vol. 9, no. 2, pp. 457–464, 2016.

[2] B. M. Muir and N. Moray, "Trust in automation. part ii. experimental studies of trust and human intervention in a process control simulation," *Ergonomics*, vol. 39, no. 3, pp. 429–460, 1996.

[3] G. Hoffman, "Evaluating fluency in human–robot collaboration," *IEEE Transactions on Human-Machine Systems*, vol. 49, no. 3, pp. 209–218, 2019.

[4] F. Vicentini, "Collaborative Robotics: A Survey," *Journal of Mechanical Design*, vol. 143, no. 4, p. 040802, 10 2020.

[5] S. El Zaatari, M. Marei, W. Li, and Z. Usman, "Cobot programming for collaborative industrial tasks: An overview," *Robotics and Autonomous Systems*, vol. 116, pp. 162–180, 2019.

[6] N. Mitsunaga, C. Smith, T. Kanda, H. Ishiguro, and N. Hagita, "Adapting robot behavior for human–robot interaction," *IEEE Transactions on Robotics*, vol. 24, no. 4, pp. 911–916, 2008.

[7] H. Nemlekar, N. Dhanaraj, A. Guan, S. K. Gupta, and S. Nikolaidis, "Transfer learning of human preferences for proactive robot assistance in assembly tasks," in *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 2023, p. 575–583.

[8] W. Wang, R. Li, Y. Chen, Z. M. Diekel, and Y. Jia, "Facilitating human–robot collaborative tasks by teaching-learning-collaboration from human demonstrations," *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 2, pp. 640–653, 2019.

[9] E. C. Grigore, A. Roncone, O. Mangin, and B. Scassellati, "Preference-based assistance prediction for human-robot collaboration tasks," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 4441–4448.

[10] A. Khamis, A. Hussein, and A. Elmogy, *Multi-robot Task Allocation: A Review of the State-of-the-Art*. Cham: Springer International Publishing, 2015, pp. 31–51.

[11] C. Schmidbauer, S. Zafari, B. Hader, and S. Schlund, "An empirical study on workers' preferences in human–robot task assignment in industrial assembly systems," *IEEE Transactions on Human-Machine Systems*, vol. 53, no. 2, pp. 293–302, 2023.

[12] A. Pupa, W. Van Dijk, C. Brekelmans, and C. Secchi, "A resilient and effective task scheduling approach for industrial human-robot collaboration," *Sensors*, vol. 22, no. 13, 2022.

[13] Y. Cheng, L. Sun, and M. Tomizuka, "Human-aware robot task planning based on a hierarchical task model," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1136–1143, 2021.

[14] E. Lamon, A. De Franco, L. Peternel, and A. Ajoudani, "A capability-aware role allocation approach to industrial assembly tasks," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3378–3385, 2019.

[15] K. Darvish, E. Simetti, F. Mastrogiovanni, and G. Casalino, "A hierarchical architecture for human–robot cooperation processes," *IEEE Transactions on Robotics*, vol. 37, no. 2, pp. 567–586, 2021.

[16] A. Noormohammadi, A. Dahiya, A. M. Aroyo, S. L. Smith, and K. Dautenhahn, "The effect of robot decision making on human perception of a robot in a collaborative task - a remote study," in *Proceedings of the 9th International Conference on Human-Agent Interaction*, 2021, p. 423–427.

[17] A. Noormohammadi-Asl, A. Ayub, S. L. Smith, and K. Dautenhahn, "Task selection and planning in human-robot collaborative processes: To be a leader or a follower?" in *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2022, pp. 1244–1251.

[18] A. Noormohammadi-Asl, S. L. Smith, and K. Dautenhahn, "To lead or to follow? adaptive robot task planning in human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01483

[19] A. Noormohammadi-Asl, A. Ayub, S. L. Smith, and K. Dautenhahn, "Adapting to human preferences to lead or follow in human-robot collaboration: A system evaluation," in *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2023, pp. 1851–1858.

[20] A. Noormohammadi-Asl, K. Fan, S. L. Smith, and K. Dautenhahn, "Human leading or following preferences: Effects on human perception of the robot and the human-robot collaboration," 2024. [Online]. Available: https://arxiv.org/abs/2401.01466

[21] J. E. Colgate and M. A. Peshkin, "Cobots," Sep. 14 1999, uS Patent 5,952,796.

[22] P. Akella, M. Peshkin, E. Colgate, W. Wannasuphoprasit, N. Nagesh, J. Wells, S. Holland, T. Pearson, and B. Peacock, "Cobots for the automobile assembly line," in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No. 99CH36288C)*, vol. 1. IEEE, 1999, pp. 728–733.

