



Delft University of Technology

Human-Machine Co-Learning Anticipating, Identifying and Sharing Emergent Collaboration Patterns

van Zoelen, E.M.

DOI

[10.4233/uuid:97d5bf1a-3051-43e4-ac18-1db930845c7c](https://doi.org/10.4233/uuid:97d5bf1a-3051-43e4-ac18-1db930845c7c)

Publication date

2025

Document Version

Final published version

Citation (APA)

van Zoelen, E. M. (2025). *Human-Machine Co-Learning: Anticipating, Identifying and Sharing Emergent Collaboration Patterns*. [Dissertation (TU Delft), Delft University of Technology].
<https://doi.org/10.4233/uuid:97d5bf1a-3051-43e4-ac18-1db930845c7c>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.


Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Human-Machine Co-Learning

Anticipating, Identifying and Sharing
Emergent Collaboration Patterns

Emma M. van Zoelen



Propositions

accompanying the dissertation

HUMAN-MACHINE CO-LEARNING

ANTICIPATING, IDENTIFYING AND SHARING EMERGENT COLLABORATION PATTERNS

by

Emma Maxime VAN ZOELEN

-  1. Human-machine co-adaptation is not a gradual change of collaborative behavior, but is characterized by the emergence of stable strategies (Collaboration Patterns) and sudden adaptations (Chapter 2, 3 and 4).
-  2. Achieving successful human-machine co-learning requires a common, grounded representation of collaborative behavior (Chapter 6).
-  3. An AI-system does not need a formal model of its human team partner to properly adapt its behavior to that partner (Chapter 3 and 4).
-  4. Humans like to control their machines as tools, whereas making machines true team partners requires humans to accept machine superiority on certain aspects (Chapter 6, (Bosch et al., 2025)).
5. The average human does not exist: in human-machine interaction design and research, we need to focus on the uniqueness and individuality of each human rather than categorizing them.
6. The lack of a common vocabulary in human-machine collaboration research makes it impossible to develop a shared research paradigm.
7. AI-powered team partners should not be designed to mimic human behavior, but should instead be designed to reflect their own capabilities.
8. Being social does not require facial expressions or verbal communication.
9. Publishing the lessons learned in the process of designing research environments is a valid scientific contribution and should be a requirement of any HCI, HRI and HMC study, to prevent redundancy in scientific work.
10. The abundance of utopian and dystopian AI futures in the public discourse hinders the development of human-centered and responsible AI.

These propositions are regarded as opposable and defensible, and have been approved as such by the promotor prof. dr. M. A. Neerincx, the copromotor prof. dr. D. A. Abbink and external adviser dr. K. van den Bosch.

 Pertains to this dissertation.

HUMAN-MACHINE CO-LEARNING

ANTICIPATING, IDENTIFYING AND SHARING EMERGENT
COLLABORATION PATTERNS

EDITED BY VON ZÜRICH

HUMAN-MACHINE CO-LEARNING

ANTICIPATING, IDENTIFYING AND SHARING EMERGENT
COLLABORATION PATTERNS

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology,
by the authority of Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates,
to be defended publicly on
Monday 24 November 2025 at 12:30 o'clock

by

Emma Maxime VAN ZOELLEN

Master of Science in Artificial Intelligence,
Utrecht University, the Netherlands,
Master of Science in Industrial Design,
Eindhoven University of Technology, the Netherlands,
born in Wester-Koggenland, The Netherlands.

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus,	Chairperson
Prof. dr. M.A. Neerincx,	Delft University of Technology, <i>promotor</i>
Prof. dr. ir. D.A. Abbink	Delft University of technology, <i>promotor</i>
Dr. K. van den Bosch,	TNO, <i>external adviser</i>

Independent members:

Prof. dr. ir. J.C.F. de Winter	Delft University of Technology
Prof. dr. A. ten Teije	VU Amsterdam
Prof.dr. F. Flemisch	RWTH Aachen University, Germany
Prof. dr. ir. A. Bozzon	Delft University of Technology
Prof. dr. ir. W.P. Brinkman	Delft University of Technology, <i>reserve member</i>

This thesis was a collaboration between the Netherlands Organisation of Applied Scientific Research (TNO) and Delft University of Technology, funded by TNO under contract “4099 DO AIO Fonds”.



Keywords: Human-Agent Teaming, Co-Learning, Co-Adaptation, Human-Robot Collaboration, Emergent Interactions, Collaboration Patterns
Printed by: Proefschriftspecialist | <https://www.proefschriftspecialist.nl/>
Cover by: Emma Maxime van Zoelen

Copyright © 2025 by E.M. van Zoelen

ISBN 978-94-6518-167-7

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

CONTENTS

Summary	xi
Samenvatting	xiii
1 Introduction	1
1.1 Mutual Adaptation and Learning in Human-Machine Collaboration: A Research Challenge	3
1.2 Background	4
1.2.1 Adaptation and Learning in Human Collaborators	4
1.2.2 Human-Machine Teams	5
1.3 Co-Learning Defined	7
1.4 Research Approach/Framework	8
1.5 Research Questions and Contributions	9
1.5.1 Human Adaptation	9
1.5.2 Human and Machine Co-Adaptation with a Virtual Robot	10
1.5.3 Human and Machine Co-Adaptation with a Physical Robot	10
1.5.4 Knowledge Structures for Shared Representations of Emergent Collaboration Patterns.	11
1.5.5 Human-Machine Co-Learning Facilitated by Shared Representations of Emergent Collaboration Patterns	12
1.6 Co-Learning: Contributions and Study Designs	13
2 Identifying Interaction Patterns of Tangible Co-Adaptations	21
2.1 Introduction	22
2.2 Co-Adaptation in Human-Robot Teams	23
2.2.1 Co-Adaptation - A Definition.	23
2.2.2 Related Work	24
2.3 Research Challenge	27
2.4 Materials and Methods	28
2.4.1 Task Environment: Search-and-Navigation	28
2.4.2 Design of Human-Robot Interaction	28
2.4.3 Experiment Setup and Initial Results.	30
2.4.4 Analyzing Behavior to Uncover Interaction Patterns	36
2.5 Results: Interaction Patterns	37
2.5.1 Observed Interactions	37
2.5.2 Interaction Patterns for Adaptive Leader-Follower Behavior	38

2.6	Conclusion and Discussion	41
2.6.1	Interaction Patterns and Team Design Patterns	43
2.6.2	Interaction Pattern Language	44
2.6.3	Relation to Existing Interaction Taxonomies	45
2.6.4	Limitations.	45
2.6.5	Final Conclusion.	46
3	Becoming Team Members	51
3.1	Introduction	52
3.2	Co-Learning: Background and Definition	53
3.3	Research Challenges	55
3.4	Research Environment: Designing Task, Agent, Context	56
3.4.1	Context	56
3.4.2	Task Implementation	57
3.4.3	Learning Agent.	58
3.4.4	Claims: Expected Observations	60
3.5	Method of Study for Identifying Interaction Patterns	61
3.5.1	Participants	62
3.5.2	Design and Materials	63
3.5.3	Procedure	63
3.5.4	Data Collection and Analysis.	64
3.6	Results	65
3.6.1	Subjective Collaboration Fluency and Confidence Score	65
3.6.2	Interaction Pattern Analysis	66
3.6.3	Collaborative Learning.	69
3.7	Discussion	74
3.7.1	Interaction Patterns that Drive Co-Learning	74
3.7.2	Validating the Research Environment	74
3.7.3	Future Steps for Studying Human-Robot Co-Learning	75
3.8	Conclusion	77
4	Enabling Embodied Human-Robot Co-Learning	83
4.1	Introduction	84
4.2	Design Requirements	84
4.2.1	Shared Goal	85
4.2.2	Synchrony	85
4.2.3	Interdependence.	86
4.2.4	Adaptability	86
4.2.5	Transparency	87
4.3	Methods	87
4.3.1	Task Design	87
4.3.2	Robot Reinforcement Learning Algorithm	90
4.3.3	Evaluation	92

4.4	Results	93
4.4.1	Shared Goal (R1)	94
4.4.2	Synchrony (R2)	94
4.4.3	Interdependence (R3)	96
4.4.4	Adaptability (R4)	97
4.4.5	Transparency (R5)	97
4.5	Discussion	97
4.5.1	Facilitating Co-learning	97
4.5.2	Limitations and Future work	98
5	Ontology-Based Reflective Communication	105
5.1	Introduction	106
5.2	Ontology: Requirements and Background	106
5.2.1	Requirements for Collaboration Pattern Ontology	106
5.2.2	Ontologies in Human-Agent Teaming	108
5.2.3	Frameworks for Describing Patterns of Collaboration	108
5.3	Collaboration Pattern Ontology	109
5.4	Ontology in Practice: Translation to Context	109
5.4.1	Evaluation Method	109
5.4.2	Results: Correctness of Collaboration Patterns Recognized	111
5.4.3	Results: Patterns of Descriptions	111
5.5	Discussion	112
5.6	Conclusion	113
6	Shared Representations of Emergent Collaboration Patterns	117
6.1	Introduction	118
6.1.1	Research Questions and Hypotheses	119
6.2	Design – Task Environment, Collaboration Book, Agent Behavior	121
6.2.1	Task Environment	121
6.2.2	Collaboration Book	122
6.2.3	AI Agent Design	124
6.3	Experimental Study	126
6.3.1	Design	126
6.3.2	Participants	126
6.3.3	Dependent Variables	127
6.3.4	Procedure	127
6.3.5	Data Analysis	129
6.4	Results	131
6.4.1	Awareness of Collaboration Patterns (RQ1)	131
6.4.2	Reuse of CPs (RQ2)	131
6.4.3	Effects on Team Performance (RQ3)	133
6.4.4	Discovered Collaboration Patterns (RQ4)	141
6.4.5	Locus of Control (RQ5)	143

6.5	Discussion	143
6.5.1	RQ1 Awareness of Collaboration Patterns	144
6.5.2	RQ2 Reusing CPs.	145
6.5.3	RQ3 Influence of Collaboration Book on Team Performance.	145
6.5.4	RQ4 Discovered Collaboration Patterns	146
6.5.5	RQ5 Locus of Control	147
6.5.6	Reflection on Human-Machine Co-Learning Design.	147
6.6	Conclusion	148
7	Conclusion	153
7.1	Research Questions Addressed	153
7.2	Main conclusions	155
7.2.1	Co-Adaptation is characterized by emergent stable strategies and sudden adaptations	155
7.2.2	Collaboration Patterns emerge as a result of co-adaptation	156
7.2.3	Humans are able to recognize and formalize emergent CPs	157
7.2.4	Humans don't easily accept machines as team partners	158
7.2.5	Defining co-learning for the study of complex human-machine collaboration dynamics	158
7.3	Limitations	159
7.4	Recommendations for studying Co-Learning	160
7.4.1	Task Design	161
7.4.2	Algorithm Design	161
7.4.3	Assessing Co-Learning.	162
7.5	Future Work.	163
7.5.1	Task Environments and Real-Life Settings	163
7.5.2	Modality of Interactions on Collaboration Patterns	164
7.5.3	Algorithms for Machines Learning in Synchrony with Humans	164
7.5.4	Machine Initiative: AI Recognition of Emergent CPs	165
7.6	Final Remarks.	165
A	List of observed behaviors in Chapter 2	171
B	Flowcharts of Agent Behavior in Chapter 3	175
C	Questions asked during the experiment in Chapter 3	179
C.1	Collaboration Fluency Questionnaire	179
C.2	Interview Questions.	179
D	List of observed behaviors and interaction patterns in Chapter 3	181
E	Coding of emergent CPs in Chapter 6	187
F	Description of awareness of CPs in Chapter 6	191
G	Description of reuse of CPs in Chapter 6	197
H	CP clusters in Chapter 6	199
I	Results of Linear Mixed Models in Chapter 6	201

J	Discovered CPs in Chapter 6	213
K	Questionnaire used in Chapter 6	215
	Acknowledgements	217
	Curriculum Vitæ	221
	List of Publications	223
	List of Relevant Datasets	225

SUMMARY

Intelligent machines (in the form of physically embodied robots or virtual agents) are increasingly able to perform tasks in collaboration with humans. However, learning to become a good team takes time, especially when dynamic tasks require the team to constantly adapt to new situations. Over time, both human and machine need to not only learn how to execute the task, but also how their team partner behaves in the task, as well as how to improve their collaboration over time by attuning their behavior to each other. Existing research on human-machine collaboration often does not sufficiently address adaptation and learning. Work that does study adaptation and learning tends to focus on either machine learning and adaptation or human learning and adaptation, thereby not addressing the interaction of these learning processes that would be present in a co-learning situation.

The aim of this thesis is to understand how collaborative behaviors develop when an adaptive human and an adaptive machine (either physical or virtual) co-adapt in an interdependent task, and to find out how we can facilitate that this process yields fluent and successful collaboration. More specifically, we investigate whether providing human and machine with the means to develop and maintain a shared mental model of successful emergent collaborative behaviors supports the team in learning to collaboratively perform a dynamic task (co-learn).

We studied the development of collaborative behaviors as a result of co-adaptation in a setting with a Wizard-of-Oz physically embodied robot (**Chapter 2**), in a setting with a virtual Reinforcement Learning (RL) powered virtual agent (**Chapter 3**), and in a setting with a physically embodied RL-powered robot (**Chapter 4**). We found that co-adaptation usually does not happen gradually, but rather that different strategies (Collaboration Patterns) emerge during interaction. The team usually engages in each of these Collaboration Patterns for a while. Events in the task environment can trigger the team to suddenly shift to another Collaboration Pattern. In the studies with RL-powered machines (**Chapter 3 and 4**), we observed that these emergent Collaboration Patterns were followed by both the human and the machine; the humans would initiate actions following the Collaboration Patterns, while we could see it represented in the machine's learned policy. The human and the machine were however not always able to successfully use a Collaboration Pattern for a prolonged period of time. Therefore, teams were sometimes unable to achieve good performance.

In **Chapter 2, 3 and 4**, adaptation and learning happened in an implicit manner, and there was no communication about the emergent CPs, which is possibly one of the reasons why the human and/or machine would often deviate from a CP. Therefore, we developed an ontology model of Collaboration Patterns that could serve as a shared mental model in which CPs could be formalized. In addition, we created a user interface to provide a way for the human to share formalized CPs with the machine (**Chapter 5**). This way, once humans recognized an emerging Collaboration Pattern as fit for reuse by the

team, they could use the user interface to formalize the pattern, thereby ensuring that the machine could execute it when applicable. In a human-subjects study, we evaluated whether the ontology model and user interface improved human-machine team collaboration and performance (**Chapter 6**), to ultimately lead to co-learning. Results showed that teams using it were able to improve their performance over time to a larger extent than teams without the ontology model. This effect was, however, not maintained for more complex task situations. Participants rated the quality of the collaboration much more positively when using the ontology and interface, but it motivated humans to take control by directing the machine around, thereby breaking the co-adaptive dynamic that existed in teams without the ontology model and user interface.

In conclusion, adaptive humans and adaptive machines that need to collaborate are able to develop emergent Collaboration Patterns, and humans are able to identify and formalize these CPs when provided with the tools to do so. To ensure that this leads to good team performance, they need to be able to communicate about the CPs to their machine team partner, although there is a risk that they will use it to direct their machine team partner around. These insights ask for a way of studying human-machine collaboration that is more process-centered, to ensure that learning algorithms are able to follow the ways in which humans naturally engage in adaptation and learning with machines, while machines should also be sufficiently proactive and transparent about their contribution to the collaboration. Our designed ontology and user interface provide a possible way to achieve communication about emergent Collaboration Patterns, which can be used and extended by designers of human-machine collaboration to support shared mental models between humans and machines.

SAMENVATTING

Intelligente machines (fysieke robots of virtuele agenten) voeren steeds vaker taken uit in samenwerking met mensen. Leren om een goed team te worden kost tijd, zeker als een dynamische taak ervoor zorgt dat een mens-machine team zich constant moet aanpassen aan nieuwe situaties. Zowel mens als machine moet niet alleen leren hoe de taak uitgevoerd moet worden, maar ook hoe hun teampartner de taak uitvoert, en hoe ze hun samenwerking kunnen verbeteren door hun gedrag steeds meer af te stemmen. Onderzoek naar mens-machine samenwerking besteedt meestal niet voldoende aandacht aan adaptatie en leren. Studies die wel adaptatie en leren onderzoeken leggen meestal de nadruk op alleen de machine of alleen de mens, waardoor ze geen aandacht besteden aan de interactie tussen de leerprocessen.

Het doel van dit proefschrift is om te begrijpen hoe samenwerkingsgedragingen zich ontwikkelen wanneer een adaptieve mens en een adaptieve machine (fysiek of virtueel) samen leren (co-adapt) terwijl ze samenwerken aan een taak waarbinnen ze onderling afhankelijk zijn. Daarnaast onderzoekt het proefschrift hoe gefaciliteerd kan worden dat dit proces van samen leren leidt tot succesvolle en vloeiende samenwerking. We onderzoeken specifiek of het aanbieden van een manier om een gedeeld mentaal model van succesvol emergent samenwerkingsgedrag op te bouwen en te onderhouden het team ondersteunt in het samen leren (co-learn) om een dynamische taak uit te voeren.

We hebben de ontwikkeling van samenwerkingsgedrag als resultaat van samen leren (co-adaptation) bestudeerd in een situatie met een fysieke Wizard-of-Oz robot (**Hoofdstuk 2**), in een situatie met een virtuele Reinforcement Learning (RL)-gestuurde virtuele agent (**Hoofdstuk 3**), en in een situatie met een fysieke RL-gestuurde robot (**Hoofdstuk 4**). We ondervonden dat samen leren (co-adaptation) meestal niet gradueel verloopt, maar dat verschillende strategieën (Collaboration Patterns of patronen van samenwerking) ontstaan door interactie tussen mens en machine. Het team voert zo'n samenwerkingspatroon meestal gedurende een bepaalde tijd uit. Gebeurtenissen in de taakomgeving kunnen ervoor zorgen dat het team plotseling wisselt naar een ander samenwerkingspatroon. In de studies met RL-gestuurde machines (**Hoofdstuk 3 en 4**), observeerden we dat deze emergente samenwerkingspatronen gevolgd werden door zowel de mens als de machine; de mensen initieerden acties die pasten binnen de samenwerkingspatronen, en we konden zien dat de patronen gerepresenteerd werden in het geleerde gedrag van de machine. Het lukte de mens en machine echter niet altijd om een samenwerkingspatroon een langere periode vast te houden en te gebruiken. Daardoor lukte het teams soms niet om een goede score te behalen voor de gezamenlijke taak.



In **Hoofdstuk 2, 3 en 4** verliep adaptatie en leren impliciet, en was er geen communicatie over emergente samenwerkingspatronen. Dit is mogelijk één van de redenen waardoor de mens en/of machine vaak afweken van een samenwerkingspatroon. We hebben daarom een ontologiemodel van samenwerkingspatronen ontwikkeld, wat kan functioneren als gedeeld mentaal model waarin samenwerkingspatronen kunnen worden ge-

formaliseerd. Daarnaast hebben we een gebruikersinterface ontworpen waarmee de mens geformaliseerde samenwerkingspatronen kan delen met de machine (**Hoofdstuk 5**). Mensen kunnen via de gebruikersinterface samenwerkingspatronen, die zij geschikt vinden voor hergebruik, formaliseren om ervoor te zorgen dat de machine deze ook uit kan voeren wanneer het samenwerkingspatroon toepasbaar is. We hebben onderzocht of het ontologiemodel en de gebruikersinterface de samenwerking en prestatie van een mens-machine team verbeterden (om uiteindelijk tot co-learning te komen) in een proefpersonenonderzoek (**Hoofdstuk 6**). De resultaten lieten zien dat teams die gebruik maakten van het ontologiemodel en de gebruikersinterface in staat waren om hun teamprestatie over tijd sterker te verbeteren dan teams die het ontologiemodel niet konden gebruiken. Dit effect hield echter geen stand in complexere taaksituaties. Proefpersonen beoordeelden de kwaliteit van de samenwerking veel positiever wanneer ze het ontologiemodel en de gebruikersinterface konden gebruiken, maar het zorgde er ook voor dat mensen meer controle uitoefenden op de machine door specifieke taken voor de machine te formuleren. De wederzijdse adaptieve dynamiek die wel bestond bij de teams die geen gebruik konden maken van het ontologiemodel en de gebruikersinterface werd hierdoor doorbroken.

We concluderen dat adaptieve mensen en adaptieve machines die moeten samenwerken in staat zijn om emergente samenwerkingspatronen te ontwikkelen, en dat mensen in staat zijn om deze samenwerkingspatronen te identificeren en formaliseren wanneer zij de juiste hulpmiddelen hiervoor aangereikt krijgen. Om ervoor te zorgen dat dit leidt tot een betere teamprestatie moeten ze de mogelijkheid hebben om over de samenwerkingspatronen te communiceren met hun team partner (machine), hoewel er een risico is dat ze het zullen gebruiken om volledig te bepalen wat de machine moet doen. Deze inzichten vragen om een manier van bestuderen van mens-machine samenwerking die meer nadruk legt op het proces van samenwerking, om ervoor te zorgen dat leeralgoritmes in staat worden gesteld om de manieren waarop mensen van nature adapteren en leren met machines te volgen. Tegelijkertijd moeten machines voldoende proactief en transparant zijn over hun bijdrage aan de samenwerking. Ons ontologiemodel en gebruikersinterface bieden een mogelijke manier om communicatie over emergente samenwerkingspatronen teweeg te brengen. Deze kunnen gebruikt en uitgebreid worden door ontwerpers van mens-machine samenwerking om gedeelde mentale modellen tussen mens en machine te ondersteunen.

1

INTRODUCTION



Imagine a situation in which you need to collaborate with people you have not met before. For example, there has been an earthquake in a mountainous area, and you are part of a foreign rescue team that needs to collaborate with local rescuers in finding and saving people trapped underneath collapsed buildings. The local rescuers have a different cultural background, and therefore have different manners and customs than what you are used to. Communication might be limited to gestures and basic words. Although collaborating with local rescuers will not be an easy endeavor, it will not be impossible. An implicit shared understanding of what actions need to be done serves as a frame of reference to try to understand each other. Over time, possibly quite quickly, you will start experimenting with adapting your behavior, while the others adapt as well. As a result, you will develop ways of working together: sequences of joint or distributed actions that work well enough to be successful. Along the way you might sometimes get frustrated and feel tempted to give up and disengage. There will be hiccups and misunderstandings, but if you persevere, it is quite likely that you will be able to compensate for each other's limitations by adapting your behavior and making it work. Along the way, you and your partners will learn about each other in different ways; you will learn when (and when not to) rely on each other's actions and thus when to delegate or take initiative. Due to this collaborative learning process, collaboration can become more fluent over time. Especially if you and your partners manage to develop ways to communicate about the planning, execution, and evaluation of collaborative actions, whether through gestures, new words, or other means, this could very well be the start of a fruitful longer-term collaboration.

Now imagine that the partner you need to collaborate with is not human. You may

think of a guide dog that helps you navigate through busy traffic (Lagerstedt and Thill, 2020), or a horse you are riding, avoiding trees while speeding through a forest (Flemisch et al., 2003). To some extent, the process of learning to collaborate with an animal will be similar to collaborating with a human: both the animal as well as you will have their own understanding of what needs to be done. This understanding needs to be aligned over time by reciprocally adapting actions. Learning to communicate about these actions is more challenging, because you will not know what the animal can and cannot understand when it comes to gestures and words, and because your understanding of what the other knows about the task and about you will be more limited compared to a human partner. Throughout history, humans have collaborated with animals in many different ways. These collaborations are successful because the humans adapted their ways of working to the needs of the animals, while the animals adapted their behavior to the wishes of the humans, to develop new practices and solutions for tasks (Darling, 2021). But what if the partner is not an animal, but an intelligent machine, like a robot or an algorithmic system? As intelligent machines enter our lives at work and in private and public spaces, they will interact with humans more intensely, making humans and machines collaborators (Ajoudani et al., 2018; Akata et al., 2020). While inspiration can be taken from human-human and human-animal collaboration, collaborating with a machine could be a fundamentally different form of collaboration.

1.1. MUTUAL ADAPTATION AND LEARNING IN HUMAN-MACHINE COLLABORATION: A RESEARCH CHALLENGE

ACHIEVING successful collaboration between a human and a machine is an important research challenge in the fields of artificial intelligence, human-computer interaction and human-robot interaction (Ajoudani et al., 2018; Akata et al., 2020; Klein et al., 2004; Mackay, 2023). Adaptation and learning are key concepts to be considered in this challenge. Becoming a good team takes time; especially if dynamic tasks require the team to constantly adapt to new situations (Burke et al., 2006). In such situations, human and machine should adapt to and learn from each other in the context of their task, and develop ways to orchestrate and synchronize their actions to collaborate successfully (van den Bosch et al., 2019).

When a learning and adaptive human collaborates with a learning and adaptive machine, this can lead to the emergence of new practices and new ways of collaborating, inspired by how this happens among human collaborators. For example, they can develop different roles (e.g. leader-follower behavior, or completely new machine-specific roles) (Mörzl et al., 2012; Zoelen et al., 2020) and different ways of anticipating and taking initiative (Jiang and Odom, 2018). Mutual adaptation can support a collaborating human and machine to develop a high level of implicit knowledge about each other, by providing them with the opportunity to gather knowledge about how the other behaves in different situations. With implicit knowledge I mean that they can anticipate and respond to the other's actions (van den Bosch et al., 2019), without having explicit awareness of the reasons for such anticipatory or responsive actions. To ensure that team members learn sustainably, it is then necessary that they become explicitly aware of the learned knowledge such that they can build up a shared mental model of their collaboration (Andrews

et al., 2023; Uitdewilligen et al., 2013).

There are many machine-focused studies on how machines adapt and learn when collaborating with a human, mostly evaluating different kinds of Machine Learning (ML) algorithms (e.g. Buschmeier and Kopp, 2013; Ehrlich and Cheng, 2018; Gao et al., 2017). There are also human-focused studies into how humans adapt and learn when collaborating with a machine (e.g. Kumar, 2024; Mohammad and Toyoaki Nishida, 2008). It is however vital to also study how humans and machine learn together, simultaneously. Ensuring that both human and machine can successfully adapt and learn concurrently, presents an important scientific challenge (Ansari et al., 2018; Holstein et al., 2020; Pepe and Hutchison, 2022; van den Bosch et al., 2019; Wenskovich and North, 2020; Xu et al., 2012).

This thesis presents research on the process of mutual adaptation and learning between a human and a machine, for which I use the term *Human-Machine Co-Learning*. Co-learning is the process that facilitates human and machines to adapt their behavior to the situation, that enables the emergence of effective Collaboration Patterns (CPs), and that facilitates that these CPs can be consolidated. The aim of this thesis is twofold:

1. To better understand the characteristics and behavioral dynamics of co-learning between a collaborating human and machine;
2. To investigate how to facilitate that co-learning leads to fluent and successful collaboration. More specifically, to study how human and machine can be provided with the means to develop shared representations of (and thereby consolidate) their emergent patterns of collaboration while performing the task .

The following sections discuss literature on Human-Machine Teams and Hybrid Intelligence, specifically work related to adaptation and learning (human learning, machine learning, and co-learning). Moreover, I will provide insights from the literature on adaptation and learning in human-human collaborators, and relate this to human-machine co-learning. After that, I will explain the research approach and address the relevant research questions and contributions of this thesis.

1.2. BACKGROUND

1.2.1. ADAPTATION AND LEARNING IN HUMAN COLLABORATORS

ANY real-life task is bound to exist in a dynamic environment, as the world is ever changing. Therefore, humans who collaborate as team partners on real-life tasks will need to be able to adapt to changes, whether those appear in the task itself or in their team members. In organizational psychology, this process is described by the term ‘Team Adaptation’ (Burke et al., 2006). While exact definitions of Team Adaptation vary (Maynard et al., 2015), it is generally defined as a process within which a team changes team structures and actions as a result of certain cues or triggers, that ultimately leads to changes in team performance. Team adaptation is widely considered to have positive effects on team performance, although unsuccessful adaptation can have negative effects on team outcome (Maynard et al., 2015). An important aspect in ensuring sustained success of team adaptation is the creation of shared mental models and joint mental model updating (Uitdewilligen et al., 2013). Joint mental model updating is the process

through which team partners update and align their mental representations of the task and their collaborative structures (Andrews et al., 2023; Uitdewilligen et al., 2013). The process of updating mental models is sometimes distinguished from Team Adaptation using the term ‘Team Reflection’ (Wiedow and Konradt, 2011) or ‘Team Learning’, indicating a difference between the more implicit and ad hoc adaptation and the more explicit and cognitive reflection or learning. Literature on Team Learning emphasizes that it contains activities such as sharing, co-construction and constructive conflict (Decuyper et al., 2010); all activities within which team members explicitly reflect on and discuss their team knowledge and behavior that may have been implicit before. It can therefore be concluded that while implicit behavioral adaptations are at the basis of how human-only teams learn, reflecting and communicating about their learned behaviors is key to achieving successful learning.

In this thesis, I explore how implicit adaptation as well as explicit communication about learned behaviors develop in human-machine collaboration. While this is a relatively natural process for collaborating humans (although not always easy), it is important to study mutual adaptation and learning for collaborating humans and machines. Humans and machines can have different ways of learning and adapting, which will influence how they adapt and learn together. For example, machines usually learn by finding patterns in large amounts of data, whereas humans often learn from example or explanation in only a few trials. Moreover, humans and machines often lack a basic shared level of common ground or frame of reference, a necessity for building shared mental models. I study how the process of mutual adaptation leads to the emergence of new patterns of collaborative behavior, which I call Collaboration Patterns. In addition, I study how emergent Collaboration Patterns can be made explicit and communicated among the collaborators to create a shared mental model, and whether this affects the collaboration and team performance.

1.2.2. HUMAN-MACHINE TEAMS

IN recent decades, research in the interaction between humans and machines has increasingly shifted towards using a hybrid intelligence approach, in which humans and AI systems complement and empower each other (Akata et al., 2020). Due to increases in AI capabilities, intelligent machines perform more and more tasks alongside humans in various domains, such as manufacturing (Kolbeinsson et al., 2019; Wang et al., 2020), healthcare (Neerincx et al., 2019; Weerathna et al., 2023), and search-and-rescue (Kruijff et al., 2012; Nourbakhsh et al., 2005). This has resulted in new types of teams, in which adaptive humans and adaptive machines operate together in dynamic environments. Many new research challenges result from these new team situations (Klein et al., 2004), which require a new paradigm for how human-machine interaction research is being done; a collaborative paradigm. Different research communities have used different terms to describe this, even though they study similar concepts such as trust, interdependence, shared autonomy, transparency. The main conferences of the different communities (e.g. AAMAS or IVA for virtual agents, IROS, HRI or RO-MAN for robotic systems) show little cross-fertilization. Some of the terms used are for example:

1. Human-agent teaming (multi-agent systems community, studying software agents, e.g. Bradshaw et al., 2003; Diggelen et al., 2019; S. Li et al., 2016);

2. Human-robot collaboration (human-robot interaction community, studying embodied robots, sometimes controlled by AI algorithms, e.g. Ajoudani et al., 2018; Nikolaidis, Hsu, and Srinivasa, 2017; Ogenyi et al., 2019).

The researchers within these communities use the terms ‘teaming’ and ‘collaboration’ to distinguish machines that function as team mates from machines that are used as tools (Krüger et al., 2017), to do justice to the increasing capabilities of machines, mainly increased initiative, autonomy and adaptivity. To study how humans and machines can function as team members, researchers can design tasks and teams based upon principles of co-active design and the interdependence framework (Johnson et al., 2014). These principles support the design of tasks and teams that contain both *hard dependencies* (ways in which human and machine need each other to complete a task) as well as *soft dependencies* (ways in which human and machine can support each other in completing a task). Regarding machines as team partners or collaborators is sometimes considered controversial, as some argue that it is an unnecessary anthropomorphization that leads to humans being less in control or evading responsibility (Evans et al., 2023). However, it has also opened the door to studying themes like shared autonomy (Nikolaidis, Zhu, et al., 2017), role shifting (Mörtl et al., 2012), and trust calibration (de Visser et al., 2020); reciprocal social dynamics that people naturally engage in when interacting with machines (Nass and Moon, 2000).

In this thesis, I use a collaborative perspective to study situations in which a machine has complementary skills or knowledge to the human, and might therefore sometimes be better suited to perform a (sub) task, but in which it will not always be clear to the human team partner when this is the case. Such situations require humans to critically reflect on their own abilities and those of the machine, to ensure that they do not overuse or underuse the capabilities of the machine. As machines might not always be aware of their own strengths and weaknesses either, I believe that it is valuable to create an interaction dynamic in which human and machine reciprocally explore how to best act and collaborate. The aim is to increase the human’s knowledge of the machine through this process, as well as to increase the machine’s knowledge of the human partner.

ADAPTATION AND LEARNING IN HUMAN-MACHINE COLLABORATION

Within the literature on human-machine collaboration, adaptation and learning has been a growing topic of study (Semeraro et al., 2023). There are several aspects that distinguish the work presented in this thesis from the existing literature. First of all, most existing work focuses on AI or robot adaptation. In such research, AI or robotics algorithms are designed to adapt to a human collaborator, usually by pre-learning preferences using a large amount of data previously obtained from people (e.g. Buschmeier and Kopp, 2013; Ehrlich and Cheng, 2018; Gao et al., 2017; Sordani et al., 2015). Sometimes the AI system or robot does learn in real-time (e.g. in Sasagawa et al., 2020), but generally human learning or adaptation is not taken into account; in a broad literature research on (real-time) Machine Learning in HRC, all papers investigated the learning of an optimal solution or policy while seemingly assuming that the human’s strategy is static (Semeraro et al., 2023). Most of this work focuses on adaptivity, meaning that it is studied how the AI system or robot adapts by itself, while there is some work in which the authority to control when the AI system or robot adapts can be with the human (Flemisch et al.,

2012), also called adaptability (Calhoun, 2021). In studies that investigate learning or adaptation by the human, usually little attention is paid to the machine adaptation (e.g. Kumar, 2024). Secondly, existing work that does explicitly take a reciprocal approach usually studies adaptation or learning outcomes in the form of performance (e.g. Nikolaidis, Hsu, and Srinivasa, 2017; Shafti et al., 2020). The process underlying performance is often not empirically addressed. It has been argued that studying the effectiveness of the interaction over time is vital for creating successful human-machine partnerships (Mackay, 2023; Wiltshire et al., 2024).

In this thesis, I present a process-centered approach. I study both human and machine adaptation and learning at the same time, investigating how the individual adaptive processes influence each other and how they lead to adaptation and learning at the team level. I focus on adaptivity rather than adaptability, although later chapters provide the human with more control over the behavior of the machine, thereby also touching upon adaptability. A more in-depth description of the research framework is described in Section 1.4.

1.3. CO-LEARNING DEFINED

IN this thesis, Human-Machine Co-Learning is defined as the process in which the human and machine concurrently adapt their behavior while collaborating, and in which successful adaptations can ultimately be consolidated, shared and sustained. This can lead to improved team performance over time. I distinguish two components of human-machine co-learning:

1. Co-Adaptation: implicit behavioral adaptations as a consequence of changes in the task or collaborating partners, leading to the emergence of Collaboration Patterns;
2. Shared Mental Model Updating: the development of shared representations of the emergent Collaboration Patterns by communicating about those patterns.

A visualization of Human-Machine Co-Adaptation (component 1) and Co-Learning (component 1 and 2 combined) can be seen in Figure 1.1. The figure shows how individual and joint adaptive processes lead to the emergence of new Collaboration Patterns. Once at least one of the collaborating partners recognizes a Collaboration Pattern, they can inform the partner about this by defining a formalization of the pattern and storing it in the shared mental model of the team. In future task situations, the collaborators can then use these patterns to align actions with their partner. This will help sustain successful behavior over time and across context. Within human-machine collaboration literature, visions for dynamic behavioral adaptation in human-machine collaboration have been published before (e.g. Flemisch et al., 2012), while attention for co-learning and related processes (co-adaptation, co-evolution) has specifically risen in the past years (Döppner et al., 2019; Nikolaidis, Hsu, and Srinivasa, 2017; Pepe and Hutchison, 2022; van den Bosch et al., 2019). A more extensive definition and positioning of the term against relevant literature is provided in Chapter 3.




Figure 1.1: A visualization of human-machine co-learning, consisting of co-adaptation (component 1) and shared mental model updating (component 2).

1.4. RESEARCH APPROACH/FRAMWORK

THE work in this thesis takes a *process-centered* approach to study emergent Collaboration Patterns, which means that I study the co-learning process over time to better understand how it leads to collaboration and performance outcomes. When assessing the *process*, I use the concept of Collaboration Fluency (Hoffman, 2019) to assess the quality of the collaboration. In addition, I use qualitative methods to analyze the observed behavior of human-machine collaborators. When I do assess *outcomes*, I use a mixed-methods approach consisting of concrete performance metrics such as task completion time, and of analyses of responses to open questions (e.g. addressing whether participants are aware of emergent collaboration patterns).

To create the empirical conditions in which Co-Learning can take place, I use the interdependence framework (Johnson et al., 2014) to analyze and define the roles of a machine and a human in their joint endeavors. Moreover, I have applied this framework in a symmetrical manner; in any task used, I design situations in which the machine has to rely on the human and vice versa (hard dependence), and I create situations for the machine to support the human and vice versa (soft dependence). I apply these principles in different research environments, using virtual robots (in a simulated environment) as

well as physically embodied robots.

1.5. RESEARCH QUESTIONS AND CONTRIBUTIONS

THE following sections briefly describe the five main research questions that will be addressed in the chapters of this thesis.

1.5.1. HUMAN ADAPTATION

WHILE adaptation mechanisms are thoroughly studied for AI agents collaborating among each other (e.g. Baker et al., 2020; Foerster et al., 2016; Iqbal and Sha, 2019), and for robots adapting to humans (Semeraro et al., 2023), the adaptive interactions that humans and AI agents or robots engage in receive little attention of research. These subtle and implicit behavioral adaptations might however determine whether performance outcomes are successful. Therefore, it is relevant to study how humans engage in the adaptive process while collaborating with a machine. The adaptation of leader-follower roles is a well-studied topic in literature (see e.g. Evrard and Kheddar, 2009; Y. Li et al., 2015; Mörtl et al., 2012) that was chosen as the context for studying human-machine co-adaptation. The research question is as follows:

RQ1 How do a human and a Wizard-of-Oz robot co-adapt leader-follower behavior when performing a collaborative task with no clear optimal solution?

To this end, I designed and implemented an experimental paradigm, in which human participants needed to perform a navigation task with a physical Wizard-of-Oz robot on a leash. Within this task, specific interdependencies between human and robot were designed such that participants continuously needed to choose to lead or follow the robot, thereby adapting their behavior to the context of the task. To detect and describe the adaptive interactions present within this setting, a fitting taxonomy was necessary. There are existing taxonomies of human-machine collaboration that describe the nature of the interactions, but these are often designed in a top-down manner (Jarrassé et al., 2012) or focus on isolated actions rather than on how interactions follow each other (Madan et al., 2015). I therefore developed a new taxonomy that contains concepts specifically focused on adaptive interactions, created in a bottom-up manner from observations in the experiment using Grounded Theory (Charmaz and Belgrave, 2007). The design and results of this study are presented in Chapter 2.

CONTRIBUTION

The developed method for qualitative analysis of human-machine collaboration as used in Chapter 2 provides insight into the dynamics of co-adaptation, specifically the individual interactions that humans engage in while collaborating with robot. It gives an overview of possible emergent Collaboration Patterns, categorized in a taxonomy of adaptive interactions that can be used as a language to describe co-adaptive behaviors. Some categories in the taxonomy are for example stable situations (sequences of actions that are repeated for a period of time), sudden adaptations and gradual adaptations, as I found that the adaptive process is built up of sequences of stable actions and sudden adaptations, rather than one gradual adaptation. These insights can be used by designers of collaborative machines, to align machine adaptive mechanisms with those of hu-

mans, to ultimately create fluent and successful human-machine collaborations. Moreover, having a better understanding of the dynamics of co-adaptation on the side of the human can support the design of more complex studies with actual adaptive machines.

1.5.2. HUMAN AND MACHINE CO-ADAPTATION WITH A VIRTUAL ROBOT

WHILE it is relevant to understand the behavioral dynamics of co-adaptation from the human's side, to study human-machine co-adaptation it is necessary to also study the behavioral dynamics in contexts with an algorithm controlled learning machine. Artificially Intelligent machines are not often studied behaviorally, but understanding how they behave when interacting and collaborating with humans is vital for making sure that desired outcomes are achieved (Rahwan et al., 2019). While there are some papers that touch upon the behavioral dynamics of co-adaptation (e.g. Kumar, 2024; Shafiq et al., 2020), the analyses done are often not of the joint human-machine behavior but rather of the separate agents, and relatively informal and exploratory. This leads to the following research question:

RQ2 How do a human and an ML-powered virtual robot co-adapt when performing a collaborative task with no clear optimal solution?

This research question demands the design of a task environment and learning algorithm to create the circumstances for human-machine co-adaptation. Moreover, it requires an experimental setup with which it can be studied whether co-adaptive human-machine collaborators are able to develop a successful strategy for their task. I defined requirements for all components of this research environment, after which I designed a task and algorithm for enabling human-machine co-adaptation. The task is inspired by an Urban-Search-and-Rescue scenario, a dynamic task that inherently requires ad hoc adaptation. Within the task, several interdependencies were implemented to ensure that the human and machine needed to collaborate to be successful. A variation of a Q-learning algorithm was used to enable the machine to reciprocally adapt its behavior to a human participant. The details of this design as well as the results of a study with human participants within the designed task environment is presented in Chapter 3.

CONTRIBUTION

The study provides requirements of experimental tasks and learning algorithms for enabling co-adaptation, such that the associated process dynamic can be studied in lab environments. I also provide a method for qualitative analysis of reciprocal co-adaptation, revealing how machine learning outcomes were influenced by human adaptation and vice versa. For example, participants that showed similar behavioral strategies were clustered, and compared the learned policies of the machines across the clusters to find commonalities and differences. This combined analysis of human behavioral clustering and machine learning model outputs provides new ways in which joint human-machine behavior can be studied.

1.5.3. HUMAN AND MACHINE CO-ADAPTATION WITH A PHYSICAL ROBOT

RESEARCH into human-machine collaboration is carried out by different research communities: some use virtual agents; some use physically embodied robots. The choice

often depends on the respective use case and on the origin of the research community. As it is well-known that the embodiment of an intelligent machine can have a large impact on the human-machine interaction (Herrera Perez and Barakova, 2020), I wanted to explore whether physical embodiment makes a difference in the co-adaptation dynamic as well. This leads to the following research question:

RQ3 How do a human and an embodied, ML-powered robot co-adapt when performing a collaborative task with no clear optimal solution?

Additionally, I wanted to investigate how it could be systematically assessed whether co-adaptation took place within a human-machine dyad.

I designed a task and algorithm intended to enable co-adaptation within a team consisting of a human and a physically embodied robot. The task is a handover task for which subtle physical coordination is vital. Moreover, the thesis presents a set of criteria with which it can be evaluated if co-adaptation has taken place in a specific human-machine dyad. Results of a pilot study within the designed task environment as well as a systematic evaluation of co-adaptation against the set of criteria are presented in Chapter 4.

CONTRIBUTION

The study provides insights into how a co-adaptive dynamic can be achieved in the collaboration between a human and a physical robot. It presents a systematic evaluation of criteria used to assess co-adaptation. Moreover, I provide a method for qualitative analysis of co-adaptation strategies, which was used to provide an overview of how several human-machine dyads co-adapted in context of the designed task. For example, I used a comparison of emergent strategies and the extent to which the human or the robot was responsible for failures to determine whether the human and robot co-adapted in a synchronous manner (and thereby met the criterion ‘synchrony’). Metrics and criteria like these can be used to improve design and evaluation of co-adaptation in future studies.

1.5.4. KNOWLEDGE STRUCTURES FOR SHARED REPRESENTATIONS OF EMERGENT COLLABORATION PATTERNS

IN this thesis co-learning is regarded as a two-step-process, consisting of co-adaptation and shared mental model updating. The studies discussed so far address the first step: co-adaptation. If the process of co-adaptation allows a human-machine team to develop successful emergent Collaboration Patterns, the team needs to become aware of those patterns and consolidate them. This can be achieved by jointly formalizing and updating emergent CPs in their shared mental model. In human-only teams, humans will usually do this by discussing and reflecting on their collaboration, assuming that they have a basic level of common ground. In human-machine teams, this cannot be assumed, and the team needs to be provided with an infrastructure that gives them an actionable shared representation. Such an infrastructure needs to fulfill the following two functions:

1. A model of team behaviors in which emergent Collaboration Patterns can be represented. The concepts in the model need to be understandable for human and machine team partners;

2. An interface that enables human and machine team partners to access the model, and to provide input regarding emergent and existing Collaboration Patterns.

There are existing frameworks that can be used to describe patterns of collaboration between partners (e.g. Plays (Kasmier et al., 2021; Miller, 2005) or Social Practices (Dignum, 2018)). In these frameworks, the Collaboration Patterns are generally predefined. Patterns that emerge while collaborating cannot be (easily) accommodated by these frameworks. Therefore a new framework is needed. This leads to the following research question:

RQ4 How can we facilitate co-adaptive human-robot collaborators with an infrastructure to formalize and share emergent collaboration patterns?

The thesis presents the design of an ontology model and an accompanying Graphical User Interface, which will be called the Collaboration Book. The purpose of this system is to enable human and machine partners to formalize emergent CPs in a format that is understandable for both the human and the machine. The Collaboration Book also allows partners to modify existing CPs in the ontology. The qualities of the Collaboration Book are empirically evaluated in a human subjects study. This is reported in Chapter 5. In this study, human participants watched videos of situations in which Collaboration Patterns emerged as a result of co-adaptation. Participants were given the instruction to report recognized CPs and to formalize these in the ontology model using the GUI.

CONTRIBUTION

The presented design of the Collaboration Book and its evaluation contributes to the design of models intended to enable human-machine team members to formalize and share successful collaborative human-machine team behavior. A knowledge structure like the Collaboration Book is especially valuable for teamwork in contexts where team behavior needs to evolve via co-adaptation and co-learning.

1.5.5. HUMAN-MACHINE CO-LEARNING FACILITATED BY SHARED REPRESENTATIONS OF EMERGENT COLLABORATION PATTERNS

THE studies of this thesis focused on what human-machine co-learning is (1.5.1 and 1.5.2), what requirements need to be present in a task in order to allow for co-learning (1.5.2), and whether co-learning can actually be observed in a human-agent team when a task is presented that meets these requirements (1.5.2). In study 1.5.3 this was again investigated, using a physical robot as team mate of the human (rather than an agent). An ontology and accompanying GUI (together called the Collaboration Book) were developed, intended to enable human and machine partners to formalize emergent CPs in a format that is understandable to both, thus providing the basis for a shared representation of successful team behaviors (1.5.4). A pilot study showed that humans were able to use the Collaboration Book successfully (1.5.4). In the final study of this thesis the effects of the system on human-machine collaboration and on team performance is empirically and experimentally investigated. The aim was to study component 1 (Co-Adaptation) and component 2 (Shared Mental Model updating) in conjunction. Given that emergent Collaboration Patterns as a result of co-adaptation are not always stable, I theorized that

providing human-machine teams with an infrastructure to formalize and share representations of these CPs (the Collaboration Book) would support co-learning and would subsequently lead to improved team performance. There is theoretical work on shared mental models in human-machine teams, but although there are AI agents that are capable of modeling human team members, there are no empirical evaluations of the effects that shared mental models in human-machine teams have on behavior, understanding and performance (Andrews et al., 2023). In the present study it was investigated whether the developed system for formalizing and sharing CPs facilitates human team members to develop a stronger awareness of emergent Collaboration Patterns. Furthermore, it was investigated whether the system supported the reuse of CPs. Finally, it was studied whether the system helped human-machine teams to improve their performance. This leads to the following research question:

RQ5 Does the developed Collaboration Book (Chapter 5) support human-machine team members in becoming aware of and reusing the CPS, as well as in improving their team performance?

The thesis presents an experiment in which I compared human-machine teams making use of the designed Collaboration Book to teams that could not make use of this infrastructure in Chapter 6.

CONTRIBUTION

The presented experiment provides insights into how access to a Shared Mental Model of emergent Collaboration Patterns influences task performance, Collaboration Fluency and the Collaboration Patterns itself within a co-learning human-machine team. The designed research environment with the addition of the ontology model and GUI can be adapted and used to study other topics related to co-learning in a systematic manner, such as the influence of memory or other machine learning algorithms.

1.6. CO-LEARNING: CONTRIBUTIONS AND STUDY DESIGNS

HUMAN-machine collaboration is often studied in static environments, and when learning is involved, often only human or machine learning is taken into account. With this thesis, I provide an elaborate study of co-learning, viewing it as an ongoing dynamic process and focusing on emergent behaviors and the consolidation of successful Collaboration Patterns, rather than regarding co-learning as a team's gradual convergence to a single optimal strategy. I provide a definition of the co-learning process, thereby unifying the language used about similar processes in human-agent teaming and human-robot collaboration literature. This can help to make co-learning a more outlined area of study within these fields. Moreover, I provide insights into how dynamic human-machine interaction can be studied, including requirements for study design as well as methods for analysis. In addition, I study how collaborating humans and machines can be facilitated in developing shared mental models about their collaboration. This can aid the design of human-machine teams in which partners can improve the understanding of each other, both implicitly and explicitly.

By taking a process-centered rather than an outcome-centered perspective we can learn about how humans actually engage in collaborating with a machine, for exam-

ple in how they solve conflict, as well as about how ML-controlled machines behave when confronted with an actively learning and adapting human. The contributions in this thesis towards a better understanding of complex human-machine collaboration will support the design of interactions and interventions that ensure fruitful long-term human-machine collaborations.

BIBLIOGRAPHY

- Ajoudani, A., Zanchettin, A., Ivaldi, S., Albu-Schäffer, A., Kosuge, K., & Khatib, O. (2018). Progress and Prospects of the Human-Robot Collaboration. *Autonomous Robots*, 42(5), 957–975. <https://doi.org/10.1007/s10514-017-9677-2>
- Akata, Z., Balliet, D., Rijke, M. d., Dignum, F., Dignum, V., Eiben, G., Fokkens, A., Grossi, D., Hindriks, K., Hoos, H., Hung, H., Jonker, C., Monz, C., Neerinx, M., Oliehoek, F., Prakken, H., Schlobach, S., Gaag, L. v. d., Harmelen, F. v., ... Welling, M. (2020). A Research Agenda for Hybrid Intelligence: Augmenting Human Intellect With Collaborative, Adaptive, Responsible, and Explainable Artificial Intelligence [Conference Name: Computer]. *Computer*, 53(8), 18–28. <https://doi.org/10.1109/MC.2020.2996587>
- Andrews, R. W., Lilly, J. M., Srivastava, D., & Feigh, K. M. (2023). The role of shared mental models in human-AI teams: A theoretical review. *Theoretical Issues in Ergonomics Science*, 24(2), 129–175. <https://doi.org/10.1080/1463922X.2022.2061080>
- Ansari, F., Erol, S., & Sihm, W. (2018). Rethinking Human-Machine Learning in Industry 4.0: How Does the Paradigm Shift Treat the Role of Human Learning? *Procedia Manufacturing*, 23, 117–122. <https://doi.org/10.1016/j.promfg.2018.04.003>
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., & Mordatch, I. (2020). Emergent Tool Use From Multi-Agent Autocurricula [arXiv:1909.07528]. *arXiv:1909.07528 [cs, stat]*. Retrieved December 10, 2020, from <http://arxiv.org/abs/1909.07528>
- Bradshaw, J. M., Sierhuis, M., Acquisti, A., Feltovich, P., Hoffman, R., Jeffers, R., Prescott, D., Suri, N., Uszok, A., & Van Hoof, R. (2003). Adjustable Autonomy and Human-Agent Teamwork in Practice: An Interim Report on Space Applications. In G. Weiss, H. Hexmoor, C. Castelfranchi, & R. Falcone (Eds.), *Agent Autonomy* (pp. 243–280, Vol. 7). Kluwer Academic Press. https://doi.org/10.1007/978-1-4419-9198-0_11
- Burke, S., Stagl, K., Salas, E., Pierce, L., & Kendall, D. (2006). Understanding Team Adaptation: A Conceptual Analysis and Model. *The Journal of applied psychology*, 91, 1189–207. <https://doi.org/10.1037/0021-9010.91.6.1189>
- Buschmeier, H., & Kopp, S. (2013). Co-constructing Grounded Symbols - Feedback and Incremental Adaptation in Human-Agent Dialogue. *KI - Künstliche Intelligenz*, 27, 137–143. <https://doi.org/10.1007/s13218-013-0241-8>
- Calhoun, G. (2021). Adaptable (not adaptive) automation: Forefront of human-automation teaming. *Human factors*, 64(2), 269–277.
- Charmaz, K., & Belgrave, L. L. (2007). *Grounded Theory*. Wiley & Sons Ltd.
- Darling, K. (2021, April). *The New Breed: What Our History with Animals Reveals about Our Future with Robots* [Google-Books-ID: oTrGDwAAQBAJ]. Henry Holt; Company.

- de Visser, E. J., Peeters, M. M. M., Jung, M. F., Kohn, S., Shaw, T. H., Pak, R., & Neerincx, M. A. (2020). Towards a Theory of Longitudinal Trust Calibration in Human–Robot Teams. *International Journal of Social Robotics*, 12(2), 459–478. <https://doi.org/10.1007/s12369-019-00596-x>
- Decuyper, S., Dochy, F., & Van den Bossche, P. (2010). Grasping the dynamic complexity of team learning: An integrative model for effective team learning in organisations. *Educational Research Review*, 5(2), 111–133. <https://doi.org/10.1016/j.edurev.2010.02.002>
- Diggelen, J. v., Neerincx, M., Peeters, M., & Schraagen, J. M. (2019). Developing Effective and Resilient Human-Agent Teamwork Using Team Design Patterns [Conference Name: IEEE Intelligent Systems]. *IEEE Intelligent Systems*, 34(2), 15–24. <https://doi.org/10.1109/MIS.2018.2886671>
- Dignum, F. (2018). Interactions as Social Practices: Towards a formalization [arXiv: 1809.08751]. *arXiv:1809.08751 [cs]*. Retrieved July 1, 2021, from <http://arxiv.org/abs/1809.08751>
- Döppner, D. A., Derckx, P., & Schoder, D. (2019). Symbiotic Co-Evolution in Collaborative Human-Machine Decision Making: Exploration of a Multi-Year Design Science Research Project in the Air Cargo Industry. *Proceedings of the 52nd Hawaii International Conference on Systems Science*, 264–273. <https://doi.org/10.24251/HICSS.2019.033>
- Ehrlich, S. K., & Cheng, G. (2018). Human-agent co-adaptation using error-related potentials. *Journal of Neural Engineering*, 15(6), 066014. <https://doi.org/10.1088/1741-2552/aae069>
- Evans, K. D., Robbins, S. A., & Bryson, J. J. (2023). Do We Collaborate With What We Design? [eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/tops.12682>]. *Topics in Cognitive Science*, 00(n/a), 1–20. <https://doi.org/10.1111/tops.12682>
- Evrard, P., & Kheddar, A. (2009). Homotopy switching model for dyad haptic interaction in physical collaborative tasks. *World Haptics 2009 - Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 45–50. <https://doi.org/10.1109/WHC.2009.4810879>
- Flemisch, F., Adams, C. A., Conway, S. R., Goodrich, K. H., Palmer, M. T., & Schutte, P. C. (2003). The H-Metaphor as a guideline for vehicle automation and interaction. NASA.
- Flemisch, F., Heesen, M., Hesse, T., Kelsch, J., Schieben, A., & Beller, J. (2012). Towards a dynamic balance between humans and automation: Authority, ability, responsibility and control in shared and cooperative control situations. *Cognition, Technology & Work*, 14(1), 3–18.
- Foerster, J., Assael, I. A., de Freitas, N., & Whiteson, S. (2016). Learning to Communicate with Deep Multi-Agent Reinforcement Learning. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 29* (pp. 2137–2145). Curran Associates, Inc. Retrieved February 12, 2019, from <http://papers.nips.cc/paper/6042-learning-to-communicate-with-deep-multi-agent-reinforcement-learning.pdf>

- Gao, Y., Barendregt, W., & Castellano, G. (2017). Personalised Human-Robot Co-Adaptation in Instructional Settings using Reinforcement Learning. *Proceedings of the Persuasive Embodied Agents for Behavior Change (PEACH2017) Workshop*.
- Herrera Perez, C., & Barakova, E. I. (2020). Expressivity Comes First, Movement Follows: Embodied Interaction as Intrinsically Expressive Driver of Robot Behaviour. In N. Noceti, A. Sciutti, & F. Rea (Eds.), *Modelling Human Motion: From Human Perception to Robot Design* (pp. 299–313). Springer International Publishing. https://doi.org/10.1007/978-3-030-46732-6_14
- Holstein, K., Alevan, V., & Rummel, N. (2020). A Conceptual Framework for Human–AI Hybrid Adaptivity in Education. In I. I. Bittencourt, M. Cukurova, K. Muldner, R. Luckin, & E. Millán (Eds.), *Artificial Intelligence in Education* (pp. 240–254). Springer International Publishing. https://doi.org/10.1007/978-3-030-52237-7_20
- Iqbal, S., & Sha, F. (2019). Actor-Attention-Critic for Multi-Agent Reinforcement Learning. *Proceedings of the 36th International Conference on Machine Learning*, 10.
- Jarrassé, N., Charalambous, T., & Burdet, E. (2012). A Framework to Describe, Analyze and Generate Interactive Motor Behaviors (M. O. Ernst, Ed.). *PLoS ONE*, 7(11), e49945. <https://doi.org/10.1371/journal.pone.0049945>
- Jiang, S. D., & Odom, J. (2018). Toward Initiative Decision-Making for Distributed Human-Robot Teams. *Proceedings of the 6th International Conference on Human-Agent Interaction - HAI '18*, 286–292. <https://doi.org/10.1145/3284432.3284467>
- Johnson, M., Bradshaw, J. M., Feltoovich, P. J., Jonker, C. M., Van Riemsdijk, M. B., & Sierhuis, M. (2014). Coactive Design: Designing Support for Interdependence in Joint Activity. *Journal of Human-Robot Interaction*, 3(1), 43. <https://doi.org/10.5898/JHRI.3.1.Johnson>
- Kasmier, D., Merrell, E., Kelly, R., Smith, B., Heisey, C., Maki, E., Brittain, M., Ankner, R., & Bush, K. (2021). Ontology of Plays for Autonomous Teaming and Collaboration. *Proceedings of the XIV Seminar on Ontology Research in Brazil (ONTOBRAS)*, 9–22.
- Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltoovich, P. J. (2004). Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems*, 19(6), 91–95. <https://doi.org/10.1109/MIS.2004.74>
- Kolbeinsson, A., Lagerstedt, E., & Lindblom, J. (2019). Foundation for a classification of collaboration levels for human-robot cooperation in manufacturing. *Production & Manufacturing Research*, 7(1), 448–471. <https://doi.org/10.1080/21693277.2019.1645628>
- Krüger, M., Wiebel, C. B., & Wersing, H. (2017). From Tools Towards Cooperative Assistants. *Proceedings of the 5th International Conference on Human Agent Interaction - HAI '17*, 287–294. <https://doi.org/10.1145/3125739.3125753>
- Kruijff, G. J. M., Janićek, M., Keshavdas, S., Larochelle, B., & Zender, H. (2012). Experience in System Design for Human-Robot Teaming in Urban Search & Rescue. *8th International Conference on Field and Service Robotics*, 1–14.
- Kumar, R. (2024). Unveiling the Dynamics of Human Decision-Making: From Strategies to False Beliefs in Collaborative Human-Robot Co-Learning Tasks.

- Lagerstedt, E., & Thill, S. (2020). Benchmarks for evaluating human-robot interaction: Lessons learned from human-animal interactions [ISSN: 1944-9437]. *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 137–143. <https://doi.org/10.1109/RO-MAN47096.2020.9223347>
- Li, S., Sun, W., & Miller, T. (2016). Communication in Human-Agent Teams for Tasks with Joint Action. In V. Dignum, P. Noriega, M. Sensoy, & J. S. Sichman (Eds.), *Coordination, Organizations, Institutions, and Norms in Agent Systems XI* (pp. 224–241, Vol. 9628). Springer International Publishing. https://doi.org/10.1007/978-3-319-42691-4_13
- Li, Y., Tee, K. P., Chan, W. L., Yan, R., Chua, Y., & Limbu, D. K. (2015). Continuous Role Adaptation for Human–Robot Shared Control. *IEEE Transactions on Robotics*, 31(3), 672–681. <https://doi.org/10.1109/TRO.2015.2419873>
- Mackay, W. E. (2023). Creating Human-Computer Partnerships. *CHIRA 2023 - 7th International Computer-Human Interaction Research and Applications, CCIS-1996*, 3–17. https://doi.org/10.1007/978-3-031-49425-3_1
- Madan, C. E., Kucukyilmaz, A., Sezgin, T. M., & Basdogan, C. (2015). Recognition of Haptic Interaction Patterns in Dyadic Joint Object Manipulation [Conference Name: IEEE Transactions on Haptics]. *IEEE Transactions on Haptics*, 8(1), 54–66. <https://doi.org/10.1109/TOH.2014.2384049>
- Maynard, M. T., Kennedy, D. M., & Sommer, S. A. (2015). Team adaptation: A fifteen-year synthesis (1998–2013) and framework for how this literature needs to “adapt” going forward. *European Journal of Work and Organizational Psychology*, 24(5), 652–677. <https://doi.org/10.1080/1359432X.2014.1001376>
- Miller, C. A. (2005). Delegation Architectures: Playbooks and Policy for Keeping Operators in Charge, 28–35.
- Mohammad, Y., & Toyoaki Nishida. (2008). Human adaptation to a miniature robot: Precursors of mutual adaptation. *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, 124–129. <https://doi.org/10.1109/ROMAN.2008.4600654>
- Mörzl, A., Lawitzky, M., Kucukyilmaz, A., Sezgin, M., Basdogan, C., & Hirche, S. (2012). The role of roles: Physical cooperation between humans and robots. *The International Journal of Robotics Research*, 31(13), 1656–1674. <https://doi.org/10.1177/0278364912455366>
- Nass, C., & Moon, Y. (2000). Machines and Mindlessness: Social Responses to Computers [eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/0022-4537.00153>]. *Journal of Social Issues*, 56(1), 81–103. <https://doi.org/10.1111/0022-4537.00153>
- Neerinx, M. A., van Vught, W., Blanson Henkemans, O., Oleari, E., Broekens, J., Peters, R., Kaptein, F., Demiris, Y., Kiefer, B., Fumagalli, D., & Bierman, B. (2019). Socio-Cognitive Engineering of a Robotic Partner for Child’s Diabetes Self-Management. *Frontiers in Robotics and AI*, 6. Retrieved February 10, 2023, from <https://www.frontiersin.org/articles/10.3389/frobt.2019.00118>
- Nikolaïdis, S., Hsu, D., & Srinivasa, S. S. (2017). Human-robot mutual adaptation in collaborative tasks: Models and experiments. *I. J. Robotics Res.*, 36, 618–634. <https://doi.org/10.1177/0278364917690593>

- Nikolaïdis, S., Zhu, Y. X., Hsu, D., & Srinivasa, S. (2017). Human-Robot Mutual Adaptation in Shared Autonomy. *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI '17*, 294–302. <https://doi.org/10.1145/2909824.3020252>
- Nourbakhsh, I., Sycara, K., Koes, M., Yong, M., Lewis, M., & Burion, S. (2005). Human-Robot Teaming for Search and Rescue. *IEEE Pervasive Computing*, 4(1), 72–78. <https://doi.org/10.1109/MPRV.2005.13>
- Ogenyi, U. E., Liu, J., Yang, C., Ju, Z., & Liu, H. (2019). Physical Human-Robot Collaboration: Robotic Systems, Learning Methods, Collaborative Strategies, Sensors, and Actuators. *IEEE Transactions on Cybernetics*, 1–14. <https://doi.org/10.1109/TCYB.2019.2947532>
- Pepe, K., & Hutchison, N. (2022). AI4SE and SE4AI: Setting the Roadmap toward Human-Machine Co-Learning [eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/inst.12417>]. *INSIGHT*, 25(4), 80–84. <https://doi.org/10.1002/inst.12417>
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J. W., Christiakis, N. A., Couzin, I. D., Jackson, M. O., Jennings, N. R., Kamar, E., Kloumann, I. M., Larochelle, H., Lazer, D., McElreath, R., Mislove, R., Parkes, D. C., Pentland, A. ', . . . Wellman, M. (2019). Machine behaviour. *Nature*, (568), 477–486. <https://doi.org/10.1038/s41586-019-1138-y>
- Sasagawa, A., Fujimoto, K., Sakaino, S., & Tsuji, T. (2020). Imitation Learning Based on Bilateral Control for Human–Robot Cooperation [Conference Name: IEEE Robotics and Automation Letters]. *IEEE Robotics and Automation Letters*, 5(4), 6169–6176. <https://doi.org/10.1109/LRA.2020.3011353>
- Semeraro, F., Griffiths, A., & Cangelosi, A. (2023). Human–robot collaboration and machine learning: A systematic review of recent research. *Robotics and Computer-Integrated Manufacturing*, 79, 102432. <https://doi.org/10.1016/j.rcim.2022.102432>
- Shafti, A., Tjomsland, J., Dudley, W., & Faisal, A. A. (2020). Real-World Human-Robot Collaborative Reinforcement Learning [arXiv: 2003.01156]. *arXiv:2003.01156 [cs]*. Retrieved March 8, 2021, from <http://arxiv.org/abs/2003.01156>
- Sordani, A., Galley, M., Auli, M., Brockett, C., Ji, Y., Mitchell, M., Nie, J.-Y., Gao, J., & Dolan, B. (2015). A Neural Network Approach to Context-Sensitive Generation of Conversational Responses [arXiv: 1506.06714]. *arXiv:1506.06714 [cs]*. Retrieved February 12, 2019, from <http://arxiv.org/abs/1506.06714>
- Uitdewilligen, S., Waller, M. J., & Pitariu, A. H. (2013). Mental Model Updating and Team Adaptation [Publisher: SAGE Publications Inc]. *Small Group Research*, 44(2), 127–158. <https://doi.org/10.1177/1046496413478205>
- van den Bosch, K., Schoonderwoerd, T., Blankendaal, R., & Neerinx, M. (2019). Six Challenges for Human-AI Co-learning. *Lecture Notes in Computer Science*, 11597, 572–589. https://doi.org/10.1007/978-3-030-22341-0_45
- Wang, L., Liu, S., Liu, H., & Wang, X. V. (2020). Overview of Human-Robot Collaboration in Manufacturing. In L. Wang, V. D. Majstorovic, D. Mourtzis, E. Carpanzano, G. Moroni, & L. M. Galantucci (Eds.), *Proceedings of 5th International Conference on the Industry 4.0 Model for Advanced Manufacturing* (pp. 15–58). Springer International Publishing. https://doi.org/10.1007/978-3-030-46212-3_2

- Weerarathna, I. N., Raymond, D., & Luharia, A. (2023). Human-Robot Collaboration for Healthcare: A Narrative Review. *Cureus*, *15*(11), e49210. <https://doi.org/10.7759/cureus.49210>
- Wenskovitch, J., & North, C. (2020). Interactive AI: Designing for the “Two Black Boxes” Problem. *Hybrid Human-Artificial Intelligence Special Issue*, 1–10.
- Wiedow, A., & Konradt, U. (2011). Two-Dimensional Structure of Team Process Improvement: Team Reflection and Team Adaptation [Publisher: SAGE Publications Inc]. *Small Group Research*, *42*(1), 32–54. <https://doi.org/10.1177/1046496410377358>
- Wiltshire, T. J., Van Eijndhoven, K., Halgas, E., & Gevers, J. M. (2024). Prospects for augmenting team interactions with real-time coordination-based measures in human-autonomy teams. *Topics in Cognitive Science*, *16*(3), 391–429.
- Xu, Y., Ohmoto, Y., Okada, S., Ueda, K., Komatsu, T., Okadome, T., Kamei, K., Sumi, Y., & Nishida, T. (2012). Formation conditions of mutual adaptation in human-agent collaborative interaction. *Applied Intelligence*, *36*(1), 208–228. <https://doi.org/10.1007/s10489-010-0255-y>
- Zoelen, E. M. v., Barakova, E. I., & Rauterberg, M. (2020). Adaptive Leader-Follower Behavior in Human-Robot Collaboration [ISSN: 1944-9437]. *29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 1259–1265. <https://doi.org/10.1109/RO-MAN47096.2020.9223548>

2

IDENTIFYING INTERACTION PATTERNS OF TANGIBLE CO-ADAPTATIONS IN HUMAN-ROBOT TEAM BEHAVIORS

As robots become more ubiquitous, they will increasingly need to behave as our team partners and smoothly adapt to the (adaptive) human team behaviors to establish successful patterns of collaboration over time. A substantial amount of adaptations present themselves through subtle and unconscious interactions, which are difficult to observe. Our research aims to bring about awareness of co-adaptation that enables team learning. This paper presents an experimental paradigm that uses a physical human-robot collaborative task environment to explore emergent human-robot co-adaptations and derive the interaction patterns (i.e., the targeted awareness of co-adaptation). The paradigm provides a tangible human-robot interaction (i.e., a leash) that facilitates the expression of unconscious adaptations, such as “leading” (e.g., pulling the leash) and “following” (e.g., letting go of the leash) in a search-and-navigation task. The task was executed by 18 participants, after which we systematically annotated videos of their behavior. We discovered that their interactions could be described by four types of adaptive interactions: stable situations, sudden adaptations, gradual adaptations and active negotiations. From these types of interactions we have created a language of interaction patterns that can be used to describe tacit co-adaptation in human-robot collaborative contexts. This language can be used to enable communication between collaborating humans and robots in future studies, to let them share what they learned and support them in becoming aware of their implicit adaptations.

 **Emma M. van Zoelen**, Karel van den Bosch, Matthias Rauterberg, Emilia I. Barakova, Mark Neerincx. 2021. Identifying interaction patterns of tangible co-adaptations in human-robot team behaviors. In *Frontiers in Psychology*, 12, 645545.

2.1. INTRODUCTION

WITH AI being increasingly used in social robotics (Breazeal et al., 2016), there is a growing number of possible applications in which artificially intelligent robots need to interact and collaborate with humans in the physical space. Creating AI for the physical world comes with many challenges, one of which is ensuring that a robot does not only execute its own task, but instead behaves as a team partner, to enable human and robot to become one well-functioning unit of collaboration. One of the mechanisms that can be used to enable this, is a process of co-adaptation, where both human and robot, through (physical) interaction, adapt their behavior to develop successful patterns of collaboration over time (Chauncey et al., 2017).

To define what we mean by co-adaptation, we can think of how humans adapt their behavior in a reciprocal manner when they collaborate with other humans: the kind of adaptive interactions they use to achieve a fruitful collaboration. It is known that a human team's ability to adapt to new circumstances is vital for its performance, and team members tend to rapidly develop updated interaction patterns that fit with new situations (Burke et al., 2006; Uitdewilligen et al., 2013). Humans have the ability to intuitively interpret body language of their team members and to send signals when initiating adaptations (Sacheli et al., 2013). This kind of non-verbal interaction is not obvious when a team member is a robot. While we might be able to interact with a robot using language, collaborative interactions are generally multimodal and contain many subtle and implicit non-verbal interaction cues that help us to create tacit knowledge. The focus of this paper is on these non-verbal interactions, and specifically those that are connected to physical contact.

Two classic examples of non-verbal interactions for co-adaptation in a human-non-human collaborative context can be found in human-animal interaction:

1. The interaction between a horse and its rider (Flemisch et al., 2008);
2. The interaction between a guide dog and a blind person (Lagerstedt and Thill, 2020).

When a human rides a horse, they start off as two separate entities with their own goals. As they interact for a longer period of time, they gradually start to better understand the other, adapting their interaction concurrently, until they become one joint system acting toward a common goal through subtle and implicit interactions. Another example is the interaction between blind people and their guide dogs: blind people truly need to trust and follow the choices of the guide dog, whereas in horse riding, the human makes most of the decisions. When guide dogs and blind people learn to navigate together, the human needs to learn to assess when to adapt its behavior to follow the dog, and when to give the dog directions about their route. The dog must learn to understand what the human is and isn't comfortable with and adapt its behavior to that. All this learning and adapting takes place through subtle physical interactions.

Mechanisms of adaptation have been studied in intelligent agents, more specifically in the field of multi-agent systems (e.g. Foerster et al., 2016; Iqbal and Sha, 2019). Research addresses learning algorithms, such as different types of Reinforcement Learning,

and investigates their effects on agent performance or team performance. Little to no attention is paid to the interactions that the agents engage in, which bring about the adaptations (except for some examples such as Baker et al., 2020). Even when mechanisms of adaptation are studied in human-robot interaction contexts such as in Nikolaidis, Hsu, and Srinivasa, 2017, the effects on performance are studied. We believe that research should also address the interactions that bring about successful adaptation, to come closer to the fluency and naturalness of the above-mentioned human-animal examples.

There is a need for further study of the specific interactions and interaction patterns that bring about co-adaptation when humans and robots collaborate. A deeper insight in these interactions and patterns can help researchers and designers to study and create more natural and fluent human-robot collaborations that take the limitations and affordances of the physical world into account. In addition, such insights can support the collaborating human and robot to become more aware of their implicit adaptations and communicate about them, to further improve their collaboration. In Section “Co-Adaptation in Human-Robot Teams,” we define co-adaptation in a human-robot collaboration context, and we explain the relevance of embodiment in this process in Section “Research Challenge.” We describe an experimental paradigm that we designed and implemented to conduct an empirical study into co-adaptation and how it emerges from interactions. This human-robot team task was presented to human participants, after which we analyzed the team behavior in terms of interactions and interaction patterns. The resulting interaction pattern vocabulary and language provides a thorough analysis of co-adaptive interactions surrounding leadership roles in human-robot teams.

2.2. CO-ADAPTATION IN HUMAN-ROBOT TEAMS

2.2.1. CO-ADAPTATION - A DEFINITION

IN human-only teams, the term ‘team adaptation’ is used to describe the changes that occur in team behavior and performance. More specifically, Burke et al., 2006 define team adaptation as “a change in team performance, in response to a salient cue or cue stream, that leads to a functional outcome for the entire team”. They describe that it “is manifested in the innovation of new or modification of existing structures, capacities, and/or behavioral or cognitive goal-directed actions” (p. 1190). On top of that, it is argued that an important aspect in this is that the team members update their mental models according to changes in the task situation (Uitdewilligen et al., 2013).

We use the term co-adaptation instead of team adaptation, as we study the adaptive interactions at the level of the individual actors: team adaptation is a result of adaptive behavior exhibited by the individual team members. Also, co-adaptation is used more often in the context of (physical) human-robot interaction. We define co-adaptation in human-robot teams as follows:

A process in which at least two parties change their behavior and/or mental models concurrently as a consequence to changes in task or team situation while collaborating with each other.

This concurrent changing of behavior and/or mental models is relevant for the team, as smooth collaboration requires partners to adapt to each other over time. Since humans are adaptive creatures by nature, and artificially intelligent systems are becoming

more and more adaptive, there is an opportunity to study how they adapt together as they collaborate.

Co-adaptation is a process which generally takes place over a short period of time, e.g., over the course of several seconds or minutes; this timespan is generally considered in the study of co-adaptive behaviors (e.g., in Nikolaidis, Hsu, and Srinivasa, 2017, see also Section “Related Work”). It is not necessarily a deliberate process: adaptation happens as a consequence of interactions and an implicit or explicit drive to improve performance or experience. The resulting behaviors or mental models in both adapting partners do not necessarily persist over time and contexts, as new contexts and influences may cause the co-adaptation to continue. We used the above definition to describe co-adaptation as a design pattern in Table 2.1 (according to the template specified in Diggelen et al., 2019). This table provides a detailed explanation of the possible positive and negative effects of co-adaptation, as well as an overview of the kind of contexts in which it is relevant to develop or apply co-adaptation.

2.2.2. RELATED WORK

IN the sections below, we discuss related work on co-adaptation in human-robot or human-AI collaborative contexts, as we are studying interactions that bring about co-adaptation. Since we are specifically interested in analyzing and categorizing interactions and interaction patterns, we also looked at literature on interaction taxonomies within collaborative contexts. There is a body of research on dynamic role switching in human-robot collaboration, which has many similarities with how we described co-adaptation in terms of interactions. However, the existing literature (e.g. Evrard and Kheddar, 2009; Li et al., 2015; Mörtl et al., 2012) focuses on computational approaches to enable a robot to dynamically switch roles in an attempt to optimize performance of a human-robot team. While the existing studies evaluate the impact of the robot strategies on human factors, they do not study the natural interactions between the human and robot that arise as a consequence of the necessity for role switching. Therefore, we do not go into further depth on these papers.

CO-ADAPTATION

Most work on human-agent co-adaptation focuses on making the agent adaptive to the human, using information on different properties of the human (e.g. Buschmeier and Kopp, 2013; Ehrlich and Cheng, 2018; Gao et al., 2017; Yamada and Yamaguchi, 2002). There have been studies that investigated how a human adapts in situations when collaborating with an intelligent agent, using the team's performance to determine the impact of co-adaptive collaborations (e.g. Mohammad and Toyooki Nishida, 2008; Nikolaidis, Hsu, and Srinivasa, 2017; Nikolaidis, Zhu, et al., 2017; Youssef et al., 2014). In addition to determining the effects of co-adaptation on performance, it is also necessary to study the kind of interactions that emerge throughout the co-adaptation process and support team members in the process of developing a fluent collaboration. A better understanding of these processes will help to initiate and maintain co-adaptation in human-agent teams.

Xu et al., 2012 have outlined requirements for co-adaptation to occur in human-robot teams. First, they argue that in order to achieve a common purpose, both agents

Proposed Design Pattern for Co-Adaptation	
Behavior patterns	Team members engage in collaboratively solving a task. While they do this, they observe each other's actions and adapt their behavior in an attempt to make the collaboration more fluent and effective.
Potential positive effect	The performance on the collaborative task increases. Both partners will be able to work more efficiently, as there is less idle time, fewer mistakes and more understanding between the partners.
Potential negative effect	In the process of adapting, there is a risk of making mistakes. In addition, it takes time to adapt to a working strategy, which might have negative effects on the immediate performance.
Use when	Team partners need to collaborate but don't know the best strategy to complete the task. At the same time, the task and capabilities of the team members contain many implicit aspects that are hard to explicitly communicate or make agreements about.
Example	A human and a robot arm have to collaboratively assemble a product. There are different parts that either of them can assemble, and some parts need to be jointly assembled; e.g., the robot needs to hold up a heavy part while the human adjusts the bottom. If the human has to constantly provide the robot with instructions, this will slow them down, so it is useful to let the robot move autonomously and to synchronize their actions. When they start collaborating, the human might not trust the robot enough to adjust the bottom of a part that the robot holds up, in fear of being crushed underneath the part. The robot might see the hesitation and move the part upside down, such that the human can reach the object more easily. In turn, the human will have to adjust their workflow to do their task, but the fact that the robot adapted might increase the trust and understanding between the partners, which can in turn improve future team performance. While adapting, however, the human might make the mistake of trusting the robot too much, and think they can climb on top of the heavy part whereas the robot is unable to hold that weight. The co-adaptive process, if done too quickly or inconsiderate, therefore has the risk of making mistakes that hamper immediate performance.
Design rationale	A process of mutual adaptation helps to establish and maintain common ground, one of the main aspects of necessary for enabling collaboration between humans and machines (Klein et al., 2004; Sciutti et al., 2018). This might also be called mutual understanding, meaning that both parties are able to predict and/or explain each other's actions, leading to trust and eventually smooth collaboration (Azevedo et al., 2017). In human-only teams, co-adaptation leads to team adaptation (Burke et al., 2006), which has shown to be an essential aspect of successful teams.
Type	Collective

Table 2.1: Proposed Design Pattern for Co-adaptation.

Task requirements	
Mixed initiative	Both parties can take the initiative for an interaction at any point in time (see Xu et al., 2012).
Interaction symmetry	Interaction modalities should have a certain level of symmetry, meaning that there is at least some overlap in the way the two parties can interact with the other, to enable imitation. Interaction symmetry thereby contributes to the common ground.
Performance improvement	By adapting their individual behavior, team members can support an improvement in team performance.
Collaborative advantage	It must be easier to be successful at the task when collaborating, as opposed to doing it on your own.
Common ground	There must be a common ground between the collaborating partners. In our case this comes from the physical nature of the task.

Table 2.2: Task requirements for a collaborative, co-adaptive task environment.

need to be prepared to adapt their behavior to their partner, should actively and dynamically estimate the partner's intention, and develop options of how to adapt their own behavior in response. Another requirement is that the agents need to be able to receive and appreciate feedback or reward from the other, to express their internal state to their partner in a comprehensible manner, and to establish with their partner a common protocol for interaction. Third, the authors outline several inducing conditions, derived from experimental work, that can be used to ensure a mutually adaptive process will start, for example that both agents should be able to take initiative.

We formulated our own requirements for a task environment that would fit with our research goals of studying co-adaptive interactions, which include the mixed-initiative requirement as well as the requirements for dynamic and adaptive behavior (which we connected to an improvement in team performance). Moreover, we added two requirements that relate to the presence of a common ground (general common ground as well as interaction symmetry). Common ground is considered to be necessary for any collaboration (Klein et al., 2004), while interaction symmetry is often used to provide the possibility for imitation, which can create initial common ground (e.g. in Sasagawa et al., 2020). A full description of these requirements is given in Table 2.2.

INTERACTION TAXONOMIES

The literature reports two important existing studies into interaction taxonomies that describe interactions in collaborative tasks. One of those papers describes a top-down approach of describing different types of interactive behaviors based on game theory (Jarrassé et al., 2012); the other describes a bottom-up approach where interaction behaviors were identified from empirical observations (Madan et al., 2015). Both taxonomies were validated on their applicability by successfully classifying behaviors in different HRI scenarios. Although useful to describe collaborative behavior, the top-down approach (as used in Jarrassé et al., 2012) resulted in a taxonomy that describes interaction at a high level of abstraction (distinguishing for example between competitive versus collaborative behavior). Such a taxonomy can be used to describe the overall behavior in a

task, but it does not provide insights into (atomic) interactions that drive adaptation. The taxonomy presented in the other paper (Madan et al., 2015) presents interactions at various levels of detail, where the highest level of detail describes categories of interactions (i.e., harmonious, conflicting or neutral). The lower level interactions are more closely related to what we are interested in. They describe interaction patterns such as harmonious translation, persistent conflict and passive agreement. These interaction patterns focus on interactions related to collaborative object manipulation and were observed in a specific controlled task environment. This leaves room to study interactions in other contexts, to investigate a wider range of possible interaction patterns. Moreover, they do not provide information on how the different interaction patterns relate to each other; how they follow each other or how one pattern leads to a specific other pattern. We believe that the relations between interaction patterns are especially important when looking at adaptation. In our study, we take a bottom-up approach to identifying interaction patterns, which is similar to the work of Madan et al., 2015. This means that we do not predefine or design interactions, but that we set up a task that allows participants to behave as naturally as possible, and treat the data collection and analysis as an ethnographic study. Since such an approach requires us to have as little assumptions about behavior which will be observed as possible, we do not use the existing taxonomies when identifying interaction patterns. In our analysis, we focus specifically on adaptive interactions, as well as on how the different observed behaviors relate to each other. We will reflect on how our findings overlap or differ to the existing work in Section “Relation to Existing Interaction Taxonomies,” to understand how they might complement or complete each other.

2.3. RESEARCH CHALLENGE

THE goal of this study is to empirically investigate the interactions between humans and robots that underly their co-adaptation when jointly performing a task. A challenge is that the adaptive intentions and outcomes of interactions are often not directly clear and observable. Partners may themselves not be aware that their behavior is an adaptation to the developments, and may be a response to subtle cues, possibly processed unconsciously. In order to nevertheless investigate how such important processes take place in a human-robot team, the approach of observing and analyzing embodied human-robot team behavior was taken. Expressivity and intentionality of behavior plays a large role in embodied interaction (Herrera Perez and Barakova, 2020). It is believed that in such a setting the subtle and perhaps unconsciously executed adaptations will be expressed by means of physical, embodied interactions, hence being accessible for observation and analysis.

The literature on using embodied intelligence when studying human-robot interaction shows two main lines of research:

- A line of research that focuses on human cognition: investigating how computers or robotic interfaces can be used to understand and extend human cognition and behavior. An example of this is extending human cognition through prostheses or sensory substitution (e.g., as in Bach-y-Rita and W. Kercel, 2003; Kaczmarek et al., 1991; Nagel et al., 2005);

- A line of research that focuses on using embodiment to create more intelligent computers (or machines or robots). The robot's intelligence is 'grounded' by a body with which it can interact with its environment (for example as described in Duffy and Joue, 2000; Kiela et al., 2016).

2

Our research approach is not directed at the intelligence of one particular partner of the team, but at the intelligence emerging from the interaction of partners. In the first described line of research (extending human cognition, e.g., using prosthetics), one of the main aims is to create a unity between the human and the added technology, such that for the human the artificial parts feels as though it is part of themselves. We research the unity of human and technology jointly forming a team, with both having a certain level of autonomy, and sharing a common goal. This approach distinguishes between cognition on an individual level (per agent) and collective cognition, at the team level. In our work, we focus on the team as a dyad formed by one human and one robot.

It is believed that studying embodied human-robot interaction, while they collaboratively perform a task, is an approach that can help us to discover and understand how a team adapts to the dynamics of the context, and how this adaptation emerges from the interactions between the team members.

2.4. MATERIALS AND METHODS

2.4.1. TASK ENVIRONMENT: SEARCH-AND-NAVIGATION

WE have developed a task in which a human and a robot jointly navigate through a space while searching for objects to collect additional points. The conditions and interdependencies described in Table 2.2 were implemented in this task.

The team of human and robot are given the task of navigating between two points in space. The team's assignment is to reach the goal location with as many points as possible. They start with 60 points, and lose a point each second until they reach the goal location. Virtual objects were hidden in the task area: some close to the shortest route to the goal location; others further to the side. Picking up a virtual object yields the team 10 points. These scores were chosen after trying out the task several times, such that solely focusing on the goal would yield approximately the same score as solely focusing on the objects, while combining the capabilities of both team members could potentially result in a higher score than either of the extreme strategies, ensuring a trade-off between the two. The partners have complementary capabilities: only the robot knows where the objects are; only the human can oversee the route and distance to the goal (see Figure 2.1 for an image of the field used). A sound cue is given when an object is picked up.

2.4.2. DESIGN OF HUMAN-ROBOT INTERACTION

WE designed and implemented a remotely controlled robot with a leash (Figure 2.2). An ambiguous form was selected for the robot, without anthropomorphic features. This was chosen on purpose, to allow humans interacting with it to focus on the interaction, not on its form.

The leash was designed to be the only direct communication channel between the robot and a participant, to ensure specific evaluation of the interaction through the leash



Figure 2.1: The field on which the task was executed. Participants moved from the goal on the left to the goal on the right (where the robot is stationed).

without too much noise of other interaction modalities. On top of that, the leash interaction allows for subtle and implicit interactions as both the participant and the robot can pull the leash more or less. The robot was explicitly made to be quite large and heavy, to allow it to pull the participant in a direction as well.

For our study, the robot was remotely controlled by a human operator (i.e. the experimenter). It is usually preferable that the operator is hidden from the participant, however, due to technical limitations this was not possible. The human operator was therefore on the field together with the participant and the robot during the experiment. A small pilot with two participants showed that participants only paid attention to the human operator in the first few seconds of the experiment, after which they directed their attention to the robot only. Therefore, and for practical reasons, we decided that it did not pose a problem for our study goals that the human operator was visible. The human operator controlled the robot according to a set of pre-defined rules: to direct the robot to the closest virtual object (following a default route as much as possible, as specified in Figure 2.3) if the leash was held loose by the participant (the operator, in contrast to the participant, knew the locations of all hidden virtual objects). If the participant kept a tight leash, the operator directed the robot to give in and to move toward the participant until the leash was no longer stretched. A detailed description of these rules is provided in Table 2.3. The human operator made decisions based on visual cues: they carefully watched the leash to see whether it was stretched. Human response time to visual cues is known to be on average 0.25 s, therefore, we can assume that the robot responded to participant behavior with a delay of 0.25 s.

The task and robot were designed such that both partners had their own knowledge,



Figure 2.2: Two participants interacting with the robot showing a situation with a stretched leash and thus in a leading role (top) and situation with a loose leash and thus a following role (bottom).

enabling them to initiate actions that their partner cannot initiate. The knowledge of both partners was relevant for the task, making collaboration beneficial and enabling the partners to learn how to use their knowledge in the best possible way. All communication and coordination between the human and the robot took place through the leash, which ensured that interactions are physically grounded, and allowed for subtle and implicit interactions.