[23] "Collaborative robot (cobot) market," A report from MarketsandMarkets, Avaiable at https://www.marketsandmarkets.com/Market-Reports/collaborative-robot-market-194541294.html, Mar 2020, (Last accessed date: 30 January, 2024).

[24] "Modus V, Fully automated, robotic digital microscopy with 3D visualization and voice-activated control," Synaptive Medical Company, https://www.synaptivemedical.com/products/modus-v/, (Last accessed date: 30 January, 2024).

[25] "CARLO (Cold Ablation Robot-guided Laser Osteotome)," Advanced Osteotomy Tools Company, https://aot.swiss/en/carlo/, (Last accessed date: 30 January, 2024).

[26] M. Ureel, M. Augello, D. Holzinger, T. Wilken, B.-I. Berg, H.-F. Zeilhofer, G. Millesi, P. Juergens, and A. A. Mueller, "Cold Ablation Robot-Guided Laser Osteotome (CARLO®): From bench to bedside," *Journal of clinical medicine*, vol. 10, no. 3, p. 450, 2021.

[27] "Cafe X, Fast, advanced fully automated café system," Cafe X Company, https://cafexapp.com/, (Last accessed date: 30 January, 2024).

[28] "FLIPPY, The world's first autonomous robotic kitchen assistant," MISO Robotics Company, https://misorobotics.com/flippy/, (Last accessed date: 30 January, 2024).

[29] "ALIAS, The Aircrew Labor In-Cockpit Automation System," Aurora Flight Sciences, a Boeing company, https://www.aurora.aero/wp-content/uploads/2016/10/ALIAS-Brochure_X8.pdf, (Last accessed date: 30 January, 2024).

[30] S. Haddadin and E. Croft, "Physical human–robot interaction," in *Springer handbook of robotics*. Springer, 2016, pp. 1835–1874.

[31] A. Cesta, A. Orlandini, G. Bernardi, and A. Umbrico, "Towards a planning-based framework for symbiotic human-robot collaboration," in *2016 IEEE 21st international conference on emerging technologies and factory automation (ETFA)*. IEEE, 2016, pp. 1–8.

[32] "Demystifying collaborative industrial robots," Published by International Federation of Robotics (IFR), December 2018.

[33] E. Helms, R. D. Schraft, and M. Hagele, "rob@ work: Robot assistant in industrial environments," in *Proceedings. 11th IEEE International Workshop on Robot and Human Interactive Communication*. IEEE, 2002, pp. 399–404.

[34] W. Bauer, M. Bender, M. Braun, P. Rally, and O. Scholtz, "Lightweight robots in manual assembly—best to start simply," *Examining companies' initial experiences with lightweight robots, Stuttgart*, pp. 1–32, 2016.

[35] G. Charalambous, S. Fletcher, and P. Webb, "The development of a scale to evaluate trust in industrial human-robot collaboration," *International Journal of Social Robotics*, vol. 8, no. 2, pp. 193–209, 2016.

[36] J. Kildal, A. Tellaeche, I. Fernández, and I. Maurtua, "Potential users' key concerns and expectations for the adoption of cobots," *Procedia CIRP*, vol. 72, pp. 21–26, 2018.

[37] A. C. Simões, A. L. Soares, and A. C. Barros, "Factors influencing the intention of managers to adopt collaborative robots (cobots) in manufacturing organizations," *Journal of Engineering and Technology Management*, vol. 57, p. 101574, 2020.

[38] "Executive summary world robotics 2020 industrial robots," Published by International Federation of Robotics (IFR), September 2020.

[39] "UR10 cobots offer aging workforce solution and reduce relief worker costs for global car manufacturer," https://www.universal-robots.com/case-stories/nissan-motor-company/, (Last accessed date: 30 January, 2024).

[40] "UR10 cobots optimize the assembly line at Ford Romania," https://www.universal-robots.com/case-stories/ford-motor-company/, (Last accessed date: 30 January, 2024).

[41] "UR5 allows employees at Schneider Electric to focus on creative tasks through automating repetitive processes," https://www.universal-robots.com/case-stories/schneider-electric/, (Last accessed date: 30 January, 2024).

[42] "PT JVC Electronics Indonesia improved productivity and performance quality while reducing yearly operational costs by USD 80,000 with Universal Robots UR3," https://www.universal-robots.com/case-stories/pt-jvc-electronics-indonesia/, (Last accessed date: 30 January, 2024).

[43] "Ur cobots deliver cost reduction and stable production while freeing up workers for more important jobs," https://www.universal-robots.com/case-stories/yokota-corporation/, (Last accessed date: 30 January, 2024).

[44] "Universal Robots ensures "Quality Made in The Netherlands"," https://www.universal-robots.com/case-stories/fluidics-instruments-bv/, (Last accessed date: 30 January, 2024).

[45] "Franke increases productivity with collaborative robots while boosting employee job satisfaction." https://www.universal-robots.com/case-stories/franke/, (Last accessed date: 30 January, 2024).

[46] "42 ur robots in machine tending application improve price-competitive production," https://www.universal-robots.com/case-stories/trelleborg-sealing-solutions/, (Last accessed date: 30 January, 2024).

[47] "Glidewell laboratories optimizes its production by 9 hours with collaborative robot," https://www.universal-robots.com/case-stories/glidewell-laboratories/, (Last accessed date: 30 January, 2024).

[48] "Atria packages food with user-friendly robots," https://www.universal-robots.com/case-stories/atria/, (Last accessed date: 30 January, 2024).

[49] "The UR5 makes all employees robot experts," https://www.universal-robots.com/case-stories/nordic-sugar/, (Last accessed date: 30 January, 2024).

[50] "The specialist in modern building technology deploys 8 ur cobots working in close collaboration with employees," https://www.universal-robots.com/case-stories/albrecht-jung-gmbh/, (Last accessed date: 30 January, 2024).

[51] "Rnb cosméticos integrates collaborative robots into its workforce, boosting employee talent," https://www.universal-robots.com/case-stories/rnb-cosmeticos/, (Last accessed date: 30 January, 2024).

[52] "Cobots vs. covid-19," https://www.universal-robots.com/blog/cobots-vs-covid-19/, (Last accessed date: 30 January, 2024).

[53] A. Hentout, M. Aouache, A. Maoudj, and I. Akli, "Human–robot interaction in industrial collaborative robotics: a literature review of the decade 2008–2017," *Advanced Robotics*, vol. 33, no. 15-16, pp. 764–799, 2019.

[54] A. Ajoudani, A. M. Zanchettin, S. Ivaldi, A. Albu-Schäffer, K. Kosuge, and O. Khatib, "Progress and prospects of the human–robot collaboration," *Autonomous Robots*, vol. 42, no. 5, pp. 957–975, 2018.

[55] A. Rosenfeld, N. Agmon, O. Maksimov, and S. Kraus, "Intelligent agent supporting human–multi-robot team collaboration," *Artificial Intelligence*, vol. 252, pp. 211–231, 2017.

[56] A. Khasawneh, H. Rogers, J. Bertrand, K. C. Madathil, and A. Gramopadhye, "Human adaptation to latency in teleoperated multi-robot human-agent search and rescue teams," *Automation in Construction*, vol. 99, pp. 265–277, 2019.

[57] M. F. Mahani, L. Jiang, and Y. Wang, "A bayesian trust inference model for human-multi-robot teams," *International Journal of Social Robotics*, pp. 1–15, 2020.

[58] A. Rosenfeld, A. Noa, O. Maksimov, and S. Kraus, "Human-multi-robot team collaboration for efficent warehouse operation," *Autonomous Robots and Multirobot Systems (ARMS)*, 2016.

[59] H. Yedidsion, J. Deans, C. Sheehan, M. Chillara, J. Hart, P. Stone, and R. J. Mooney, "Optimal use of verbal instructions for multi-robot human navigation guidance," in *International Conference on Social Robotics*. Springer, 2019, pp. 133–143.

[60] M. Lippi and A. Marino, "Safety in human-multi robot collaborative scenarios: a trajectory scaling approach," *IFAC-PapersOnLine*, vol. 51, no. 22, pp. 190–196, 2018.

[61] S. Pourmehr, J. Bruce, J. Wawerla, and R. Vaughan, "A sensor fusion framework for finding an HRI partner in crowd," in *IEEE Int. Conf. on Intelligent Robots and Systems, Workshop on Designing and Evaluating Social Robots for Public Settings*, 2015.

[62] H. Claure, Y. Chen, J. Modi, M. Jung, and S. Nikolaidis, "Multi-armed bandits with fairness constraints for distributing resources to human teammates," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 2020, pp. 299–308.

[63] S. Kumar and A. Sekmen, "Single robot–multiple human interaction via intelligent user interfaces," *Knowledge-Based Systems*, vol. 21, no. 6, pp. 458–465, 2008.

[64] Y. Back, Y. Zak, Y. Parmet, and T. Oron-Gilad, "Combining cognitive work analysis and empirical evaluations to understand map use by operators of small carry-on unmanned aerial systems," *Applied Ergonomics*, vol. 90, p. 103218, 2021.