2.4.3. EXPERIMENT SETUP AND INITIAL RESULTS

THE experimental paradigm described above was previously used to study leadership shifts and its influence on subjective Collaboration Fluency in human-robot teams. This section will explain the experimental protocol used as well as results obtained in that study. For the current study, we have re-analyzed the data obtained in the original study to research specifically what interactions and interaction patterns bring about




Figure 2.3: The four predefined maps with the locations of the objects (red circles), including a line indicating the default route of the robot. The bottom of the field is the starting point.

Situation	Resulting robot behavior
The leash is stretched	Follow the human in the direction that the leash is pulled in
The leash is loose AND the human-robot team is on the predefined route	Follow the predefined route
The leash is loose AND the human-robot team is not on the predefined route AND there are virtual objects that have not been picked up	Move toward the nearest virtual object that has not yet been picked up
The leash is loose AND the human-robot team is not on the predefined route AND all virtual objects have been picked up	Move towards the goal

Table 2.3: The protocol used by the human operator to control the behavior of the robot.

co-adaptation in such a task. In Section “Analyzing Behavior to Uncover Interaction Patterns” and “Data Analysis: Extracting Interaction Patterns” we will describe in detail how that analysis was done.

EXPERIMENTAL PROTOCOL

Participants were told that they had to perform a collaborative task together with an intelligent robot, while holding the leash of the robot. They were presented with the described task and human-controlled robot, and were given instructions about how they could score points.

Before the start of the experiment, the participants were given the possibility to walk from one end of the field to the other with the robot. This was done to give participants an indication of the speed of the robot. After that, the first round started. The task was performed four rounds per participant. The locations of the virtual objects were different for every round. Four predefined maps with specified locations of the virtual objects were created for the human operator (Figure 2.3). Each of these maps were used for

each participant during one of the rounds. The order of the maps was randomized for each participant to make sure that the observed behavior would not be influenced by the specific maps. After each round, participants were asked three interview questions:

2

1. Can you explain the behavior of the robot?
2. What was your strategy for completing the task?
3. How did you experience the collaboration?

An overview of the answers given to these questions was given in van Zoelen et al., 2020. For the analyses described in the following sections, the answers to these interview questions were used to support the researchers in interpreting participant behavior.

PARTICIPANTS

A total of 18 people participated in the experiment (9 male, 9 female), consisting of students from different programs within Eindhoven University of Technology, with an average age of 23 (SD = 3.9). The participants were told that the person with the highest number of points on a single run would receive a gift voucher of €10 to motivate them to perform to the best of their abilities. Before the start of the experiment, participants gave their consent after carefully reading the consent form that explained all details of the experiment except for the focus of the research (evolving leadership shifts) and the specific behavior of the robot. After the experiment, they were debriefed on the exact purpose of the experiment.

DATA ANALYSIS: CODING PROCESS

While performing the task, a camera placed in a corner of the field recorded the behavior of the participants. These videos were thoroughly analyzed through a process based on Grounded Theory (Charmaz, 2014), using different stages of open coding, closed coding, and categorizing. All videos were coded using an open coding process at first, to get a view on the different kinds of behavior present among participants as well as on events that triggered participants to switch between a more leading and a more following role. Using the results from the open coding, a coding scheme for closed coding was developed that contained codes describing task events, robot movement, participant movement, leash activity and the participant's location relative to the robot. Each code was characterized as a leading, following or neutral behavior (see Table 2.4).

All videos were then coded again using a closed coding process using The Observer XT (Noldus, 2019). This was done in an iterative manner, where each video was watched and coded again for each code category as specified in Table 2.4. Codes of different categories therefore could exist in parallel (e.g., codes for leash activity and codes for participant movement), while codes within a category (e.g., 'loose' and 'stretched') could not exist in parallel. An exception were the 'task event' codes; these were used to record how long it took participants to finish the task and to be able to see whether behavior lined up with task events. This left us with an overview of whether the participant was in a leading, following or neutral position across the three variables of leash activity, participant movement and participant location at each moment during the task. Combined with

Code Category	Code
Task events	Task is running
	Object sound
Robot movement	Standing still
	Moving toward object in goal direction
	Moving toward object away from goal
	Moving toward object across field
	Moving with participant
	Moving in goal direction
Participant movement	Standing still/waiting
	Moving around robot
	Moving in goal direction*
	Moving in robot direction
	Moving across field*
Leash activity	Loose
	Stretched*
	Pulled in direction*
	Loosening/stretching
Participant location relative to robot	Behind
	In front of*
	Next to

Table 2.4: The coding scheme that was developed to analyze the behavior of participants and the robot in the experiment. Codes marked with a * were considered leading behavior by the participant. The presence of these codes was taken as an indication for leading behavior in all further analyses.

the visualization tool in The Observer XT, this enabled us to visually analyze the (development of) different behaviors across rounds simultaneously as well as to quantify the amount of leading behaviors present in each run. Intercoder reliability for the duration of sequences with another coder for 5.6% of the data (videos of 4 runs) was found to be 97.55%.

PREVIOUS RESULTS

The task environment presented above has previously been described in Zoelen et al., 2020. The main findings focused on three aspects:

- Interactions that trigger people to reconsider leadership roles;
- How leader/follower behavior changes over time;
- The interplay between subjective Collaboration Fluency and shifting leader/follower roles.

As the current paper builds upon and greatly extends the results presented in the previously published work, we will summarize these findings in the sections below.

Switch Triggers An open coding process revealed six types of situations that typically triggered participants to reconsider whether they should behave in a more leading or following way. The first of those situations is at the start of the task, where participants express their initial idea about the role they should take on. The other five triggers are the following:

1. Sound indicating a virtual object;
2. A leash pull by the robot;
3. The robot deviating from the route that leads to the final goal without clear leash pull;
4. Getting close to the goal;
5. The robot standing still.

Of these five triggers, numbers 1 and 2 are explicitly visible and clear moments in time, while numbers 3 and 4 are more implicit, slowly emerge and are harder to observe. Number 5 is a special case, as the robot standing still was sometimes clearly linked to the collection of a virtual object, but sometimes emerged more implicitly from the interactions in the task. Besides grouping them in explicit versus implicit triggers, they can also be grouped into task feedback (1 and 4) and partner feedback (2 and 3).

Leading Behavior Development Apart from direct triggers for reconsidering leadership roles, we looked at how the level of leadership that participants expressed developed over the four different rounds in which the task was executed. Three different dimensions of behavior (leash activity, participant location relative to the robot, participant movement) were looked at separately. We found that for all these dimensions, six types of leadership behavior development could be observed, namely:

- Mostly following (a);
- Start off following, leading in the middle, following at the end (b);
- Start off following, increase of leading over time (c);
- Start off leading, increase of following over time (d);
- Start off leading, following in the middle, leading at the end (e);
- Mostly leading (f).

We categorized each dimension of behavior (leash activity, participant movement and participant location) into one of those types of leadership behavior development for each participant. This resulted in a very wide distribution of behavior, showing that participants engaged in highly personal ways of dealing with leadership roles and shifts in the context of the task. While many participants could be categorized in the same type of behavior for at least two of the dimensions (meaning that participants themselves behaved relatively consistently), the pattern of combined dimensions was unique

	Leash activity	Participant location relative to robot	Participant movement
Mostly following (a)	4 (n = 1)	4, 15 (n = 2)	15, 11 (n = 2)
Start off following, leading in the middle, following at the end (b)	13 (n = 1)	13 (n = 1)	18, 13 (n = 2)
Start off following, increase of leading over time (c)	14, 1 (n = 2)	7, 14, 10 (n = 3)	2, 12, 14, 1, 10 (n = 5)
Start off leading, increase of following over time (d)	3, 16, 7, 18, 15, 11 (n = 6)	3, 18 (n = 2)	8, 6, 3, 16, 7, 4 (n = 6)
Start off leading, following in the middle, leading at the end (e)	5, 12 (n = 2)	5, 12, 16, 1 (n = 4)	5 (n = 1)
Mostly leading (f)	9, 17, 2, 8, 6, 10 (n = 6)	9, 17, 2, 8, 6, 11 (n = 6)	9, 17 (n = 2)

Table 2.5: An overview of the distribution of participants across all six behavior development types for each behavior dimension (leash activity, participant location and participant movement). Each number represents a participant.

for almost every participant. For a distribution of participants across the behavior development types, see Table 2.5. To understand how these types of behavior relate to task performance, we created a boxplot of the task performance related to each category of behavior development, using the categorization based on leash activity (Figure 2.4). Given the small number of participants, it is impossible to draw any hard conclusions from this (especially about category (a) and (b), as only one participant was categorized in either of those). Realistically, only (d) and (f) provide relevant information since both these categories contain 6 participants; it is interesting to see that in this case, the category that is more balanced (d) indeed scores better than the category in which participants were strongly leading all the time (f).

Subjective Collaboration Fluency and Leadership Roles Besides behavioral data, subjective Collaboration Fluency was also measured after each round of performing the task using a questionnaire, based on (Hoffman, 2019). We found that the score on this questionnaire increases significantly over time. This effect was visible within three runs of performing the task. This means that regardless of how people behave, the way in which participants interacted with the robot enabled them to develop a more fluent collaboration over time.

We also found that there was a weak (but significant) negative correlation between the Collaboration Fluency score and the amount of leading behavior people expressed through the leash and movement. This means that when participants were less willing to follow the robot, they also regarded the robot as less cooperative.




Figure 2.4: An overview of the task performance of participants per category. For each participant, the average score of the four rounds was calculated. The categorization is based on which category participants were in when looking at their leash activity only: (a) mostly following, (b) start off following, leading in the middle, following at the end, (c) start off following, increase of leading over time, (d) start off leading, increase of following over time, (e) start off leading, following in the middle, leading at the end, (f) mostly leading.

2.4.4. ANALYZING BEHAVIOR TO UNCOVER INTERACTION PATTERNS

THE fact that participants were able to develop a fluent collaboration with the robot, while still showing a wide variety of behaviors, prompted us to have a closer look at the specific adaptive behaviors and interactions that emerged in this task. In the following sections, we will explain in more detail how we approached this further analysis as well as the results.

DATA ANALYSIS: EXTRACTING INTERACTION PATTERNS

Using visualizations of the video coding, we studied the videos in more detail, paying specific attention to moments at which adaptations took place (moments at which the codes switched from a leading to a following code for example). We described the specific interactions that we observed at such moments, as well as the interactions of what happened in between those moments. In this process, we tried to focus on the smallest relevant unit of interaction (we will refer to these as unit interactions later in the paper). If it was unclear what a participant was doing at a specific instance, we looked at transcriptions of the interview questions to be able to reliably interpret the intention behind their actions.

The resulting list of interactions was categorized by manually clustering them, after which we described all these different interactions using more general concepts. This process can again be seen as another iteration of open coding: we carefully read each observed interaction and created a code (or sometimes a few codes) to describe the interaction. Within this process we tried to use similar words as much as possible, to keep

the list of codes as short as possible. With this process, we aimed to make the interactions less dependent on the specific context of the task executed in the experiment, and more generally applicable. Such more generally applicable interactions are usually called patterns in literature (Diggelen et al., 2019), and are often used as reference for designing human-technology interactions across different contexts. Important to note here is that patterns are not completely generalizable; they are part of a category of concepts that are called ‘intermediate-level knowledge’ (Höök and Löwgren, 2012). They are more abstracted than a single instance, but are not as generalizable as a theory. Their value comes specifically from the fact that they are relatively close to an actual context and task, while being applicable to a range of task and contexts. We will call the more generalized versions of the observed interactions *interaction patterns*. Besides a specification of these interaction patterns themselves, we have tried to combine them into sequences to create larger interaction patterns. Also, we have specified how certain interaction patterns related to others. The combination of the set of interaction patterns (the interaction vocabulary) and the details on how they can be combined and relate to each other will be referred to as the *pattern language*.

2.5. RESULTS: INTERACTION PATTERNS

BELOW, we will describe in detail the outcomes of the analysis of the interactions and interaction patterns. As mentioned above, we will describe exactly what interactions were extracted from the video data, how they were categorized and generalized into interaction patterns and how they can be combined into larger sequences.

2.5.1. OBSERVED INTERACTIONS

BY analyzing the videos of the collaboration between the human participants and the robot, a list of 34 types of interactions could be distinguished. They are the unit interactions: the smallest relevant co-adaptive interactions than can be described within the context of the experiment. These interactions were categorized in four types:

- Stable situations (10 interactions): interactions observed in-between adaptations, such as the interaction of the human leading and the interaction of the robot leading.
- Sudden adaptations (17 interactions): interactions in which the human and/or robot adapted their leader or follower role, therefore starting a transition from one stable situation to another. The adaptation happens in a single moment, over a short period of time, often in response to an event in the task or partner.
- Gradual adaptations (5 interactions): interactions in which the human and/or robot adapted their leader or follower role, therefore starting a transition from one stable situation to another. The adaptation happens gradually over a longer period of time, often in response to a newly hypothesized or discovered property of the partner’s behavior.
- Active negotiations (2 interactions): interactions in which there was a sequence of several short adaptations that eventually also lead to a transition between stable situations.

The full list of these observed interactions and their categorization can be found in Table A.1, but some examples are the following:

- Stable situations: ‘Human speeds forward dragging the robot along’, or ‘robot is in the lead but human actively runs around robot taking into account the route’.
- Sudden adaptations: ‘Human changes direction, thereby loosening the leash, setting the robot free’, or ‘human pulls the leash and moves to the goal when getting close to the goal’.
- Gradual adaptations: ‘Gradually the robot leads more’.
- Active negotiations: ‘Alternating pulling the robot in a specific direction, waiting for the robot to go, then following the robot’.

The behavior of all participants in the experiment can be described as sequences of these unit interactions, thereby generating larger and higher level interactions.

2.5.2. INTERACTION PATTERNS FOR ADAPTIVE LEADER-FOLLOWER BEHAVIOR

THE above described interactions are specifically related to the experimental task. In order to be able to apply them to other contexts, it is necessary to describe them in more general terms. Therefore, we formulated them into general interaction patterns that can appear in any human-robot collaboration where leader-follower dynamics are relevant. Appendix A shows how the observed interactions were described with interaction patterns. Important to note here is that some of the more complex interactions were described using two or three interaction patterns, while some interaction patterns were also used to describe more than one observed interaction. Table 2.6 presents a list of the resulting more generalized interaction patterns, including their category and a short description.

The relatively long list of sudden adaptations contains a diversity of interaction patterns. Some of them are triggers for adaptation (e.g., ‘unexpected action by a team member’), while others are outcomes (e.g., ‘team member stops with what they’re doing, waits’). After a closer look at the list we believe that four components can be distinguished within these sudden adaptations:

- External trigger: an event outside of the partner (e.g., in the task, environment or other partner) triggers an adaptation to a new stable situation;
- Internal trigger: an event inside of the partner (e.g., a specific expectation or change of mind) triggers an adaptation to a new stable situation;
- Outcome: a specific action that is preceded by an internal or external trigger, that will gradually develop into a new stable situation afterward;
- In-between-situation: a specific action that is preceded by an internal or external trigger, that serves as a new trigger for adapting to a new stable situation afterward.

To understand how combinations of these components constitute an interaction pattern, each interaction pattern has been described using the above components in Table 2.7.

Using the extended description of the interaction patterns, we can create sequences of them to describe and analyze behavior that participants showed in the experiment. Examples of those are shown in Figure 2.5. The sequences shown in the figure all represent behavior that participants showed at a specific point in the task. For example, the top sequence is behavior shown by participant 14 in round 2. They were following the robot to pick up the object (stable situation, following). At some point, they were approaching the goal (the robot was also moving toward the goal), which triggered the participant to try to take over the robot's task by further exploring the field for objects (sudden adaptation, being close to finishing the task and trying to finish the other's task when the other is done). To urge the robot to follow, the participant pulled the leash in short intervals, but as the robot had already collected all objects, it would continue to move to the goal when the leash was loose (active negotiation, executing leading in short intervals). This resulted in the participant giving in and they again followed the robot (stable situation, following). Another interesting example is the sequence from participant 5, shown in round 4. The participant was focused on reaching the goal (stable situation, leading), when the robot drove over the participant's feet in an attempt to move with the participant (sudden adaptation, partner-interfering mistake). This caused the participant to immediately take over the robot's task by exploring the field for objects themselves (stable situation, taking over the other's task).

From these examples, it can be seen that sometimes different stable situations can exist at the same time to form more complex behavior. Also, different adaptations can happen after each other before a new stable situation is reached. This usually happens when a sudden adaptation is described as an outcome or an in-between-situation. Using sequences of interaction patterns of varying lengths, we can look at the dynamics of co-adaptation at different levels of complexity. This allows us to analyze the effect that small, short-term adaptations have on the overall development of leader-follower roles, but also to dissect large sequences of observed behavior into small units. An explanatory overview of how the observations translate into sequences of interaction patterns is given in a video in Supplementary Material B of van Zoelen et al., 2021.

Category	Concept	Description
Stable situation	Human following	Human lets the robot do its task
	Human actively on top of things, actively supervising	Human is constantly in touch with the robot
	Active observation by human	Human is actively observing what the robot is doing
	Human leading	Human leads the robot
	Human executing the robot's task	Human executes the task of the robot

Sudden adaptation

Proactive following by human	Human actively predicts and observes what the robot will do, following their course of action
Human dragging the robot along while doing all the work, the robot is a burden	Human ignores the robot as much as possible while focusing on completing the task
Human focuses on their own task, but leaving time for the robot to catch up	Human executes their own task while leaving space for the robot to follow them in that course of action
Unexpected action by a robot team member	The robot does something the human did not expect, possibly triggering a leadership shift
Human waiting for the robot to finish their task	The human waits for the robot to finish their task, and decides on a leadership role after that
Human trying to finish the robot's task when the robot is done	When the robot has finished their task, the human takes over the task to see if it can be improved upon
Partner-interfering mistake	The robot makes a mistake that directly and strongly interferes with the human's course of action
Human losing contact with the robot due to focus on own task	The human focuses very much on their own task, therefore lose contact with the robot
Being close to finishing the task	The team is very close to finishing the task, which possibly triggers a leadership shift
Human actively making up for the robot's limitations	The human foresees a limitation of the robot will hinder their performance, therefore undertakes action to avoid that
Task achievement	A task achievement is reached, possibly triggering a leadership shift
Human urging the robot to be more active, 'come on'	The robot is relatively passive, causing the human to actively urge the robot to be more active
Human stops with what they're doing, waits	The human suddenly stops with what they are doing to wait, after which a new leadership role is chosen
Repeating previous behavior patterns	The human recognizes a situation similar to an earlier situation, and repeats the behavior previously executed

	Human recognizing the autonomy of the robot	The human recognizes the autonomous capabilities of the robot, possibly triggering a leadership shift
	Quick response to leadership shifts due to continuous connection	Due to continuous contact between the team members, a leadership shift initiated by one team member is quickly and smoothly followed by the other
	Robot becomes active after being inactive	After a period of waiting of being inactive, the robot suddenly becomes active again, possibly triggering a leadership shift
Gradual adaptation	Human gradually letting the robot do more	The human gradually lets the robot do more over time
	Human learning to predict the robot's behavior	Over time, the human gradually gains insight into the robot's behavior, thereby enabling them to better predict their behavior
	Human trying to regain control in different ways until eventually taking the lead	Over time, the human attempts to take the lead and regain control in different ways, to eventually find a way to keep taking the lead
Active negotiation	Human executing leading in short intervals	The human takes the lead several times in short intervals, observing what the robot does in the following intervals, to actively search for and negotiate a new stable situation

Table 2.6: The interaction patterns identified from the behavioral data, including a description of what they entail.

2.6. CONCLUSION AND DISCUSSION

W^E We have studied the process of co-adaptation within the context of human-robot collaboration. We focused on the adaptations that emerge within the team as a result of interactions around dynamic leadership roles and complementary capabilities. An embodied approach was taken to study subtle and unconscious interactions that manifest themselves in observable physical behavior. We believe that the design of our experiment provides a different way of looking at HRI; one imposes little assumptions about interactions on the design, and that allows for natural interactions based on affordances. In the sections below, we will go into more detail on how the different aspects of our results can be of use for future HRI research and design.

Interaction Pattern	Type of sudden adaptation
Unexpected action by a robot team member	External trigger
Human waiting for the robot to finish their task	In-between situation, preceded by trigger of the other partner working on a specific subtask, succeeded by a new stable situation
Human trying to finish the robot's task when the robot is done	External trigger and outcome
Partner-interfering mistake	External trigger
Human losing contact with the robot due to focus on own task	Internal trigger and outcome
Being close to finishing the task	External trigger, followed by any outcome
Human actively making up for the robot's limitations	Internal trigger (expectations) and outcome
Task achievement	External trigger
Human urging the robot to be more active, 'come on'	Outcome, preceded by trigger of the other being inactive
Human stops with what they're doing, waits	Outcome, preceded by any trigger
Repeating previous behavior patterns	Outcome, preceded by internal trigger
Human recognizing the autonomy of the robot	In-between-situation, preceded by external trigger (behavior of the other), succeeded by a new stable situation
Quick response to leadership shifts due to continuous connection	In-between-situation, preceded by any trigger, succeeded by a new stable situation
Robot becomes active after being inactive	Outcome and internal trigger

Table 2.7: The interaction patterns that fall in the category of sudden adaptations described in more detail.




Figure 2.5: Several example sequences of interaction patterns as they appeared in the experiment.

2.6.1. INTERACTION PATTERNS AND TEAM DESIGN PATTERNS

WE have extracted a list of interaction patterns from observed human-robot team behavior. The idea of describing human-robot or human-agent team behavior with patterns has been explored before, such as in Diggelen et al., 2019; van Diggelen and Johnson, 2019; van der Waa et al., 2020, under the name ‘Team Design Patterns.’ In existing research, it is described how these patterns can be useful for designers of human-robot teams, as well as for the actual team members to recognize what activities they are engaged in. These existing pattern languages are generally created in a top-down approach. While Diggelen et al., 2019 mention that Team Design Patterns can emerge from interactions between the human(s) and agent(s) in the team, the pattern languages described in van Diggelen and Johnson, 2019; van der Waa et al., 2020 are still designed by the authors of the paper, although the design process is not described

in detail. We deliberately use a different name to describe the patterns in our pattern language (interaction patterns instead of team design patterns), because our interaction patterns have not been designed. Rather, they were extracted from existing observations, while they emerged naturally from the context of the human-robot team task. While the Team Design Patterns are very useful, we believe that it is important to also study interactions in human-robot teams in a bottom-up manner, to represent the processes that occur naturally within teams when members collaborate in the real world. The embodied approach of our study enabled us to generate a new interaction pattern language that is based completely on empirical data. It describes the interaction patterns as an emergent feature, while we attempted to keep our own projections of human-only team interactions out of the analysis. Therefore, they can be used as a library of existing natural interactions when designing human-robot interactions; they provide pointers for what natural and fluent co-adaptive HRI can look like.

The approach of studying embodied interactions in a natural setting, and the development of a language to interpret the observed interactions, enabled us to identify the interaction patterns that underly the co-adaptation processes taking place within a team. The interaction patterns can be used in other contexts, other tasks, and other teams, due to our efforts to describe them in a way that is as context-independent as possible. This positions our work as an addition to the work of other researchers (van Diggelen and Johnson, 2019; van der Waa et al., 2020) who study team behavior at a higher level of abstraction, that is more focused on team composition and task division. In the design of and research into human-robot or human-agent teams, both types of pattern languages can be used in different stages of the process. The high-level pattern languages can be used for deciding on team composition and general collaborative interactions, while the elements in the lower level pattern language we describe can serve as pointers for designing the specific detailed interactions between the team members that elicit or support effective team behavior.

2.6.2. INTERACTION PATTERN LANGUAGE

THE interaction patterns that we have described show that leader-follower dynamics can be described using the concepts of stable situations, sudden adaptations, gradual adaptations, and active negotiations. They give us a better understanding of the subtleties in leader-follower dynamics: very often it is not so much a matter of leading or following, but a bit of both: leadership roles constantly shift, and very often leadership is divided across the team members. The complete pattern language, consisting of interaction patterns as the vocabulary and the connections between them as grammar, provides a framework for analyzing co-adaptive interactions in human-robot collaborations, also in contexts different from the one used in our experiment. Using our pattern language to describe interactions can make it easier to understand why specific role divisions emerge and what can be done to change them.

Moreover, the pattern language can be used by collaborating humans and robots for when they want to communicate about the interactions they are engaged in. The different concepts described by the pattern language can for example be used in a knowledge base for the robot (e.g., in the form of an ontology). This can support the team members in becoming aware of their current leadership roles and possible developments in those

roles, to give them more agency in making strategic decisions about the collaboration.

2.6.3. RELATION TO EXISTING INTERACTION TAXONOMIES

OUR pattern language shows similarities to the interaction taxonomy described in Madan et al., 2015. More specifically, their description of harmonious interactions is similar to what we consider stable situations, while their description of conflicting interactions has overlap with our sudden adaptations and active negotiations. Our pattern language therefore partly confirms, but also extends their interaction taxonomy. We provide a more detailed description and categorization of their conflicting interactions, by expressing the difference between sudden adaptations and active negotiations, and by also adding gradual adaptations. Related to this, we feel that the term adaptation is more encompassing than conflict, as not all adaptive interactions within these categories come from a directly observable conflict. Moreover, we provide a detailed and task-independent description of the different types of sudden adaptations. The extensions originate from the fact that we explicitly focused on interactions that drive co-adaptation, rather than collaborative interaction in general. Moreover, through our extended description of sudden adaptations, we provide information on how different interaction patterns relate to each other (i.e. the ‘grammar’ of our pattern language), where in the work of Madan et al., 2015 only the taxonomy is provided (i.e. the ‘vocabulary’). Our interaction patterns are also more detailed than those presented in the existing literature. They are described in such a way that they can also be used to design interactions, rather than to just analyze them.

In terms of the lower level interaction patterns, both the work of Madan et al., 2015 and our work are to some extent related to the task used to obtain them. Their interaction patterns were generated in the context of collaborative object manipulation, while ours were generated within a collaborative navigation context. We, however, explicitly formulated the interaction patterns in such a way that they are generally applicable outside of this initial context. To understand the extent of their generalizability, further evaluation in other task contexts will be useful.

2.6.4. LIMITATIONS

WHILE the list of interaction patterns is quite extensive, it is probably not complete. The specific task context that we used in our experiment of course limits the kind of interactions possible. Also, while the analysis of the data was done in a systematic manner, it is bounded by the frame of reference of the researcher. In order to obtain evidence for the relevance of the proposed language, it is important to attempt to apply the analysis of interaction patterns used here to other tasks. That will provide more insights into the extent of the generalizability of the pattern language, as well as into necessary extensions or adjustments.

Furthermore, there are a few limitations forthcoming from the manner in which the task in the experiment was executed. We claim to study human-robot teamwork, but in our experiment a human operator controlled the robot following pre-configured rules. It may be that the robot behaved different from how a real robot would behave. Moreover, the participants were aware of the fact that the robot was controlled by a human operator, and even though a pilot study showed us that participants did not pay much

attention to the operator, it may have still influenced the interactions that emerged. The task was also defined with a relatively low level of agency of the robot, causing the robot to initiate few adaptive behaviors. It is likely that participants noticed this, therefore it might have influenced their initiative to take or delegate leadership. Moreover, we studied a human-robot team in the form of a dyad, whereas the dynamics of team interactions can be very different for other (larger) team compositions. This again stresses the importance of testing the results of the present study in other tasks and contexts and, if possible, with real robots and different team compositions. Outcomes of such studies will help to elaborate and refine the interaction pattern language, eventually enabling a better understanding of co-adaptation in human-robot teams. This, in turn, will support the design of adaptive human-robot teams that are able to operate successfully in the complexity of the real world.

2.6.5. FINAL CONCLUSION

BY observing embodied interactions within a human-robot team, we have extracted an interaction pattern language that can be used to describe co-adaptive behavior. This pattern language consists of a list of interaction patterns (the vocabulary) that together make up the different elements of co-adaptation. The interaction patterns can be categorized into stable situations, sudden adaptations, gradual adaptations and active negotiations. Furthermore, the sudden adaptations are built up of external triggers, internal triggers, outcomes and in-between-situations. These categorizations and concepts can be used to link different interaction patterns together, to make sequences of co-adaptive behavior. They can therefore be seen as the grammar of our pattern language.

In future studies, we will use the pattern language to analyze co-adaptive behavior in different tasks and contexts. We will analyze how the presence of certain interaction patterns influences team behavior and performance, to validate how useful the different patterns are in creating successful human-robot teams that make use of fluent co-adaptations.

BIBLIOGRAPHY

- Azevedo, C. R. B., Raizer, K., & Souza, R. (2017). A vision for human-machine mutual understanding, trust establishment, and collaboration. *2017 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA)*, 1–3. <https://doi.org/10.1109/COGSIMA.2017.7929606>
- Bach-y-Rita, P., & W. Kercel, S. (2003). Sensory substitution and the human-machine interface. *Trends in Cognitive Sciences*, 7(12), 541–546. <https://doi.org/10.1016/j.tics.2003.10.013>
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., & Mordatch, I. (2020). Emergent Tool Use From Multi-Agent Autocurricula [arXiv: 1909.07528]. *arXiv:1909.07528 [cs, stat]*. Retrieved December 10, 2020, from <http://arxiv.org/abs/1909.07528>
- Breazeal, C., Dautenhahn, K., & Kanda, T. (2016). Social Robotics. In B. Siciliano & O. Khatib (Eds.), *Springer Handbook of Robotics* (pp. 1935–1972). Springer International Publishing. https://doi.org/10.1007/978-3-319-32552-1_72
- Burke, S., Stagl, K., Salas, E., Pierce, L., & Kendall, D. (2006). Understanding Team Adaptation: A Conceptual Analysis and Model. *The Journal of applied psychology*, 91, 1189–207. <https://doi.org/10.1037/0021-9010.91.6.1189>
- Buschmeier, H., & Kopp, S. (2013). Co-constructing Grounded Symbols - Feedback and Incremental Adaptation in Human-Agent Dialogue. *KI - Künstliche Intelligenz*, 27, 137–143. <https://doi.org/10.1007/s13218-013-0241-8>
- Charmaz, K. (2014, March). *Constructing Grounded Theory* [Google-Books-ID: v_GGAwAAQBAJ]. SAGE.
- Chauncey, K., Harriott, C., Prasov, Z., & Cunha, M. (2017). A Framework for Co-adaptive Human-Robot Interaction Metrics. *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 123–124.
- Diggelen, J. v., Neerincx, M., Peeters, M., & Schraagen, J. M. (2019). Developing Effective and Resilient Human-Agent Teamwork Using Team Design Patterns [Conference Name: IEEE Intelligent Systems]. *IEEE Intelligent Systems*, 34(2), 15–24. <https://doi.org/10.1109/MIS.2018.2886671>
- Duffy, B., & Joue, G. (2000). Intelligent Robots: The question of embodiment. Retrieved December 11, 2020, from [/paper/INTELLIGENT-ROBOTS%3A-THE-QUESTION-OF-EMBODIMENT-Duffy-Joue/bf3584b68801a88943be738519abf81227aebc32](https://arxiv.org/abs/1909.07528)
- Ehrlich, S. K., & Cheng, G. (2018). Human-agent co-adaptation using error-related potentials. *Journal of Neural Engineering*, 15(6), 066014. <https://doi.org/10.1088/1741-2552/aae069>
- Evrard, P., & Kheddar, A. (2009). Homotopy switching model for dyad haptic interaction in physical collaborative tasks. *World Haptics 2009 - Third Joint EuroHaptics conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 45–50. <https://doi.org/10.1109/WHC.2009.4810879>

- Flemisch, F., Kelsch, J., Löper, C., Schieben, A., & Schindler, J. (2008). Automation spectrum, inner / outer compatibility and other potentially useful human factors concepts for assistance and automation. *Human Factors for assistance and automation*, 1–16.
- Foerster, J., Assael, I. A., de Freitas, N., & Whiteson, S. (2016). Learning to Communicate with Deep Multi-Agent Reinforcement Learning. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 29* (pp. 2137–2145). Curran Associates, Inc. Retrieved February 12, 2019, from <http://papers.nips.cc/paper/6042-learning-to-communicate-with-deep-multi-agent-reinforcement-learning.pdf>
- Gao, Y., Barendregt, W., & Castellano, G. (2017). Personalised Human-Robot Co-Adaptation in Instructional Settings using Reinforcement Learning. *Proceedings of the Persuasive Embodied Agents for Behavior Change (PEACH2017) Workshop*.
- Herrera Perez, C., & Barakova, E. I. (2020). Expressivity Comes First, Movement Follows: Embodied Interaction as Intrinsically Expressive Driver of Robot Behaviour. In N. Noceti, A. Sciutti, & F. Rea (Eds.), *Modelling Human Motion: From Human Perception to Robot Design* (pp. 299–313). Springer International Publishing. https://doi.org/10.1007/978-3-030-46732-6_14
- Hoffman, G. (2019). Evaluating Fluency in Human–Robot Collaboration. *IEEE Transactions on Human-Machine Systems*, 49(3), 209–218. <https://doi.org/10.1109/THMS.2019.2904558>
- Höök, K., & Löwgren, J. (2012). Strong concepts: Intermediate-level knowledge in interaction design research. *ACM Transactions on Computer-Human Interaction*, 19(3), 18. Retrieved April 19, 2021, from <https://dl.acm.org/tudelft.idm.oclc.org/doi/abs/10.1145/2362364.2362371>
- Iqbal, S., & Sha, F. (2019). Actor-Attention-Critic for Multi-Agent Reinforcement Learning. *Proceedings of the 36th International Conference on Machine Learning*, 10.
- Jarrassé, N., Charalambous, T., & Burdet, E. (2012). A Framework to Describe, Analyze and Generate Interactive Motor Behaviors (M. O. Ernst, Ed.). *PLoS ONE*, 7(11), e49945. <https://doi.org/10.1371/journal.pone.0049945>
- Kaczmarek, K. A., Webster, J. G., Bach-y-Rita, P., & Tompkins, W. J. (1991). Electrotactile and vibrotactile displays for sensory substitution systems. *IEEE Transactions on Biomedical Engineering*, 38(1), 1–16. <https://doi.org/10.1109/10.68204>
- Kiela, D., Bulat, L., Vero, A. L., & Clark, S. (2016). Virtual Embodiment: A Scalable Long-Term Strategy for Artificial Intelligence Research [arXiv: 1610.07432]. *arXiv:1610.07432 [cs]*. Retrieved December 11, 2020, from <http://arxiv.org/abs/1610.07432>
- Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems*, 19(6), 91–95. <https://doi.org/10.1109/MIS.2004.74>
- Lagerstedt, E., & Thill, S. (2020). Benchmarks for evaluating human-robot interaction: Lessons learned from human-animal interactions [ISSN: 1944-9437]. *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 137–143. <https://doi.org/10.1109/RO-MAN47096.2020.9223347>

- Li, Y., Tee, K. P., Chan, W. L., Yan, R., Chua, Y., & Limbu, D. K. (2015). Continuous Role Adaptation for Human–Robot Shared Control. *IEEE Transactions on Robotics*, 31(3), 672–681. <https://doi.org/10.1109/TRO.2015.2419873>
- Madan, C. E., Kucukyilmaz, A., Sezgin, T. M., & Basdogan, C. (2015). Recognition of Haptic Interaction Patterns in Dyadic Joint Object Manipulation [Conference Name: IEEE Transactions on Haptics]. *IEEE Transactions on Haptics*, 8(1), 54–66. <https://doi.org/10.1109/TOH.2014.2384049>
- Mohammad, Y., & Toyoaki Nishida. (2008). Human adaptation to a miniature robot: Precursors of mutual adaptation. *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, 124–129. <https://doi.org/10.1109/ROMAN.2008.4600654>
- Mörthl, A., Lawitzky, M., Kucukyilmaz, A., Sezgin, M., Basdogan, C., & Hirche, S. (2012). The role of roles: Physical cooperation between humans and robots. *The International Journal of Robotics Research*, 31(13), 1656–1674. <https://doi.org/10.1177/0278364912455366>
- Nagel, S. K., Carl, C., Kringe, T., Martin, R., & König, P. (2005). Beyond sensory substitution—learning the sixth sense. *Journal of Neural Engineering*, 2(4), R13–R26. <https://doi.org/10.1088/1741-2560/2/4/R02>
- Nikolaïdis, S., Hsu, D., & Srinivasa, S. S. (2017). Human-robot mutual adaptation in collaborative tasks: Models and experiments. *I. J. Robotics Res.*, 36, 618–634. <https://doi.org/10.1177/0278364917690593>
- Nikolaïdis, S., Zhu, Y. X., Hsu, D., & Srinivasa, S. (2017). Human-Robot Mutual Adaptation in Shared Autonomy. *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI '17*, 294–302. <https://doi.org/10.1145/2909824.3020252>
- Noldus, L. P. J. J. (2019, November). The Observer XT. <https://www.noldus.com/observer-xt>
- Sacheli, L. M., Tidoni, E., Pavone, E. F., Aglioti, S. M., & Candidi, M. (2013). Kinematics fingerprints of leader and follower role-taking during cooperative joint actions | SpringerLink. *Experimental Brain Research*, 226(4), 473–486. Retrieved June 12, 2019, from <https://link.springer.com/article/10.1007/s00221-013-3459-7>
- Sasagawa, A., Fujimoto, K., Sakaino, S., & Tsuji, T. (2020). Imitation Learning Based on Bilateral Control for Human–Robot Cooperation [Conference Name: IEEE Robotics and Automation Letters]. *IEEE Robotics and Automation Letters*, 5(4), 6169–6176. <https://doi.org/10.1109/LRA.2020.3011353>
- Sciutti, A., Mara, M., Tagliasco, V., & Sandini, G. (2018). Humanizing Human-Robot Interaction: On the Importance of Mutual Understanding. *IEEE Technology and Society Magazine*, 37(1), 22–29. <https://doi.org/10.1109/MTS.2018.2795095>
- Uitdewilligen, S., Waller, M. J., & Pitariu, A. H. (2013). Mental Model Updating and Team Adaptation [Publisher: SAGE Publications Inc]. *Small Group Research*, 44(2), 127–158. <https://doi.org/10.1177/1046496413478205>
- van Diggelen, J., & Johnson, M. (2019). Team Design Patterns. *Proceedings of the 7th International Conference on Human-Agent Interaction*, 118–126. <https://doi.org/10.1145/3349537.3351892>

- van Zoelen, E. M., van den Bosch, K., Rauterberg, M., Barakova, E., & Neerincx, M. (2021). Identifying Interaction Patterns of Tangible Co-Adaptations in Human-Robot Team Behaviors. *Frontiers in Psychology*, 12, 645545. <https://doi.org/10.3389/fpsyg.2021.645545>
- van der Waa, J., van Diggelen, J., Cavalcante Siebert, L., Neerincx, M., & Jonker, C. (2020). Allocation of Moral Decision-Making in Human-Agent Teams: A Pattern Approach. In D. Harris & W.-C. Li (Eds.), *Engineering Psychology and Cognitive Ergonomics. Cognition and Design* (pp. 203–220). Springer International Publishing. https://doi.org/10.1007/978-3-030-49183-3_16
- Xu, Y., Ohmoto, Y., Okada, S., Ueda, K., Komatsu, T., Okadome, T., Kamei, K., Sumi, Y., & Nishida, T. (2012). Formation conditions of mutual adaptation in human-agent collaborative interaction. *Applied Intelligence*, 36(1), 208–228. <https://doi.org/10.1007/s10489-010-0255-y>
- Yamada, S., & Yamaguchi, T. (2002). Mutual adaptation to mind mapping in human-agent interaction. *Proceedings. 11th IEEE International Workshop on Robot and Human Interactive Communication*, 105–110. <https://doi.org/10.1109/ROMAN.2002.1045606>
- Youssef, K., De Silva, P. R. S., & Okada, M. (2014). Sociable Dining Table: Incremental Meaning Acquisition Based on Mutual Adaptation Process. In M. Beetz, B. Johnston, & M.-A. Williams (Eds.), *Social Robotics* (pp. 206–216, Vol. 8755). Springer International Publishing. https://doi.org/10.1007/978-3-319-11973-1_21
- Zoelen, E. M. v., Barakova, E. I., & Rauterberg, M. (2020). Adaptive Leader-Follower Behavior in Human-Robot Collaboration [ISSN: 1944-9437]. *29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 1259–1265. <https://doi.org/10.1109/RO-MAN47096.2020.9223548>

3

BECOMING TEAM MEMBERS: IDENTIFYING INTERACTION PATTERNS OF MUTUAL ADAPTATION FOR HUMAN-ROBOT CO-LEARNING

Becoming a well-functioning team requires continuous collaborative learning by all team members. This is called co-learning, conceptualized in this paper as comprising two alternating iterative stages: partners adapting their behavior to the task and to each other (co-adaptation), and partners sustaining successful behavior through communication. This paper focuses on the first stage in human-robot teams, aiming at a method for the identification of recurring behaviors that indicate co-learning. Studying this requires a task context that allows for behavioral adaptation to emerge from the interactions between human and robot. We address the requirements for conducting research into co-adaptation by a human-robot team, and designed a simplified computer simulation of an urban search and rescue task accordingly. A human participant and a virtual robot were instructed to discover how to collaboratively free victims from the rubbles of an earthquake. The virtual robot was designed to be able to real-time learn which actions best contributed to good team performance. The interactions between human participants and robots were recorded. The observations revealed patterns of interaction used by human and robot in order to adapt their behavior to the task and to one another. Results therefore show that our task environment enables us to study co-learning, and suggest that more participant adaptation improved robot learning and thus team level learning. The identified interaction patterns can emerge in similar task contexts, forming a first description and analysis method for co-learning. Moreover, the identification of interaction patterns support awareness among team members, providing the foundation for human-robot communication about the co-adaptation (i.e., the second stage of co-learning). Future research will focus on these human-robot communication processes for co-learning.

 **Emma M. van Zoelen**, Karel van Den Bosch, Mark Neerinx. 2021. Becoming team members: Identifying interaction patterns of mutual adaptation for human-robot co-learning. In *Frontiers in Robotics and AI*, 8, 692811.

3.1. INTRODUCTION

WHEN people collaborate in teams, it is of key importance that all team members get to know each other, explore how they can best work together, and eventually adapt to each other and learn to make their collaboration as fluent as possible. While humans do this naturally (Burke et al., 2006), it is not self-evident for robots that are intended to function as team partners in human-robot collaborations. It is well known that robotic team partners should be transparent, predictable, and explainable, but it is often overlooked that human team partners become predictable and explainable through a process of exploration and mutual learning.

We call the above mentioned process co-learning (van den Bosch et al., 2019). While existing work on human-robot collaboration and mutual adaptivity often focuses on short-term single interactions, we believe it is necessary to also look at repeated interactions to study co-learning as a mechanism for building fluent human-robot collaborations. We conceptualize co-learning as comprising two alternating iterative stages. In the *first stage*, partners observe each other and adapt their behavior to the other, leading to successful emergent team behaviors. Such adaptation can be done deliberately but often occurs implicitly and unconsciously. In the *second stage*, partners communicate about their adaptations and give each other feedback, thereby giving meaning to and becoming aware of the learned behavior. Especially this second stage of creating awareness of what has been learned helps to sustain the behavioral adaptations over time and across contexts.

We regard co-learning to be vital for creating successful human-robot collaborations. However, the term “co-learning” is relatively new in human-robot interaction literature, and it is not yet precisely defined what human-robot co-learning looks like in practice and how it should be studied. Emergent behavior can only be investigated through empirical studies; to investigate human-robot co-learning it is therefore necessary that both partners can learn in real-time while collaborating with each other. In this study, we therefore chose to empirically study co-learning with a human participant and a Reinforcement Learning (RL) virtual robot. For the investigation of emerging co-adaptive behaviors, we distinguish four main research questions:

1. How to identify and classify recurring behaviors that indicate co-learning in a human-robot team?
2. Which recurring sequences of these behaviors (co-learning patterns) can be identified, such that they can be used by the team partners to communicate about their adaptations?
3. How does the robot's learning, emerging from the interactions, affect the human's behavior and learning?
4. How does the human's learning, emerging from the interactions, affect the robot's behavior and Reinforcement Learning?