[65] R. Williams, "Collaborative multi-robot multi-human teams in search and rescue." in *Proceedings of the International ISCRAM Conference*, vol. 17, 2020.

[66] Z. Talebpour and A. Martinoli, "Adaptive risk-based replanning for human-aware multi-robot task allocation with local perception," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3790–3797, 2019.

[67] J. Patel and C. Pinciroli, "Improving human performance using mixed granularity of control in multi-human multi-robot interaction," in *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 2019, pp. 1135–1142.

[68] B. Busch, G. Maeda, Y. Mollard, M. Demangeat, and M. Lopes, "Postural optimization for an ergonomic human-robot interaction," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 2778–2785.

[69] L. Peternel, N. Tsagarakis, D. Caldwell, and A. Ajoudani, "Adaptation of robot physical behaviour to human fatigue in human-robot co-manipulation," in *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2016, pp. 489–494.

[70] T. Law and M. Scheutz, "Trust: Recent concepts and evaluations in human-robot interaction," in *Trust in Human-Robot Interaction*. Elsevier, 2021, pp. 27–57.

[71] Z. Wang, K. Mülling, M. P. Deisenroth, H. Ben Amor, D. Vogt, B. Schölkopf, and J. Peters, "Probabilistic movement modeling for intention inference in human–robot interaction," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 841–858, 2013.

[72] M. Khoramshahi and A. Billard, "A dynamical system approach to task-adaptation in physical human–robot interaction," *Autonomous Robots*, vol. 43, no. 4, pp. 927–946, 2019.

[73] S. Devin and R. Alami, "An implemented theory of mind to improve human-robot shared plans execution," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2016, pp. 319–326.

[74] J. Baraglia, M. Cakmak, Y. Nagai, R. Rao, and M. Asada, "Initiative in robot assistance during collaborative task execution," in *2016 11th ACM/IEEE international conference on human-robot interaction (HRI)*. IEEE, 2016, pp. 67–74.

[75] S. Nikolaidis, R. Ramakrishnan, K. Gu, and J. Shah, "Efficient model learning from joint-action demonstrations for human-robot collaborative tasks," in *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2015, pp. 189–196.

[76] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa, "Planning with trust for human-robot collaboration," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, p. 307–315.

[77] K. P. Hawkins, N. Vo, S. Bansal, and A. F. Bobick, "Probabilistic human action prediction and wait-sensitive planning for responsive human-robot collaboration," in *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 2013, pp. 499–506.

[78] L. Johannsmeier and S. Haddadin, "A hierarchical human-robot interaction-planning framework for task allocation in collaborative industrial assembly processes," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 41–48, 2016.

[79] K. Bogner, U. Pferschy, R. Unterberger, and H. Zeiner, "Optimised scheduling in human–robot collaboration–a use case in the assembly of printed circuit boards," *International Journal of Production Research*, vol. 56, no. 16, pp. 5522–5540, 2018.

[80] S. Nikolaidis, D. Hsu, and S. Srinivasa, "Human-robot mutual adaptation in collaborative tasks: Models and experiments," *The International Journal of Robotics Research*, vol. 36, no. 5-7, pp. 618–634, 2017.

[81] V. Gabler, T. Stahl, G. Huber, O. Oguz, and D. Wollherr, "A game-theoretic approach for adaptive action selection in close proximity human-robot-collaboration," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 2897–2903.

[82] Y. Li, K. P. Tee, W. L. Chan, R. Yan, Y. Chua, and D. K. Limbu, "Role adaptation of human and robot in collaborative tasks," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 5602–5607.

[83] M. Biswas and J. Murray, "Robots that refuse to admit losing–a case study in game playing using self-serving bias in the humanoid robot marc," in *International Conference on Intelligent Robotics and Applications*. Springer, 2016, pp. 538–548.

[84] P. Kaniarasu and A. M. Steinfeld, "Effects of blame on trust in human robot interaction," in *The 23rd IEEE international symposium on robot and human interactive communication*. IEEE, 2014, pp. 850–855.

[85] D. P. Van der Hoorn, A. Neerincx, and M. M. de Graaf, "" i think you are doing a bad job!" the effect of blame attribution by a robot in human-robot collaboration," in *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 2021, pp. 140–148.

[86] F. Correia, S. F. Mascarenhas, S. Gomes, P. Arriaga, I. Leite, R. Prada, F. S. Melo, and A. Paiva, "Exploring prosociality in human-robot teams," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 143–151.