The literature on human-robot interaction and learning contains a large body of research on personalized robot tutors, in which a robot tutor learns to personalize its interactions to support the learning process of a human student, focused on classroom or

training related contexts (a few examples are Baxter et al., 2017; Belpaeme et al., 2018; Gao et al., 2017; Vignolo et al., 2019). Formal training is important to initiate learning and to steer development in the right direction. However, it is important to realize that learning continues after training has been completed. Every new experience in the task provides an opportunity for human-robot teams to learn from their collaboration. An important aspect of co-learning in actual task contexts is developing and refining (shared) mental models about team members and about the task at hand, to increase mutual understanding of the best way to perform the task (Klein et al., 2004). Therefore, we specifically attempt to answer the above mentioned questions in a task context where learning happens during task execution.

In this paper, we present a behavioral study of how a human and a virtual robot, which uses a Reinforcement Learning algorithm to adapt and optimize its actions, adapt their behavior to collaboratively solve a task. We are interested in how the behavior of the human and the behavior of the robot changes as a result of this process, making our study fit within a new area of research in which both human and machine behavior are assessed (with their mutual dependencies; cf. Rahwan et al., 2019). We first provide an in-depth elaboration of the concept “co-learning” within the context of human-robot collaboration, resulting in a definition of co-learning and the related concepts of co-adaptation and co-evolution. Based on this, we identify the requirements for empirical research into co-adaptation and co-learning, and present the design of an environment for studying co-learning. This environment has been built and used to conduct an empirical study into human-robot co-learning. From the analysis of the observed human-robot interactions, a list of patterns of adaptive interactions, and the switching between these patterns over time, were identified. These patterns can emerge in similar task contexts, forming a first description method for co-learning analyses. Moreover, it supports the creation of awareness, providing the foundation (concepts) for human-robot communication about the co-adaptation (i.e., the second stage of co-learning).

3.2. CO-LEARNING: BACKGROUND AND DEFINITION

COLLABORATIVE learning is a widely studied mechanism in human-only contexts, and it was Dillenbourg (Dillenbourg et al., 1996) who suggested that collaborative learning can also take place between humans and computers. Collaborative learning in the context of Dillenbourg’s work means that learning (the acquisition of new knowledge, skills, behavior, etc.) results from collaborative activities between team partners. If we look at human-robot interaction literature, several terms are used that describe a similar process in which two parties or systems change their behavior and/or mental states concurrently while interacting with each other. Co-adaptation (Chauncey et al., 2017; Nikolaidis, Hsu, and Srinivasa, 2017; Xu et al., 2012) and co-learning (van den Bosch et al., 2019) are two of them, but we also encounter co-evolution (Döppner et al., 2019), in which “co” stands for collaborative, also meaning “mutual.”

Co-adaptation and co-learning are often used interchangeably, making it difficult to understand what they stand for. There are several vision papers explaining the importance of both co-adaptation (e.g. Ansari et al., 2018; Xu et al., 2012) as well as co-learning (e.g. Holstein et al., 2020; van den Bosch et al., 2019; Wenskovitch and North, 2020), but a clear distinction between the two, or a definition specifically for co-learning, is miss-

ing from these papers. When looking at the experimental literature, however, it seems that there are subtle differences. Experimental studies on co-adaptation often focus on making the agent or robot adaptive to the human, using different kinds of information about the human (e.g. Buschmeier and Kopp, 2013; Ehrlich and Cheng, 2018; Sordoni et al., 2015; Yamada and Yamaguchi, 2002). Some studies have investigated how a human adapts in situations in which they collaborate with an intelligent agent or robot. These studies mostly focus on the performance of the human and their resulting subtle behavior change in short experiments (e.g. Mohammad and Toyooki Nishida, 2008; Nikolaidis, Hsu, and Srinivasa, 2017; Nikolaidis, Zhu, et al., 2017). The studies that use “co-learning” tend to take a more symmetrical approach by looking at agent or robot learning as well as human learning, and pay more attention to the learning process and changing strategies of the human as well, often looking at many repetitions of a task (C. Lee et al., 2018; C.-S. Lee et al., 2020; Ramakrishnan et al., 2017; Shafti et al., 2020). Studies on co-evolution, on the other hand, monitor a long-term real-world application in which behavior of the human as well as the robot subtly changes over time (Döppner et al., 2019).

Following these differences, we propose to distinguish the terms using three dimensions, namely 1) the time over which the development takes place, 2) the persistence of the resulting behavior/mental state over time and across contexts, and 3) the intention of the development. Table 3.1 shows the proposed definitions in detail. Within our research, we focus on co-learning as defined here

In a human-robot co-learning process, a human and a robot collaborate on a given task. In order to do well on the task, they need to learn all kinds of implicit and explicit knowledge related to both the task itself as well as the collaboration and interaction between them. Related to the task they can, for example, learn the technical details of how the task should be executed. Related to the collaboration, they can, for example, learn social collaboration skills. Related to both, they can learn about their own role and the role of the other in the task and the consequences of their own and their partner’s actions and mental state on the task (how to collaborate in context of the task). Ultimately, learning this should help them to together perform well on the task, to build understanding of each other in context of the task and to calibrate the trust that the human and the robot have in each other. We focus our work on this last type of combined task and collaboration learning.

We further define co-learning to be comprised of two stages that follow each other in continuous iterative cycles, namely 1) co-adaptation, and 2) a communication process. Part (a) is therefore a process in which team members (sometimes unconsciously) adapt to each other and the task, thereby changing and developing their behavior as a consequence of interactions and an implicit or explicit drive to improve performance or experience (see again Table 3.1). Part (b) is a process in which these implicitly developed behaviors are shared and discussed through direct communications or interactions between team members, thereby making the team members aware of the implicit adaptations. This combination ensures that learned strategies are grounded in the context and task and can be strategically used in new contexts.

	Co-adaptation	Co-learning	Co-evolution
Timespan	Short (seconds — hours)	Medium (hours — weeks)	Long (weeks — years)
Persistence	Developed behavior/mental state does not necessarily persist over time, and probably not at all across contexts	Developed behavior/mental state persists over time and possibly across contexts	Developed behavior/mental state might persist for a while but possibly continues to evolve, similar to the development of habituation
Intention	Changes and developments happen as a consequence of interactions and an implicit or explicit drive to improve performance or experience	Explicitly goal-driven: Attempts to improve performance or experience; learning is an explicit goal	Changes and developments happen as a consequence of interactions and possibly an implicit drive to improve performance or experience

Table 3.1: The concepts co-adaptation, co-learning and co-evolution defined in terms of timespan in which they occur, persistence and intention.

3.3. RESEARCH CHALLENGES

MANY research challenges follow from the conceptualization of co-learning, due to the fact that both human and robot are non-static. They are both constantly developing, changing and adapting, and they influence each other in the process. This means that it is not possible to study only one of the team partners; it is necessary to take a symmetrical approach, where both human and robot are studied through the interactions between them. Moreover, co-learning in dynamic tasks is a continuous process in which new task situations that appear dynamically require new learning over and over again. Therefore, focusing on one specific interaction, or on team performance as end result, does not offer a complete picture. We need to study all interactions that contribute to this process. These specific dynamic properties need to be taken into account in the design of experiments, as well as in the analysis and discussion of results. Following from this, and to provide a broader view on the specific study that we present in this paper, we have defined three research directions that need to be addressed in the study of human-robot co-learning:

1. **Research into enabling and assessing co-learning:** to understand the dynamics of co-learning, we need to investigate what kind of behaviors and interactions drive co-adaptation and co-learning, and how learning processes of human and robot team members influence each other.
2. **Research into interaction patterns that make team partners explicitly aware of learned behavior, such that behavior can be sustained over time and context:** in

order to create sustainable team behavior, human and robot need to communicate about learned behavior to ensure that they are aware of useful learned behavior. It is important to investigate what kind of communication interaction patterns enable this specific type of communication.

3. **Research into a dynamic team mental model that takes into account naturally occurring changes in interaction patterns, and how such a model can support the robot in its learning process:** as humans, we are able to anticipate on the fact that our team members learn and change. It is important to investigate how a dynamic team model can enable robots to also anticipate the fact that their human team member is continuously changing.

The study presented in this paper focuses on the first research direction; the research questions presented in the introduction have been derived from it. More specifically, in the experiment that we describe in the following sections, we have chosen to focus on the first stage of the co-learning process: co-adaptation as a precursor for co-learning. We do not yet address questions concerning communication about learned behaviors (research direction 2), but focus on the implicit behavioral adaptations that occur within a relatively short time span. It is expected that results of the present study will provide pointers for how to investigate the issues associated with research challenges two and three above.

3.4. RESEARCH ENVIRONMENT: DESIGNING TASK, AGENT, CONTEXT

3.4.1. CONTEXT

TO study co-learning in human-robot teams, a suitable task context needs to be designed. We identify the following requirements for such a task context in which we can study co-adaptation according to the definition in Table 3.1:

1. It should be possible for the team to improve its performance by making effective use of the capabilities of the individual team members (as this is necessary to make it a team task (Johnson et al., 2014));
2. There should be possibilities for implicit adaptation and learning for both human and robot team members;
3. It should accommodate different emergent collaborative strategies for solving the task;
4. For this first study, the task and team work should be simple enough for a Reinforcement Learning agent to learn new behavior in a short number of rounds, such that we can study co-adaptation in relatively short experimental sessions;
5. To ensure societal relevance of this research, the task should be based on a real-life domain in which there is a need for autonomous robots that function as team partners.

As a general context for defining a task, we chose Urban Search and Rescue (USAR). A lot of research on human-agent teaming is done in USAR-related tasks (Lematta et al., 2019), because the safety-critical nature makes the application of human-robot collaboration very useful; there are ongoing initiatives which aim to use robots in real USAR teams (requirement 5). Moreover, it is a dynamic task context with many possible subtasks and possibilities for the introduction of threats, safety risks and changing information.

We developed an earthquake scenario for our human-robot USAR team, with the team's task to remove rubble and debris from a victim. To get a better understanding of the task, and of the knowledge and capabilities it requires from partners, we created a storyboard (Figure 3.1). The storyboard shows a possible scenario in which the robot picks up a large rock to clear it away, not realizing that the action may lead to a small rock falling on the head of the victim (Figure 3.1C). When the human notices this, this provides an opportunity to jump in, and to prevent the rock from crushing the victim (Figure 3.1D). This event facilitates the human learning that the robot apparently does not understand the risks of falling rocks. It provides the robot with the opportunity to learn that it made a mistake. The event furthermore provides an opportunity for team members to communicate about the event, their actions, and to plan how they will manage such situations together in the future. This storyboard illustrates that using the unique capabilities of both the human (insight into strategic choices) and robot (physical strength) can be exploited to achieve better task performance (requirement 1). The task of removing rubble from a victim allows for a great diversity in task planning and execution, and for the development of individual strategies (requirement 2 and 3), as the different debris can differ in shape, size and location, while enabling simple basic actions to create strategies for solving the task (requirement 4).

3.4.2. TASK IMPLEMENTATION

WE developed a digital task simulation of the described USAR context using Python and the MATRX package ("MATRX Software", 2021). MATRX is a package for rapid prototyping of human-agent team environments, which supports easy generation of an environment, object and agents. Figure 3.2 shows a screenshot of the simulation. The scene involves three characters: a victim buried underneath a pile of rocks (shown in the middle), an explorer (avatar on the left, played by a human participant) and a Reinforcement Learning robot agent (avatar on the right). The goal of the task is to free the victim by clearing away all rocks that are in front of the victim, as well as to create a pathway to the victim from either the left or the right side. In order to score well, this must happen as quickly as possible, and no additional rocks (or as little as possible) should fall on top of the victim, as that will cause extra harm. Both the human and the virtual robot each have a set of actions they can perform, such as picking up rocks and dropping them somewhere else. However, the extent to which they can perform actions differs: the robot can pick up large and small objects, and break large objects into pieces. Humans can only pick up small objects. Humans however have a better insight in certain aspects of the task that dictate which actions are useful to do, such as how rocks will fall when other rocks are removed or replaced. This insight stems from the fact that humans have "common sense," which helps us understand the probable consequences of actions. In order to complete the task successfully participants must collaborate with the virtual




Figure 3.1: A storyboard describing how a human-robot team might free an earthquake victim from underneath a pile of rocks. In this storyboard, we can see how the robot picks up a large rock, unaware that this will cause another rock to fall on the head of the victim (panel C). The human notices the issue, and steps in to prevent the rock from falling (panel D). This event can help the robot learn about the task, and that it apparently made a mistake. The human can learn about the capabilities of the robot, namely that it didn't understand how the rocks would fall and that it would cause harm.

robot (requirement 1), while managing their actions in such a way that the robot does not accidentally drop rocks on the victim's head. Since it is not clear at the start what the best strategy would be to solve the task quickly, both partners need to learn and adapt as they go (requirement 2). The levels are designed such that there are different possible ways to solve the task (requirement 3), and it is a discrete environment build on a simple state machine, making it possible to design a Reinforcement Learning agent that can process the environment (requirement 4).

3.4.3. LEARNING AGENT

To be able to empirically study how a human and a robot co-adapt while collaborating on a task, the robot should be able to try out and evaluate different actions, to be able to choose the policy that best fits the goals of the team given the adaptations done by the human team member. We chose to use Reinforcement Learning to enable the agent to learn, for three main reasons:

1. The robot had to be able to learn real-time, on the basis of rewards: as the behavior of the human team partner is adaptive and unpredictable, we cannot determine optimal behavior before the start of the task. This means that the best way to find the optimal strategy for solving the task collaboratively would be to get feedback




Figure 3.2: The USAR task environment programmed in MATRX. It shows a victim underneath a pile of rocks, and a human and a robot representing the team members. The dashed red square (above the human's head) represents the hand of the human that can be moved to pick up rocks. The dashed blue square represents the hand of the robot. Scene (A) was used as level 1 in the experiment, while scene (B) was used as level 2.

or rewards on performance.

2. The described task can be solved by a sequence of actions that manipulate the state of the world with each action, therefore, the learning algorithm had to be able to learn a sequence of actions given different sequential task situations.
3. The described task is a human-robot scenario, but it can also conceptually be seen as a multi-agent scenario as both the robot and the human are autonomous and learning agents within the collaboration. Reinforcement Learning is an often used and widely studied mechanism in multi-agent scenarios, for reasons related to reasons one and two above as well (see e.g. Foerster et al., 2016; Kapetanakis and Kudenko, 2002).

Reinforcement Learning has been designed for learning sequences of actions in tasks that can be modeled as Markov Decision Processes (van Otterlo and Wiering, 2012), in which the transitions between states are unknown. In contexts where agents collaborate and learn with a human, these transitions are unknown since it is unknown what the human will do; this is also the case in our task. While such a human-agent collaborative context poses many challenges (e.g. large state spaces, long convergence times and random behavior in the beginning) (Dulac-Arnold et al., 2019), earlier work has shown that RL can be used successfully for learning behavior in real time when interacting with a human, provided that the learning problem is simple enough (Weber et al., 2018). Since we used RL mostly as a tool to ensure that the agent could adapt over time, and not as a goal in itself, we created a RL mechanism that is much simpler than the current state-of-the-art, but that would provide the basic learning that is sufficient for our research goals. We simplified the task by modeling it as a semi-Markov Decision Process (Sutton et al., 1999), which means that the task is divided into several “phases,” which serve as the states in the RL algorithm. Normally states last one timestep, whereas in a semi-Markov Decision Process, these phases can last variable amounts of time. Our state definition describes the state of the environment based on the amount of rocks present in the area around the victim. The state space is defined by $S = (\text{Phase 1, Phase$

Phase	Description
Phase 1	The starting phase: Describes the state of the task environment when no rocks have been moved
Phase 2	The heights of all piles of rocks added up is now at least 10 rocks lower than in phase 1
Phase 3	Phase 2 has been reached, and the heights of all piles of rocks added up is now at least 20 rocks lower than in phase 1
Phase 4	Phase 2 and 3 have been reached, and either there are no more rocks directly on top of the victim, OR one of the sides of the task field is cleared from rocks, meaning there is an access route to the victim from either the left or right side
Goal phase	Phase 2, 3 and 4 have been reached, and there are no more rocks directly on top of the victim, AND one of the sides of the task field is cleared from rocks, meaning there is a free route from either the left or right side to the victim. The task terminates when this phase is reached

Table 3.2: The task conditions specified for each Phase Variable used in the state space of the Reinforcement Learning algorithm.

2, Phase 3, Phase 4, Goal Phase). Table 3.2 describes the details of the individual states. We chose to not explicitly represent learning about collaboration in the learning agent, since we wanted to focus on implicit behavioral adaptations (as explained in Research Challenges). We combined this with a system inspired by the Options Framework (Stolle and Precup, 2002) and a basic greedy Q-learning algorithm. In the Options Framework, agents use RL to learn a meta-policy as well as several “sub-policies.” These sub-policies can also be seen as macro-actions; they are combinations of atomic actions that are used together to solve parts of the task. Usually, these macro-actions are learned in parallel with the meta-policy, but sometimes they are pretrained, such as for example in Illanes et al., 2019. To further simplify the learning problem, we chose to predefine three rule-based macro-actions; the agent could choose from these macro-actions in each phase of the task (a description of each macro-action is given in Figure B.1, Figure B.2, and Figure 3.5). The rewards for the RL algorithm are based on two factors: 1) the time it took the team to move to the next phase, and 2) the amount of additional harm done to the victim. The agent would receive this reward when transitioning into a new phase, or when the task terminates due to becoming unsolvable or due to a timeout. The height of the rewards was made such that the total reward given was always negative. With initial Q-values of 0, this ensured that in the first three runs of the experiment, the agent would try out all three macro-strategies in order, to enforce initial exploration. A visual overview of the learning problem is provided in Figure 3.6, after we have explained more details about the experimental method.

3.4.4. CLAIMS: EXPECTED OBSERVATIONS

WE expect to observe several behaviors within this task environment, given that it was designed to study co-learning behavior. We have formulated these expected




Figure 3.3: A flowchart showing the rule-based decision making the agent would go through when using Macro-Action 1. See Appendix B for a full page version of this Figure.




Figure 3.4: A flowchart showing the rule-based decision making the agent would go through when using Macro-Action 2. See Appendix B for a full page version of this Figure.

observations as the following claims:

- Different participants develop different ways of performing the task;
- The agent learns different sequences of macro-actions for different participants;
- Different teams converge to different ways of performing the task;
- The agent converges to a specific sequence of macro-actions for most participants;
- The human converges to a specific strategy within the experiment.

In the Discussion (Discussion), we use the results of our experiment to critically evaluate whether we have been able to study co-adaptation as a precursor for co-learning with our methods by verifying to what extent these claims hold.

3.5. METHOD OF STUDY FOR IDENTIFYING INTERACTION PATTERNS

THE experimental setup and procedure described below was approved by the Human Research Ethics Committee at Delft University of Technology on August 17th, 2020 (reference number: 1261).




Figure 3.5: A flowchart showing the rule-based decision making the agent would go through when using Macro-Action 3.




Figure 3.6: An overview of the representation of the learning problem embedded in the experiment. It shows the different runs that a participant went through (5 runs for level 1, 3 runs for level 2), as well as how the runs were separated into 4 phases defined by the Phase Variables. The colors show how in R1.1, R1.2 and R1.3, the robot usually used *O1*—picking up all, *O2*—passive large rocks and *O3*—breaking respectively in each phase. From R1.4 onwards, the robot would choose a Macro-action based on the learned Q-values. The Future Run portrays the behavior that the robot would engage in if there were another run, based on the Q-values after R2.3.

3.5.1. PARTICIPANTS

A total of 24 people participated in the experiment (17 female, seven male), recruited through personal connections on LinkedIn, within the university, from a Slack community on AI and Design and from interns at TNO. The average age among the participants was 24.8 (Std = 2.47). All of the participants had a university degree in a STEM field. Most of them had little to some experience with gaming ($n = 7$ for “little experience,” $n = 10$ for “some experience”). Also, most of them expressed that they had no to little experience with human-robot collaboration ($n = 11$ for “no experience,” $n = 7$ for “little experience”) or human-robot collaboration research ($n = 11$ for “no experience,” $n = 5$ for “little experience”).

Due to a few problems in the data collection, some participants were excluded from all of the analyses or some of the analyses. Two participants (one female, one male) were excluded from all analyses, because there were significant connectivity issues during the execution of the experiment and/or data collection went wrong on more than one factor.

One participant (female) was excluded from the questionnaire analyses, because their data was not properly saved, and one participant (male) was excluded from the robot behavioral analyses, because the log data was not properly saved.

3.5.2. DESIGN AND MATERIALS

PARTICIPANTS were divided over two conditions: 1) a condition in which participants were instructed to think aloud and 2) a condition in which they were asked to perform the task in silence. Since we study learning processes, and since it is known that thinking aloud can have an effect on learning, the two conditions ensured that we had control over any possible effect.

We presented the task environment described in *Task Implementation* to all participants, in the form of two different levels. The first level was designed to be relatively easy, as it could be solved by simply clearing away all rocks (Figure 3.2A). A complicating factor was that breaking rocks would easily hurt the victim, which would therefore need to be avoided. The second level was designed to be more challenging: it contained a brown rock that could not be picked up at all (Figure 3.2B). This means that if the brown rock would fall on top of the victim, it would no longer be possible to save the victim and finish the task.

Participants played the first level five times, as it was estimated from pilot runs that five times would provide ample opportunity for both the participant and the robot to learn a working strategy. Participants then played the second level three times, to give the team the opportunity to adapt to the new situation. The repetition allowed for within-subject analyses, in which the behavior of participants could be compared between rounds, as well as between-subjects analyses of learning. For an overview of how the definition of the task in the Reinforcement Learning algorithm combines with this setup, Figure 3.6.

3.5.3. PROCEDURE

THE experiment was conducted through a video call between the experimenter and each individual participant, while both were located in their own home for the course of the experiment. Participants were given access to the experimental task using Parsec, which is a screen-sharing platform made for collaborative gaming (“Parsec”, n.d.). This ensured that participants had control over the task environment, while allowing the experimenter to observe their behavior.

All participants went through the following steps:

1. Participants were seated in front of their own computer at home;
2. They read the instruction, signed the consent form and provided some demographic information as well as information on their experience with video games, human-robot collaboration and human-robot collaboration research;
3. Participants had the opportunity to do a short test scenario of the task without the virtual robot, to familiarize them with the task environment and the controls;
4. Participants were presented with the first pre-specified level. After five runs, the new level was presented to the participant, which they played three times;

- (a) The participants in condition A were asked to think aloud during the execution of the levels;
- (b) After each level, participants completed a selection of the questionnaire on Subjective Collaboration Fluency (taken from Hoffman, 2019, see Appendix C for the questions used). In addition, they were asked to rate how confident they were that their strategy was a good strategy for solving the task on a scale of 1–10;
- (c) After the first five runs and at the end of the experiment the participants were interviewed about their experiences.

3.5.4. DATA COLLECTION AND ANALYSIS

SEVERAL Several types of data were collected in order to answer our research questions:

1. Screen captures and notes of behavior in the MATRX environment during the execution of the experiment
2. Voice recordings of the participants in condition A while they are thinking aloud during the execution of the experiment
3. Voice recordings of short interviews (see Appendix C for the questions asked)
4. Collaboration Fluency scores
5. Confidence of Strategy scores
6. Q-table as learned by the robot and log of how it changes

We will explain in more detail how this data was collected and how it relates to our research questions in the sections below.

BEHAVIOR

In order to identify what interaction patterns drive co-adaptation and co-evolution, we wanted to look at how the behavior and strategy of the team changed over time, and which interactions were used in that process. We used data types 1, 2, 3 and 6 for this. The screen captures and notes (data type 1) serve mainly as data on the human behavior, while the thinking aloud output and the interviews (data types 2 and 3) help to explain why humans behave in a certain way. The Q-tables (data type 6) serve to see what strategy the robot chose in each phase of the task. Normally, behavior of a robot driven by RL is assessed by looking at the cumulative rewards. As we are not necessarily interested in performance, but in the behavior resulting from the learning process (as prescribed in Rahwan et al., 2019), we chose to look at the development of the Q-tables, to understand what macro-action the robot learned to choose in each phase.

A Grounded Theory (Charmaz, 2014) process was used to identify recurring adaptive behaviors from the screen captures and notes. This means that we went through a process of open coding first, while constantly writing short memos of observed patterns. After that, we collected all codes and categorized and clustered them until reaching the desired level of detail.

We will explain how the behavioral data and Q-tables were used to answer our research questions in more detail in Section Results.

SUBJECTIVE COLLABORATION FLUENCY AND CONFIDENCE SCORE

Within the task that we designed, it is quite difficult to keep track of task performance due to the possibility for large differences in strategies, as well as because the task can become unsolvable. To still keep track of how the human-robot team performed over the course of the experiment, we have chosen two measures for tracking subjective task performance: subjective collaboration fluency and confidence score (data types 4 and 5). These measures helped us to validate that our experiment setup actually allowed for learning and improvement.

For subjective collaboration fluency, we used a short version of an existing questionnaire (Hoffman, 2019). To measure participants' confidence in their strategy, we asked them to rate confidence on a scale from 1 to 10 with the following question: "How confident are you that your strategy is the right strategy?"

3.6. RESULTS

3.6.1. SUBJECTIVE COLLABORATION FLUENCY AND CONFIDENCE SCORE

WE have created a box plot of the Subjective Collaboration Fluency scores (Figure 3.7). The Confidence scores followed a very similar pattern, therefore we do not go into further detail about those. Both scores follow a pattern with scores starting off relatively high in run one, after which they drop for run two and three, move up again for four and five, drop again for six and then move up for the last two runs. To test whether thinking aloud and the number of the run affected participants' experience of collaboration with the robot, the Subjective Collaboration Fluency score was entered in a one-way repeated measures ANOVA with Thinking Aloud (yes/no) as between-subjects factor, and run number (1–8) as a within-subjects repeated measure factor. Results show that there was no significant difference between the participants who were instructed to think aloud (Mean = 40.56, Std = 22.78) and those who were not (Mean = 45.12, Std = 25.75) ($F = 0.81$, $p = 0.38$), while there was a significant effect on the run ($F = 5.97$, $p < 0.0001$). When looking at Figure 3.7, we expected this significant difference to exist between round one and round two, round three and round four, round five and round six and round six and round seven (scores went down after round one, up after round three, down again after round five and up after round six). To test whether these differences between rounds were significant, we did a post-hoc analysis using a Tukey HSD test, which mostly confirmed the differences visible from the plot: R1.1 and R1.2 are significantly different ($p = 0.006$), R1.3 and R1.4 are significantly different ($p = 0.006$), R1.5 and R2.1 are almost significantly different ($p = 0.058$) but R1.4 and R2.1 are ($p = 0.003$). R2.1 and R2.2 did not differ statistically, but R2.1 and R2.3 do, although not significantly ($p = 0.114$).

The pattern of scores on both fluency and confidence over runs are probably caused by the setup of our experiment. In the first run, the robot would use macro-action 1 for the whole task, which is the easiest to work with from a participant perspective. Therefore, participants may have been inclined to assign high scores in the beginning. In run 2 and 3, the way the RL algorithm is implemented causes the robot to use macro-action




Figure 3.7: The Collaboration Fluency scores per run in the experiment for all participants.

2 and 3 respectively for the whole task, which are quite hard to understand from a participant perspective, arguably leading to a lower level of experienced fluency. In run 4 and 5, the robot would start picking its macro-action based on previous performance, while the participant would have learned to work with the robot a bit more. It is likely that this made the fluency scores go up again. Run 6, however, introduced the new scenario in which previously learned strategies often did not work anymore. In runs seven and eight the human-robot team would then learn to perform better at this second scenario, inducing participants to give higher scores on experienced collaboration fluency. Following this explanation of the scores, these results suggest that our experiment design indeed allowed for a learning process of the human-robot team as we anticipated.

3.6.2. INTERACTION PATTERN ANALYSIS

THE open coding process of behavioral summaries, based upon the information from videos, interviews and notes, yielded a list of 52 different behaviors. These behaviors consist of task-related actions by the participant; interactions between the participant and the virtual robot; learning (participant learns something about the task or the collaboration); strategies (combinations of actions executed over longer periods of time); team performance and participant emotional responses. After excluding behaviors that did not relate to adaptation in specific (e.g. actions such as “picking up top rocks”) and behaviors that were more of an assessment of the quality of a behavior rather than a description (e.g. performance factors such as “not understanding the link between waiting and robot action”), a list of 38 behaviors was left.

These 38 behaviors were categorized in the following two categories (based on the categorizations made in van Zoelen et al., 2021):

- **Stable situations (9 behaviors):** behaviors observed in-between adaptations, such as the behavior of the participant alternating acting and waiting for the robot.
- **Sudden adaptations (29 behaviors):** behaviors in which the human and/or robot adapted their actions, thus starting a transition from one stable situation to another. The adaptation happens in a single moment or over a short period of time, often in response to a newly hypothesized or discovered property of the partner's behavior.

The full list and categorization can be found in Appendix D. The behaviors listed in the Appendix are closely tied to the experimental task. The descriptions of the behaviors were processed to fit co-adaptation in general (such that we can call them interaction patterns). For this purpose, some of the behaviors were combined into one descriptive interaction pattern. The resulting list, consisting of 23 interaction patterns (five stable situations, 18 sudden adaptations), is presented in Table 3.3.

The biggest group is that of “sudden adaptations,” the patterns that often arise in response to a discovery, an expectation, or a surprise of one of the partners. In order to better understand this important group of adaptive interaction patterns, we explored in more detail the nature of the triggers that initiate them, what characterizes the execution of these patterns, and what they bring about in the human-robot collaboration. Again following the approach taken in van Zoelen et al., 2021, we used the following terms to describe the sudden adaptations:

- **External trigger:** an event outside of the partner (e.g. in the task, environment or other partner) triggers an adaptation to a new stable situation;
- **Internal trigger:** an event inside of the partner (e.g. a specific expectation or change of mind) triggers an adaptation to a new stable situation;
- **Outcome:** a specific action that is preceded by an internal or external trigger, that will gradually develop into a new stable situation afterward;
- **In-between-situation:** a specific action that is preceded by an internal or external trigger, that serves as a new trigger for adapting to a new stable situation afterward.

The results of this can be found in Table D.3.

Category	Concept	Description
Stable situation	Actively synchronizing actions with a team member	Human understands the capabilities of another team member and actively uses their own actions to make optimal use of the combined capabilities

Sudden adaptation

Alternating actively working on the task and waiting for a team member	Human switches between performing their own task for a while, then waiting for a team member to perform their task, and so on
Being generally passive and letting a team member do most of the work	Human is overall passive and lets the other team member do the work
Damage control: Prevent damage caused by a team member	Human performs actions that prevent their team member from causing intentional or unintentional harm or damage
Focusing on own task	Human performs their own task without paying much attention to their team member
Avoiding communication with a team member	One of the team member actively avoids the other team member to avoid unwanted communication interpretations
Being confused by non-human-like behavior	A human team member is confused by non-human-like behavior performed by a team member
Being confused by unexpected behavior (negative)	One of the team members is confused or frustrated by behavior performed by their team member that they did not expect
Being happy that a team member does as expected	One of the team members is happy that their team member performs the kind of behavior that they expect and hoped for
Being surprised by unexpected behavior (positive)	One of the team members is positively surprised by behavior performed by their team member that they did not expect
Coming into action when a team member comes into action	A team member starts to actively perform their task after a period of inaction, when their team member also starts to actively perform their task after a period of inaction
Doing useless or harmful actions because there is nothing else to do	A team member is unable to perform useful actions, therefore starts performing useless or harmful actions
Feeling alone, as if team member does not help	A human team member feels left alone

Following a team member's action	A team member follows or copies the action performed by another team member
Learning about behavioral cues	A team member gains insight into specific behavior performed by another team member
Learning about own capabilities	A team member gains insight into their own capabilities
Learning about team member's capabilities or strategy	A team member gains insight into the capabilities or strategy of another team member
Moving around different task components	A team member moves around different task components without actually performing any task
Team member changes strategy, which is visible by a behavioral cue	A team member observes that another team member changes strategy by a behavioral cue
Team member performs an action that makes no sense	A team member performs a useless action
Trying to communicate by interacting with a team partner	A team member attempts to communicate with another team member by directly interacting with them, for example by coming close to them
Trying to communicate by signaling task actions	A team member attempts to communicate with another team member by trying out different actions that they want their team member to perform
Waiting for a team member to start acting	A team member waits for another team member to start performing their task

Table 3.3: The interaction patterns identified from the behavioral data, including a description of what they entail.

3.6.3. COLLABORATIVE LEARNING

IN addition to developing a comprehensive description of adaptive interaction patterns, we further explored how human behavior, and specifically human behavior adaptation, influenced learning by the virtual robot. We analyzed and coded human adaptive behavior at a detailed level by identifying the interaction patterns as described above, but in order to analyze how the development of robot behavior and human learning depend on each other, a different level of detail was necessary. We looked at three aspects of the data:

- For each participant, we looked at the Q-tables of the virtual robot at the end of

each run in the experiment, to see which of the three macro-actions received the highest expected reward in the different phases and runs of the task;

- For each participant, we identified the main behavioral strategy used by the human per run, as well as during the whole experiment;
- We analyzed how the chosen macro-actions of the virtual robot can be associated with specific behavioral strategies of the participants.

We will describe our process for all three of these aspects in more detail below.

VIRTUAL ROBOT Q-TABLES

The Reinforcement Learning algorithm addressed learning when to apply which of three macro-actions or options (as described in Figure B.1, Figure B.2, and Figure 3.5; we will call them *O1—picking up all*, *O2—passive large rocks* and *O3—breaking* from here onwards) and used four phase variables to identify states. Figure 3.8 shows an overview of how often the robot learned to pick specific macro-actions in each phase (Figure 3.8A) and each run (Figure 3.8B), based on the macro-action with the highest expected reward.

As the robot's choice for macro-options is not clearly related to the phases in the task (especially phase 2, 3 and 4 are very similar, as can be seen in Figure 3.8A), we looked mostly at Figure 3.88B to understand how the robot's behavior developed. The figure shows that in the first three runs the robot mostly tried out all macro-strategies one by one, as determined by how the algorithm was programmed. Small deviations from this are likely caused by some participants going back and forth between phases in the task, rather than moving through them linearly as we initially expected. The robot generally learned to select *O1—picking up all* most of the time for most participants over the course of the next few runs, which fits with how level 1 of the experiment was designed. From run 6 onwards, when the second level was introduced, the robot learned to choose *O2—passive large rocks* and *O3—breaking* more often. This shows that the robot is able to generally learn what works best for the task.

PARTICIPANT BEHAVIORAL CLUSTERING

To better understand how the behavior of the participants developed over time, we performed a manual clustering of participant behavior per run. Based on the behavior observations as described in Section Interaction Pattern Analysis, we defined the following behavioral clusters:

- Just focus on own behavior efficiently
- Balancing acting and waiting
- Exploring how the robot works by observing and trying to communicate
- Actively using *O3—breaking*
- Actively using *O2—passive large rocks*




Figure 3.8: An overview of how often certain Macro-actions were chosen by the robot across all participants per phase (A) and per run (B).

The result of this clustering can be seen in Figure 3.9. It is difficult to find detailed insights from this figure, apart from the fact that more participants showed more adaptive behavior in the later runs, as indicated by the red and olive green bars in the figure. Participants’ strategies did not develop linearly, and it also did not converge to one specific type of behavior consistently within our experiment. To be able to see whether human learning had an influence on robot learning, we chose to remove the dimension of time (runs) from our participant data, and focus on whether and how a participant adapted over the whole experiment. We created the following clusters based on participant adaptation over the whole experiment:

- Does not adapt: participant shows no signs of adapting to strategies employed by the robot; participant either focuses on their own task, or constantly switches between behavior strategies as they focus too much on the robot.




Figure 3.9: An overview of how many participants used specific behavioral strategies per run.

Cluster	Participants
Does not adapt	2, 6, 9, 15, 21, 22, 23, 24 (n = 8)
Adapts by balancing passively waiting and acting	12, 13, 14, 16, 27, 28 (n = 6)
Adapts by actively using O2 or O3	3, 8, 10, 17, 19, 20, 26 (n = 7)

Table 3.4: The clusters resulting from manually clustering participants based on whether they adapted to the robot across the whole experiment.

- Adapts by balancing waiting and acting: participant shows signs that they adapt by waiting for the robot to act, and to use that robot behavior to determine their own response. It suggests that the participant understands that being passive for a while may cause the robot to act.
- Adapts by actively using O2—*passive large rocks* or O3—*breaking*: participant visibly adapts as they actively guide the robot to pick up or break rocks by waiting on top of those rocks.

This clustering of participants according to their dominant strategy resulted in three clusters with a similar number of participants per cluster, as shown in Table 3.4.

COMBINING PARTICIPANT ADAPTATION AND ROBOT LEARNING

In order to explore whether these different types of adaptation employed by participants affected robot learning, and whether differences occur between clusters, we plotted the




Figure 3.10: An overview of how often certain Macro-actions were chosen by the robot across all participants per run, split up by the level adaptation the participant showed: (A) shows participants who adapted by actively using O2—passive large rocks and/or O3—breaking, (B) shows participants who adapted by balancing waiting and acting, and (C) shows participants who did not adapt.

robot strategies per human adaptation cluster, as shown in Figure 3.10. These figures present the bar graphs of how often the different macro-actions were chosen across the group, in an attempt to more closely evaluate any possible differences between the three clusters. The figures suggest a response in the robots’ behavior to the participants’ actions, especially later on, in the final runs. When we compare the figures for “adapts by balancing passively waiting and acting” (Figure 3.10B) and “does not adapt” (Figure 3.10C), the former shows the trend of using O1—picking up all in the first level (first five runs) and using other strategies in the second level (run 6–8) more strongly. The figures for “adapts by actively using O2 or O3” (Figure 3.10A), however, shows that the robot already learned to use O2—passive large rocks and O3—breaking before being introduced with the second level, and actually moving back to O1—picking up all a little more toward the final runs.

This suggests that if participants learned more about the robot behavior and adapted their own behavior more strongly (by actively guiding the robot with O2—passive large rocks and O3—breaking), the robot was also able to learn to use those strategies more often. This combined learning effect therefore can be seen as learning at the team level.

3.7. DISCUSSION

3.7.1. INTERACTION PATTERNS THAT DRIVE CO-LEARNING

WE set out to investigate what interaction patterns between humans and robots drive co-adaptation as a precursor for co-learning. From the behaviors observed in our experiment, we identified a list of interactions and a set of interaction patterns. It should be noted, however, that (interactive) behavior is very much determined by the specific context. This means that our list should not be considered as a complete list of all possible co-adaptive interactions. It should rather be seen as a collection of interaction patterns that are likely to appear in contexts similar to ours, where co-adaptation is centered around mutual observation and harmonizing actions when collaborating.

The collection of interaction patterns can be used as a language for recording, analyzing and coding co-adaptive behavior. By describing observed behavior with such interaction patterns, complex behavioral observations of human-robot co-adaptive strategies can more easily be compared. The interaction patterns can also be useful as a vocabulary for a human-robot team itself to discuss the adaptations that they are engaged in, to help them elaborate and sustain successful collaborations over time.

3.7.2. VALIDATING THE RESEARCH ENVIRONMENT

AN important objective of this study is to improve our understanding of how human-robot co-learning develops, as well as how the adaptive processes of both partners interact. We defined claims for the experimental environment we designed; if these claims are justified, it means that it enabled us to study co-adaptation, the process that we consider to be a precursor for co-learning. Table 3.5 shows the claims described in Claims: Expected Observations and annotated conclusions as to whether we were able to justify the claims in the present study.

As can be seen in the table, only one of the claims was justified completely. Fortunately, many of the other claims were partly justified. For the claims that we did not realize, the results provide cues for how to design a research environment that better fits the claims. We regard these findings as an important step toward studying and revealing the processes involved in human-robot co-learning. Aspects that should be improved upon or need further work center around a few problems that we will elaborate on below:

1. Behavior strategy: Convergence vs. flexible adaptations
2. Statistical analysis of complex behavioral data
3. Behavior of the individual team member vs. behavior of the team
4. Task effects vs. participant effects

In our claims in Table 3.5, we mentioned convergence several times. This stems from the principle that a Reinforcement Learning algorithm should aim for convergence toward an optimal solution. However, when studying co-learning, we specifically use dynamic task environments that have no fixed optimal solution, and in which unpredicted events can require strategy changes. In such environments, convergence is not a good criterion for performance, as agents (human as well as robot) are required to continuously learn and adapt. For human-robot co-learning it can be argued that it is better to make the

algorithm learn certain repeated subsequences of interactions (or interaction patterns), and to store those in a rule-based manner. Once a pattern of interaction has proven to be successful in multiple instances of task situations, it can be applied, combined and if necessary revised in similar but other task situations. We therefore believe that future research into co-learning should not take convergence as a criterion for the robot's behavior, but to focus on the emergence and sustainability of successful interaction patterns (aspect 1).

The results show that the robot had a similar learning process across all participants despite the high variety between individual participants. However, the behavioral data of both the robot and the participant is quite complex. Sometimes there are radical changes in behavior between one run and the next, and even within one run participants sometimes quite radically changed their behavior. It is a challenge to analyze such data as it is often difficult to clarify the origin of the behavior from the data. Our qualitative analysis and clustering is able to deal with this complexity and provides many useful insights, therefore we would advise future research into co-learning to include similar qualitative analyses. When further investigating co-learning, it will, however, also be relevant and interesting to verify insights statistically. This will require different design considerations. The current complexity in behavioral data is partly due to the interaction between two adaptive systems, and probably an inherent property of co-learning. Moreover, the human and robot can approach the task in many different ways by design. This property is a strength of our experiment, as it allows participants to behave relatively freely and naturally, but it also contributes to the complexity of the resulting co-adaptation. For future research, it will be important to explicitly take these properties into account when designing an environment for experimentation (aspect 2), in such a way that insights can be verified statistically. For example, in our design, the strategies of the robot were separate, nominal actions, but if we can design a learning agent such that their learned behavior is ordinal (e.g. by using more or less of a certain behavior), it might be easier to apply statistical methods. Moreover, we can look into data analysis methods used in complexity science, to see if they can be applicable to co-learning scenarios.

Lastly, it is currently a challenge to determine which aspects of the final team behavior are caused by adaptations of individual team members, and which by interactions between them. Similarly, we cannot yet conclusively determine which aspects of the learned strategy are caused by the task, and which by the individuality of a participant (e.g. what does the robot learn just because of the task, and what does the robot learn because a certain participant behaved a certain way). To solve this problem, we need to find ways to separate the different effects, for example by creating relevant baseline results. Letting the robot perform a task and learn by itself is not an option in the context of team tasks, as the nature of such task dictates dependencies between the team members. A possibility might be to create a simulated, possibly rule-based human agent for the robot agent to collaborate with.

3.7.3. FUTURE STEPS FOR STUDYING HUMAN-ROBOT CO-LEARNING

THE discussion regarding the interactions that underlie co-learning provide several pointers and suggestions for improving the design of our research methods. Besides

Claim	Justified	Explanation
Different participants develop different ways of performing the task	Yes	When looking at the different interaction patterns that people engage in, and categorizations of their adaptive behavior, we can see that different people indeed performed the task in a variety of ways
The agent learns different sequences of strategy options for different participants	Partly	The results showed that not all agents learned the same model on an individual level. However, the models had much in common, suggesting that all agents learned similar behavior. When splitting this up in groups based on human adaptive behavior, there seems to be a difference in learned agent behavior between the different groups. Currently, however, we did not do any statistical analysis to test whether this is a significant result
Different teams converge to different ways of performing the task	Partly	When looking at the different interaction patterns that participants engaged in with their robot team partner, different teams solved the task in a variety of ways (see H1). However, it is unclear to what extent the robot contributed to this. Moreover, while participants generally gained more confidence in their strategy and expressed to experience a greater subjective collaboration fluency toward the end of the experiment, it is unclear to what extent the strategy of the team really converged to a stable one
The agent converges to a specific sequence of strategy options for most participants	No	While we did observe a logical development of the Q-values on a population level, this does not count for all of the individual agents. Moreover, it is not clear to what extent the agents really converged to a stable set of actions
The human converges to a specific strategy within the experiment	Partly	The categorizations of participant behavior show that participants settle on a stable strategy more and more over the course of the experiment. This is also shown by the development of the confidence scores and subjective collaboration fluency. True convergence to a stable strategy, however, is not clearly visible within the 8 runs of the experiment

Table 3.5: The claims as presented in Claims: Expected Observations that need to be justified in a co-adaptation experiment, including whether they were validated and an explanation of that conclusion.

the suggestions above, however, there are several other directions in which we believe that co-learning research should develop. In this paper we developed an approach for studying co-learning. Since we were still defining what it means to study co-learning, the scope of the task, learning algorithm and opportunities for interaction had to be limited. Eventually, if we want to enable and study co-learning in full-fledged teams, it will be necessary to use more complex task environments, more intelligent agents or robots and more elaborate interaction and communication between the human and the robot. In the following section we will therefore further outline the two research directions mentioned in Research Challenges that we did not further address in this paper (research direction two and three).