[87] A. Chanseau, K. Dautenhahn, K. L. Koay, M. L. Walters, G. Lakatos, and M. Salem, "How does peoples' perception of control depend on the criticality of a task performed by a robot," *Paladyn, Journal of Behavioral Robotics*, vol. 10, no. 1, pp. 380–400, 2019.

[88] A. Pasparakis, J. De Vries, and M. De Koster, "In control or under control? human-robot collaboration in warehouse order picking," *Human-Robot Collaboration in Warehouse Order Picking (March 31, 2021)*, 2021.

[89] System-directed picking empower associates to empower associates to be more productive. (Last accessed date: 30 January, 2024). [Online]. Available: https://6river.com/directed-picking/

[90] How robotic picking is revolutionizing warehouse productivity. (Last accessed date: 30 January, 2024). [Online]. Available: https://www.dhl.com/global-en/delivered/digitalization/locus-robotics-robotic-picking.html

[91] C. Schmidbauer, S. Schlund, T. B. Ionescu, and B. Hader, "Adaptive task sharing in human-robot interaction in assembly," in *2020 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, 2020, pp. 546–550.

[92] R. Müller, M. Vette, and O. Mailahn, "Process-oriented task assignment for assembly processes with human-robot interaction," *Procedia CIRP*, vol. 44, pp. 210–215, 2016, 6th CIRP Conference on Assembly Technologies and Systems (CATS).

[93] P. Tsarouchi, G. Michalos, S. Makris, T. Athanasatos, K. Dimoulas, and G. Chryssolouris, "On a human–robot workplace design and task allocation system," *International Journal of Computer Integrated Manufacturing*, vol. 30, no. 12, pp. 1272–1279, 2017.

[94] G. Michalos, J. Spiliotopoulos, S. Makris, and G. Chryssolouris, "A method for planning human robot shared tasks," *CIRP Journal of Manufacturing Science and Technology*, vol. 22, pp. 76–90, 2018.

[95] M.-L. Lee, S. Behdad, X. Liang, and M. Zheng, "Task allocation and planning for product disassembly with human–robot collaboration," *Robotics and Computer-Integrated Manufacturing*, vol. 76, p. 102306, 2022.

[96] M. Faroni, A. Umbrico, M. Beschi, A. Orlandini, A. Cesta, and N. Pedrocchi, "Optimal task and motion planning and execution for multiagent systems in dynamic environments," *IEEE Transactions on Cybernetics*, pp. 1–12, 2023.

[97] S. Alirezazadeh and L. A. Alexandre, "Dynamic task scheduling for human-robot collaboration," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8699–8704, 2022.

[98] H. Nemlekar, J. Modi, S. K. Gupta, and S. Nikolaidis, "Two-stage clustering of human preferences for action prediction in assembly tasks," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 3487–3494.

[99] O. C. Görür, B. Rosman, F. Sivrikaya, and S. Albayrak, "FABRIC: A framework for the design and evaluation of collaborative robots with extended human adaptation," *J. Hum.-Robot Interact.*, vol. 12, no. 3, may 2023.

[100] M. Fiore, A. Clodic, and R. Alami, *On Planning and Task Achievement Modalities for Human-Robot Collaboration.* Springer International Publishing, 2016, pp. 293–306.

[101] A. Tausch and A. Kluge, "The best task allocation process is to decide on one's own: effects of the allocation agent in human–robot interaction on perceived work characteristics and satisfaction," *Cognition, Technology & Work*, vol. 24, no. 1, pp. 39–55, 2022.

[102] M. Gombolay, A. Bair, C. Huang, and J. Shah, "Computational design of mixed-initiative human–robot teaming that considers human factors: situational awareness, workload, and workflow preferences," *The International Journal of Robotics Research*, vol. 36, no. 5-7, pp. 597–617, 2017.

[103] B. Akgun, M. Cakmak, K. Jiang, and A. L. Thomaz, "Keyframe-based learning from demonstration: Method and evaluation," *International Journal of Social Robotics*, vol. 4, pp. 343–355, 2012.

[104] J. Huang, D. Fox, and M. Cakmak, "Synthesizing robot manipulation programs from a single observed human demonstration," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4585–4592.

[105] V. V. Unhelkar, "Effective information sharing for human-robot collaboration," Ph.D. dissertation, MIT, 2020.

[106] J. Ayoub, L. Avetisian, X. J. Yang, and F. Zhou, "Real-time trust prediction in conditionally automated driving using physiological measures," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–9, 2023.

[107] H. Soh, Y. Xie, M. Chen, and D. Hsu, "Multi-task trust transfer for human–robot interaction," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 233–249, 2020.

[108] C. Huang, W. Luo, and R. Liu, "Meta preference learning for fast user adaptation in human-supervisory multi-robot deployments," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 5851–5856.

[109] R. Luo, R. Hayne, and D. Berenson, "Unsupervised early prediction of human reaching for human–robot collaboration in shared workspaces," *Autonomous Robots*, vol. 42, pp. 631–648, 2018.

[110] S. Reddy, S. Levine, and A. Dragan, "First contact: Unsupervised human-machine co-adaptation via mutual information maximization," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35. Curran Associates, Inc., 2022, pp. 31 542–31 556.

[111] A. Kanazawa, J. Kinugawa, and K. Kosuge, "Adaptive motion planning for a collaborative robot based on prediction uncertainty to enhance human safety and work efficiency," *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 817–832, 2019.

[112] A. Bajcsy, D. P. Losey, M. K. O'Malley, and A. D. Dragan, "Learning from physical human corrections, one feature at a time," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, p. 141–149.

[113] Y.-S. Tung, K. Bishop, B. Hayes, and A. Roncone, "Bilevel optimization for just-in-time robotic kitting and delivery via adaptive task segmentation and scheduling," in *2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2022, pp. 524–531.

[114] V. Krueger, F. Rovida, B. Grossmann, R. Petrick, M. Crosby, A. Charzoule, G. Martin Garcia, S. Behnke, C. Toscano, and G. Veiga, "Testing the vertical and cyber-physical integration of cognitive robots in manufacturing," *Robotics and Computer-Integrated Manufacturing*, vol. 57, pp. 213–229, 2019.

[115] M. Wise, M. Ferguson, D. King, E. Diehr, and D. Dymesich, "Fetch and freight: Standard platforms for service robot applications," in *Workshop on autonomous mobile service robots*, 2016, pp. 1–6.

[116] T. Sanders, K. E. Oleson, D. R. Billings, J. Y. Chen, and P. A. Hancock, "A model of human-robot trust: Theoretical model development," in *Proceedings of the human factors and ergonomics society annual meeting*, vol. 55, no. 1. SAGE Publications Sage CA: Los Angeles, CA, 2011, pp. 1432–1436.

[117] Z. R. Khavas, S. R. Ahmadzadeh, and P. Robinette, "Modeling trust in human-robot interaction: A survey," in *International conference on social robotics*. Springer, 2020, pp. 529–541.

[118] J. D. Lee and N. Moray, "Trust, self-confidence, and operators' adaptation to automation," *International journal of human-computer studies*, vol. 40, no. 1, pp. 153–184, 1994.

[119] R. Wiczorek and J. Meyer, "Effects of trust, self-confidence, and feedback on the use of decision automation," *Frontiers in Psychology*, vol. 10, 2019.

[120] K. J. Williams, M. S. Yuh, and N. Jain, "A computational model of coupled human trust and self-confidence dynamics," *ACM Transactions on Human-Robot Interaction*, vol. 12, no. 3, 2023.

[121] S. Miller, "Workload measures," *National Advanced Driving Simulator. Iowa City, United States*, 2001.

[122] S. G. Hart and L. E. Staveland, "Development of NASA-TLX (task load index): Results of empirical and theoretical research," in *Advances in psychology*. Elsevier, 1988, vol. 52, pp. 139–183.

[123] R. G. Freedman, S. J. Levine, B. C. Williams, and S. Zilberstein, "Helpfulness as a key metric of human-robot collaboration," *arXiv preprint arXiv:2010.04914*, 2020.

[124] B. Shneiderman and P. Maes, "Direct manipulation vs. interface agents," *interactions*, vol. 4, no. 6, pp. 42–61, 1997.

[125] G. Dove, K. Halskov, J. Forlizzi, and J. Zimmerman, "Ux design innovation: Challenges for working with machine learning as a design material," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 2017, p. 278–288.

[126] J. Heer, "Agency plus automation: Designing artificial intelligence into interactive systems," *Proceedings of the National Academy of Sciences*, vol. 116, no. 6, pp. 1844–1850, 2019.

[127] Q. Roy, F. Zhang, and D. Vogel, "Automation accuracy is good, but high controllability may be better," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, p. 1–8.

[128] P. G. Northouse, *Introduction to leadership: Concepts and practice*. Sage, 2014.

[129] R. E. Kelley, *The power of followership: How to create leaders people want to follow, and followers who lead themselves*. Broadway Business, 1992.

[130] B. Laugwitz, T. Held, and M. Schrepp, "Construction and evaluation of a user experience questionnaire," in *HCI and Usability for Education and Work*, A. Holzinger, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 63–76.