The first direction is aimed at enabling the communication between partners within the team, especially communication about adaptations. We believe that in order for team members to produce successful sequences of interactive behavior, that can be used strategically across contexts, it is necessary that the team members can communicate with one another. The interaction patterns that we have identified might be used as a start for a vocabulary for such communication interactions, but the specific timing, modality and details of the interactions will have to be designed and studied.

The second direction is aimed at making the agent or robot more intelligent in terms of its abilities for co-learning. In the experiment we presented, our agent only learned based on task-related rewards. It makes sense to also explicitly reason about or take into account the human's behavior and preferences in learning. There is a large body of research on personalizing robot behavior, e.g. by making the robot develop a user model of its partner, but these models often do not explicitly take into account that the human continuously learns and adapts. We therefore believe that there is a need for user models and team models that specifically accommodate the adaptive interactions as described in this paper. A team mental model that is able to represent the interactions within a team will support the partners in developing and sustaining successful adaptations and to synchronize and align their actions and learning processes.

3.8. CONCLUSION

CO-LEARNING is an important mechanism for building successful human-robot teams. However, there is no general understanding of what co-learning means in the context of human-robot teams. In this paper, we defined the concept of co-learning based on related literature, and positioned it in relation to co-adaptation and co-evolution. From this definition, it is clear that adaptive interactions between humans and robots play a central role in co-learning. We defined requirements for studying how bi-lateral adaptation emerges from the interactions between humans and robots.

From these requirements, we developed an experimental task environment based on a real-life Urban Search and Rescue task. The task was designed such that it allowed human participants to behave relatively naturally and freely, enabling us to record and analyze emerging adaptive interactions between a human and a robot. A bottom-up coding process based on Grounded Theory allowed us to identify recurring interaction patterns. The resulting list of interaction patterns describe stable situations (repeating subsequences of stable behavior) as well as sudden adaptations (changes in behavior happening over short periods of time). These patterns can emerge in similar task con-

texts, thereby forming a description and analysis method of co-adaptive behavior in human-robot teams.

Over the course of the experiment, the robot learned similar strategies for most participants. However, the results show that the learned strategies were slightly different depending on whether a human participant adapted their behavior to the robot. This suggests that human learning affected the robot's Reinforcement Learning. More specifically, the human learning about and adapting to a specific strategy of the robot enabled the robot to learn to stick to that strategy. This shows how learning on the individual level lead to team level learning in our experiment.

This paper presents a theoretical framework and a methodological approach for studying the processes that underlie co-learning in human-robot teams. The strengths as well as the shortcomings of our approach provide ample directions for future research into this important process that ultimately defines the quality of human-robot teaming.

BIBLIOGRAPHY

- Ansari, F., Hold, P., Mayrhofer, W., Schlund, S., & Sihm, W. (2018). AUTODIDACT: INTRODUCING THE CONCEPT OF MUTUAL LEARNING INTO A SMART FACTORY INDUSTRY 4.0, 8.
- Baxter, P., Ashurst, E., Read, R., Kennedy, J., & Belpaeme, T. (2017). Robot education peers in a situated primary school study: Personalisation promotes child learning (N. Reich-Stiebert, Ed.). *PLoS ONE*, 12(5), e0178126. <https://doi.org/10.1371/journal.pone.0178126>
- Belpaeme, T., Vogt, P., van den Berghe, R., Bergmann, K., Göksun, T., de Haas, M., Kanero, J., Kennedy, J., Küntay, A. C., Oudgenoeg-Paz, O., Papadopoulos, F., Schodde, T., Verhagen, J., Wallbridge, C. D., Willemsen, B., de Wit, J., Geçkin, V., Hoffmann, L., Kopp, S., ... Pandey, A. K. (2018). Guidelines for Designing Social Robots as Second Language Tutors. *International Journal of Social Robotics*, 10(3), 325–341. <https://doi.org/10.1007/s12369-018-0467-6>
- Burke, S., Stagl, K., Salas, E., Pierce, L., & Kendall, D. (2006). Understanding Team Adaptation: A Conceptual Analysis and Model. *The Journal of applied psychology*, 91, 1189–207. <https://doi.org/10.1037/0021-9010.91.6.1189>
- Buschmeier, H., & Kopp, S. (2013). Co-constructing Grounded Symbols - Feedback and Incremental Adaptation in Human-Agent Dialogue. *KI - Künstliche Intelligenz*, 27, 137–143. <https://doi.org/10.1007/s13218-013-0241-8>
- Charmaz, K. (2014, March). *Constructing Grounded Theory* [Google-Books-ID: v_GGAwAAQBAJ]. SAGE.
- Chauncey, K., Harriott, C., Prasov, Z., & Cunha, M. (2017). A Framework for Co-adaptive Human-Robot Interaction Metrics. *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 123–124.
- Dillenbourg, P., Baker, M., Blaye, A., & O'Malley, C. (1996). The evolution of research on collaborative learning. In *Learning in Humans and Machine: Towards an interdisciplinary learning science* (pp. 189–211). Elsevier.
- Döppner, D. A., Derckx, P., & Schoder, D. (2019). Symbiotic Co-Evolution in Collaborative Human-Machine Decision Making: Exploration of a Multi-Year Design Science Research Project in the Air Cargo Industry. *Proceedings of the 52nd Hawaii International Conference on Systems Science*, 264–273. <https://doi.org/10.24251/HICSS.2019.033>
- Dulac-Arnold, G., Mankowitz, D., & Hester, T. (2019). Challenges of Real-World Reinforcement Learning [arXiv: 1904.12901]. *arXiv:1904.12901 [cs, stat]*. Retrieved February 12, 2021, from <http://arxiv.org/abs/1904.12901>
- Ehrlich, S. K., & Cheng, G. (2018). Human-agent co-adaptation using error-related potentials. *Journal of Neural Engineering*, 15(6), 066014. <https://doi.org/10.1088/1741-2552/aae069>

- Foerster, J., Assael, I. A., de Freitas, N., & Whiteson, S. (2016). Learning to Communicate with Deep Multi-Agent Reinforcement Learning. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* 29 (pp. 2137–2145). Curran Associates, Inc. Retrieved February 12, 2019, from <http://papers.nips.cc/paper/6042-learning-to-communicate-with-deep-multi-agent-reinforcement-learning.pdf>
- Gao, Y., Barendregt, W., & Castellano, G. (2017). Personalised Human-Robot Co-Adaptation in Instructional Settings using Reinforcement Learning. *Proceedings of the Persuasive Embodied Agents for Behavior Change (PEACH2017) Workshop*.
- Hoffman, G. (2019). Evaluating Fluency in Human–Robot Collaboration. *IEEE Transactions on Human-Machine Systems*, 49(3), 209–218. <https://doi.org/10.1109/THMS.2019.2904558>
- Holstein, K., Aleven, V., & Rummel, N. (2020). A Conceptual Framework for Human–AI Hybrid Adaptivity in Education. In I. I. Bittencourt, M. Cukurova, K. Muldner, R. Luckin, & E. Millán (Eds.), *Artificial Intelligence in Education* (pp. 240–254). Springer International Publishing. https://doi.org/10.1007/978-3-030-52237-7_20
- Illanes, L., Yan, X., Toro Icarte, R., & McIlraith, S. A. (2019). Symbolic Planning and Model-Free Reinforcement Learning - Training Taskable Agents. Retrieved May 31, 2021, from https://drive.google.com/file/d/15b4w56ODJa4PZpSNbhho5q1s3pl_DlyP/view?usp=embed_facebook
- Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., Van Riemsdijk, M. B., & Sierhuis, M. (2014). Coactive Design: Designing Support for Interdependence in Joint Activity. *Journal of Human-Robot Interaction*, 3(1), 43. <https://doi.org/10.5898/JHRI.3.1.Johnson>
- Kapetanakis, S., & Kudenko, D. (2002). Reinforcement Learning of Coordination in Cooperative Multi-Agent Systems. *Proceedings of the Nineteenth National Conference on Artificial Intelligence*, 326–331.
- Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems*, 19(6), 91–95. <https://doi.org/10.1109/MIS.2004.74>
- Lee, C., Wang, M., Huang, T., Chen, L., Huang, Y., Yang, S., Tseng, C., Hung, P., & Kubota, N. (2018). Ontology-based Fuzzy Markup Language Agent for Student and Robot Co-Learning. *2018 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, 8–13. <https://doi.org/10.1109/FUZZ-IEEE.2018.8491610>
- Lee, C.-S., Wang, M.-H., Tsai, Y.-L., Ko, L.-W., Tsai, B.-Y., Hung, P.-H., Lin, L.-A., & Kubota, N. (2020). Intelligent agent for real-world applications on robotic edutainment and humanized co-learning. *Journal of Ambient Intelligence and Humanized Computing*, 11(8), 3121–3139. <https://doi.org/10.1007/s12652-019-01454-4>
- Lematta, G. J., Coleman, P. B., Bhatti, S. A., Chiou, E. K., McNeese, N. J., Demir, M., & Cooke, N. J. (2019). Developing Human-Robot Team Interdependence in a Synthetic Task Environment [Publisher: SAGE Publications Inc]. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 63(1), 1503–1507. <https://doi.org/10.1177/1071181319631433>

- MATRX Software [original-date: 2020-02-27T12:34:36Z]. (2021, January). Retrieved February 5, 2021, from <https://github.com/matrix-software/matrix>
- Mohammad, Y., & Toyoaki Nishida. (2008). Human adaptation to a miniature robot: Precursors of mutual adaptation. *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, 124–129. <https://doi.org/10.1109/ROMAN.2008.4600654>
- Nikolaïdis, S., Hsu, D., & Srinivasa, S. S. (2017). Human-robot mutual adaptation in collaborative tasks: Models and experiments. *I. J. Robotics Res.*, 36, 618–634. <https://doi.org/10.1177/0278364917690593>
- Nikolaïdis, S., Zhu, Y. X., Hsu, D., & Srinivasa, S. (2017). Human-Robot Mutual Adaptation in Shared Autonomy. *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction - HRI '17*, 294–302. <https://doi.org/10.1145/2909824.3020252>
- Parsec. (n.d.). Retrieved February 12, 2021, from <https://parsec.app>
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J. W., Christiakis, N. A., Couzin, I. D., Jackson, M. O., Jennings, N. R., Kamar, E., Kloumann, I. M., Larochele, H., Lazer, D., McElreath, R., Mislove, R., Parkes, D. C., Pentland, A. ', ... Wellman, M. (2019). Machine behaviour. *Nature*, (568), 477–486. <https://doi.org/10.1038/s41586-019-1138-y>
- Ramakrishnan, R., Zhang, C., & Shah, J. (2017). Perturbation Training for Human-Robot Teams. *Journal of Artificial Intelligence Research*, 59, 495–541.
- Shafti, A., Tjomsland, J., Dudley, W., & Faisal, A. A. (2020). Real-World Human-Robot Collaborative Reinforcement Learning [arXiv: 2003.01156]. *arXiv:2003.01156 [cs]*. Retrieved March 8, 2021, from <http://arxiv.org/abs/2003.01156>
- Sordoni, A., Galley, M., Auli, M., Brockett, C., Ji, Y., Mitchell, M., Nie, J.-Y., Gao, J., & Dolan, B. (2015). A Neural Network Approach to Context-Sensitive Generation of Conversational Responses [arXiv: 1506.06714]. *arXiv:1506.06714 [cs]*. Retrieved February 12, 2019, from <http://arxiv.org/abs/1506.06714>
- Stolle, M., & Precup, D. (2002). Learning Options in Reinforcement Learning. In S. Koenig & R. C. Holte (Eds.), *Abstraction, Reformulation, and Approximation* (pp. 212–223). Springer. https://doi.org/10.1007/3-540-45622-8_16
- Sutton, R. S., Precup, D., & Singh, S. (1999). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1-2), 181–211. [https://doi.org/10.1016/S0004-3702\(99\)00052-1](https://doi.org/10.1016/S0004-3702(99)00052-1)
- van Otterlo, M., & Wiering, M. (2012). Reinforcement Learning and Markov Decision Processes [Series Title: Adaptation, Learning, and Optimization]. In M. Wiering & M. van Otterlo (Eds.), *Reinforcement Learning* (pp. 3–42, Vol. 12). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-27645-3_1
- van Zoelen, E. M., van den Bosch, K., Rauterberg, M., Barakova, E., & Neerincx, M. (2021). Identifying Interaction Patterns of Tangible Co-Adaptations in Human-Robot Team Behaviors. *Frontiers in Psychology*, 12, 645545. <https://doi.org/10.3389/fpsyg.2021.645545>
- van den Bosch, K., Schoonderwoerd, T., Blankendaal, R., & Neerincx, M. (2019). Six Challenges for Human-AI Co-learning. *Lecture Notes in Computer Science*, 11597, 572–589. https://doi.org/10.1007/978-3-030-22341-0_45

- Vignolo, A., Powell, H., McEllin, L., Rea, F., Sciutti, A., & Michael, J. (2019). An adaptive robot teacher boosts a human partner's learning performance in joint action [ISSN: 1944-9437]. *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. <https://doi.org/10.1109/RO-MAN46459.2019.8956455>
- Weber, K., Ritschel, H., Lingenfeller, F., & André, E. (2018). Real-Time Adaptation of a Robotic Joke Teller Based on Human Social Signals, 3.
- Wenskovitch, J., & North, C. (2020). Interactive AI: Designing for the “Two Black Boxes” Problem. *Hybrid Human-Artificial Intelligence Special Issue*, 1–10.
- Xu, Y., Ohmoto, Y., Okada, S., Ueda, K., Komatsu, T., Okadome, T., Kamei, K., Sumi, Y., & Nishida, T. (2012). Formation conditions of mutual adaptation in human-agent collaborative interaction. *Applied Intelligence*, 36(1), 208–228. <https://doi.org/10.1007/s10489-010-0255-y>
- Yamada, S., & Yamaguchi, T. (2002). Mutual adaptation to mind mapping in human-agent interaction. *Proceedings. 11th IEEE International Workshop on Robot and Human Interactive Communication*, 105–110. <https://doi.org/10.1109/ROMAN.2002.1045606>

4

ENABLING EMBODIED HUMAN-ROBOT CO-LEARNING: REQUIREMENTS, METHOD, AND TEST WITH HANDOVER TASK

Despite a large body of research on robot learning, it has not yet been thoroughly studied how collaborating humans and robots learn reciprocally. In such situations, both humans and robots continuously learn about each other and the task through interaction. This paper addresses the research question: “How can human-robot co-learning be facilitated in physically embodied collaborative tasks?”. First, we derived five requirements for successful human-robot co-learning from literature: shared goal, synchrony, interdependence, adaptability, and transparency. Based on these requirements, we designed a collaborative human-robot handover task and a robot Q-learning method. In an evaluation with six human participants co-learning was indeed found to emerge in the hand-over task. Particularly, for three of the human-robot dyads, our designed setup proved to facilitate co-learning in a way that met all five requirements. The task and robot learning method presented in this paper demonstrate how human-robot co-learning can be enabled in physically embodied tasks.

 **Emma M. van Zoelen**, Hugo Veldman-Loopik, Karel van den Bosch, Mark Neerincx, David A. Abbink, Luka Peternel. 2024. Enabling Embodied Human-Robot Co-Learning: Requirements, Method, and Test With Handover Task. In *IEEE Robotics and Automation Letters*, 1425 - 1432.

4.1. INTRODUCTION

HUMAN-ROBOT collaboration research has rapidly evolved in the last decade (Ajoudani et al., 2018). Collaborative robots are being used in various industries and domains, performing an increasing amount of tasks side by side with humans. To ensure that humans and robots collaborate effectively, it is essential that they learn about each other and the task, to improve their collaboration over time (van den Bosch et al., 2019). Especially in dynamic, real-world environments, this learning will partly need to take place on the job, while humans and robots are collaborating.

We call this continuous collaborative learning process Co-Learning. When a human and robot are Co-Learning, both agents simultaneously learn how to collaborate effectively as a team by adapting their behavior to the other (Li et al., 2023; van Zoelen, van den Bosch, and Neerincx, 2021; van Zoelen, van den Bosch, Rauterberg, et al., 2021). As a result of this reciprocal adaptation, new patterns of collaboration emerge. Human-robot Co-Learning can be used to improve performance and personalize robot behavior to their human collaborator (Shafti et al., 2020). Recent research has explored the use of Co-Learning to improve Collaboration Fluency and task performance (Johnson et al., 2014; van den Bosch et al., 2019; van Zoelen, van den Bosch, and Neerincx, 2021; van Zoelen, van den Bosch, Rauterberg, et al., 2021).

Co-Learning is a relatively new and unstudied topic within human-robot collaboration. The dynamics of Co-Learning have previously been explored in virtual environments (van Zoelen, van den Bosch, Rauterberg, et al., 2021) and with Wizard-of-Oz setups (van Zoelen, van den Bosch, and Neerincx, 2021). These initial studies have shown that humans and robots collaborating can develop successful patterns of collaboration as a result of reciprocal adaptations. There is however limited research on Co-Learning in physically embodied environments with robots whose actions are governed by Machine Learning algorithms. There are some relevant studies on human-robot mutual adaptation and collaborative learning (e.g. Amirshirzad et al., 2019; Ikemoto et al., 2012; Nikolaidis et al., 2017; Peternel et al., 2016), within which there is often a strong focus on task performance improvement. There are only a few exploratory studies on the process of Co-Learning and the patterns of collaboration that emerge as a result (Kumar et al., 2024; Shafti et al., 2020). As the existing studies pay little attention to how Co-Learning can be facilitated, our research was guided by the following research question: *“How can human-robot co-learning be facilitated in physically embodied collaborative tasks?”*

We studied this question for a team consisting of a human-robot dyad, and focused on Reinforcement Learning (RL) as the learning method for the robot. This paper provides a set of core requirements for human-robot co-learning, and presents the design and evaluation of a human and a robot collaborating on a handover task based on these requirements. We use a qualitative approach to get an in-depth understanding of the co-learning process in relation to our requirements and design, as well as to provide a basis for future research on co-learning.

4.2. DESIGN REQUIREMENTS

WE have defined five design requirements for human-robot Co-Learning in physically embodied tasks based on literature research (Loopik, 2022): shared goal, syn-




Figure 4.1: Experiment setup for the human-robot co-learning of an object handover task. The robot consists of the *KUKA LBR iiwa7 800* robotic arm with the *qb-softhand* attached. The *Optitrack* motion tracking system is used to track the pose of the human hand via a sensorized glove. The human is also performing a secondary task, as explained in Section 4.3.1

chrony, interdependence, adaptability, and transparency. We examine each of them in the following subsections.

4.2.1. SHARED GOAL

ENSURING that both team members have the same goal is crucial for them to converge to congruent strategies (Johnson et al., 2014; Katz-Navon and Erez, 2005; van Zoelen, van den Bosch, and Neerincx, 2021). This can be done by rewarding both team members based on the joint task goal (e.g., Rijgersberg-Peters et al., 2023). This leads to the first requirement:

- R1** Both the human and the robot are rewarded similarly, based on their collaborative performance.

4.2.2. SYNCHRONY

CO-LEARNING is most likely to succeed when both agents learn synchronously to enable continuity, reciprocity and complementarity in their learning process (cf., Hendrikse, 2024; Mörtl et al., 2014). If team members' learning is "disconnected", the mo-

tivation for collaboration can be lost due to uncertainty about the other's progress and contribution. One of the collaborators being ahead in learning could also cause a hierarchy in the team that could be harmful for interdependence (Katz-Navon and Erez, 2005). Our second design requirement therefore states:

R2 The robot has the ability to learn in synchrony with the human team member.

4.2.3. INTERDEPENDENCE

TO enable Co-Learning, all team members should be able to meaningfully contribute to the task by complementing and supporting each other. Such a team relationship is described by *interdependence*, which is a requirement for collaboration (Doorewaard et al., 2002; Johnson et al., 2014; Katz-Navon and Erez, 2005) and therefore for Co-Learning (Burke et al., 2006; van den Bosch et al., 2019; van Zoelen, van den Bosch, and Neerinx, 2021; van Zoelen, van den Bosch, Rauterberg, et al., 2021). Interdependence is often used to describe team and task designs in studies on team collaboration (Katz-Navon and Erez, 2005), collaborative performance (Burke et al., 2006; Katz-Navon and Erez, 2005), team task design (Johnson et al., 2014; van Zoelen, van den Bosch, Rauterberg, et al., 2021) and team learning (van den Bosch et al., 2019; van Zoelen, van den Bosch, and Neerinx, 2021).

Interdependence is built up of two types of dependence: (1) *hard dependence*, in which team members can only complete a task together, and (2) *soft dependence*, when team members do not strictly need each other to achieve the group goal, but have opportunities to collaborate to perform better as a team. This can lead to team members proactively adapting to and supporting each other, which is a vital part of the Co-Learning process. Soft dependencies have a recursive nature; when interdependence is established, soft dependencies can arise, retaining and strengthening the interdependent relationship. To enable Co-Learning, we therefore need to facilitate the formation of an interdependent relationship between the human and the robot. This is done by ensuring hard dependencies between the human and the robot and by creating opportunities for soft dependencies to emerge. The third requirement is as follows:

R3 The task design ensures hard dependencies and allows for soft dependencies between the human and the robot, in both directions.

4.2.4. ADAPTABILITY

IN Co-Learning the robot algorithm must remain adaptable to change, because the human team member is learning at the same time and might therefore change its behavior later on (van den Bosch et al., 2019). This can cause certain state-action pairs, that were previously discarded by the robot algorithm, to now be preferred due to changes in the policy of the human. We therefore defined a requirement that ensures that the robot always keeps exploring:

R4 The RL algorithm can continuously adapt its behavior during all stages of the learning process.

4.2.5. TRANSPARENCY

MUTUAL transparency is crucial for the understanding of each others' contribution to the joint task performance (Haresamudram et al., 2023; Verhagen et al., 2022; Vössing et al., 2022). Team members' behaviors and decisions must be observable, predictable and directable in a collaboration (Johnson et al., 2014; van den Bosch et al., 2019). Both team members should be able to observe the state and actions of the other team member, to allow them to adapt to each other to develop patterns of collaboration. Mutual transparency helps to avoid hierarchical inequalities within the team and ensures that both team members are able to properly adapt. This leads to the fifth requirement:

- R5** The human and the robot are able to observe and understand each other's state and actions.

4.3. METHODS

4.3.1. TASK DESIGN

WE designed a human-robot handover task, a common task found in physical human-robot collaboration (Ortenzi et al., 2021). To coordinate the specific moment in which the responsibility of not dropping the handover object switches from one team member to the other, the team members must collaborate to successfully complete the task. This ensures that a symmetrical hard dependency is embedded in the task. Moreover, passing an object involves multiple elements in which soft dependencies can arise. For instance, the position and orientation at which the object is handed over needs to be predicted or learned, thereby allowing for proactively reciprocating the strategy of the other team member. It also allows for different strategies (e.g. the robot drops the object while the human holds its hand up, or the robot conveys the object close to the human until the human seizes it). Therefore, there is space for a human-robot team to explore and learn what works well for their team. The presence of both hard and soft dependencies follows **R3**.

Additionally, the task of handing over an object is relatively short and can either succeed or fail. It is ideal for rewarding the team based on their collaborative performance (**R1**), and, as it is a short task, the team can rehearse the task often in a short amount of time. Therefore, the robot is rewarded on a regular basis, allowing it to continuously update its policy. This contributes to requirement **R2**, as it allows the robot to learn at a human timescale.

To accommodate **R3** more strongly, the task was designed such that responsibilities are divided over both agents, creating dependencies between the human and the robot. We describe the capabilities of the robot, defined by the State-Action space of the RL algorithm, in Section 4.3.1. We explain how we established a fixed set of capabilities for the human by creating a secondary task that limits the human ability to act and observe the environment in Section 4.3.1. Figure 4.1 shows an overview of the whole setup.

STATE-ACTION SPACE (ROBOT CAPABILITIES)

To meet **R2**, the state-action space used by the RL algorithm should be designed such that it enables a sufficient learning pace. RL algorithms usually require a great amount

of training iterations, which is not possible if they need to learn alongside their human team partner. We have reduced the number of training iterations necessary by modeling the task through a state-action space that is as small as possible. This makes it possible to quickly explore all possible state-action pairs, also enabling relearning and therefore **R4**.

The actions modeled are a set of seven predetermined movements, and the states are determined by a set of four binary state factors, visualized in Figure 4.2. The state factors mostly contain information about the human team member. Therefore, they provide the robot with some transparency of the human (**R5**). Moreover, the handover task is broken up into three phases. In each phase, only one or two of the state factors are taken into consideration, thereby effectively breaking the task up into three separate learning problems to further reduce the complexity. In each phase, the robot has different actions that it can choose from, as shown in Figure 4.2.

Phase 1 describes the start of the handover, during which the robot needs to learn when to start handing over the object. The robot can observe whether the hand of the human team member is in the robot's workspace. The robot has two available actions: waiting until the state changes, or moving the object towards the human with the action *Go to human*.

When the robot takes the action (*Go to human*), it moves to Phase 2, during which the robot is moving towards the human. While moving, it can decide on the orientation it will use to hand over the object. The state factor that can be observed is the orientation of the human hand. The robot can choose between two predetermined orientations: the palm of the robot hand facing up (*Serve*), so that the human can take it out, or the palm facing down (*Drop*), to drop the object into the human hand.

After either action in Phase 2, the robot will move to Phase 3, in which the focus is on the handover. The robot needs to learn when and how far to open its hand, while the human needs to grasp or catch the object to prevent it from falling. The human can influence the robot's behavior by pulling on the object and displacing the end-effector. An additional state factor describes whether the robot hand is still fully closed or partially opened. This combination of capabilities presents opportunities for multiple strategies and soft dependencies to emerge. For example, the robot can learn to wait until the end-effector is displaced before opening its hand to ensure that the human is already holding the object when the robot lets go, or to open its hand enough to allow the human to take the object out without dropping it.

SECONDARY TASK (HUMAN CAPABILITIES)

A secondary task was introduced for the human in the form of a game-like task, that could only be completed if the human successfully received the handover object. The introduction of a secondary task gives the human incentive to complete the task by rewarding the human for the collaborative performance, to ensure a shared goal (**R1**). It also creates a motive for the human to get the object handed over from the robot, thereby creating a hard dependency (**R3**). Moreover, the secondary task could compensate for the superior ability of the human to observe task, to prevent a transparency inequality (**R5**) and give control over the capabilities of the human.

The secondary task we designed required the human to track an asteroid on a screen. For that, they needed to continuously have one hand on a mouse and constantly have




Figure 4.2: A flow diagram that shows the capabilities of the robot throughout three phases of the task. The figure shows the binary state factors (rectangles) and the possible actions (ellipsoids). Actions are red if they do not affect the environment, and blue if they result in the robot advancing to the next phase. The yellow action influences the yellow state, as shown with the yellow arrow.

their eyes on the screen. They were instructed to deflect the asteroid to complete the game, and that for this to be possible, they needed a physical object that served as a projectile. They could not get up and get the object themselves, as they needed to keep tracking the asteroid on the screen. The task consisted of two stages. During the first

stage, the human needed to keep their second hand on a button, until a loading bar was filled. As soon as the bar was full, they could let go of the button, while a timer started to count down, starting the second stage. In this stage, the human had twenty seconds to receive the object from the robot with their free hand, while still tracking the target on the screen with the other. When the team succeeded, the human was rewarded with the same score as the robot. If the task failed, both were rewarded negatively. Auditory feedback was provided in addition to reward screens, to engage the human.

4.3.2. ROBOT REINFORCEMENT LEARNING ALGORITHM

AFTER comparison of multiple RL algorithms (Akalın and Loutfi, 2021; Haarnoja et al., 2018; Ramstedt and Pal, 2019) and their suitability for embodied human-robot co-learning applications, we chose to extend and adapt a Q-learning algorithm (Watkins and Dayan, 1992). Q-learning is a simple and robust RL technique, that is often used in the domain of social robotics (Hemminahaus and Kopp, 2017; Moro et al., 2018; Papaioannou et al., 2017) and specifically co-learning (van Zoelen, van den Bosch, and Neerinx, 2021). We adapted the Q-learning algorithm using decomposition techniques based on MAXQ value decomposition (Dietterich, 2000) and extended it with eligibility traces (Sutton, 1984) to specifically meet our design requirements.

DECOMPOSITION

Hierarchical RL with MAXQ Value Function Decomposition (Dietterich et al., 1998) decomposes the learning problem into multiple smaller problems with a hierarchical structure, resulting in faster learning (Chan and Nejat, 2012). Splitting the problem into smaller problems can also increase adaptability (Dietterich, 2000), as the policy of one phase of the learning problem can change without affecting the policies of other phases.

The idea of decomposing the problem is based on the concept that not every state variable is important in every phase of the task. The three phases in our task (see Figure 4.2) are however sequential instead of hierarchical, meaning that they can not be decomposed using Dietterich's hierarchical value decomposition (Dietterich et al., 1998). Therefore, we instead decomposed the learning problem into three sequential Q-learning problems, each with their own Q-table, creating the same effect of decreasing the amount of Q-values without affecting the amount of actions and state variables. The Q-values are thus a function of state (s), action (a) as well as phase (ϕ):

$$Q_*(\phi, s, a) = E[R(\phi', s') + \gamma \max_{a'} Q_*(\phi', s', a')]. \quad (4.1)$$

By decomposing the task, we provided the robot with information about the importance of state variables in different phases of the task. This significantly reduces the number of states and thereby decreases the scale of the learning problem, increasing the overall learning pace and adaptability of the RL agent, ensuring **R2** and **R4**.

REWARD FUNCTION

Design requirement **R1** states that both agents get rewarded based on performance, and that both agents get rewarded similarly. Both agents therefore received either positive or negative feedback at the end of each episode. This reward was based on whether the handover was completed successfully without dropping the object, as well as the time

left to do so. The robot received a positive reward of (+10) if the task was completed successfully, and a negative of (−10) if the task failed. Additionally, if the task succeeded, the amount of seconds left to complete the task was added to the positive reward. As the team was given 20 seconds at the start of the task, the positive reward was always between +10 and +30. The reward function was extended with a small punishment (−1) for each action necessary to prevent a policy where the robot gets stuck in a loop. To accommodate **R1**, the human would see this same reward as a score given for the completion of the task.

ELIGIBILITY TRACES

Rewarding the Q-learning algorithm at the end of each episode creates two problems. First of all, most actions get a delayed reward (Sutton, 1992), making the learning pace slow. Second, due to the decomposition that we implemented, the Markov property is not satisfied across the whole task (as the task is decomposed into separate learning problems). Therefore, regular backpropagation does not work as effectively as it would otherwise. We solved these problems with eligibility traces (Sutton, 1984).

An eligibility trace is a trace of all previously visited Q-values. These traces are stored in a table for each state-action pair in each phase $S(\phi, s, a)$. Using eligibility traces, the algorithm tracks all state-action pairs reached during the episode. When a reward is received at the end of an episode, it updates all corresponding Q-values based on this reward. This not only speeds up the learning process, but it also ensures that mistakes made in early phases of the task get rewarded negatively in case of an unsuccessful episode (Singh and Sutton, 1996).

With eligibility-traces, all Q-values are updated after every action. To do so, we first calculate what would have been the updated Q-value for the last phase-state-action combination $\hat{Q}(\phi, s, a)$ shown in (4.2a) using the decomposed Bellmann equation (4.1). Then we use \hat{Q} to calculate the update-value Δ_Q (4.2b):

$$\hat{Q}(\phi, s, a) = R(\phi', s') + \gamma \max_{a'} Q(\phi, s', a'), \quad (4.2a)$$

$$\Delta_Q = Q(\phi, s, a) - \hat{Q}(\phi, s, a). \quad (4.2b)$$

This update-value (Δ_Q) is then used to update all Q-values based on their eligibility. As shown in (4.3):

$$Q(\phi, s, a) = Q(\phi, s, a) + \alpha \Delta_Q S(\phi, s, a) \quad \forall S(\phi, s, a). \quad (4.3)$$

The learning rate α is a value between 1 and 0. It is used to determine to what extent new experiences override what has been learned already.

EPSILON DECAY

The algorithm uses epsilon decay to balance exploration and exploitation. Usually, methods for balancing exploration and exploitation are designed to converge to a greedy policy. This is not beneficial for adaptability in the later stages of the learning process. In our algorithm, ϵ never completely decays to zero. This enhances **R4** as the system must never stop exploring to stay adaptable during all stages of the learning process.

The reward (R) and the epsilon decay rate (γ_ϵ) are used to update ϵ as follows:

$$\epsilon = \begin{cases} \max(\gamma_\epsilon \epsilon, 0.2) & \text{if } R < 0 \vee \epsilon > 0.5 \\ \min(\frac{1}{\gamma_\epsilon} \epsilon, 0.5) & \text{if } R > 0 \end{cases} \quad (4.4)$$

Epsilon starts at a value of 1 to guarantee exploration when no policy is learned yet. Epsilon then decays until it reaches a 50% change of exploration ($\gamma_\epsilon = 0.9$). During the rest of the episodes, ϵ decays further when the team has a high success rate, so the robot has a higher chance to exploit its current policy. When the team experiences more failure, the chance of exploring grows.

4.3.3. EVALUATION

To test whether the designed task and robot algorithm would allow for co-learning, we evaluated the setup with human participants using a qualitative, case-by-case analysis. Co-learning is an open-ended process in which good task performance can manifest in many different ways, due to the multiple possible strategies that can be taken by the team. A qualitative approach can provide 1) an understanding of how individual co-learning behaviors evolve over time, and 2) insight into how our proposed new co-learning task provides an environment to hypothesize and quantify co-learning processes in the future. We aim to contribute to existing work on Co-Learning that also takes a qualitative approach (Kumar et al., 2024; Shafti et al., 2020). Six human participants (students from Master programs at Delft University of Technology) each performed the task for four sessions of ten minutes (40 minutes of co-learning per dyad in total). This resulted in approximately 20 to 30 handover attempts per session. Participants received written information about the procedure beforehand, and were further instructed verbally. We used six measurements for our evaluation: performance, subjective Collaboration Fluency, behavioral strategies, relative liability, action preference from Q-values, and answers to interview questions. The procedure was approved by the Human Research Ethics Committee at Delft University of Technology.

PERFORMANCE

To track performance over time, we stored whether each attempt of the task was successful or not, and calculated the percentage of successful attempts per every ten-minute session.

SUBJECTIVE COLLABORATION FLUENCY

To measure how the human participants experienced the collaboration, they were asked to complete a survey on human-robot Collaboration Fluency (Hoffman, 2019) after each ten-minute session.

BEHAVIORAL STRATEGIES

As described earlier in the paper, there are several possible strategies that all lead to a successful handover. Considering the task design, we identified three possible strategies:

- S1** The robot lets the object go, trusting the human will catch it.

S2 The human pulls on the object, letting the robot know it can let go.

S3 The robot opens its hand partially, letting the human take the object.

We recorded videos of the participants, such that we could qualitatively assess which strategies were followed in successful attempts.

RELATIVE LIABILITY

Relative liability describes the proportion in which team members caused episodes to fail in each 10-minute session. It portrays the relative learning pace of both agents, since when the learning pace is similar, it should stay the same for both agents over time. If one team member learns faster than the other, there is a shift in relative liability because the proportion of mistakes made by the superior agent goes down. We determined relative liability by checking which agent made the mistake that caused the episode to fail in the case of a failed episode.

We considered the robot to be responsible for failure when the object was not passed within the allocated time when the human did try to signal the robot, or when the robot dropped the object without the human touching it. For any other reason of failure, we considered the human liable.

ACTION PREFERENCE FROM Q-VALUES

Action preference describes the specific policy of the robot in different phases of the task. We have evaluated the Q-tables after each episode, to track which action the robot preferred in each state and whether and how this changed as the experiment progressed. This gave us insight into the behavior learned by the robot, as well as how adaptable this behavior is (how much it changes over time).

INTERVIEW

We conducted a short interview with each of the participants after the last learning, in which we asked the following three questions:

Q1 Please indicate what your objective was during the learning process.

Q2 Describe the different strategies that you used, and how did this change over time.

Q3 Did you rely on a specific strategy of the robot?

The first question was asked to investigate whether the goal of the human corresponded to the goal of the robot, to test whether the team had a shared goal **R1**. The second question was used to find if the human explored different strategies during the learning process, and more specifically whether it converged towards preferring one strategy over other strategies. With the last question, we intended to find whether the human experienced soft dependencies.

4.4. RESULTS

We present the results of the evaluation by analyzing whether we succeeded in facilitating the design requirements for each human-robot dyad (summarized in Table 4.1).

Table 4.1: Overview of each requirement and whether it was met during the experiment in each team. Additionally, the bottom row shows whether the results indicate that co-learning took place during the experiment. The content of the bottom row is discussed in Section 4.5.

Teams	A	B	C	D	E	F
R1 - Shared Goal	X	✓	✓	-	✓	✓
R2 - Synchrony	✓	✓	-	X	✓	✓
R3 - Interdependence	✓	✓	✓	✓	-	✓
R4 - Adaptability	✓	✓	-	✓	-	✓
R5 - Transparency	✓	✓	X	-	✓	✓
Co-learning	✓	✓	-	-	-	✓

4

4.4.1. SHARED GOAL (R1)

PARTICIPANTS B, C, E, and F indicated in the interview that their goal was to complete each episode without dropping the object. Participants B and F even indicated that they had a secondary goal of improving the time in which they succeeded, to optimize their score. Therefore, it is shown in Table 4.1 that requirement **R1** is met in these teams.

Participants A and D indicated that their main goal was not to succeed at the task, but to train the robot to follow their preferred strategy. Participant A said that this was their main objective during the whole experiment. For instance, they never let the task succeed if the robot did not let go of the object. This is why strategy **S3** is never seen in team A in Figure 4.3. Additionally, it can be seen that the human's perception of Fluency is relatively high, while the team's performance is low. This can be explained by the fact that the human met their objective of influencing the robot's behavior, at the expense of the robot's goal of succeeding at the task. This resulted in both agents perceiving different rewards.

Participant D indicated that they changed their objective between session three and four. First, it matched the robot's objective, while during the last session their goal was only to train the robot to their preferred strategy. Participant D explained that their goal changed due to the realization that the robot used trial-and-error learning. They realized that they could influence the robot's behavior by rewarding desired behavior and punishing undesired behavior. Participant D therefore suddenly changed their behavior, resulting in the performance drop in the 4th session (Figure 4.3) and the human deliberately failing the task to train the robot (Figure 4.4). This behavior resulted in the team having a shared goal for only a part of the task, leaving **R1** inconclusive in team D.

4.4.2. SYNCHRONY (R2)

FIGURE 4.4 shows that relative liability is relatively constant for teams A and E, which means that the learning of the human and the robot was synchronous. For teams B and E, there is an initial large shift towards the human being responsible for a large part of the failures, but this recovers to the middle over subsequent sessions. While the robot learned faster in the beginning, the human was able to catch up, leading to synchronous learning overall.

In team D we can observe a shift in relative liability in the first 3 sessions, as the robot




Figure 4.3: The distribution of the three different strategies in phase 3 that could lead to a successful handover, combined with the performance and Collaboration Fluency score. The figure shows how the preference for different strategies changes over time. The three strategies are as follows:

- S1:** The robot lets go of the object, trusting the human will catch it.
- S2:** The human pulls on the object, letting the robot know it can let go.
- S3:** The robot opens its hand partially, letting the human take the object.

could not keep up with the learning of the human. Therefore, this synchrony requirement was not met. Moreover, the human still learned faster than the robot during the fourth session, even though Figure 4.4 seems to imply that the robot made a recovery. The reason this proportion drops back towards 50%, however, is that the human started to deliberately fail the task to actively train the robot to prefer strategy **S2**, as mentioned by the participant in the interview. This can also be seen by the sudden decrease in performance rate during this session in Figure 4.3.

In team C, a shift in relative liability can be seen in Figure 4.4. In this case, the robot




Figure 4.4: The percentage of how many times each agent was responsible for failing an episode during each 10-minute session is shown for each team (relative liability). The rate of failure of the robot can be read on the left axis, while the human failure rate is displayed on the right axis.

improved its policy faster than the human. Combined with the fact that the team barely improved their performance during the four sessions (see Figure 4.3), we can deduce that the human did not improve its policy at all. Therefore, no conclusion could be drawn about this requirement for this team.

4.4.3. INTERDEPENDENCE (R3)

DIFFERENT individuals prefer different strategies. Figure 4.3 shows that the method enables different teams to learn different strategies. Team A, for instance, converges completely to strategy **S1**, while teams B and D learned that this strategy did not work for them.

Post-experiment interviews revealed some important underlying insights about the development of the strategies during co-learning. Participants A and C stated that they did not want to take the object from the robot without its permission, in an attempt to maintain the trust of the robot. This complies with the action preferences from the Q-values, visualized in Figure 4.3 that shows that strategy **S3** was not preferred in these teams. By actively not choosing this strategy, the human depends on the robot to open its hand completely for the task to be completed. This shows the establishment of a soft dependency between the human and the robot, which is beneficial for the team's relationship.

All three observed strategies contain their own similar soft dependencies. This means that soft dependencies arise during the learning process, when a team converges to pre-

ferring one specific strategy. It can be seen in Figure 4.3 that in all teams multiple strategies were explored. In teams A, B, D, and F, there was convergence to one specific strategy during the experiment. Thus, soft dependencies emerged in these teams. Team E is the only team that kept executing all strategies until the end of the experiment. This means that both team members never fully committed to being dependent on the other, making it the only team for which it is inconclusive whether design requirement R3 is met.

4.4.4. ADAPTABILITY (R4)

FIGURE 4.3 clearly shows that teams A, B, D, and F made a relatively drastic change in preferred strategy towards the end of the experiment. This shows that the RL algorithm was able to adapt its policy to accommodate strategy changes.

In teams C and E, no large change in the policy of the robot occurred during the experiment. However, this does not necessarily mean that the robot had no adaptability, as the result could have also been caused by behavior of the human. Therefore, these cells are inconclusive in Table 4.1.

4.4.5. TRANSPARENCY (R5)

THE method allowed both agents to observe each other by design, as explained earlier in the paper. We attempted to avoid an imbalance in learning pace by ensuring that both human and robot had access to a similar level of limited information about the other.

Figure 4.4 shows that there was indeed no imbalance in learning pace in teams A, B, E, and F, as explained in Section 4.4.2. The unequal learning pace in team C, however, was caused by the fact that the human was not able to understand the policy of the robot. This was a result of the human being too occupied by the secondary task, resulting in unbalanced transparency. There is no indication that the unequal learning pace in team D was caused by the same issue.

While the secondary task prevented the human from constantly looking at the robot, as explained in Section 4.3.1, Figure 4.3 shows that in teams B, D, and F, the human preferred to rely on tactile sensing to know where to grasp the object, as they do not follow strategy S1. They were therefore able to compensate for their lack of visual observation. Further investigation of the video recordings of the experiment showed that participants A and E also relied on tactile sensing to locate the object, they just did it in a subtle manner, so as to not displace the robot.

In short, in teams A, B, E, and F we can state that both agents had transparency and that no unwanted imbalance appeared. This means the requirement is met. In team C, the secondary task over-hindered the human's ability to visually observe the robot, causing this requirement not to be met, while in team D the results are inconclusive.

4.5. DISCUSSION

4.5.1. FACILITATING CO-LEARNING

IN summary, in three out of six teams we managed to create the circumstances for Co-Learning. Table 4.1 shows that all the requirements are met in team B and team F, meaning that these teams demonstrated successful Co-Learning. Nevertheless, partially

fulfilling the requirements can still mean that there was some degree of Co-Learning, as shown by the development of interesting and useful human-robot collaboration patterns. For example, even though the human and the robot did not have the same goal (R1) in team A, they still managed to improve their collaboration by co-learning a joint strategy. The reason R1 was not met in team A, is that the participant chose to train the robot, while the goal of the robot was to succeed at the task. In practice, however, there are still multiple congruent strategies that reach both goals. Moreover, Figure 4.3 shows an increase in performance over time for team A, as well as a growth in the participant's perception of fluency in the team. The better team performance may be the result of effective collaboration patterns, and the improved fluency suggests that soft dependencies emerged in team A.

In team C, we observed that the participant struggled to understand how to do the task and seemed unable to learn this within the given time. Still, the team developed some collaboration patterns and soft dependencies. However, by not meeting R2, R4 and R5 we cannot claim that this team was able to achieve a full Co-Learning process.

In team D, the human learned faster than the robot, which led to an imbalance in contribution over time (Figure 4.4). This imbalance may have caused the human to change their motivation over time. Even though there were indications for co-learning during the first sessions of the experiment, it was not sustained during the last session. Therefore, in team D, multiple design requirements were left inconclusive or were not met.