[131] M. Schrepp, J. Thomaschewski, and A. Hinderks, "Design and evaluation of a short version of the user experience questionnaire (ueq-s)," *International Journal of Interactive Multimedia and Artificial Intelligence*, vol. 4, no. 6, pp. 103–108, 12/2017 2017.

APPENDICES

Appendix A

Leadership and Followership Styles Questionnaires

A.1 Leadership Style

5-point Likert scale questionnaire [128]:
1–Strongly Disagree, 2–Disagree, 3–Neutral, 4–Agree, 5–Strongly Agree

For each of the statements below, select the option that indicates the degree to which you agree or disagree. Give your immediate impressions. There are no right or wrong answers.

1. Employees need to be supervised closely, or they are not likely to do their work.
2. Employees want to be a part of the decision-making process.
3. In complex situations, leaders should let subordinates work problems out on their own.
4. It is fair to say that most employees in the general population are lazy.
5. Providing guidance without pressure is the key to being a good leader.
6. Leadership requires staying out of the way of subordinates as they do their work.

7. As a rule, employees must be given rewards or punishments to motivate them to achieve organizational objectives.
8. Most workers want frequent and supportive communication from their leaders.
9. As a rule, leaders should allow subordinates to appraise their own work.
10. Most employees feel insecure about their work and need direction.
11. Leaders need to help subordinates accept responsibility for completing their work.
12. Leaders should give subordinates complete freedom to solve problems on their own.
13. The leader is the chief judge of the achievements of the members of the group.
14. It is the leader's job to help subordinates find their "passion."
15. In most situations, workers prefer little input from the leader.
16. Effective leaders give orders and clarify procedures.
17. People are basically competent, and if given a task, will do a good job.
18. In general, it is best to leave subordinates alone.

A.2 Followership Style

7-point Likert scale questionnaire [129]

For each statement, think of how would respond when facing a specific "followership" situation. Select a number from 0 to 6 to indicate the extent to which the statement describes you, where 0 indicates never applies and 6 indicates virtually always applies.

1. Does your work help you fulfill some societal goal or personal dream that is important to you?
2. Are your personal work goals aligned with your organization's priority goals?
3. Are you highly committed to and energized by our work and organization, giving them your best ideas and performance?
4. Does your enthusiasm also spread to and energize your coworkers?
5. Instead of waiting for or merely accepting what your organizational leader tells you, do you personally identify which activities are most critical for achieving the organization's most important goals?
6. Do you actively develop a distinctive competence in those critical activities so that you become more valuable to your leader and your organization?
7. When starting a new task, do you promptly build a record of successes in tasks that are important to your organizational leader?
8. Can your organizational leader give you a difficult assignment without the benefit of much supervision, knowing that you will meet your deadline with highest-quality work and that you will "fill in the cracks" if need be?
9. Do you take the initiative to seek out and successfully complete assignments that go above and beyond your job?
10. When you are not the leader of a group project, do you still contribute at a high level, often doing more than your share?
11. Do you independently think up and champion new ideas that will contribute significantly to the leader's or the organization's goals?

12. Do you try to solve tough problems (technical or organizational) rather than look to the leader to do it for you?

13. Do you help out other co-workers, making them look good, even when you don't get any credit?

14. Do you help the leader or group see both the upside potential and the downside risks of ideas or plan, planning the devil's advocate if need be?

15. Do you understand the leader's needs, goals, and constraints and then work hard to help meet the leader's needs and goals and work within the leader's constraints?

16. Do you actively and honestly admit to your strengths and weaknesses rather than delay evaluation?

17. Do you make a habit of internally questioning the wisdom of the leader's decisions rather than just doing what you are told?

18. When the leader asks you to do something that runs contrary to your professional or personal preferences, do you say "no" rather than "yes"?

19. Do you act on your own ethical standards rather than the leaders or the group's standards?

20. Do you assert your views on important issues, even though it might conflict with your group or reprisals from your leader?

Appendix B

Ethics Approval

UNIVERSITY OF WATERLOO

Notification of Ethics Clearance to Conduct Research with Human Participants

---

Principal Investigator: Kerstin Dautenhahn (Electrical and Computer Engineering)

Student investigator: Ali Noormohammadi Asl (Systems Design Engineering)

Collaborator: Stephen Smith (Electrical and Computer Engineering)

Collaborator: Ali Ayub (Electrical and Computer Engineering)

Collaborator: Kevin Fan (Electrical and Computer Engineering)

File #: 44549

Title: Planning for human-robot collaboration

The Human Research Ethics Board is pleased to inform you this study has been reviewed and given ethics clearance.