In Team E, we did not see soft dependencies arise during the experiment. Additionally, Figure 4.3 does not show an increase in performance or fluency. However, changes in preferred collaboration patterns over time can still be observed in Figure 4.3. They are not substantial enough to prove that R3 or R4 were met, but it does suggest that a longer collaboration could have resulted in the requirements being met. Furthermore, Figure 4.4 shows a balanced learning pace between the two agents. Overall, it seems that the team was still exploring after the four sessions, and full Co-Learning might have emerged in a longer collaboration.

4.5.2. LIMITATIONS AND FUTURE WORK

WHILE the results showed that the developed task-, algorithm- and interaction design enabled co-learning, the short-term performance was not increased for most teams (Figure 4.3). This can be explained by the fact that the designed method and evaluation focused on the Co-Learning process. The essence of designing for Co-Learning is to design the conditions in which collaborators can learn the behavior needed to develop smooth and effective collaboration, which can facilitate long-term performance improvement. This means that Co-Learning can be present without an immediate performance increase. We expect that when a similar experiment is done for a longer duration of time, with more participants, an increase in performance should be measurable in teams where Co-Learning is identified. The work presented in this paper can serve as a basis for future larger, quantitative studies into long-term effects of Co-Learning in embodied human-robot teams.

This study focused on facilitating human-robot Co-Learning through five requirements. As we were able to identify co-learning in at least one of the teams (team A)

despite one of the requirements not being met, the individual effect and weight of each design requirement requires further research.

BIBLIOGRAPHY

- Ajoudani, A., Zanchettin, A. M., Ivaldi, S., Albu-Schäffer, A., Kosuge, K., & Khatib, O. (2018). Progress and prospects of the human–robot collaboration. *Autonomous Robots*, 42(5), 957–975.
- Akalin, N., & Loutfi, A. (2021). Reinforcement learning approaches in social robotics. *Sensors*, 21(4), 1292.
- Amirshirzad, N., Kumru, A., & Oztop, E. (2019). Human adaptation to human–robot shared control. *IEEE Transactions on Human-Machine Systems*, 49(2), 126–136.
- Burke, C. S., Stagl, K. C., Salas, E., Pierce, L., & Kendall, D. (2006). Understanding team adaptation: A conceptual analysis and model. *Journal of Applied Psychology*, 91(6), 1189.
- Chan, J., & Nejat, G. (2012). Social intelligence for a robot engaging people in cognitive training activities. *International Journal of Advanced Robotic Systems*, 9, 1. <https://doi.org/10.5772/51171>
- Dietterich, T. G., et al. (1998). The maxq method for hierarchical reinforcement learning. *ICML*, 98, 118–126.
- Dietterich, T. G. (2000). Hierarchical reinforcement learning with the maxq value function decomposition. *Journal of artificial intelligence research*, 13, 227–303.
- Doorewaard, H., Van Hootegem, G., & Huys, R. (2002). Team responsibility structure and team performance. *Personnel review*, 31(3), 356–370. <https://doi.org/10.1108/00483480210422750>
- Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *International conference on machine learning*, 1861–1870.
- Haresamudram, K., Larsson, S., & Heintz, F. (2023). Three levels of ai transparency. *Computer*, 56(2), 93–100.
- Hemminahaus, J., & Kopp, S. (2017). Towards adaptive social behavior generation for assistive robots using reinforcement learning. *2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 332–340.
- Hendrikse, S. C. F. (2024). *Interpersonal synchrony: Analyzing and modeling social interaction dynamics* (tech. rep.) (PhD Thesis). Vrije Universiteit Amsterdam.
- Hoffman, G. (2019). Evaluating fluency in human–robot collaboration. *IEEE Transactions on Human-Machine Systems*, 49(3), 209–218.
- Ikemoto, S., Amor, H. B., Minato, T., Jung, B., & Ishiguro, H. (2012). Physical human-robot interaction: Mutual learning and adaptation. *IEEE robotics & automation magazine*, 19(4), 24–35.
- Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., van Riemsdijk, M. B., & Sierhuis, M. (2014). Coactive design: Designing support for interdependence in joint activity. 3(1), 43–69. <https://doi.org/10.5898/JHRI.3.1.Johnson>

- Katz-Navon, T. Y., & Erez, M. (2005). When collective- and self-efficacy affect team performance: The role of task interdependence. *Small Group Research*, 36(4), 437–465. <https://doi.org/10.1177/1046496405275233>
- Kumar, R., Huan, Y., & Yao, N. (2024). Unveiling the dynamics of human decision-making: From strategies to false beliefs in collaborative human-robot co-learning tasks. *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 632–636. <https://doi.org/10.1145/3610978.3640743>
- Li, Y., Zhang, Z., & Zhang, F. (2023). Hybrid co-learning for proximate human-robot teaming. *20th International Conference on Ubiquitous Robots (UR)*, 239–244. <https://doi.org/10.1109/UR57808.2023.10202393>
- Loopik, H. (2022). *Comparing reinforcement learning techniques for embodied co-learning in a human-robot team* (tech. rep.) (MSc Literature Study). Technische Universiteit Delft.
- Moro, C., Nejat, G., & Mihailidis, A. (2018). Learning and personalizing socially assistive robot behaviors to aid with activities of daily living. *J. Hum.-Robot Interact.*, 7(2). <https://doi.org/10.1145/3277903>
- Mörtl, A., Lorenz, T., & Hirche, S. (2014). Rhythm patterns interaction-synchronization behavior for human-robot joint action. *PLoS one*, 9(4), e95195.
- Nikolaïdis, S., Hsu, D., & Srinivasa, S. (2017). Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research*, 36(5-7), 618–634.
- Ortenzi, V., Cosgun, A., Pardi, T., Chan, W. P., Croft, E., & Kulic, D. (2021). Object handovers: A review for robotics. *IEEE Transactions on Robotics*, 37(6), 1855–1873.
- Papaioannou, I., Dondrup, C., Novikova, J., & Lemon, O. (2017). Hybrid chat and task dialogue for more engaging hri using reinforcement learning. *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 593–598.
- Peternel, L., Oztop, E., & Babič, J. (2016). A shared control method for online human-in-the-loop robot learning based on locally weighted regression. *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3900–3906.
- Ramstedt, S., & Pal, C. J. (2019). Real-time reinforcement learning. *CoRR*, abs/1911.04448.
- Rijgersberg-Peters, R., van Vught, W., Broekens, J., & Neerincx, M. A. (2023). Goal ontology for personalised learning and its implementation in child's health self-management support. *IEEE Transactions on Learning Technologies*.
- Shafti, A., Tjomslund, J., Dudley, W., & Faisal, A. A. (2020). Real-world human-robot collaborative reinforcement learning. *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 11161–11166. <https://doi.org/10.1109/IROS45743.2020.9341473>
- Singh, S. P., & Sutton, R. S. (1996). Reinforcement learning with replacing eligibility traces. *Machine learning*, 22(1-3), 123–158.
- Sutton, R. S. (1984). Temporal credit assignment in reinforcement learning [Technical Report].
- Sutton, R. S. (1992). Introduction: The challenge of reinforcement learning. *Reinforcement learning*, 1–3.

- van den Bosch, K., Schoonderwoerd, T., Blankendaal, R., & Neerincx, M. (2019). *Six challenges for human-ai co-learning* (Vol. 11597 LNCS).
- van Zoelen, E. M., van den Bosch, K., & Neerincx, M. (2021). Becoming team members: Identifying interaction patterns of mutual adaptation for human-robot co-learning. *Frontiers in Robotics and AI*, 8. <https://doi.org/10.3389/frobt.2021.692811>
- van Zoelen, E. M., van den Bosch, K., Rauterberg, M., Barakova, E., & Neerincx, M. (2021). Identifying interaction patterns of tangible co-adaptations in human-robot team behaviors. *Frontiers in Psychology*, 12. <https://doi.org/10.3389/fpsyg.2021.645545>
- Verhagen, R. S., Neerincx, M. A., & Tielman, M. L. (2022). The influence of interdependence and a transparent or explainable communication style on human-robot teamwork. *Frontiers in Robotics and AI*, 9, 993997.
- Vössing, M., Kühl, N., Lind, M., & Satzger, G. (2022). Designing transparency for effective human-ai collaboration. *Information Systems Frontiers*, 24(3), 877–895.
- Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3), 279–292. <https://doi.org/10.1007/BF00992698>

5

ONTOLOGY-BASED REFLECTIVE COMMUNICATION FOR SHARED HUMAN-AI RECOGNITION OF EMERGENT COLLABORATION PATTERNS

When humans and AI-agents collaborate, they need to continuously learn about each other and the task. We propose a Team Design Pattern that utilizes adaptivity in the behavior of human and agent team partners, causing new Collaboration Patterns to emerge. Human-AI Co-Learning takes place when partners can formalize recognized patterns of collaboration in a commonly shared language, and can communicate with each other about these patterns. For this, we developed an ontology of Collaboration Patterns. An accompanying Graphical User Interface (GUI) enables partners to formalize and refine Collaboration Patterns, which can then be communicated to the partner. The ontology was evaluated empirically with human participants who viewed video recordings of joint human-agent activities. Participants were requested to identify Collaboration Patterns in the footage, and to formalize patterns by using the ontology's GUI. Results show that the ontology supports humans to recognize and define Collaboration Patterns successfully. To improve the ontology, it is suggested to include pre- and post-conditions of tasks, as well as parallel actions of team members.

 **Emma M. van Zoelen**, Karel van den Bosch, David A. Abbink, Mark A. Neerinx. 2022. Ontology-Based Reflective Communication for Shared Human-AI Recognition of Emergent Collaboration Patterns. In *International Conference on Principles and Practice of Multi-Agent Systems*, 621-629. Springer International Publishing.

5.1. INTRODUCTION

A growing body of research on human-agent teaming (Schneider et al., 2022; van Diggelen et al., 2021) and human-robot collaboration (Ajoudani et al., 2018) studies how to make optimal use of the qualities of both humans and AI agents by making them team partners. An important aspect of becoming successful team members is to continuously learn about each other and the task, to make sure the team becomes a fluently functioning unit; a process called co-learning (Schoonderwoerd et al., 2022).

To us humans, adapting to and learning with our fellow human team members often comes natural. In a hybrid human-agent team, successful adaptation and learning is not self-evident, as humans and AI agents differ in the way in which they learn and adapt. Still, implicit co-adaptation is bound to occur as the learning processes of human and agent will influence each other while they collaborate on a task. As a result, new team behaviors will emerge (van Zoelen et al., 2021); successful emergent team behaviors (coordination and cooperation as defined by Engeström, 1992) can be specified as ‘Collaboration Patterns’¹. To make sure that the team can successfully co-learn by achieving reflective communication (the highest quality collaboration), it is necessary for team members to consolidate and organize their collaborative efforts by developing a shared model of the collaboration (Engeström, 1992).

Based on exploratory work done by van Zoelen et al., 2021, this process can be expressed as a Team Design Pattern (Table 5.1) that describes how human and agent team members co-learn by communicating about emergent Collaboration Patterns. In this paper we address the following research question:

What kind of model and communication interface enables a human-agent team to establish shared recognitions of emergent Collaboration Patterns?

Existing frameworks of collaboration usually predefine Collaboration Patterns (e.g. as Plays (Miller, 2005) or Social Practices (Dignum, 2018)), and do not study how Patterns can be created or updated through communication during human-agent collaboration. We propose a model in the form of an ontology and an accompanying communication interface that can be used for communication about and formalization of emergent Collaboration Patterns. An advantage of ontologies is that they provide shared univocal conceptualizations that can be used for communication and reasoning (Fensel, 2001). The concepts and structure of the ontology and its communication interface were evaluated empirically, with human participants. This ontology-based reflective communication enables human-agent co-learning of successful coordination and collaboration behaviors.

5.2. ONTOLOGY: REQUIREMENTS AND BACKGROUND

5.2.1. REQUIREMENTS FOR COLLABORATION PATTERN ONTOLOGY

To enable an agent to reason and communicate about patterns of collaboration, the agent needs a representation of relevant concepts that underlie the collaboration. Relevant concepts are, for example, the entities (e.g. actors, objects) that take part in the collaboration, details about the actions that should be executed (e.g. which actions happen when, and in what order) and the context in which the Collaboration Pattern

¹Previously called ‘Interaction Patterns’ in Chapter 2 and 3

Name	Human-AI Co-Learning
Description	When human and an adaptive AI agent collaborate as team partners, they both adapt their behavior constantly to dynamically changing requirements of the task. When doing so, Collaboration Patterns emerge. A team member recognizes a CP as valuable to the task, and communicates this to its team member. By jointly reflecting on the CP they can refine or adjust the CP until they agree on its value and use. The CP is defined and stored in the shared ontology; now both team partners are explicitly aware of this CP, and can use it when relevant.
Structure	

Table 5.1: Team Design Pattern for co-learning in a human-agent team. The pattern supports definition of emergent Collaboration Patterns in a shared ontology, enabling partners to communicate about them. In this paper we focus on the dashed arrows on the human side.

takes place (e.g. when does it start and end). Together these concepts should support the agent to connect a particular pattern of events (e.g., a series of collaborative actions, hence an 'Collaboration Pattern') to a particular instance of a context, based on specific contextual conditions. In addition, the model should allow the agent to determine the success of a Collaboration Pattern, in terms of its contribution (or harm) to the team's task.

As we want to use the model for communication during collaboration, and for storing and updating newly emergent CPs, it is necessary that it can be dynamically updated. Given that the ontology should function as a shared model between team members, the structure and concepts of the ontology should be fixed, while instances of specific Collaboration Patterns can be updated. This is analogous to a frame-based approach, in which an Upper Ontology describes concepts and relations in a generic manner, while a Lower Ontology describes unique instances of the concepts and relations (Fensel, 2001).

To summarize, the requirements for the ontology are as follows: (1) it should store and specify the structure of patterns of collaboration; (2) it should contain a model of context at a level that is understandable by humans; (3) it should support the agent to reason about the appropriateness of patterns in specific contexts; (4) it should allow live updating; (5) there should be a distinction between high level concepts (Upper Ontology) that provide the structure of the model, and low level instances (Lower Ontology) that can be used directly in a task by the team members.

5.2.2. ONTOLOGIES IN HUMAN-AGENT TEAMING

THERE are several existing ontologies in the areas of human-agent teaming and human-robot collaboration. Some of these used for human-agent teaming represent knowledge about team configuration (Diggelen et al., 2019; Pico-Valencia et al., 2019), but as they do not represent information about collaborative actions, these are not directly useful for our purposes. Some ontologies that describe actions done by the team members focus on high-level tasks and the accompanying goals (e.g. Madni and Madni, 2018), while some more recent papers touch upon complex aspects of team behavior, such as the integration of intent (e.g. Schneider et al., 2022), or a combination of tasks, goals and intent by introducing Plays in their ontology (e.g. Kasmier et al., 2021; van Diggelen et al., 2021). A play is a set of instructions that tells actors in a team how to act. This concept is similar to our notion of Collaboration Patterns, although a play is a predefined set of instructions, whereas CPs emerge and develop during collaboration. In Kasmier et al., 2021, the ontology contains the concept ‘Play’, of which there are several domain-specific sub-plays. The content and structure of these plays are defined outside of the ontology. In van Diggelen et al., 2021 the authors use predefined plays as part of their ontology as well. While the concept of plays is relevant for our work, the reported ontologies that represent plays do not provide information on the structure of a play and therefore do not meet requirement 1.

For human-robot collaboration, many ontologies exist that enable robots to behave autonomously in a certain practical task (Olivares-Alarcos et al., 2019). They contain concepts such as object, task, actor, etc., and have a large overlap with task models (such as Welie et al., 1998), but add aspects that are necessary for robots, such as hardware knowledge (e.g. about sensors and actuators that the robot is equipped with). These ontologies sometimes support communication to humans, but the focus of that communication is on context-dependent information sharing. The task model related aspects of these ontologies are reusable for the description of context (requirement 2).

In conclusion, existing ontologies that use task models as described by Olivares-Alarcos et al., 2019; Welie et al., 1998 can provide a basis for formalizing context within an ontology. We used these as a starting point for designing the context part of our ontology (requirement 2).

5.2.3. FRAMEWORKS FOR DESCRIBING PATTERNS OF COLLABORATION

WE have looked at frameworks that describe patterns of collaboration between partners, for example from sociology, that have not been formalized in an ontology. As we aim to represent CPs that are defined while collaborating, we are looking for a structure of the concepts and relations that build up a CP, to incorporate that in the ontology. ‘Plays’ were introduced and are used mostly in human-agent teaming applications in the military (Miller, 2005). They are inspired on playbooks used in sports, to ensure communication about goals and plans between a supervisor and a group of actors. Key is that the supervisor provides instructions and constraints, while the group of actors can act autonomously within the given constraints. A play is similar to our concept of CPs. It is, however, unclear whether there is a formal structure that defines a play.

Another framework is Social Practices, which originates from sociology, although it has been formalized for agent reasoning (Dignum, 2018). Social Practices are described

as ways of doing things that are shared between actors. They contain patterns of behavior that are strongly tied to the specific contexts in which they are executed. Social Practices emerge as humans interact with each other in an environment. Although Social practices are predefined in the formalization by Dignum, 2018, the formalization is detailed and contains the elements that build up a Social Practice. The formalization by Dignum, 2018 uses concepts such as Actors, Roles, and Places to describe context, and also includes concepts such as Possible Actions and Strategies to describe expected sequences of behavior.

We used the Social Practice formalization as a starting point for designing the part of the ontology that represents Collaboration Patterns (requirement 1).

5.3. COLLABORATION PATTERN ONTOLOGY

USING the frameworks mentioned as a starting point, we chose a minimal set of concepts and relations, that can be extended for use in specific domains. This minimal set of concepts and relations serves as the Upper Ontology. In choosing the concepts and relations, we have based ourselves on previous work; we wanted to ensure that our ontology could at least describe CPs found by van Zoelen et al., 2021. An example of such a CP is 'Alternating actively working on the task and waiting for a team member'; in this CP, a human team member clears away some small rocks from a pile, then waits for their agent team member to clear away large rocks from that pile. This CP would take place when there is a rock pile that contains both large and small rocks. In this example, we need tasks (clear away, wait), actors (human, robot) and objects (small and large rocks), as well as a way to describe the order in which tasks happen, and a condition to choose the CP. Figure 5.1 presents a graphical overview of the ontology.

5.4. ONTOLOGY IN PRACTICE: TRANSLATION TO CONTEXT

TO evaluate whether the chosen approach supports humans in expressing Collaboration Patterns to agents, we chose to use a USAR task in which a human-agent team collaborated in saving an earthquake victim from underneath a pile of rocks (see van Zoelen et al., 2021). We created a context-dependent set of specifications for several concepts, as well as a communication interface, based on use of the ontology within this task. The context-dependent set of specifications are translations of the concepts specified in Figure 5.1.

To support a human team partner in defining and communicating a Collaboration Pattern and the context-dependent conditions for its application to an AI-agent, a drag-and-drop graphical user interface (GUI) was developed (see Figure 2). The GUI consisted of predefined 'building blocks' based on the context-dependent concepts.

5.4.1. EVALUATION METHOD

BY evaluating the ontology, we wanted to find out if human participants were able to find emergent Collaboration Patterns in video footage, and what the differences and similarities were between patterns they detect and those previously identified by the researchers. Moreover, we wanted to learn whether participants were able to define these CPs with the ontology, the context dependent specification and the designed GUI,




Figure 5.1: An overview of the Collaboration Pattern ontology. The red items are the concepts, the yellow items are the relations. Dashed lines represent relations that enable an agent to reason about whether CPs fit a certain context.


as well as what combinations of the concepts they used to create their descriptions.

A series of short video clips of human-agent collaboration were presented to participants. These clips were taken from a previous experiment (van Zoelen et al., 2021). All clips contained a fragment of human-agent collaboration in which a Collaboration Pattern was previously identified by the experimenter. The experiment was conducted by video call with each individual participant. It consisted of two parts: (1) participants practiced with the task of the experiment described by van Zoelen et al., 2021, and (2) participants watched the videos and defined Collaboration Patterns they distinguished using the GUI as well as verbally. Ten students participated (7 M, 3 F). All participants were in the final stages of AI-related Master programs. The procedure was approved by the Human Research Ethics Committee at Delft University of Technology on November 12th, 2021.

On top of the Collaboration Patterns that participants described using the GUI, we collected audio recordings of verbal descriptions and answers to interview questions. The GUI descriptions of the Collaboration Patterns were coded in two different ways: (1) correct, incorrect or semi-correct compared to descriptions made by the researchers, and (2) open coded based on how the descriptions were constructed. The verbal de-

Situation

The new situation contained:



What we do

In this new situation, we did this:



+

Figure 5.2: An example Collaboration Pattern and accompanying Pattern Condition in the GUI. The user selects context-factors from colored blocks and drags them into the left grey area ("Situation") to define the contextual conditions of the CP, and into the right grey area ("What we do") to define the actions constituting the CP. The description portrayed is the example mentioned in Section 3, 'Alternating actively working on the task and waiting for a team member'.

5

descriptions of the Collaboration Patterns and the interview data were used to clarify what participants meant in case their GUI descriptions were unclear and for additional exploration of what aspects of the approach could be improved.

5.4.2. RESULTS: CORRECTNESS OF COLLABORATION PATTERNS RECOGNIZED

ONE participant (participant 13) had a hard time with the syntax of the GUI, and described Collaboration Patterns in an extremely minimal manner. As no feed-back was given on the descriptions during the experiment, there was no opportunity for them to learn. Across participants, video 5 was impossible to understand, which was understandable; the CP shown in the video was mostly dependent on the implicit intentions of the human showing the behavior. Participants became better at recognizing and describing that the human was directing the robot as the experiment progressed.

Interestingly, anything that required the human to wait for the robot was ignored in the descriptions. On the other hand, waiting behavior by the robot was expressed often. Moreover, some participants decided to decouple the human and robot behavior completely, describing both separately and therefore ignoring any interaction or coordination between the two team members. Most participants did this for some descriptions.

5.4.3. RESULTS: PATTERNS OF DESCRIPTIONS

THERE were several ways in which people indicated that the human was directing the robot: by standing still, moving back and forth, 'Robot move to Human' and by putting the human actions in the situation description. Several participants also indicated that they had difficulty expressing this. Participants felt the need to describe

causality, sometimes in the form of a pre-condition or trigger for an action, but sometimes also for the post-condition or consequence of an action.

Information (such as locations, objects or agents) was often left out of actions when it was deemed self-evident (e.g. the information is already in the situation description, or two consecutive actions that are done by the same agent). Some-times, participants added extra information to make their descriptions more specific, for example by using double location specifications to describe a more specific location.

About half of all participants took a modular approach in describing the Collaboration Patterns. They used several small CPs to describe behavior observed in one video. Some participants created separate CPs for the human and the robot. These participants often attempted to create a complete set of CPs that would describe all possible behavior sequences observed; therefore, many of these small CPs were reused in several videos.

5.5. DISCUSSION

5

OUR Our work presents an approach for enabling a human-agent team to co-learn, and to identify and share emergent Collaboration Patterns. The ontology has been designed to be generic; it can be used in or easily adapted to other contexts. It should be seen as a first iteration for creating a model to formalize emergent CPs, that can be built upon. To enable a human-agent team to use the ontology of CPs, we designed an interface. We chose to design a drag-and-drop GUI (rather than, e.g., a natural language interface), to restrict formalization of and communication about the CPs to concepts and relations present in the ontology. The GUI allowed us to perform an experiment with human participants. The results show that it supports people in formalizing CPs within the boundaries of the ontology. Research into what interfaces are suitable for communicating about CPs in different task contexts will be valuable.

Results of the evaluation can be used to improve the design of the ontology and the GUI. Participants tended to not explicitly formalize information that was described earlier in the same CP; they assumed that the agent is capable of inferring that this information still holds. To ensure that this does not hamper the process of developing shared representations of CPs, we might want to equip the agent with inferencing capabilities, or expand the interface to guide the human in checking whether assumptions are met. Several participants also formalized human and agent behavior separately. The GUI should support the user more extensively in formalizing the interactions as a CP, rather than a sequential series of actions. This requires elaboration of the ontology, by allowing parallel tasks, and/or the coupling of actions through pre- and post-conditions.

The evaluation was done offline, through video recordings, instead of on-task, based on experienced collaboration. Therefore, it does not address how descriptions might change over time due to behavior portrayed by the agent team partner. A challenge is to implement the additional design requirements, and to investigate whether they support human-agent co-learning on the job. This will help to improve our understanding of how participants' understanding of the task, agent and CPs evolves over time.

5.6. CONCLUSION

THE Ontology in this paper enables a human-agent team to represent emergent Collaboration Patterns explicitly, thus making them available for use in future task situations. The formal representation in the ontology enables partners to correct and refine CPs, based upon new experiences or reflections. This way, the ontology system supports co-learning within the team. The drag-and-drop graphical user interface that we presented provides a common language for the team members by translating high-level concepts from the ontology to more contextualized concepts. Evaluation with human participants showed that people are able to identify relevant CPs from videos, and that they were able to formalize them using the GUI. CPs in which the human directs the agent proved difficult to describe. It is therefore considered necessary to expand the ontology system with a function to represent behaviors conducted in parallel by human and agent, as well as to explicitly provide ways of representing pre- and post-conditions.

Our work contributes to developing representations of emergent Collaboration Patterns that support co-learning by humans and agents. Further research is needed to evaluate, expand and refine the proposed ontology system.

BIBLIOGRAPHY

- Ajoudani, A., Zanchettin, A., Ivaldi, S., Albu-Schäffer, A., Kosuge, K., & Khatib, O. (2018). Progress and Prospects of the Human-Robot Collaboration. *Autonomous Robots*, 42(5), 957–975. <https://doi.org/10.1007/s10514-017-9677-2>
- Diggelen, J. v., Neerincx, M., Peeters, M., & Schraagen, J. M. (2019). Developing Effective and Resilient Human-Agent Teamwork Using Team Design Patterns [Conference Name: IEEE Intelligent Systems]. *IEEE Intelligent Systems*, 34(2), 15–24. <https://doi.org/10.1109/MIS.2018.2886671>
- Dignum, F. (2018). Interactions as Social Practices: Towards a formalization [arXiv: 1809.08751]. *arXiv:1809.08751 [cs]*. Retrieved July 1, 2021, from <http://arxiv.org/abs/1809.08751>
- Engeström, Y. (1992). *Interactive Expertise: Studies in Distributed Working Intelligence. Research Bulletin 83* [ISSN: 0359-5749]. Department of Education, University of Helsinki, Bulevardi 18, SF-00120 Helsinki, Finland. Retrieved June 1, 2022, from <https://eric.ed.gov/?id=ED349956>
- Fensel, D. (2001). *Ontologies: A Silver Bullet for Knowledge Management and Electronic Commerce*. Retrieved June 1, 2022, from <https://link-springer-com.tudelft.idm.oclc.org/book/10.1007/978-3-662-04396-7>
- Kasmier, D., Merrell, E., Kelly, R., Smith, B., Heisey, C., Maki, E., Brittain, M., Ankner, R., & Bush, K. (2021). Ontology of Plays for Autonomous Teaming and Collaboration. *Proceedings of the XIV Seminar on Ontology Research in Brazil (ONTOBRAS)*, 9–22.
- Madni, A. M., & Madni, C. C. (2018). Architectural Framework for Exploring Adaptive Human-Machine Teaming Options in Simulated Dynamic Environments [Number: 4 Publisher: Multidisciplinary Digital Publishing Institute]. *Systems*, 6(4), 44. <https://doi.org/10.3390/systems6040044>
- Miller, C. A. (2005). Delegation Architectures: Playbooks and Policy for Keeping Operators in Charge, 28–35.
- Olivares-Alarcos, A., Beßler, D., Khamis, A., Goncalves, P., Habib, M. K., Bermejo-Alonso, J., Barreto, M., Diab, M., Rosell, J., Quintas, J., Olszewska, J., Nakawala, H., Pignaton, E., Gyrard, A., Borgo, S., Alenyà, G., Beetz, M., & Li, H. (2019). A review and comparison of ontology-based approaches to robot autonomy. *The Knowledge Engineering Review*, 34(0), 1–29. <https://doi.org/10.1017/S0269888919000237>
- Pico-Valencia, P., Holgado-Terriza, J. A., & Sierra Martínez, L. M. (2019). A Preliminary Ontology for Human-Agent Collectives. In F. De La Prieta, A. González-Briones, P. Pawleski, D. Calvaresi, E. Del Val, F. Lopes, V. Julian, E. Osaba, & R. Sánchez-Iborra (Eds.), *Highlights of Practical Applications of Survivable Agents and Multi-Agent Systems. The PAAMS Collection* (pp. 176–187). Springer International Publishing. https://doi.org/10.1007/978-3-030-24299-2_15

- Schneider, M. F., Miller, M. E., Ford, T. C., Peterson, G., & Jacques, D. (2022). Intent integration for human-agent teaming [eprint: [https://onlinelibrary.wiley.com/doi/pdf/10.1002/SystemsEngineering.n/a\(n/a\)](https://onlinelibrary.wiley.com/doi/pdf/10.1002/SystemsEngineering.n/a(n/a), 1-13)], 1–13. <https://doi.org/10.1002/sys.21616>
- Schoonderwoerd, T. A. J., Zoelen, E. M. v., Bosch, K. v. d., & Neerincx, M. A. (2022). Design patterns for human-AI co-learning: A wizard-of-Oz evaluation in an urban-search-and-rescue task. *International Journal of Human-Computer Studies*, 164, 102831. <https://doi.org/10.1016/j.ijhcs.2022.102831>
- van Diggelen, J., Barnhoorn, J., Post, R., Sijs, J., van der Stap, N., & van der Waa, J. (2021). Delegation in Human-Machine Teaming: Progress, Challenges and Prospects. In D. Russo, T. Ahram, W. Karwowski, G. Di Bucchianico, & R. Tair (Eds.), *Intelligent Human Systems Integration 2021* (pp. 10–16). Springer International Publishing. https://doi.org/10.1007/978-3-030-68017-6_2
- van Zoelen, E. M., van den Bosch, K., & Neerincx, M. (2021). Becoming Team Members: Identifying Interaction Patterns of Mutual Adaptation for Human-Robot Co-Learning. *Frontiers in Robotics and AI*, 8, 200. <https://doi.org/10.3389/frobt.2021.692811>
- Welie, M. V., Veer, G. C. V. D., & Eliëns, A. (1998). An Ontology for Task World Models. *Proceedings of DSV-IS98, Abingdon*, 3–5.

6

SHARED REPRESENTATIONS OF EMERGENT COLLABORATION PATTERNS

We aim to understand how human and AI team members adapt and learn together, and how they can be supported in developing successful team behaviors. We call the process of mutual adaptation and learning co-learning. In previous work, we studied how Collaboration Patterns emerge spontaneously from adaptations of human and AI team members to each other and the task. We expect that formalizing emergent CPs can support human-AI teams in reusing and refining the CPs. This paper reports on the improvement of a previously developed ontology model and user interface that enable a human-AI team to formalize and communicate about emergent Collaboration Patterns. The experiment presented in this paper investigates whether humans can successfully use the ontology to formalize emergent Collaboration Patterns for sharing with their AI-partner, and whether this improves subsequent learning and performance of the team. Human participants (n=63) were coupled with a virtual robot to form a two-member rescue team operating in a simulated earthquake area. Effects on team performance were measured during and between exercises. Results show that teams using the ontology model improve their team performance over time when compared to teams that did not have access to the ontology model, while they also scored higher on several subjective measures such as Collaboration Fluency, self-efficacy, and subjective performance. At the same time, the introduction of the ontology model provided human team partners with more control over the robot behavior, leading to most participants directing the robot around rather than collaborating as a team partner. Our study shows that having a shared representation of emergent Collaboration Patterns supports human-AI teams in co-learning, but that maintaining a fluent co-adaptive team dynamic is not guaranteed when support is provided in such a way that the human has more control and initiative than the machine.

6.1. INTRODUCTION

RESEARCH on human-AI teaming, in which humans and AI-systems perform tasks together as team members, is becoming more relevant as AI systems increasingly enter our work environments (Akata et al., 2020; van Diggelen et al., 2021; van der Waa et al., 2020). In dynamic real-world task environments that change constantly, any human-AI team will need some time to develop a successful collaboration. To deal with this, humans continuously adapt (ad-hoc short term behavior changes, often implicitly and unconsciously) and learn (long term, sustainable behavior changes), to improve collaboration with their team members (Burke et al., 2006). With AI-systems that exhibit self-learning capabilities, the two adaptation and learning processes will interact over time. Therefore, it is relevant to understand how human-AI team members adapt and learn together and how they can be supported in developing successful team behaviors (van den Bosch et al., 2019).

This study builds on previous work in our group (E. M. van Zoelen, van den Bosch, and Neerinx, 2021), in which the process of co-learning was defined as consisting of two phases: (1) implicit behavioral adaptation and (2) communication and reflection about emergent Collaboration Patterns (E. M. van Zoelen, van den Bosch, and Neerinx, 2021), with the goal of building a shared mental model about the collaboration. Phase 1 (the process that naturally occurs as a result of the individual adaptive capabilities of the team members) results in the development of relatively stable patterns of collaborative behavior, so-called Collaboration Patterns (CPs) (E. M. van Zoelen, van den Bosch, Rauterberg, et al., 2021).

Sustaining successful Collaboration Patterns throughout the collaboration, however, requires team members to be aware of those patterns, to reflect on their usefulness, and to communicate about them (Phase 2 of the Co-Learning process). They need to know the actions that both team members need to engage in, and they need to understand how these actions depend upon each other and the environment. A shared mental model of emergent Collaboration Patterns can support shared understanding of the collaboration, and can facilitate further learning and communication (Andrews et al., 2023). To be able to assess whether CPs are beneficial to team performance, and to make sure that successful CPs can be agreed upon by the team and maintained, team members should therefore be equipped with:

1. A shared mental model of emergent Collaboration Patterns in which patterns can be formalized;
2. A way to communicate about these emergent Collaboration Patterns.

Existing work on Co-Learning tends to focus on implicit behavioral adaptations, the phase we call co-adaptation (Kumar, 2024; Shafti et al., 2020). Studies that focus on mental models tend to focus solely on knowledge learning rather than behavioral learning (Schoonderwoerd et al., 2022). In this paper, we bring the two phases of Co-Learning together by facilitating a human-AI team in formalizing emergent Collaboration Patterns, following the above two requirements. Previously, we have designed an ontology model and Graphical User Interface (GUI), together called "Collaboration Book", to document, share and re-use collaboration patterns in a human-agent team (E. M. van Zoelen et al.,

2022). This Collaboration Books aims to enhance team's meta-cognition, i.e. help team-members to reflect on their own behavioral adaptations. The previous study (Chapter 5) has shown that the ontology model and GUI enables human participants to formalize Collaboration Patterns from videos if they are able to recognize them as such (E. M. van Zoelen et al., 2022).

The aim of the study presented in this paper is to investigate whether the Collaboration Book supports human-AI team members in their collaborative task performance. We have investigated this by comparing human-AI teams that can make use of the Collaboration Book with teams that do not have access to this Book. The Collaboration Book provides the teams with extensive support for their collaboration (such as support for documenting and reflecting on the collaboration). A part of this is that it also serves as a form of communication between the team members. Additional communication could possibly improve the collaboration even if the other aspects of the Collaboration Book would not be present. To be able to control for the addition of this communication, we introduced two conditions besides the condition in which participants can use the Collaboration Book: a baseline condition consisting of teams that do not have access to any support or language-based communication (similar to the study presented in E. M. van Zoelen, van den Bosch, and Neerincx, 2021) and a condition in which teams have access to support in the form of basic communication consisting of chat messages (but not to other possible advantages of the Collaboration Book). We will provide a more detailed overview of our research questions in Section 6.1.1. Section 6.2 provides a description of the designed task environment, ontology model, GUI and AI agent behavior. Section 6.3 contains the experimental design and setup of the empirical study, of which the results are presented in Section 6.4.

6.1.1. RESEARCH QUESTIONS AND HYPOTHESES

WE introduced the concept co-learning as consisting of two phases, of which the first one is an implicit process. Achieving step two requires collaborators to be aware of their implicit behavior adaptations to the extent that they can recognize emergent collaboration patterns and communicate about them. In the present study, the AI agent is programmed to follow any CP formalized in the Collaboration Book, whereas the human can use it as a tool to formalize CPs they recognize as such. We are therefore interested in whether the provided ontology model and GUI supports humans in becoming aware of emergent Collaboration Patterns:

RQ1 Do human team members that can use the Collaboration Book in a human-AI team setting become more aware of emergent Collaboration Patterns compared to humans that do not have access to it?

We expect teams that make use of the Collaboration Book to have a higher awareness of the Collaboration Patterns they engage in throughout the collaboration. We expect this to be the case regardless of whether the teams without Collaboration Book can communicate with each other.

Secondly, we investigate whether the Collaboration Book supports teams to use the awareness to perform Collaboration Patterns more consistently; i.e., whether the Collaboration Book help teams to reuse emergent Collaboration Patterns over time, as well

as across context (such that it can be applied in slightly different task situations):

RQ2 Do team members that can use the Collaboration Book reuse emergent Collaboration Patterns more often than teams that do not have access to it?

We want to know whether the use of the Collaboration Book has an effect on task performance and on the fluency of the collaboration. Specifically, we also want to know if the Collaboration Book has an influence on learning, and therefore if task performance and fluency of the collaboration increase over time more strongly if the Collaboration Book is being used. We measure task performance through task metrics like completion time, as well as through self-reports. Collaboration fluency is measured in-task and through self-reports as well (Hoffman, 2019). Additionally, we measure self-efficacy, as it provides insight into participant's confidence in the collaboration. This brings us to the following set of research questions:

RQ3 Does a human-agent team become better at a collaborative task when they can make use of the Collaboration Book?

H3a. We expect to observe an increase in task performance over time.

H3b. We expect to observe an increase in objective Collaboration Fluency over time (i.e., less idle time in teams using the CB).

H3c. We expect to observe an increase in subjective Collaboration Fluency (i.e., teams using a CB having higher ratings on the Collaboration Fluency questionnaire).

H3d. We expect to observe an increase in self-efficacy during the experiment.

H3e. We expect to observe an increase in subjective performance during the experiment.

We expect that the increase in task performance over time will differ for teams with different levels of communication at their disposal. Communication provides insights into the behavior of team members which could help in attuning collaborative actions, which could therefore lead to faster and stronger learning as well as better (objective and subjective) performance. We expect that the Collaboration Book provides the best support and generates the largest increase; more than just basic communication. We expect that teams using the Collaboration Book will have an overall higher Collaboration Fluency, objective and subjective, than the other teams. Moreover, we expect that the increase will remain when the task changes such that the team needs to relearn and new CPs are necessary.

Next, we are interested in how the addition of the Collaboration Book qualitatively influences the Collaboration Patterns that emerge. That means that we want to explore what kinds of Collaboration Patterns emerge, and which of those are stored in the Collaboration Book:

RQ4 What Collaboration Patterns can we discover in human-AI teams?

4a. How are Collaboration Patterns described in the Collaboration Book by human participants?

- 4b. Which Collaboration Patterns that we can discover through observation are stored in the Collaboration Book?
- 4c. What are the similarities and differences in the Collaboration Patterns identified by the participants?

Earlier experiments have shown large individual differences in the ways in which participants adapt to and with their AI team partner, resulting in a large diversity of emergent Collaboration Patterns (E. M. van Zoelen, van den Bosch, and Neerincx, 2021; E. M. van Zoelen, van den Bosch, Rauterberg, et al., 2021). Up until now, we have not been able to find predictors for this variety. We hypothesize that locus of control might have an influence. People with an internal locus of control tend to believe that events happening to them are a consequence of their own actions. In contrast, people with an external locus of control tend to ascribe events to luck, fate, or the actions of others. Internal and external locus of control have been found to be predictors of performance-related behaviors and outcomes (Nießen et al., 2022). Specifically, it has been found that it can influence the level of trust a human has in a machine in human-machine collaborative scenarios (Chiou et al., 2021; Sharan and Romano, 2020). To investigate whether LoC affects task performance and a human's experience of the collaboration with agents, we include the following research question:

RQ5 Can locus of control predict how human participants perform at and experience the task?

We expect people with a higher internal locus of control to perform better at the task. This is based on our previous observation that it is difficult to co-learn successfully with a virtual robot for people who get frustrated about the robot not doing what they want it to do (E. M. van Zoelen, van den Bosch, and Neerincx, 2021; E. M. van Zoelen, van den Bosch, Rauterberg, et al., 2021).

6.2. DESIGN – TASK ENVIRONMENT, COLLABORATION BOOK, AGENT BEHAVIOR

6.2.1. TASK ENVIRONMENT

THE task environment used is based on the task environment presented in E. M. van Zoelen, van den Bosch, and Neerincx, 2021. This is an Urban-Search-and-Rescue task, in which a human and a virtual robot collaboratively need to save a victim from underneath a pile of rocks by removing the rocks (Figure 6.1). The human and robot have complementary capabilities; the human is assumed to have the general world knowledge needed for understanding the order in which rocks need to be picked up (e.g. not picking up a rock that would result in heavier rocks burying the victim). Furthermore, the human can pick up small rocks. The robot has limited general world knowledge, but can instead lift and remove large rocks and break large rocks into pieces to make it easier to remove them. Like the human, the robot can also pick up and carry small rocks. Both human and robot have the capability to learn, meaning that they are able to change their behavior over time to improve their performance. The details of how the robot learns are outlined in Section 6.2.3.



Figure 6.1: An overview of the USAR task environment, containing **actors** (the human operator, the robot, the victim) and **objects with different size** (small rocks, large rocks) and color (not shown here).

6.2.2. COLLABORATION BOOK

THE task is extended by the Collaboration Book: a tool given to participants with which they can represent Collaboration Patterns, that can consequently be used by the robot in the task. The Collaboration Book consists of the two elements introduced earlier: (1) an ontology model within which the CPs are represented, and (2) a Graphical User Interface that allows participants to describe patterns they deem relevant.

ONTOLOGY MODEL

To achieve an interaction dynamic in which human and AI agent can reason and communicate about Collaboration Patterns, it is necessary to provide them with a knowledge structure within which these patterns can be represented, that can serve as a shared mental model. The concepts used in this knowledge structure should be grounded for both collaborators; that means that the AI agent should be able to use them for reasoning, and that the human should have an understanding of what they mean. To achieve this, the ontology model was designed based on existing task models and frameworks for human social behavior (e.g. Dignum, 2018; Welie et al., 1998). Moreover, we used the language that participants used to describe their behavior in previous experiments as input for the design of the ontology. The full high-level ontology used in the experiment can be seen in Figure 6.2.

GRAPHICAL USER INTERFACE

To access the Collaboration Book (and thus the ontology) we designed a Graphical User Interface (GUI) (see Figure 6.3) that can be opened like a chat window while working on the task. This GUI allows participants to formalize Collaboration Patterns by dragging and dropping items for the start conditions and actions of the CP. Sequences of collaborative actions formalized via this GUI are then stored in the accompanying ontology model, from which the AI agent can retrieve them. The GUI should be seen as a transla-




Figure 6.2: Ontology Model for Collaboration Patterns. Red rectangles show concepts, yellow diamonds show relations, and blue ellipses show attributes.

tion layer between the ontology, that contains generally applicable high-level concepts, and the task environment used. It presents concrete instances of the concepts actor (e.g. human, robot), object (e.g. large rock, small rock), location (e.g. top of rock pile, bottom of rock pile) and task (e.g. pick up, drop) that are present in the specific task environment. It also provides the basic structure of the ontology in a visually easy to understand manner.

COLLABORATION BOOK IMPROVEMENTS

An earlier version of the ontology model and GUI were presented in E. M. van Zoelen et al., 2022. Based on insights gained from the study done in E. M. van Zoelen et al., 2022 as well as pilot runs for this experiments, a few improvements were made:

- Rather than presenting the actions as one single column in which participants would need to specify the actor doing the action, we now present a designated action-column for the human and for the robot. This can motivate participants to formalize the actions of both actors, as well as enable them to specify parallel and sequential actions.
- The previous study showed that people were sometimes inclined to specify the

Situation
The situation when the actions start:

You can use these types of items: object actor location

+ -

What we do
In this situation, we did this:

You can use these types of items: action object location

Human _____ Robot _____

+ -

Label

Top of rock pile Large rock

Above rock pile Small rock

Bottom of rock pile Brown rock

Left side of rock pile Robot

Left side of field Human

On top of Object Victim

Move to Object

Move back and forth in location

Stand still in location

Pick up object in location

Drop object in location

Break object in location

Figure 6.3: Graphical User Interface as presented to participants.

result of a Collaboration Pattern, but pilot runs showed that adding an ‘end conditions’ section confused participants. They were unable to specify end conditions in a useful manner; therefore we did not include it in the final version.

- In the ontology model, we specified the start conditions via the presence of context factors, as that enabled us to use the existing structure for tasks also for start conditions (in the earlier version, this was more limited).

6.2.3. AI AGENT DESIGN

THE AI agent (here called the virtual robot) was designed with the ability to learn in the task presented, while collaborating with the human. The virtual robot was equipped with a) the necessary skills for performing the task (e.g. to pick up rocks), including capabilities that the human did not have (e.g. picking up large rocks), and b) learning abilities, to be able to learn a policy based on feedback from the environment and from the human’s actions. The virtual robot would not be able to finish the task on its own; it is partly dependent on the human, yet lacks a specific strategy for executing the task together with a human partner. For example, it knows *how* to pick up rocks in general, but not *when* it should pick up a specific rock.

The virtual robot’s learning behavior resulted from a combination of machine learn-

AGENT BEHAVIOR




Figure 6.4: Visualization of the behavior of the virtual robot, based on rule-based actions (left) and machine learning (right). A decision tree decides whether it should take a rule-based action (perform a CP) or an action informed by learning (e.g. a basic behavior action or choosing an applicable CP).

ing and rule-based behavior. We use machine learning to ensure that co-learning can take place, while the rule-based behavior sufficiently structures the virtual robot's behavior to be able to study the relevant interaction aspects. The agent will execute a CP in a rule-based manner if the human participant formalized a CP applicable to a situation it encounters. If such a CP is not available, it will use Reinforcement Learning (greedy Q-learning) to choose from a set of atomic actions (e.g. 'pick up', we call this the basic behavior). If a situation occurs in which more than one CP is applicable, it uses a Contextual Bandit algorithm (Upper Confidence Bound algorithm) to choose which CP to execute. Figure 6.4 visualizes an overview of the high-level behavior tree of the virtual robot.

For both the Reinforcement Learning as well as the Contextual Bandit algorithm, an abstract state representation is used composed of three state features: 1) progress of the task (measured by percentage of rocks already removed), 2) relative contribution of the team members (measured by the number of rocks removed by each team member), and 3) whether the human is standing still (measured by a minimum of five ticks in which the human stays in the same location).

State S is built up of progress P , contribution C , human standing still HS as follows: $S = P \times C \times HS$. We define P as $P = \{1, 2, 3, 4\}$, C as $C = \{human, equal, robot\}$ and HS as $HS = \{True, False\}$.

These state factors were chosen to generate a small state space, to ensure a similar learning pace to that of the human, and to facilitate understanding of the learned behavior. The actions used in the Reinforcement Learning algorithm are 'Pick up', 'Drop', 'Break', 'Stand still' and 'Move back and forth'.

C1: Baseline (no support)	C2: Basic Communication	C3: Collaboration Book
Based on a previous experiment in E. M. van Zoelen, van den Bosch, and Neerincx, 2021, with no direct communication available. Learning about the other happens by behavioral observation only. The behavior of the virtual robot is different from the previous experiment, as described in Section 6.2.3.	The baseline with additional communication about actions: every time one of the team members executes an action, a message appears in a chat-window that communicates about that action (e.g. ‘Now executing pick up’). This communication is automated for both the human and the virtual robot.	The basic communication condition with use of the ontology and GUI added. The participant can formalize and share Collaboration Patterns with the virtual robot, so that the virtual robot uses the CPs when the appropriate conditions arise during the task. The virtual robot informs the participant (by means of a chat message) when it starts executing a CP.

Table 6.1: Overview of the levels of collaboration support used in the experiment.

6.3. EXPERIMENTAL STUDY

6

6.3.1. DESIGN

THE experimental design is a mixed design consisting of a between-subjects factor (Collaboration support) and a within-subjects factor (Round). Participants were randomly assigned to one of three conditions presented in Table 6.1. Each participant was presented with eight rounds of the task, with the level of collaboration support corresponding to their condition. Each Round presented a new variation of the task scenario that contained the same number of rocks, and the order of the different variations was counterbalanced to account for possible differences in difficulty. The first four rounds presented variations of the task scenario that were relatively similar and required a similar strategy to solve the task. Teams were expected to learn collaboration patterns that fit the specific type of situation in rounds 1-4. In rounds 5-8, a new element was introduced in the task scenarios (i.e. a brown rock that cannot be picked up) that often made the learned strategy unsuccessful. The team then had to relearn how to adapt the strategy, that is, develop new collaboration patterns. In the rest of the paper, we will refer to rounds 1-4 as Block A (Small and large rocks), and to rounds 5-8 as Block B (Brown rock).

The experimental design as well as the analysis method were preregistered before executing the experiment (E. van Zoelen et al., 2024). The analysis method was adapted after the experiment to better fit the data collected. The original plan for statistical analyses did not sufficiently account for the temporal nature of the experiment. Moreover, the answers to the open questions did not fit with the original plan for qualitative coding. A detailed description of the analyses conducted can be found in Section 6.3.5.

6.3.2. PARTICIPANTS

A total of 67 people participated in the experiment (20 in C1, 23 in C2 and 24 in C3). The allocation of participants to a condition was done randomly. In C2, one partici-

part was excluded due to hardware problems on the side of the participant. In C3, three participants were excluded: for one participant, questionnaire data was not complete, and for two participants, a significant software bug appeared that was solved afterwards. For the analysis, we will therefore have data for 63 participants (C1: 14 M, 6 F; C2: 9 M, 13 F; C3: 10 M, 11 F). The average age was 28.8 overall (STD = 7.8), 30.2 in C1 (STD = 12.1), 29.1 in C2 (STD = 5.1) and 27.4 in C3 (STD = 5.1). Participants were recruited through university PhD communities, by contacting interns at TNO, and via different forums within TU Delft and TNO. All participants were either enrolled in or completed a form of higher education.

6.3.3. DEPENDENT VARIABLES

To answer our research question, we measured the following dependent variables during the experiment:

- Subjective Collaboration Fluency: measured through the Fluency and Improvement scales from Hoffman, 2019 presented on a slider;
- Objective Collaboration Fluency: measured through robot and human idle time, the number of move actions and the number of productive actions in the task;
- Self-efficacy: measured through a self-efficacy statement with a slider;
- Subjective performance: measured through a performance statement with a slider;
- Task performance: measured through the number of simulation ticks needed to finish the task, the harm done to the victim by falling rocks, and the number of rocks remaining in case the task was not solved in time;
- Emergent Collaboration Patterns: screen captures of four randomly selected participants per condition provided behavior from which we could extract emergent Collaboration Patterns;
- CP Awareness: the extent to which participants are aware of emergent Collaboration Patterns, measured through open questions asking about their strategy;
- Formalized Collaboration Patterns: the Collaboration Patterns that participants in C3 formalized during the experiment.

6.3.4. PROCEDURE

Participants completed the experiment online, within a scheduled timeslot. They were given instructions via a video. In this video, they were told that they form a team with the robot, that the team is yet unexperienced in the victim-saving task, and that this will become clear when they start to work with the robot. In addition, participants in C3 were instructed to monitor their own performance and that of the robot, and that when they recognize a pattern of behavior in their collaboration that works well, they should record this in a “Collaboration Book” (the ontology model and accompanying GUI). These participants were then told that the team uses this Collaboration Book as a tool to improve their cooperation, and that once a Collaboration Pattern is stored in the

Collaboration Book, the robot is able to access to it, and will automatically initiate the actions that are its part of the CP. They were instructed that if the outcomes of a CP do not meet the participant's expectations, they can edit the CP to change or finetune the actions. If the participant thinks that the pattern is of no use after all, they can decide to delete the CP.

Before starting the experiment, participants had the opportunity to practice the task and the interface operations in a simple variation of the task scenario without the robot. After this practice round, they were asked to provide some demographic information and answer a few questions on their experience with gaming, human-robot interaction, and locus of control (Nießen et al., 2022). See Appendix K for the full questionnaire.

For each round, the participant went through the following steps:

- Perform the task (i.e. try to save the victim as quickly as possible) and meanwhile monitor the collaboration. For C3, at any point in time, the participant could pause the task for the purpose of describing and storing an identified Collaboration Pattern in the Collaboration Book.
- Participants in C3 received a prompt asking them to describe Collaboration Patterns if they observed any.
- Evaluate six statements (five taken from Hoffman, 2019, the Fluency and Improvement items, and one self-efficacy statement) on a slider ranging from 'strongly disagree' to 'strongly agree'.
- Participants in C3 were asked to explain for each stored Collaboration Pattern why they stored it, participants in C1 and C2 were asked for their strategies.
- Evaluate their performance subjectively on a slider ranging from 'poor' to 'excellent'.

After all eight rounds, participants were asked to answer the following open questions:

- What was the strategy of the team in the end?
- Can you explain who did what in the task?
- Did the strategy change over the course of the experiment? How?
- How fluently did you and the robot collaborate? Did this change over the course of the experiment?
- (C3 only) What was your general strategy for storing Collaboration Patterns?

For four randomly chosen participants per condition, we recorded their behavior in the experiment via screen captures. We organized a separate second meeting with these participants, in which we showed them the video captures of their task execution, and asked them to describe their own behavior.

Baseline: 20 participants * 8 rounds = 160 datapoints
 Basic Communciation: 22 participants * 8 rounds = 176 datapoints
 Collaboration Book: 21 participants * 8 rounds = 168 datapoints




Figure 6.5: Overview of data collected and analyses performed on specific subsets of the data.

6.3.5. DATA ANALYSIS

FIGURE 6.5 shows an overview of all data collected and analyses performed for specific subsets of the data.

To answer **RQ1** (Pattern Awareness) and **RQ2** (Pattern Reuse), we first manually coded from observation which Collaboration Patterns emerged in each round for a subset of four randomly selected participants per condition. As a coding scheme, we used the CPs identified in E. M. van Zoelen, van den Bosch, and Neerincx, 2021. That means that our coding scheme contained the following Collaboration Patterns:

1. Focus on own task: the participant picks up and clears away small rocks, without paying much attention to what the robot is doing;
2. Passive: the participant mostly stands still while observing what the robot is doing;
3. Alternating: the participant alternates between standing still on the side, which triggers the robot to pick up or break large rocks, and actively picking up small rocks themselves;

4. Active synchrony: the participant actively directs the robot towards large rocks that need to be picked up or broken down by hovering over them;
5. Damage control: the participant builds piles of rocks to prevent the robot from dropping rocks on top of the victim.

In addition, we used the code 'jitter' for when the participant would move around the task in a jittery manner without clear goal. For participants in C3, we used the names of the formalized CPs given by the participant in case these were used. The full result of the coding can be found in Appendix E.

For **RQ1** (Pattern Awareness), we then read through participants' answers to the open questions about strategy, and assessed whether their answers matched the emergent Collaboration Patterns that we observed. We created descriptions of their awareness using quotes from the open questions. For participants in C3, we also assessed whether CPs formalized in the Collaboration Book matched earlier emergent CPs. For **RQ2** (Pattern Reuse), we assessed to what extent CPs reappeared in the eight rounds, based on the coding in Appendix E. For both **RQ1** and **RQ2** we distinguish between 'emergent CPs', which are CPs that emerge spontaneously from co-adaptive behavior, 'formalized CPs', which are CPs that were formalized using the Collaboration Book (whether they originated from an emergent CP or not), and 'designed CPs', which are formalized CPs that do not originate from an emergent CP but are instead actively designed by the human participant.

To investigate **RQ3** (Team Performance), we generated Linear Mixed Models with *Round* (8), *Collaboration support* (3) and *Block* (2) as interacting factors, and a random intercept for Participant. We chose to use LMMs because of the double repeated measures design of the experiment with high importance of the temporal aspect (Round and Block), as well as a large variation in outcomes between individual participants. In addition to the LMMs for the full experiment, we generated separate Linear Mixed Models for Block A (Round 1-4) and Block B (Round 5-8) with *Round* (4) and *Collaboration support* (3) as interacting factors, and a random intercept for Participant. By constructing these models for the full experiment as well as for the separate blocks, we could explore interactions between the blocks as well as separate effects. For each of the constructed models, we used C3 (Collaboration Book) as the reference value. We used R and *lme4* (Bates et al., 2015) to perform these analyses, using *lmerTest* (Kuznetsova et al., 2017) to obtain p-values.

To investigate **RQ4** (Pattern Types), we:

1. performed a qualitative analysis of the formalized CPs for a subset of four participants. We attributed codes to all formalized CPs that were used in the task (using an open coding process), and compared those to emergent behavior that we, the researchers, observed and coded in Appendix E;
2. created bag-of-words feature vectors (Qader et al., 2019) of each formalized CP of all participants that was used during the experiment, using the names of the items in the GUI as vocabulary (i.e. we used 'Large rock' as a word rather than 'large' and 'rock' separately). In cases where a CP was edited, the latest edit was used for the specific round of the task in which it was used. This left 125 CPs in total, that

we clustered using a K-Means clustering algorithm, to find commonalities across all formalized CPs. We performed the clustering four times: (1) for the situation description, (2) for the human actions, (3) for the robot actions, and (4) for the full CP. We tried different numbers of clusters for each separate K-means clustering and qualitatively looked at the content of the clusters. For each subset of clusters, we created a definitive list of clusters based on a combination of the K-Means result and our qualitative judgment.

To address **RQ5** (Locus of Control), we calculated correlations (Pearson correlation) between internal and external locus of control scores and all performance metrics.

6.4. RESULTS

6.4.1. AWARENESS OF COLLABORATION PATTERNS (RQ1)

AN extensive description of our qualitative analysis process is provided in Section 6.3.5.

Most participants in C1 or C2 became aware of emergent CPs early in the experiment (participants 4006, 4007, 4016, 4055, 4056). They often explicitly mentioned the behavior that we as the researchers observed (e.g. 'by holding my hand at a big rock, the robot will pick it up', participant 4006, or 'trying to hover over the big rocks needed to be picked up', participant 4056). However, they then started to doubt themselves and tended to deny the existence of the CP later on, even if they used it throughout the experiment (e.g. 'the robot seemed to understand what I was doing but then it seems it stopped working', participant 4007). We suspect that this is caused by the robot sometimes deviating from CPs, because the RL-algorithm of the robot causes a certain level of randomness in its behavior.

Participants in C3 were sometimes (but definitely not always) able to formalize a CP that emerged spontaneously early in the experiment (participants 4071 and 4098), which they subsequently used throughout. We observed some instances in which a formalized CP matched an emergent CP observed by the researchers (e.g., participant 4098 formalized a CP to 'instruct the robot to break up large rocks', after the CP 'Active synchrony' emerged). Interestingly, in such instances participants didn't mention that they saw the CP emerge and then formalized it. Later on in the experiment, several instances were observed in which participants designed new CPs that did not emerge earlier (e.g., as participant 4071 mentioned, to 'allow for more control over the actions of the robot'). For an extensive description of the awareness per participant including more example quotes, see Appendix F.

6.4.2. REUSE OF CPs (RQ2)

FOR most of the participants selected for qualitative analysis, emergent CPs were reused throughout the experiment; generally, they would use CPs that emerged early in the task also later on in the task, even when a new task factor (the brown rock) was introduced (see Appendix E for the CPs used per participant per round). Participants in C1 and C2 did this most consistently. Even if they were unable to succeed at the task, they would stick to strategies that had worked in earlier rounds of the experiment (see for example participant 4008 in Appendix E and G, who used a CP that emerged in round 1




Figure 6.6: Outcomes of the experimental rounds sorted by condition. C1 is the baseline condition, C2 is the basic communication condition, and C3 is the Collaboration Book condition.

consistently throughout the experiment apart from some exploration in round 2 and 5). For participants in C1 and C2, different CPs emerged early on in the experiment, but the variety of CPs used gradually declined over the rounds, meaning that they converged to one or two CPs that they used consistently (see for example participants 4016 and 4057 in Appendix E and G).

Participants in C3, however, continued to use different kinds of CPs throughout the experiment, even though these were mostly designed rather than emergent (see for example participant 4099 in Appendix E and G who continued to specify new, more specific CPs to direct the robot throughout the experiment). It seems that the availability of the Collaboration Book motivated participants to explore more, as participants in C3 tended to create new CPs when the brown rock was introduced (see for example participant 4073 in Appendix E and G, who created new variations of CPs from round 6 onwards), although not always successfully. It also prompted C3-participants to create new variations of previously created CPs when they discovered that the robot did not behave as expected, thereby creating CPs for very specific situations that would not work in other situations (see again participant 4073 in Appendix E and G). The basis of their behavior, however, remained relatively constant, meaning that they usually had one or two CPs that they would continuously fall back on (see for example participant 4098 in Appendix E and G, who returned to their original formalized CP in the final round).

Appendix G contains an extensive description of the reuse of CPs throughout the experiment per participant.




Figure 6.7: Visualization of the outcome of the Linear Mixed Model for time in ticks.

Factor	Estimate	Std. Error	df	t	p
Round	-184.30	41.23	358.99	-4.470	<0.001
Round:C1	165.61	59.03	358.99	2.805	0.005
Round:C2	136.26	57.64	358.99	2.364	0.020
Round:B.Brown rock	184.11	63.75	368.02	2.888	0.004

Table 6.2: Results of the Linear Mixed Model for time in ticks, for the factors that show a significant effect. A visualization of the LMM is shown in Figure 6.7.

6.4.3. EFFECTS ON TEAM PERFORMANCE (RQ3)

IN each Round, there were at least two possible outcomes for the human-AI team: (1) successfully saving the victim, and (2) not saving the victim within the time limit (timeout). In Block B, it was possible to make a fatal mistake that would render the scenario unsolvable (when the brown rock would fall on top of the victim). Figure 6.6 shows how often each outcome occurred within each round for each condition of the experiment. Only a small number of human-AI teams was able to successfully save the victim within the time given. Fatal mistakes were common as well. In the following sections, we present an analysis of the different subjective and objective performance metrics and the extent to which Collaboration support and Round had an effect on those. The full results of the Linear Mixed Models can be found in Appendix I.

TASK PERFORMANCE

Time in ticks We analyzed the time in ticks (a tick lasted 0.05 seconds, meaning that there are 20 ticks in a second) that participants spent in the task, which could be between 0 and 3000 (when the clock ran out, after 150 seconds). We excluded rounds in which the task became unsolvable due to the brown rock falling on top of the victim from this analysis, as these would give an unrealistically low number of ticks.

The results of the Linear Mixed Model with Round, Condition and Block (see Figure 6.7) showed significant effects for the factors presented in Table 6.2 (average starting value = 3101.26). These results indicate that Time in ticks decreases over the course of the rounds for C3, as can be seen from the negative estimate for Round; each Round, Time in ticks decreases by -184.30. The values for Round:C1 and Round:C2, however, are of similar magnitude but positive, which almost cancels out this effect. This means that for participants in C1 and in C2 this effect is much smaller than for participants in C3. This can also be seen in Figure 6.7. Moreover, this effect only exists in Block A, as indicated by the value for Round:B. Brown rock, which is almost identical to the value for Round but with a positive sign.




Figure 6.8: Visualization of the outcome of the Linear Mixed Model for victim harm.

Victim harm Victim harm consists of the number of rocks falling on top of the victim multiplied by 100. The results of the Linear Mixed Model with Round, Condition and Block (see Figure 6.8) showed significant effects for the factors presented in Table 6.3 (average starting value = 752.38). These results show that Victim Harm decreases per Round, as indicated by the negative value for Round. This effect is practically non-existent for participants in C1 and C2, as indicated by the similar but positive values for

Round:C1 and Round:C2, meaning that the effect only exists in C3. Moreover, the decrease per Round also practically does not exist in Block B (as is indicated by the similar but positive value for Round:B. Brown rock).

Factor	Estimate	Std. Error	df	t	p
Round	-140.95	46.06	441.00	-3.060	0.002
Round:C1	152.95	65.95	441.00	2.319	0.021
Round:C2	147.32	64.40	441.00	2.287	0.023
Round:B. Brown rock	147.79	65.15	441.00	2.269	0.024
Round:C1:B. Brown rock	-221.79	93.27	441.00	-2.378	0.018
Round:C2:B. Brown rock	-223.25	91.08	441.00	-2.451	0.015

Table 6.3: Results of the Linear Mixed Model for victim harm, for the factors that show a significant effect. A visualization of the LMM is shown in Figure 6.8.

Remaining rocks For participants that were unable to save the victim within the time allocated, we logged their progress by the number of rocks remaining for rounds ending in a Timeout or Fatal Mistake. We did not find any effects of any of the between- and within-subject factors in our Linear Mixed Model analysis.

OBJECTIVE COLLABORATION FLUENCY AND EFFICIENCY

To measure Collaboration Fluency, we chose idle time of the human and the robot. In general, these cannot be compared against each other, because the action duration of human actions in the task environment was much shorter than the action duration of robot actions, due to the way actions were registered in the task environment (human actions were only registered at a single tick, whereas robot actions were specified to have a duration of a certain amount of ticks). Of course, the human has additional thinking time, but this is not visible in the logged data. Regardless, we can analyze both human and robot idle time separately across rounds and levels of collaboration support.

Human idle time The results of the Linear Mixed Model with Round, Condition and Block (see Figure 6.9) showed significant effects for the factors presented in Table 6.4 (average starting value = 0.923). These results indicate that the amount of relative idle time by the human decreases over the rounds, as indicated by the negative estimate for Round. This is however dependent on the block; the effect only exists in Block A, as indicated by the same but positive estimate for Round:B. Brown rock. Moreover, in C1 (Baseline), participants had less idle time overall than in C3 (Collaboration Book), as indicated by the negative estimate for C1.

Robot idle time The results of the Linear Mixed Model with Round, Condition and Block (see Figure 6.10) showed significant effects for the factors presented in Table 6.5 (average starting value = 0.311). These results indicate that the relative robot idle time increases over the rounds, as indicated by the positive estimate for Round.




Figure 6.9: Visualization of the outcome of the Linear Mixed Model for human idle time (idle ticks) per tick in the task.

Factor	Estimate	Std. Error	df	t	p
Round	-0.009	0.003	441.00	-3.076	0.002
C1	-0.036	0.015	264.78	-2.357	0.019
Round:B. Brown rock	0.009	0.004	441.00	1.965	0.050

Table 6.4: Results of the Linear Mixed Model for human idle time (idle ticks) per tick in the task, for the factors that show a significant effect. A visualization of the LMM is shown in Figure 6.9.

Human move actions The results of the Linear Mixed Model with Round, Condition and Block (see Figure 6.11) showed significant effects for the factors presented in Table 6.6 (average starting value = 0.074). These results indicate that the relative number of move actions done by the human increased over the rounds, as indicated by the positive estimate for Round, but that this effect is influenced by block; the effect only exists in Block A, as indicated by the similar but negative estimate for Round:B. Brown rock. Moreover, participants in C1 (Baseline) performed more move actions than participants in C3 (Collaboration Book), as indicated by the positive estimate for C1.

Robot move actions The results of the Linear Mixed Model with Round, Condition and Block (see Figure 6.12) showed significant effects for the factors presented in Table 6.7 (average starting value = 0.538). These results indicate that the number of move actions performed by the robot decreased over the rounds, as indicated by the negative estimate for Round. This effect is influenced by block, meaning that the effect is non-existent




Figure 6.10: Visualization of the outcome of the Linear Mixed Model for robot idle time (idle ticks) per tick in the task.

Factor	Estimate	Std. Error	df	t	p
Round	0.044	0.017	441.00	2.615	0.009

Table 6.5: Results of the Linear Mixed Model for robot idle time (idle ticks) per tick in the task, for the factors that show a significant effect. A visualization of the LMM is shown in Figure 6.10.

for Block B, as indicated by the similar but positive estimate for Round:B. Brown rock. Moreover, the robot performed fewer move actions overall in Block B when compared to Block A, as indicated by the negative estimate for B. Brown rock.

Human productive actions We did not find any clear effects in our Linear Mixed Model for the full experiment. The models for the separate blocks did indicate a significant effect for C1 (Baseline) in both blocks.

Robot productive actions We did not find any effects in our Linear Mixed Model. The models for the separate blocks did indicate a significant effect for Round:C1, indicating an increase over the rounds for C1 (Baseline) when compared to C3 in Block A. This is also visible in Figure 6.13.




Figure 6.11: Visualization of the outcome of the Linear Mixed Model for human move actions per tick.

Factor	Estimate	Std. Error	df	t	p
Round	0.009	0.003	441.00	3.211	0.001
C1	0.029	0.014	264.10	2.090	0.038
Round:B. Brown rock	-0.008	-0.004	441.00	-2.034	0.043

Table 6.6: Results of the Linear Mixed Model for human move actions per tick, for the factors that show a significant effect. A visualization of the LMM is shown in Figure 6.11.

Factor	Estimate	Std. Error	df	t	p
Round	-0.057	0.012	441.00	-4.665	<0.001
B. Brown rock	-0.231	0.087	441.00	-2.642	0.009
Round:B. Brown rock	0.055	0.017	441.00	3.174	0.002

Table 6.7: Results of the Linear Mixed Model for robot move actions per tick, for the factors that show a significant effect. A visualization of the LMM is shown in Figure 6.12.




Figure 6.12: Visualization of the outcome of the Linear Mixed Model for robot move actions per tick.




Figure 6.13: Visualization of the outcome of the Linear Mixed Model for robot productive actions per tick.

SUBJECTIVE COLLABORATION FLUENCY

The results of the Linear Mixed Model with Round, Condition and Block (see Figure 6.14) showed significant effects for the factors presented in Table 6.8 (average starting value = 29.167). These results indicate that subjective Collaboration Fluency increases over the rounds, as indicated by the positive estimate for Round, but that this effect is influenced by block; in Block B, the effect is slightly opposite, but almost non-existent, as indicated by the similar but negative estimate for Round:B. Brown rock.




Figure 6.14: Visualization of the outcome of the Linear Mixed Model for subjective collaboration fluency.

Factor	Estimate	Std. Error	df	t	p
Round	6.915	1.833	441.00	3.773	<0.001
Round:B. Brown rock	-7.262	2.592	441.00	-2.802	0.005

Table 6.8: Results of the Linear Mixed Model for subjective collaboration fluency, for the factors that show a significant effect. A visualization of the LMM is shown in Figure 6.14.

SELF-EFFICACY

The results of the Linear Mixed Model with Round, Condition and Block (see Figure 6.15) showed significant effects for the factors presented in Table 6.9 (average starting value = 57.48). These results indicate that participants in C2 (Basic communication) had a lower self-efficacy overall than participants in C3 (Collaboration Book), as indicated by the negative estimate for C2. The models for the separate blocks indicated an additional significant effect for Round in Block B, indicating a decrease of self-efficacy in this block.




Figure 6.15: Visualization of the outcome of the Linear Mixed Model for self-efficacy.

Factor	Estimate	Std. Error	df	t	p
C2	-22.75	9.415	190.43	-2.416	0.017

Table 6.9: Results of the Linear Mixed Model for self-efficacy, for the factors that show a significant effect. A visualization of the LMM is shown in Figure 6.15.

SUBJECTIVE PERFORMANCE

The results of the Linear Mixed Model with Round, Condition and Block (see Figure 6.16) showed significant effects for the factors presented in Table 6.10 (average starting value = 19.429). These results indicate that subjective performance increased over the rounds, as indicated by the positive estimate for Round, but that this effect was influenced by block; the effect is non-existent in Block B, as indicated by the similar but negative estimate for Round:B. Brown rock.

6.4.4. DISCOVERED COLLABORATION PATTERNS (RQ4)

QUALITATIVE ANALYSIS OF CPS

IN Appendix J, we present a full list of the formalized CPs that the four selected participants 4071, 4073, 4098 and 4099 used in the task.

Participants 4071 and 4098 formalized a CP after it emerged from co-adaptation. Both CPs involved the human directing the robot towards breaking a large rock. Additional CPs formalized by these participants did *not* originate from emergent behavior, but were actively designed by the participant to optimize the collaborative behavior; they came up with these CPs without seeing them emerge first. Participant 4071 formal-




Figure 6.16: Visualization of the outcome of the Linear Mixed Model for subjective performance.

Factor	Estimate	Std. Error	df	t	p
Round	8.524	2.043	441.00	4.172	<0.001
Round:B. Brown rock	-7.529	2.889	441.00	-2.606	0.009

Table 6.10: Results of the Linear Mixed Model for subjective performance, for the factors that show a significant effect. A visualization of the LMM is shown in Figure 6.16.

ized three CPs, which centered around the human directing the robot to large rocks and efficiently removing small rocks. Participant (4098) formalized five CPs that describe the order in which rocks need to be picked up in more detail (e.g., three of these CPs involve breaking or removing specific rocks on top of other rocks).

For the other two participants (4073 and 4099) there was no CP that demonstrably originated from emergent behavior (we, the researchers, could not identify a CP that emerged from co-adaptive behavior that contained similar behavior), although participant 4073 behaved relatively separately from the robot, and this was represented in a CP in which the human picked up small rocks while the robot stood still. All CPs that participant 4073 formalized contained either the robot or the human standing still. Participant 4099, on the other hand, did not formalize any human behavior, but only designed the behavior of the robot.

The variety in the formalized CPs shows that different participants have different strategies with which they approach the formalization of CPs. It is noticeable that many of the designed CPs centered around the behavior of the robot and did not contain a lot of (or even any) actions for the human. When human actions were formalized, they

were often used to direct the robot to a specific rock (often via standing still), although there were a few CPs that simply described the human and the robot picking up rocks concurrently.

CLUSTERING OF CPs

A full overview of the clustering results can be found in Appendix H. There were roughly four different types of CPs stored via the Collaboration Book:

1. A CP which would apply once there was a brown rock, in which the robot was supposed to pick up or break that brown rock. This would not yield any results, as the brown rock could not be picked up or broken;
2. A CP which would apply once there was a large rock in a certain location, that the robot needed to pick up;
3. A CP which would apply once there was a large rock in a certain location, in which the human would stand still somewhere, triggering the robot to start breaking or picking up;
4. A CP which would apply once there was a small rock in a certain location, which the human and the robot would both start picking up.

From the cluster results of the separate CP parts, we could see that within these four, there is variety in how the separate parts were formalized. Some remarkable variations were for example the lack of human actions (no human actions specified at all), pick up actions followed by drop actions versus just pick up actions for the robot, and actions that instruct the robot to break or pick up a brown rock (which was not possible, and participants were told as such in the instructions). We conclude that while there are some general patterns in how people formalized CPs, there are a lot of individual differences even though the interface we designed provided participants with a limited vocabulary to perform the formalization.

6.4.5. LOCUS OF CONTROL (RQ5)

WE found one significant weak to medium correlation between internal locus of control and subjective Collaboration Fluency ($R = 0.31$, $p = 0.015$). This suggests that participants with a stronger internal locus of control generally perceive their Collaboration Fluency to be higher. The R is however very low. Moreover, as this is a subjective scale, this does not provide us with much information about whether they were actually able to collaborate more fluently with the robot; it might just be that people with larger internal LoC tend to assess Fluency more positively.

6.5. DISCUSSION

IN this paper, we present an experiment in which a human and a virtual robot collaborate on a task to save a victim from underneath a pile of rocks. Within this experiment, both human and virtual robot needed to learn how to collaborate at the task, creating a co-learning setting. We investigated whether an ontology of emergent Collaboration Patterns and an accompanying GUI, with which the human could formalize

CPs that emerged as a result of implicit co-learning, would improve human-robot team performance. This serves as an empirical study of how shared mental models impact human-AI collaborations (Andrews et al., 2023). Previous work has shown that in order to study co-learning, an approach in which the dynamics of coordination that lead to emergent Collaboration Patterns are studied is vital (Kumar, 2024; E. M. van Zoelen, van den Bosch, and Neerincx, 2021; Wiltshire et al., 2024). At the same time, to be able to understand the conditions for successful co-learning, it is necessary to design experiments in which performance improvements can be quantified to be able to make inferences about what leads to successful co-learning outcomes. Our main contribution is twofold:

1. We provide an experimental framework that combines insights into behavioral adaptive dynamics with quantitative performance metrics;
2. We study co-adaptation and shared mental model updating jointly to be able to understand the co-learning process as a whole, while existing work often focuses on either co-adaptation (e.g. Kumar, 2024; Shafti et al., 2020) or explicit knowledge learning (e.g. Schoonderwoerd et al., 2022).

Below, we will address the research questions, followed by a reflection on our study design and directions for future research.

6.5.1. RQ1 AWARENESS OF COLLABORATION PATTERNS

FIRST of all, we wanted to understand if and how the use of the Collaboration Book influenced awareness of emergent Collaboration Patterns, as it was designed as a tool for formalizing emergent behavior (or, in other words, explicating the emergent behavior, thereby aiding awareness). This relates to how we defined co-learning as consisting of (1) a co-adaptive phase in which implicit, ad hoc adaptations lead to collaborative behavior, and (2) a phase in which partners build and update a shared mental model of emergent CPs, which makes them aware of these emergent CPs. The Collaboration Book was designed to facilitate the second phase. Based on qualitative analyses of behavior and self-reports, we found that without the Collaboration Book, participants noticed emergent CPs quickly (in the first few rounds of the experiment, as can be seen in Appendix F), but that they were often unsure about the CP, as the robot would sometimes show unexpected behavior. This would often cause them to either forget or deny that the CP existed at all later in the experiment, although some participants were able to confirm the presence of certain emergent CPs towards the end of the experiment and become aware of them. Some participants using the Collaboration Book seem to become aware of CPs that emerge early in the experiment, as the first CP that they formalize in the Collaboration Book resembles emergent behavior. However, they never mention that this is why they formalize the behavior as such, which suggests that their awareness might have still been very implicit. We therefore cannot conclude that the Collaboration Book *per se* leads to higher awareness. This might however partly be due to a third process which interfered with both phases of co-learning: a process in which participants proactively tried to design successful team behaviors, possibly because they wanted more control over the robot. We suspect that the amount of attention that they spent on this third process caused them to have less time and attention for emergent CPs. We did not foresee this, and it could be relevant to further study how a desire for control interferes with

co-learning, and what this means for achieving Meaningful Human Control (Santoni de Sio and van den Hoven, 2018). For example, we could design a study which investigates whether this kind of control that manifests itself as proactive design of robot behaviors can be considered 'meaningful', or whether co-learning itself contributes to the meaningfulness of control (and therefore we should try to prevent humans from expressing this kind of control).

6.5.2. RQ2 REUSING CPs

WE set out to investigate to what extent emergent CPs were reused, and whether this was influenced by the Collaboration Book. As the learning algorithm was designed to enable continuous adaptation in order to establish a co-learning interaction, the virtual robot could sometimes behave unexpectedly. Especially when participants changed their behavior due to new insights (for example if they discovered that they could direct the robot towards specific large rocks), this might influence the rewards and thus the robot behavior, not always in a useful manner. This variability in behavior makes it a challenge for human-machine partners to develop stable collaborative strategies. We designed the Collaboration Book to facilitate agreements on CPs, such that they could be used more stably over time and possibly across context. However, when looking at the results, we can conclude that this did not exactly happen. While all participants went through a phase in which they tried out different CPs in the beginning of the experiment, participants without the Collaboration Book (C1 and C2) most clearly converged towards one or two CPs that they kept using, although this would not always lead to success. While participants with the Collaboration Book (C3) would capitalize on a certain CP by formalizing it, they would continue tweaking this CP, and continued to design new CPs all throughout the experiment. We can conclude that the availability of the Collaboration Book provided participants with the means to continue improving their strategy, while without it, participants tended to use the strategy that they knew somewhat worked. At the same time, participants in C3 were as unsuccessful as participants in the other conditions in changing their strategy to adapt to the brown rock in Block B. This means that while the CB gives the opportunity for continuous improvement as described above, it does not per se stimulate participants to be more creative in dealing with new challenging situations.

6.5.3. RQ3 INFLUENCE OF COLLABORATION BOOK ON TEAM PERFORMANCE

WE investigated whether access to the Collaboration Book would improve human-AI team performance. First of all, for both time as well as victim harm, using the Collaboration Book lead to a stronger performance improvement over time. For remaining rocks we did not find these effects; we only found an effect of Round which reveals a gradual improvement over time.

Both the human and the robot had a higher idle time and performed fewer actions when the Collaboration Book was available. For the robot, this can be explained easily; namely, the introduction of CPs that include wait actions ensure that there are no useless move actions. At the same time, CPs that do not contribute to the team task and do not work let the robot idle more unnecessarily. For the human, the higher idle time when using the CB is more difficult to explain. It might be that the formalized CPs cause the

human to behave more efficiently (meaning that they also performed fewer useless move actions, like the robot), as they know better what is required of them. In general, the teams became more efficient with the introduction of the Collaboration Book (higher performance with fewer actions).

The visualizations of the Linear Mixed Models suggest that all self-reported measures (subjective Collaboration Fluency, self-efficacy and subjective performance) are affected positively by the availability of the Collaboration Book, although this is only significant between C2 and C3 for self-efficacy. Interestingly, the visualizations (Figure 6.14, 6.15 and 6.16) suggest that this effect is partly maintained in Block B. Overall, we conclude that participants valued the Collaboration Book. In Block A, this self-reported benefit can be explained by an increase in the performance measures (see Figure 6.7 and 6.8). In Block B, however, this is not the case. This can mean two things: either the presence of the Collaboration Book makes participants view the collaboration more positively regardless of performance, or the positive effect that the performance increase in Block A caused has improved participants' trust in the virtual robot in such a way that it does not immediately drop when performance drops in Block B. In human-machine teams, it is important that humans have an appropriate level of trust in their machine team members, which they can achieve through a process called trust calibration (de Visser et al., 2020). Potentially, participants in C3 may have overtrusted the robot in Block B based on their experience in Block A, and therefore needed to recalibrate their trust. If so, this is a process that may take longer than the four rounds we provided. An interesting direction for future research would be to investigate relations between subjective performance metrics and task performance, as well as the role of trust calibration in co-learning.

6.5.4. RQ4 DISCOVERED COLLABORATION PATTERNS

As we used the same task as in previous studies, we were able to establish that the CPs emerging in the present study were similar to CPs found in previous work (E. M. van Zoelen, van den Bosch, and Neerincx, 2021). Most participants in the present study formalized the same kind of CPs, but there was a lot of variety in the way in which they formalized them. The vocabulary available to participants in the Collaboration Book did not lead to a clear understanding of how formalizations could best lead to the intended behavior, as we observed that participants edited their CPs often to get the desired result. They were not able to specify the intended behavior in one try, which means that reaching sufficient common ground was a challenge for them. Recent innovations in LLMs might enable interactions in which models can interpret free form input (natural language) more easily, and translate it to robot behavior in such a way that it comes closer to the human's intention due to better intent recognition (see e.g. Zhou et al., 2024). It will be interesting to research if that would support the creation of a shared mental model for collaborative tasks among human-AI collaborators. Even with LLMs, however, there will always be a possibility for translation errors, as the human might expect the machine to understand their intent whereas the capabilities of the machine to actually translate that into actions might be limited. An interesting direction for future research is to study how to provide the human with enough freedom to express their intent while not suggesting that the machine will be able to execute all of their intentions if they are not capable to.

6.5.5. RQ5 LOCUS OF CONTROL

UNFORTUNATELY, we could not find an effect of Locus of Control on human-AI team performance. Currently, we have focused on performance metrics, and it could be useful to study whether factors like LoC have an effect on behavioral metrics, such as the kind of Collaboration Patterns that emerge.

6.5.6. REFLECTION ON HUMAN-MACHINE CO-LEARNING DESIGN

FOR the present study, we have used a slightly adapted version of a task used in a previous study (E. M. van Zoelen, van den Bosch, and Neerincx, 2021). When comparing to the previous study, in this study we found that CPs emerged earlier in the experiment, and especially in C1 and C2 we observed that the teams converged to a main CP quickly. This was probably due to the difference in how the task was modeled for the learning algorithm. In the previous studies, states were modeled as phases in the experiment, while the robot would pick a rule-based strategy as an action. This meant that the decision for an action was only made a few times. In the experiment presented here, the actions were specified at a more atomic level; therefore, the action decision was made often, which made it easier for the robot to obtain rewards and learn quickly. This led to less learning-pressure for the human, making it easier to follow the robot and converge towards a CP. This probably made it harder for the participants to change their behavior when the brown rock was introduced, as by then they were already used to a certain way of working with the robot.

The new version of the state and action modeling also caused the behavior of the virtual robot to be less transparent and predictable to the human participants. This is probably caused by the use of more complex algorithms and state-action spaces. While these bring about an increase in performance for the algorithm (i.e. faster learning), it does not automatically lead to a better result when interacting with a human, possibly due to the mentioned decreased transparency and predictability as well as less synchrony in learning pace with the human. Providing explanations to make the robot behavior more transparent could possibly help in overcoming this problem. This could be especially useful when combined with the Collaboration Book, as that provides the machine with a model of what the human expects of the machine. Basing explanations on such a mental model has been mentioned as a prerequisite for effective explanations (Sreedharan et al., 2022).

While the use of the Collaboration Book had positive effects on task performance and participants' perception of the collaboration, the qualitative analyses show that this did not always mean that there was improved collaboration between the human participants and the virtual robot. In fact, many participants used the Collaboration Book to design the robot behavior rather than recognizing and documenting emergent CPs. The tendency to design behavior seemed to prevent participants to get a better understanding of the (principles underlying the) robot's behavior (see 6.4.1). In a way, we provided the participants with more control over the robot, which made the collaborations unbalanced (the human had significantly more power in determining the collaborative behavior than the robot), and we suspect that the Collaboration Book caused participants to view the robot as a tool rather than a team member. In future work, it will be necessary to carefully craft co-learning interventions, to reflect on the impact they will have on the

level of initiative and control team partners have within the human-AI team, as well as their general capability to contribute to Collaboration Patterns. If we truly want to maintain a collaborative dynamic in which both partners can contribute their expertise, we need to be mindful of how we design the level of control that both partners have over the task.

6.6. CONCLUSION

CO-LEARNING consists of two processes: ad hoc, implicit mutual adaptations (co-adaptation) and shared mental model updating. In this paper, we focus on the latter process by introducing an ontology model and accompanying Graphical User Interface (the Collaboration Book) which has been designed to facilitate shared mental model updating by enabling the formalization of emergent Collaboration Patterns.

The Collaboration Book provides a first attempt at the grounding of collaborative behavior in human-AI teams. We conclude that providing participants with the Collaboration Book enabled them to improve their task performance over time with the virtual robot. Participants also evaluated the collaboration with the robot more positively, even when their performance dropped later in the experiment. However, we cannot conclude that participants became more aware of emergent Collaboration Patterns, as they tended to design the CPs in a proactive manner rather than formalize spontaneously emerged CPs. Additionally, the Collaboration Book gave them more control over the robot, which caused a fundamentally different interaction dynamic in which there was less space for co-adaptation and learning about the robot partner. Moreover, in more complicated task scenarios, participants using the Collaboration Book were unable to maintain their performance increase.

Introducing support in the form of a Collaboration Book to foster co-learning in a human-AI team can help to improve team performance. However, maintaining a collaborative dynamic in which both team partners contribute equivalently to the task is a design requirement that needs to be explicitly taken into account when designing co-learning support and interventions. When designing interventions for shared mental model updating, the level of initiative that both human and AI team partners have needs to be evaluated critically.