Initial Approval Date: 10/05/23 (m/d/y)

University of Waterloo Research Ethics Boards are composed in accordance with, and carry out their functions and operate in a manner consistent with, the institution's guidelines for research with human participants, the Tri-Council Policy Statement for the Ethical Conduct for Research Involving Humans (TCPS, 2nd edition), International Conference on Harmonization: Good Clinical Practice (ICH-GCP), the Ontario Personal Health Information Protection Act (PHIPA), the applicable laws and regulations of the province of Ontario. Both Boards are registered with the U.S. Department of Health and Human Services under the Federal Wide Assurance, FWA00021410, and IRB registration number IRB00002419 (HREB) and IRB00007409 (CREB).

This study is to be conducted in accordance with the submitted application and the most recently approved versions of all supporting materials.

Expiry Date: 10/06/23 (m/d/y)

Multi-year research must be renewed at least once every 12 months unless a more frequent review has otherwise been specified. Studies will only be renewed if the renewal report is received and approved before the expiry date. Failure to submit renewal reports will result in the investigators being notified ethics clearance has been suspended and Research Finance being notified the ethics clearance is no longer valid.

Level of review: Delegated Review

Signed on behalf of the Human Research Ethics Board



Heather Dekker, Ethics Advisor, hdekker@uwaterloo.ca, 519-888-4567, ext. 41506

This above named study is to be conducted in accordance with the submitted application and the most recently approved versions of all supporting materials.

Documents reviewed and received ethics clearance for use in the study and/or received for information:

UNIVERSITY OF WATERLOO

Notification of Ethics Clearance to Conduct Research with Human Participants

---

Principal Investigator: Kerstin Dautenhahn (Electrical and Computer Engineering)

Student investigator: Ali Noormohammadi Asl (Systems Design Engineering)

Collaborator: Stephen Smith (Electrical and Computer Engineering)

Collaborator: Alexander Aroyo (Electrical and Computer Engineering)

File #: 42760

Title: Robot's Decision Making Effects on Human Trust and Satisfaction in a Collaborative Task

---

The Human Research Ethics Board is pleased to inform you this study has been reviewed and given ethics clearance.

Initial Approval Date: 04/08/21 (m/d/y)

University of Waterloo Research Ethics Boards are composed in accordance with, and carry out their functions and operate in a manner consistent with, the institution's guidelines for research with human participants, the Tri-Council Policy Statement for the Ethical Conduct for Research Involving Humans (TCPS2 2022), the Ontario Personal Health Information Protection Act (PHIPA), and all laws and regulations of the province of Ontario (as applicable). Additionally, CREB operates in a manner consistent with the International Conference for Harmonization of Technical Requirements for Pharmaceuticals for Human Use (ICH) Guidance E6(R2): Good Clinical Practice , the International Organization for Standardization of Good Clinical Practices (GCP) as set out by ISO 14155 - Clinical investigation of medical devices for human subjects, Part C, Division 5 of the Food and Drug Regulations, Part 4 of the Natural Health Products Regulations, Part 3 of the Medical Devices Regulations. Both Boards are registered with the U.S. Department of Health and Human Services under the Federal Wide Assurance, FWA00021410, and IRB registration number IRB00002419 (HREB) and IRB00007409 (CREB).

Expiry Date: 04/09/22 (m/d/y)

Multi-year research must be renewed at least once every 12 months unless a more frequent review has otherwise been specified. Studies will only be renewed if the renewal report is received and approved before the expiry date. Failure to submit renewal reports will result in the investigators being notified ethics clearance has been suspended and Research Finance being notified the ethics clearance is no longer valid.

Level of review: Delegated Review

Signed on behalf of the Human Research Ethics Board



Heather Root, Senior Manager, Ethics, heather.root@uwaterloo.ca, 519-888-4567, ext. 41469

This above named study is to be conducted in accordance with the submitted application and the most recently approved versions of all supporting materials.

Documents reviewed and received ethics clearance for use in the study and/or received for information:

Appendix C

Flyer



Eligibility **Criteria**

If you can distinguish the four colors of these stars, you are ready to work with Fetch.



CHARGE FETCH

Location & **Duration**

E5 - 5123, Approximately 1.5 + 0.5 hour (in-person & online)

Ethical **Considerations**

- This study has been reviewed by, and received ethics clearance through a University of Waterloo Research Ethics Board (ORE# 44549)

Participate now:
ali.asl@uwaterloo.ca
Ali Asl, PhD Student

SIRRL
Social and Intelligent Robotics Research Laboratory

UNIVERSITY OF WATERLOO