BIBLIOGRAPHY

- Akata, Z., Balliet, D., Rijke, M. d., Dignum, F., Dignum, V., Eiben, G., Fokkens, A., Grossi, D., Hindriks, K., Hoos, H., Hung, H., Jonker, C., Monz, C., Neerincx, M., Oliehoek, F., Prakken, H., Schlobach, S., Gaag, L. v. d., Harmelen, F. v., ... Welling, M. (2020). A Research Agenda for Hybrid Intelligence: Augmenting Human Intellect With Collaborative, Adaptive, Responsible, and Explainable Artificial Intelligence [Conference Name: Computer]. *Computer*, 53(8), 18–28. <https://doi.org/10.1109/MC.2020.2996587>
- Andrews, R. W., Lilly, J. M., Srivastava, D., & Feigh, K. M. (2023). The role of shared mental models in human-AI teams: A theoretical review. *Theoretical Issues in Ergonomics Science*, 24(2), 129–175. <https://doi.org/10.1080/1463922X.2022.2061080>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Burke, S., Stagl, K., Salas, E., Pierce, L., & Kendall, D. (2006). Understanding Team Adaptation: A Conceptual Analysis and Model. *The Journal of applied psychology*, 91, 1189–207. <https://doi.org/10.1037/0021-9010.91.6.1189>
- Chiou, M., McCabe, F., Grigoriou, M., & Stolkin, R. (2021). Trust, Shared Understanding and Locus of Control in Mixed-Initiative Robotic Systems [ISSN: 1944-9437]. *2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN)*, 684–691. <https://doi.org/10.1109/RO-MAN50785.2021.9515476>
- de Visser, E. J., Peeters, M. M. M., Jung, M. F., Kohn, S., Shaw, T. H., Pak, R., & Neerincx, M. A. (2020). Towards a Theory of Longitudinal Trust Calibration in Human–Robot Teams. *International Journal of Social Robotics*, 12(2), 459–478. <https://doi.org/10.1007/s12369-019-00596-x>
- Dignum, F. (2018). Interactions as Social Practices: Towards a formalization [arXiv: 1809.08751]. *arXiv:1809.08751 [cs]*. Retrieved July 1, 2021, from <http://arxiv.org/abs/1809.08751>
- Hoffman, G. (2019). Evaluating Fluency in Human–Robot Collaboration. *IEEE Transactions on Human-Machine Systems*, 49(3), 209–218. <https://doi.org/10.1109/THMS.2019.2904558>
- Kumar, R. (2024). Unveiling the Dynamics of Human Decision-Making: From Strategies to False Beliefs in Collaborative Human-Robot Co-Learning Tasks.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13), 1–26. <https://doi.org/10.18637/jss.v082.i13>
- Nießen, D., Schmidt, I., Groskurth, K., Rammstedt, B., & Lechner, C. M. (2022). The Internal–External Locus of Control Short Scale–4 (IE-4): A comprehensive validation

- of the English-language adaptation. *PLoS ONE*, 17(7), e0271289. <https://doi.org/10.1371/journal.pone.0271289>
- Qader, W. A., Ameen, M. M., & Ahmed, B. I. (2019). An overview of bag of words; importance, implementation, applications, and challenges. *2019 international engineering conference (IEC)*, 200–204.
- Santoni de Sio, F., & van den Hoven, J. (2018). Meaningful Human Control over Autonomous Systems: A Philosophical Account. *Frontiers in Robotics and AI*, 5(15), 1–14. <https://doi.org/10.3389/frobt.2018.00015>
- Schoonderwoerd, T. A. J., Zoelen, E. M. v., Bosch, K. v. d., & Neerincx, M. A. (2022). Design patterns for human-AI co-learning: A wizard-of-Oz evaluation in an urban-search-and-rescue task. *International Journal of Human-Computer Studies*, 164, 102831. <https://doi.org/10.1016/j.ijhcs.2022.102831>
- Shafti, A., Tjomslund, J., Dudley, W., & Faisal, A. A. (2020). Real-World Human-Robot Collaborative Reinforcement Learning [arXiv: 2003.01156]. *arXiv:2003.01156 [cs]*. Retrieved March 8, 2021, from <http://arxiv.org/abs/2003.01156>
- Sharan, N. N., & Romano, D. M. (2020). The effects of personality and locus of control on trust in humans versus artificial intelligence. *Heliyon*, 6(8), e04572. <https://doi.org/10.1016/j.heliyon.2020.e04572>
- Sreedharan, S., Kulkarni, A., & Kambhampati, S. (2022). *Explainable human-ai interaction: A planning perspective*. Springer Nature.
- van Diggelen, J., Barnhoorn, J., Post, R., Sijs, J., van der Stap, N., & van der Waa, J. (2021). Delegation in Human-Machine Teaming: Progress, Challenges and Prospects. In D. Russo, T. Ahram, W. Karwowski, G. Di Bucchianico, & R. Taiar (Eds.), *Intelligent Human Systems Integration 2021* (pp. 10–16). Springer International Publishing. https://doi.org/10.1007/978-3-030-68017-6_2
- van Zoelen, E., Neerincx, M. A., van den Bosch, K., & Abbink, D. (2024). Reflective communication for human-ai co-learning.
- van Zoelen, E. M., van den Bosch, K., Abbink, D., & Neerincx, M. (2022). Ontology-Based Reflective Communication for Shared Human-AI Recognition of Emergent Collaboration Patterns. In R. Aydoğan, N. Criado, J. Lang, V. Sanchez-Anguix, & M. Serramia (Eds.), *PRIMA 2022: Principles and Practice of Multi-Agent Systems* (pp. 621–629). Springer International Publishing. https://doi.org/10.1007/978-3-031-21203-1_40
- van Zoelen, E. M., van den Bosch, K., & Neerincx, M. (2021). Becoming Team Members: Identifying Interaction Patterns of Mutual Adaptation for Human-Robot Co-Learning. *Frontiers in Robotics and AI*, 8, 200. <https://doi.org/10.3389/frobt.2021.692811>
- van Zoelen, E. M., van den Bosch, K., Rauterberg, M., Barakova, E., & Neerincx, M. (2021). Identifying Interaction Patterns of Tangible Co-Adaptations in Human-Robot Team Behaviors. *Frontiers in Psychology*, 12, 645545. <https://doi.org/10.3389/fpsyg.2021.645545>
- van den Bosch, K., Schoonderwoerd, T., Blankendaal, R., & Neerincx, M. (2019). Six Challenges for Human-AI Co-learning. *Lecture Notes in Computer Science*, 11597, 572–589. https://doi.org/10.1007/978-3-030-22341-0_45

- van der Waa, J., van Diggelen, J., Cavalcante Siebert, L., Neerincx, M., & Jonker, C. (2020). Allocation of Moral Decision-Making in Human-Agent Teams: A Pattern Approach. In D. Harris & W.-C. Li (Eds.), *Engineering Psychology and Cognitive Ergonomics. Cognition and Design* (pp. 203–220). Springer International Publishing. https://doi.org/10.1007/978-3-030-49183-3_16
- Welie, M. V., Veer, G. C. V. D., & Eliëns, A. (1998). An Ontology for Task World Models. *Proceedings of DSV-IS98, Abingdon*, 3–5.
- Wiltshire, T. J., Van Eijndhoven, K., Halgas, E., & Gevers, J. M. (2024). Prospects for augmenting team interactions with real-time coordination-based measures in human-autonomy teams. *Topics in Cognitive Science*, *16*(3), 391–429.
- Zhou, J., Su, X., Fu, W., Lv, Y., & Liu, B. (2024). Enhancing intention prediction and interpretability in service robots with llm and kg. *Scientific Reports*, *14*(1), 26999.

7

CONCLUSION

THE aim of this thesis is to better understand co-adaptation processes that lead to emergent behavior between a collaborating human and a learning machine, both virtual (agent) and physical (robot). Additionally, the aim is to understand what factors can shape this process towards successful collaboration, specifically through developing shared representations of emergent Collaboration Patterns (to ensure co-learning). These aims were studied in the context of task environments with clear collaborative goals and dependencies between human and machine team members, but in which the way in which the goals could be reached was not obvious at the start, thus requiring adaptation and learning by the human-machine team. Below, I will discuss how the five research questions following from this general aim were addressed, after which I discuss the main takeaways of the thesis as a whole. After that, I present several recommendations for future co-learning studies.

7.1. RESEARCH QUESTIONS ADDRESSED

RQ1 How do a human and a Wizard-of-Oz robot co-adapt leader-follower behavior when performing a collaborative task with no clear optimal solution?

I designed a new experimental paradigm in which a human and a robot-on-a-leash have to perform a joint navigation task requiring the human to continuously make tradeoffs between leading and following the robot. The embodied design of the interactions (leash pulls) allowed us to observe ad hoc and implicit behavioral adaptations. Participants used strategies that allowed both human and robot to partly lead, e.g. letting the robot decide on the direction for navigation, while walking in front and determining the speed themselves. They developed stable strategies with specific leader-follower balances fit to a certain stage of the task. When the strategy proved no longer adequate (due to environmental or internal triggers) participants suddenly adapted their behavior to use another stable strategy with a different leader-follower balance. Rather than a gradual development of leader-follower roles, roles were determined on an ad hoc basis, leading to a

variety of strategies per human-robot dyad that were reused across task instances (Table 2.6, p. 41).

RQ2 How do a human and an ML-powered virtual robot co-adapt when performing a collaborative task with no clear optimal solution?

I designed an experimental setting consisting of an Urban-Search-and-Rescue task with different possible solutions. Human participants needed to save an earthquake victim by clearing away rocks, together with a virtual robot, without having complete knowledge of the robot capabilities. I used a combination of rule-based strategies and basic Q-learning for the behavior of a virtual robot, which ensured that the robot would be able to learn different behaviors while collaborating with a human participant. Participants explored different strategies for collaborating with the robot, and developed and used several strategies tuned to the behavior of the robot and the context of the task, similar to what was observed when investigating RQ1 (see Table 3.3, p. 69). The robot learned and applied a similar set of strategies for most participants. However, when the level of adaptation on the human's side would render certain strategies more appropriate than others, the robot learned to prefer those specific strategies. Thus, the robot's policy was influenced by the level of adaptation of the human; see Figure 3.8, p. 71.

RQ3 How do a human and an embodied, ML-powered robot co-adapt when performing a collaborative task with no clear optimal solution?

I designed a human-machine hand-over task with a robot arm that was equipped with a Reinforcement Learning algorithm. This setup was used to evaluate five human-machine co-learning requirements: (1) having a shared goal, (2) synchrony in learning, (3) interdependence, (4) continuous adaptability of the machine's learning algorithm, and (5) mutual transparency. Several Collaboration Patterns (CPs) emerged. Initially, participants repeatedly switched between these CPs, until they converged towards a preferred CP (see Figure 4.3, p. 95). Of the evaluated human-robot dyads, three out of six met all requirements. The results warrant the assumption that at least two of the other dyads would also have met the requirements if they would have been allowed more time to practice (see Table 4.1, p. 94). The five requirements provide a starting point for systematically evaluating co-adaptation in human-machine teams.

RQ4 How can we facilitate co-adaptive human-robot collaborators with an infrastructure to formalize and share emergent collaboration patterns?

I designed (1) an ontology for formalizing emergent Collaboration Patterns, containing a description of the task situation in which the CP is applicable, and of the constituent actions of the CP, and (2) a Graphical User Interface (GUI) that enables the user to formalize the CPs in the ontology (see Figure 5.1, p. 110 and Figure 5.2, p. 111). This is called the 'Collaboration Book'. In an evaluation with human participants, these participants were able to recognize emergent CPs from video clips (recorded during a previous experiment) of human-robot teams exploring how to collaborate on a new task. They were able to formalize most of these CPs using the Collaboration Book. However, sometimes they failed, in particular when participants attempted to express properties

of collaboration that were not accommodated by the Collaboration Book, such as parallel actions and consequences of actions. I developed an improved version of the Collaboration Book in which these properties are included.

RQ5 Does the developed Collaboration Book (Chapter 5) support human-machine team members in becoming aware of and reusing the CPs, as well as in improving their team performance?

In an experimental study with human participants (Chapter 6), it was found that the developed Collaboration Book enables participants to improve their task performance over time (Figure 6.7, p. 133, as well as that participants evaluated the collaboration with the robot more positively when they had access to the Collaboration Book (Figure 6.14, 140). No evidence that participants became more aware of the emergent CPs was found, nor that they reused the CPs more often than participants without the Collaboration Book. The freedom offered by the Collaboration Book to unlimitedly design new CPs possibly prevented the participants to reuse earlier used CPs. It was also observed that the Collaboration Book induced participants to direct the robot's behavior rather than engaging in co-adaptation until successful collaboration emerged.

7.2. MAIN CONCLUSIONS

7.2.1. CO-ADAPTATION IS CHARACTERIZED BY EMERGENT STABLE STRATEGIES AND SUDDEN ADAPTATIONS

THE tasks in this thesis were purposefully designed to allow for and study human-machine co-adaptation (defined as the first step in co-learning from Chapter 3 onwards). In particular, these tasks were characterized by 1) various possibilities to achieve satisfactory task performance, and 2) both human and machine having incomplete information about each other and/or the task.

Before the empirical studies, the co-adaptation process was believed to be a gradual process in which the behavior of both human and machine would increasingly adapt until they reached some stable way of collaborating. Yet for the tasks considered, little evidence was found that human-machine co-adaptation follows such a gradual process. In three studies, transient behavioral co-adaptations were triggered by changes in the environment and changes in the human or machine, leading to a process in which the human-machine team jumps between different emergent stable strategies (Collaboration Patterns, Chapters 2, 3 and 4). To illustrate this with an example from Chapter 2: participants would sometimes completely follow the direction chosen by the robot (stable strategy), but suddenly start pulling the robot towards the goal location after crossing the middle of the field. Rather than a gradual development of behavior towards a single strategy, human-machine teams in the studies in this thesis developed a set of strategies (Collaboration Patterns) in an ad hoc manner that they applied contextually. Moreover, the dynamics of jumping between different emergent stable strategies are similar in human-virtual robot collaborations and human-physical robot collaborations (Chapters 3 and 4).

While emergence of behavioral patterns is an important topic of study in the field of multi-agent systems (see e.g. Baker et al., 2020; Teo et al., 2013), as far as I know, it has not

been systematically studied in the context of human-machine interaction or collaboration apart from some recent exploratory studies (Kumar, 2024; Shafti et al., 2020). Within the multi-agent systems community, studies on the emergence of language between AI agents do sometimes mention the potential of studying emergence in human-machine interactions to increase mutual understanding between the human and machine (Foerster et al., 2016; Lazaridou and Baroni, 2020), but they do not perform studies with human participants nor mention concrete frameworks for how this could be studied. Some vision papers have attempted to progress research on emergence in human-machine collaboration (Brandizzi and Iocchi, 2022; Taniguchi, 2016), but generally, research on human-machine collaboration focuses on gradual convergence towards a single optimal strategy.

The empirical work in this thesis shows that emergence is a central theme in human-machine collaboration, and that it is necessary to perform studies with human participants to gain insights into emergence. It should then be kept in mind that adaptation is usually not characterized by gradual convergence, but, as mentioned before, by stable strategies that collaborating partners jump between via sudden adaptations. The concepts of stable strategies and sudden adaptations can be used as a basic coding scheme to describe collaborative behavior and specifically emergence in a systematic manner.

7.2.2. COLLABORATION PATTERNS EMERGE AS A RESULT OF CO-ADAPTATION

COLLABORATION Patterns emerge from joint adaptive processes. This means that behavior learned by one of the agents influences the partner's view on what is considered successful behavior, and vice versa. If the machine learns some kind of strategy based on rewards obtained in the task, and then the human changes their strategy, this will affect the outcome of their joint behavior and therefore also the rewards obtained by the machine. This can in turn trigger the machine to re-adapt, which may then cause the human to adapt their behavior. This process may repeat itself several times, until it eventually leads to a stable and shared strategy, here called a Collaboration Pattern. Chapters 2, 3 and 4 present evidence for how joint learning processes lead to the emergence of stable Collaboration Patterns (for the observed duration of the task). Some of these CPs will be strongly dictated by the design of the task (through certain affordances) and can therefore already be predicted by the designers of the task, while some can be surprising even to the designers of the task.

Results from Chapters 3 and 4 show that these emergent Collaboration Patterns can be followed and initiated by both the human as well as the machine, even if they do not have an explicit representation of the CPs. In both chapters, qualitative coding of the joint human-machine behavior lead to the discovery of Collaboration Patterns. The policy of the RL-powered machine (the learned Q-tables) matched with the CPs identified by qualitative coding, even though there was no explicit modeling of those CPs (see Figure 3.8, p. 71). Human participants, in turn, repeatedly initiated behavior that fit with the identified CPs. For example, if the human and the machine engaged in a CP within which the human would stand still on a large rock after which the robot would pick up that rock (Chapter 3), participants would move to and stand still on a large rock repeatedly, while the Q-tables showed that the robot preferred actions that fit with such a strategy. Participants are sometimes able to explicitly mention this behavior when asked for

their strategy after the task, but several participants denied having such a strategy while still performing the behavior.

I have proposed the term ‘Collaboration Pattern’ to describe emergent collaborative behavior in human-machine teams. Initially, I used the term ‘Interaction Pattern’, as that is a term more commonly used in literature to describe humans and machines interacting (Baltzer et al., 2019; Flemisch et al., 2022; Madan et al., 2015; Mörtl et al., 2012). However, the studies presented in this thesis specifically regard collaborative behavior, not just any type of human-machine interactions. In human-machine teaming literature, the term ‘Team Design Pattern’ is also used frequently (see e.g. van Diggelen and Johnson, 2019; van der Waa et al., 2020), but that tends to refer to the design of a team configuration (including their responsibilities) rather than the actual actions they engage in. Throughout the process of co-learning the human and machine may interact in different ways, but what emerges from their adaptations is a way to collaborate, which can therefore be called a Collaboration Pattern. All patterns categorized as ‘Stable Situations’ (Chapter 2) are, in fact, the emergent Collaboration Patterns. The other identified interaction patterns (Sudden Adaptations, Gradual Adaptations and Active Negotiations) are indeed ‘just’ Interaction Patterns. By adding a nuance between Interaction Patterns and Collaboration Patterns, I hope to contribute to a language for describing dynamic collaborations and to make emergent Collaboration Patterns a central topic in the study of human-machine co-learning.

7.2.3. HUMANS ARE ABLE TO RECOGNIZE AND FORMALIZE EMERGENT CPs

ONCE people understand the concept of CPs, they are generally able to recognize the emergence of such a pattern when they collaborate with a machine. Moreover, they are able to formalize a pattern using an ontology for documenting CPs and accompanying GUI, although the structure and details in which they formalize patterns vary widely (Chapters 5 and 6). When humans are able to share formalized CPs with their machine team partner during collaboration, this improves the human-machine team performance. Furthermore, the human will become more confident that outcomes will be successful and will generally evaluate the collaboration more positively (Chapter 6).

The Collaboration Book provides human-machine partners with a knowledge structure to formalize their emergent Collaboration Patterns. This can help them build a shared mental model of their collaborative behavior, comparable to how shared mental models function in human-only teams (Andrews et al., 2023; Uitdewilligen et al., 2013). In the literature on human-machine collaboration, many studies address the creation of common ground by conducting activities to ensure that partners agree on the meaning of a particular concept (for example by pointing at an object and giving it a name). Having common ground (a shared representation of the environment) is generally regarded as the basic requirement for collaboration (Klein et al., 2004). Literature on achieving common ground in human-machine collaboration is generally about creating models of the physical environment (e.g. grounding specific objects, see Buschmeier and Kopp, 2013; Chai et al., 2014). In this thesis I argue that it is also necessary to achieve common ground on collaborative behavior by engaging in the grounding of collaborative actions. Chapters 5 and 6 report results of a literature search on the grounding of collaborative behavior. Different structures have been proposed to formalize collaborative

behavior (e.g. Plays (Kasmier et al., 2021), sometimes provided in a Playbook (Miller et al., 2005), see also Chapter 5). In all of these structures, however, the researchers design the patterns of collaborative behavior in a top-down manner. This approach does not ensure that the patterns are embedded in the perceptions of both team members, which is necessary to achieve common ground. In such cases, it is hard to know if the partners understand the behavior in a similar manner. Only some work on Social Practices (Dignum, 2018) suggests that formalizations of collaborative behavior could be based on emergent patterns, but in the research itself, this is not implemented. By providing the Collaboration Book as a way to formalize emergent CPs (i.e. CPs that originate as a result of joint adaptation), I have provided a first attempt at enabling human-machine partners to ground collaborative behavior. The Collaboration Book can be filled and refined during task execution, and is by design machine-interpretable.

7.2.4. HUMANS DON'T EASILY ACCEPT MACHINES AS TEAM PARTNERS

IN this thesis I adopt a human-machine teaming perspective on human-machine interaction. It is clear, also in other related work not presented in this thesis (Schoonderwoerd et al., 2022), that it is difficult to create a setting in which humans accept the machine as their team member rather than as a tool. Creating team behavior and a team feeling requires frequent reciprocal interactions, unique contributions from both team members, and a mindset in the human recognizing that they themselves have limited capabilities too. Once the team dynamic is achieved, it is fragile, and it breaks easily when new task elements are introduced that make the contribution less balanced, as was shown in Chapter 6. For example, if the human is given more control over the robot, they will then immediately revert to directing the robot around as if it is a tool. This was often the case with the Collaboration Book, as participants designed the collaborative behavior based on their own preferences rather than observing the robot's behavior (see Chapter 6).

7.2.5. DEFINING CO-LEARNING FOR THE STUDY OF COMPLEX HUMAN-MACHINE COLLABORATION DYNAMICS

CO-LEARNING is discussed and defined in Chapter 3. In several related research communities (for example human-robot collaboration, human-agent teaming, multi-agent systems), the terms Co-Learning, Co-Adaptation and Co-Evolution are often used interchangeably without clear definitions (see below for some example references). I aimed to unify these communities by creating operational definitions for the different terms. The definitions provided enable the study of co-learning by making it concrete enough to directly inform experimental designs. This has led to the requirements for co-learning presented in Chapter 4 (having a shared goal, synchrony in learning, interdependence, continuous adaptability of the machine's learning algorithm, and mutual transparency), as well as the design of the Collaboration Book (Chapters 5 and 6).

Due to rising interest in hybrid intelligence and human-machine collaboration, the number of papers using the term 'co-learning' has started to grow rapidly, but most papers do not provide or follow a formal definition. Rather, they follow the broad definition of a human and AI system learning together without explicitly distinguishing co-learning from other terms in literature such as co-adaptation and co-evolution (see e.g. Kumar,

2024; Pepe and Hutchison, 2022; Shafti et al., 2020). Sometimes, the definition I proposed was cited (see e.g. Li et al., 2023). However, the operationalizations of co-learning in the mentioned papers do not follow concrete requirements, nor do they reflect on the distinction between co-adaptation and co-learning (the latter also involving Shared Mental Model Updating, as defined in Chapter 3). To support that co-learning becomes a main topic of study in human-machine collaboration, a wide discussion is needed on the definition of co-learning, as well as on its operationalization in experimental settings. With this thesis, I provide a first step in that direction, hoping that others will adopt or reflect on the definitions and operationalizations to continue improving them.

7.3. LIMITATIONS

ASIDE from the contributions emanating from the studies in this thesis, there are also some limitations. First of all, co-learning has been studied in the context of collaborating humans and machines presented as agents (robot-like machines, even when they were presented in a virtual environment); human and machine act within the same environment. Generalizing the findings to other types of machines, such as Decision Support Systems and Chat Agents, is therefore not possible. Interaction with such machines will usually involve more communication actions, but fewer behavioral actions, which will change what it means to co-adapt. The conclusions should be interpreted as applicable to machines presented as agents first and foremost, and further research is necessary to understand if they can be extended to Decision Support Systems and the like. Additionally, I have only studied dyadic collaborations (one human, one machine). Therefore, I do not know how the results scale to settings with more than two collaborators. To study co-learning in larger teams, it might be necessary to describe CPs at a more abstract level. Especially formalizing the CPs in the Collaboration Book will quickly become more complex and cumbersome the more team members are involved, creating new design requirements.

Secondly, the studies were conducted using tasks and research environments with low ecological validity. The thesis does not provide insights on how the requirements relate to real-life tasks in which co-learning could be applicable. While the tasks were designed analogous to plausible human-machine collaboration scenarios (such as USAR), the next step in co-learning research should be to move to more realistic task environments to increase ecological validity.

Thirdly, the number of task repetitions used in the experiments proved to sometimes not be sufficient to establish clear emergent Collaboration Patterns. It may very well be that the time-on-task in the lab experiments was too short (usually about an hour) to demonstrate solid learning effects. Studies with more exposure, more practice, and longer timelines will be necessary to study co-learning behavior in real-life settings.

Moreover, the algorithms used for governing machine behavior throughout the work in this thesis do not match the current State-of-the-Art of Machine Learning algorithms, because designing algorithms for learning in collaboration with humans has its own unique challenges (e.g. data scarcity). Additionally, the algorithms served mostly to facilitate the interaction dynamic necessary to study co-learning and extensively exploring different algorithms was therefore out of scope. To give co-learning research more relevance and to connect more strongly to the AI community, it will be necessary to

perform deeper investigations of the State-of-the-Art of several Machine Learning sub-fields (e.g. one-shot learning, real-time learning, interactive Machine Learning) to evaluate whether other algorithms could work in co-learning settings. Complexities in such learning algorithms might change how it adapts to human adaptation, which might in turn affect how humans adapt their behavior (e.g. they might adapt less because the machine is more adaptive).

Last, the design of the GUI presented in Chapters 5 and 6 was limited by the possibilities of the experiment environment used (the MATRX environment, see “MATRX Software”, 2021). While it was possible to build a custom interface, the connection with the code of the task and agents has some limitations on speed and form. Moreover, no exploration and evaluation of design alternatives was done. While the GUI was sufficient for the studies presented, as shown by the evaluations described in Chapter 5, it will be useful to do more extensive research and design of interaction modalities for future studies.

7.4. RECOMMENDATIONS FOR STUDYING CO-LEARNING

THROUGHOUT the thesis, several task and research environments specifically designed for co-learning have been presented. From these, recommendations for future co-learning studies became clear in three main areas: task design, algorithm design, and assessing co-learning. Before presenting these recommendations in the sections below, I will first categorize the research environments presented in this thesis following the classification scheme made in Chung et al., 2024, to support researchers in choosing the right testbed and task environment for their research and to provide the opportunity for reusing my task environments.

- **Team composition:** One-to-one
- **Task interdependence:** Reciprocal
- **Role structure:** Functional
- **Leadership structure:** Temporary (leadership can shift in all of the presented research environments)
- **Authority differentiation:** None
- **Communication structure:** Dyadic
- **Communication direction:** Bidirectional in Chapters 2, 3 and 4 (only via behavioral interactions); unidirectional in Chapters 5 and 6
- **Communication medium:** Haptic-based in Chapter 2; gesture-based in Chapter 3; a combination of gesture-based and haptic based in Chapter 4; text-based in Chapters 5 and 6
- **Team life span:** Long-term

7.4.1. TASK DESIGN

TO study co-learning, a task must be **sufficiently open-ended**. The best strategy for performing the task should not be self-evident at the start of the collaboration, to allow for exploration and learning, and preferably the task should have several ways in which it can be completed successfully. This open-endedness should be balanced with creating sufficiently controlled experimental environments to ensure that it is still possible to compare the results of different participants. The tasks presented in Chapters 2, 3 and 4 all present instances of such tasks. In Chapter 2, open-endedness is created by a continuous tradeoff between speed and points obtained by picking up objects. In Chapter 3, rocks can be picked up in many different orders to save the victim. In Chapter 4, an object handover can be achieved by several different hand and arm positions. In all of these tasks, there was a clear task objective, and a limited variety of possible strategies to achieve the objective.

Designing hard dependencies (situations in which a task can only be completed when both human and machine contribute) into a task ensures that it is necessary for human and machine to collaborate. To support co-learning, the **opportunity for soft dependencies** (situations in which a task could be completed by one of the collaborators alone, but in which performance can be improved significantly if they assist each other) is however vital. Specifically to ensure that co-adaptations can occur, human and machine need to be able to explore how both can contribute by supporting the other in a proactive and anticipatory manner. When only hard dependencies are present, this easily results in a collaboration in which both partners work on their own task without much interaction, removing opportunities for further co-adaptation. The observed CPs in Chapter 2 (Table 2.6) show different examples of how soft dependencies made the team perform better, e.g. when the human pulled the robot in a direction chosen by the robot to increase its speed. In Chapter 3, one of the main observed CPs was that the human would do 'damage control' by building piles of rocks to anticipate possible mistakes done by the virtual robot and also to nudge them towards picking up specific rocks.

The above relates to the **establishment of a team-attitude** in the human. When humans view the machine as a tool, they are inclined to either control the machine, or to focus on their own task and let the machine do their part autonomously. The results of the experiment in Chapter 6 show how a team-attitude can be facilitated through task design, and also how it can be violated by the introduction of an intervention that gives the human partner too much control. To ensure that both human and machine can learn and improve, the human needs to allow the machine some autonomy, while at the same time paying attention to its behavior. Such an attitude can be triggered by a clear need to collaborate. Of course the machine should also be capable of acting as a team partner, for example by proactively anticipating on the human's actions (soft dependence) and providing transparency on its way of working.

7.4.2. ALGORITHM DESIGN

TO study co-learning, an important challenge is to design an algorithm for the machine that generates dynamic co-adaptive behaviors. A certain level of **synchrony** (synchronized learning pace, i.e. both human and machine learn at a similar pace) is necessary, to ensure that both human and machine can adapt to adaptations done by

the other. For example, in Chapter 4, results showed that one of the participants learned much faster than the robot, which caused them to feel like the robot was dragging them down. This led to the participant starting to fail at the task on purpose to ‘train’ the robot towards a desired behavior (Figure 4.4, p. 96). State-of-the-Art algorithms often cannot achieve such synchrony, because they require large amounts of data for learning that cannot be obtained in real time. They therefore cannot adapt to human adaptations when necessary, but only after training for a period of time. A possible way to achieve synchrony is to combine abstract state representations and high level actions with relatively simple learning algorithms such as Q-learning, like it was done in this thesis.

Algorithms also need to be configured such that they can **continuously adapt and learn in new situations**, as the real world is ever changing. In Chapters 3 and 6, a brown rock was introduced in the victim evacuation task, which required a different strategy than scenarios without the brown rock. These circumstances demanded the partners to again adapt to the task and to each other, thus developing and learning new Collaboration Patterns. If co-learning is the objective, introducing such interventions is a way to foster this and to evaluate whether human and machine are capable of learning new CPs. In the studies in this thesis, it was ensured that the algorithm would always keep exploring or that the rewards would be such that the algorithm could easily change its policy.

7.4.3. ASSESSING CO-LEARNING

STUDYING the dynamics of co-learning requires a **mixed-methods approach**, to be able to capture the richness of behavioral interactions as well as make statistical inferences (see also Fragiadakis et al., 2024). It is important to combine traditional performance metrics with thorough qualitative analyses of adaptive behavior of both humans and machines, as well as of how their understanding of the task and of each other develops over time. Moreover, attention should be paid to the constructive nature of any human-machine interaction study; the fact that the behavior of the machine is designed requires reflection on how designs influence observed results (Koskinen et al., 2011).

In the studies of this thesis, observations and thematic analyses were used for measuring behavior, and coding schemes were developed to analyze the emergent Collaboration Patterns (Chapters 2, 3 and 6). These coding schemes can be adapted for use in future co-learning studies, especially for characterizing the dynamics of jumping between different emergent Collaboration Patterns. The coding schemes provide a vocabulary with which complex collaborative behavior can be described, which helps to unify results from different types of methods.

The results from these coding processes were combined with quantitative results in Chapter 6, to evaluate how behavior relates to performance. Moreover, requirements for the design of tasks and algorithms has been a main theme throughout this thesis, to reflect on their influence on the results, as well as to help future researchers in designing their own tasks.

7.5. FUTURE WORK

THE conclusions and limitations of the studies in this thesis lead to several possible lines of future work, that I will outline further in the sections below.

7.5.1. TASK ENVIRONMENTS AND REAL-LIFE SETTINGS

IN this thesis, insight into the dynamics of Co-Adaptation and Co-Learning was gained for the specific tasks using physical robots (robot on a leash, robot arm) and virtual robots. It is important to further investigate whether the finding that Co-Adaptation leads to the emergence of stable Collaboration Patterns transfers to other task environments, such as real-life human-machine collaboration settings. Addressing this question results in two directions for future research:

1. Comparing co-learning dynamics across different task environments;
2. Studying co-learning dynamics in real world tasks.

COMPARING TASK ENVIRONMENTS

It would be useful to systematically compare co-learning in different task environments. A first step would be to analyze the tasks used in earlier studies in terms of task properties, which should then be compared with existing task taxonomies. A starting point could be taken from general human-machine teaming literature; Chung et al., 2024 created an initial classification scheme, and more efforts were initiated during the Lorentz Workshop on Research Environments for Human-Machine Teaming (“Center for Scientific Workshops in All Disciplines - Research Environments for Human-Machine Teaming — lorentzcenter.nl”, n.d.). Using the resulting task categorization, it should be possible to design task environments that differ systematically on task aspects (such as modality, level of open-endedness, number of task repetitions) while adhering to the task requirements for establishing a Co-Learning dynamic as identified in this thesis. Such systematic comparisons would help deepen the understanding of how Co-Learning dynamics are influenced by different task characteristics.

REAL WORLD TASKS

Real world tasks are often dynamic, and humans are always prone to adapt and learn. In the tasks presented in this thesis, it was attempted to simulate the dynamic nature of the real world by designing them in an open-ended manner and introducing unexpected changes to the environment (e.g. the introduction of the brown rock in Chapters 3 and 6), but of course this only simulates limited aspects of the real world. The thesis presents several conclusions on co-adaptive processes (e.g. consisting of sudden adaptations, from which stable Collaboration Patterns emerge). To evaluate whether sudden adaptations and emergent stable CPs can also be observed in real-life human-machine collaborations, similar studies should be executed in real world tasks. Moreover, it should also be evaluated whether interactions with an ontology, such as via a Collaboration Book, can also be useful in real-life task environments. Studies in real world settings could for example involve first responders collaborating with drones (see for example studies like Kruijff-Korbayová et al., 2015; Mioch et al., 2021), factory workers engaging in complex

manufacturing tasks together with robots, or studies done in the context of training programs (such as for defense). Lab studies with experiment environments that have higher ecological validity could also be used to test the generalizability of the conclusions.

7.5.2. MODALITY OF INTERACTIONS ON COLLABORATION PATTERNS

THE research tasks in this thesis involved relatively limited interaction between human and machine. The initial chapters use behavioral interactions (Chapters 2, 3 and 4), and later on communication via a GUI is used (Chapters 5 and 6), to enable the sharing of emergent Collaboration Patterns. The GUI provided the human-machine team with the means to enter and access shared representations of Collaboration Patterns. The first design of the GUI was evaluated and improved to ensure that it worked well enough, but its functioning was not compared with other communication forms or modalities. It is well known that communication modalities affect the effectiveness of human-machine interactions (Bonarini, 2020). Therefore, it can be expected that it will affect the effectiveness of Shared Mental Model updating as well. Results of the experiments in Chapter 6 showed limitations of the used GUI (e.g. participants were sometimes confused about the meaning of some of the concepts and indicated that they felt restricted in what behavioral interactions they could express). It would therefore be interesting to explore whether different user interfaces and modalities can establish better communication about emergent Collaboration Patterns, possibly also in other task contexts.

To make this more concrete, it would be interesting to investigate whether natural language models can be used to bring about effective communication about emergent CPs. LLM-based tools can for example be used to create a natural language interface for this purpose (e.g. as proposed in Chin et al., 2024). This would enable people to express themselves more freely (compared to a restricted GUI or traditional chatbot), while making it easier to let the machine give feedback on the formalized CPs. To ensure grounding of concepts used in the CP formalizations, the challenge would lie in translating natural language utterances into the concepts of the ontology model.

7.5.3. ALGORITHMS FOR MACHINES LEARNING IN SYNCHRONY WITH HUMANS

TO achieve co-learning, machines should be able to adapt to and learn with humans. Creating algorithms that allow machines to do so synchronously with humans is an open research challenge. The machines used in this thesis were equipped with different variations of a basic Q-learning algorithm, using a high-level state representation to ensure a small state-space. This ensures that the machine can learn within a small amount of trials, enabling it to learn in synchrony with the human (i.e., at a human time scale). These learning mechanisms are relatively basic, they do not always produce stable policies, and are limited in what they can learn. State of the art machine learning algorithms that have more impressive learning capabilities, however, are rarely designed to learn at a human time scale, due to the need for large amounts of data.

Recently, developments in AI, specifically in Foundation Models, have triggered people to work on algorithms that allow AI agents (specifically RL agents) to adapt and learn at a human time scale, even when using large state and action spaces (Bauer et al., 2023).

The results of this work show that collaboration patterns emerged in a multi-agent setting, in few-shots and at a human timescale. The evaluations were however only done with software agents, so it remains an open question whether patterns will emerge when these agents collaborate with humans whose learning follows a different course than a learning software agent. There are other studies that attempt to leverage cognitive models for better and faster adaptation when collaborating with humans (e.g. Çelikok, 2023), compared to traditional ML models. To connect research on co-learning with high performing algorithms, an interesting direction for future research would be to evaluate promising existing algorithms such as the ones mentioned above in human-machine co-learning experiments. This can help evaluate if their performance remains high when collaborating with humans, as well as whether they are suitable for co-adaptive interactions.

7.5.4. MACHINE INITIATIVE: AI RECOGNITION OF EMERGENT CPs

IN the initial studies in this thesis, human-machine interactions were designed in a relative symmetrical way: both human and machine have capabilities that the other does not have, and communication opportunities of both are similar (Chapters 2, 3 and 4). In later studies, asymmetry was introduced in the mechanism for Shared Mental Model updating (Chapters 5 and 6). Chapter 6 showed that when one partner has more control over the collaboration (in this case through the initiation of formalized CPs), this is likely to break the collaborative dynamic between human and machine. Rather than relying on emergent learning from mutual adaptation, the human partner is then inclined to treat the robot as a tool rather than a partner. If symmetry could be ensured, e.g. by providing the machine with the means to also propose formalizations of Collaboration Patterns, it would be interesting to see whether that would lead to a more balanced co-learning process. For achieving this, it would be necessary to provide the machine with the capability to recognize emergent Collaboration Patterns out of the stream of interactions. This could possibly be done using pattern recognition algorithms (see e.g. Abbaspour et al., 2021). Moreover, the machine would need a way to assess whether emergent Collaboration Patterns are valuable or not with regards to the goal of the task. Combining pattern recognition with reward structures that have already been implemented for the Reinforcement Learning used by the machine may be of help in this matter.

7.6. FINAL REMARKS

THE aim driving the research in this thesis was to better understand co-learning of collaborating humans and machines in dynamic tasks. Another aim was to find methods to facilitate successful outcomes of human-AI co-learning. I developed a definition of co-learning and used the concept of emergence to study human-machine collaborative dynamics. Co-adaptation is not a gradual process. Instead, human-machine partners develop emergent Collaboration Patterns that they engage in for a period of time, and sudden behavioral deviations from these patterns can be triggered by a multitude of factors, such as a new insight about the task or achieving a subgoal. The co-adaptive dynamic of alternating between different emergent Collaboration Patterns exists in collaborations with a virtual robot partner, as well as with a physical robot partner. Designing

tasks and research environments that allow for a co-adaptive dynamic is, however, not an easy challenge. I have proposed several tasks and research environments to address this challenge. Moreover, I have provided extensive reflections and a set of requirements to evaluate whether an environment sufficiently allows for co-adaptation to develop, and thereby for it to be useful for conducting co-learning studies.

The ontology design and accompanying GUI showed that people are able to recognize and formalize emergent Collaboration Patterns when prompted to do so. Ensuring that this leads to formalizations in which human expectations and robot behaviors match is still an open research challenge, but the proposed design did show that the formalization of emergent Collaboration Patterns can help the team to perform better and faster. This thesis provides a foundation for future research into human-machine co-learning and related processes, and I hope that it inspires researchers to take a process-centered perspective on understanding and designing for co-learning in human-machine collaboration.

BIBLIOGRAPHY

- Abbaspour, S., Naber, A., Ortiz-Catalan, M., GholamHosseini, H., & Lindén, M. (2021). Real-Time and Offline Evaluation of Myoelectric Pattern Recognition for the Decoding of Hand Movements [Number: 16 Publisher: Multidisciplinary Digital Publishing Institute]. *Sensors*, 21(16), 5677. <https://doi.org/10.3390/s21165677>
- Andrews, R. W., Lilly, J. M., Srivastava, D., & Feigh, K. M. (2023). The role of shared mental models in human-AI teams: A theoretical review. *Theoretical Issues in Ergonomics Science*, 24(2), 129–175. <https://doi.org/10.1080/1463922X.2022.2061080>
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., & Mordatch, I. (2020). Emergent Tool Use From Multi-Agent Autocurricula [arXiv: 1909.07528]. *arXiv:1909.07528 [cs, stat]*. Retrieved December 10, 2020, from <http://arxiv.org/abs/1909.07528>
- Baltzer, M. C. A., López, D., & Flemisch, F. (2019). Towards an interaction pattern language for human machine cooperation and cooperative movement. *Cognition, Technology & Work*, 21(4), 593–606. <https://doi.org/10.1007/s10111-019-00561-8>
- Bauer, J., Baumli, K., Behbahani, F., Bhoopchand, A., Bradley-Schmieg, N., Chang, M., Clay, N., Collister, A., Dasagi, V., Gonzalez, L., Gregor, K., Hughes, E., Kashem, S., Loks-Thompson, M., Openshaw, H., Parker-Holder, J., Pathak, S., Perez-Nieves, N., Rakicevic, N., ... Zhang, L. M. (2023). Human-Timescale Adaptation in an Open-Ended Task Space [ISSN: 2640-3498]. *Proceedings of the 40th International Conference on Machine Learning*, 1887–1935. Retrieved August 27, 2024, from <https://proceedings.mlr.press/v202/bauer23a.html>
- Bonarini, A. (2020). Communication in human-robot interaction. *Current Robotics Reports*, 1(4), 279–285.
- Brandizzi, N., & Iocchi, L. (2022). Emergent communication in human-machine games. Retrieved August 29, 2024, from <https://openreview.net/forum?id=rqLgeQWCXZ9>
- Buschmeier, H., & Kopp, S. (2013). Co-constructing Grounded Symbols - Feedback and Incremental Adaptation in Human-Agent Dialogue. *KI - Künstliche Intelligenz*, 27, 137–143. <https://doi.org/10.1007/s13218-013-0241-8>
- Çelikok, M. M. (2023). *Model-based Multi-agent Reinforcement Learning for AI Assistants* [ISSN: 1799-4942 (electronic)]. Aalto University. Retrieved August 28, 2024, from <https://aaltodoc.aalto.fi/handle/123456789/120725>
- Center for Scientific Workshops in All Disciplines - Research Environments for Human-Machine Teaming — lorentzcenter.nl [[Accessed 13-04-2025]]. (n.d.).
- Chai, J. Y., She, L., Fang, R., Ottarson, S., Littley, C., Liu, C., & Hanson, K. (2014). Collaborative Effort towards Common Ground in Situated Human-Robot Dialogue. *2014 9th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 33–40.

- Chin, D., Wang, Y., & Xia, G. (2024). Human-centered llm-agent user interface: A position paper. *arXiv preprint arXiv:2405.13050*.
- Chung, H., Holder, T., Shah, J., & Yang, X. J. (2024). Developing a Team Classification Scheme for Human-Agent Teaming [Publisher: SAGE Publications Inc]. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 68(1), 1394–1399. <https://doi.org/10.1177/10711813241260387>
- Dignum, F. (2018). Interactions as Social Practices: Towards a formalization [arXiv: 1809.08751]. *arXiv:1809.08751 [cs]*. Retrieved July 1, 2021, from <http://arxiv.org/abs/1809.08751>
- Flemisch, F., Usai, M., Herzberger, N. D., Baltzer, M. C. A., Hernández, D. L., & Pacaux-Lemoine, M.-P. (2022). Human-Machine Patterns for System Design, Cooperation and Interaction in Socio-Cyber-Physical Systems: Introduction and General overview [ISSN: 2577-1655]. *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 1278–1283. <https://doi.org/10.1109/SMC53654.2022.9945181>
- Foerster, J., Assael, I. A., de Freitas, N., & Whiteson, S. (2016). Learning to Communicate with Deep Multi-Agent Reinforcement Learning. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 29* (pp. 2137–2145). Curran Associates, Inc. Retrieved February 12, 2019, from <http://papers.nips.cc/paper/6042-learning-to-communicate-with-deep-multi-agent-reinforcement-learning.pdf>
- Fragiadakis, G., Diou, C., Kousiouris, G., & Nikolaidou, M. (2024). Evaluating human-ai collaboration: A review and methodological framework. *arXiv preprint arXiv:2407.19098*.
- Kasmier, D., Merrell, E., Kelly, R., Smith, B., Heisey, C., Maki, E., Brittain, M., Ankner, R., & Bush, K. (2021). Ontology of Plays for Autonomous Teaming and Collaboration. *Proceedings of the XIV Seminar on Ontology Research in Brazil (ONTOBRAS)*, 9–22.
- Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. J. (2004). Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems*, 19(6), 91–95. <https://doi.org/10.1109/MIS.2004.74>
- Koskinen, I., Zimmerman, J., Binder, T., Redström, J., & Wensveen, S. (2011). *Design research through practice*. morgan kaufmann.
- Kruijff-Korbayová, I., Colas, F., Gianni, M., Pirri, F., de Greeff, J., Hindriks, K., Neerincx, M., Ögren, P., Svoboda, T., & Worst, R. (2015). Tradr project: Long-term human-robot teaming for robot assisted disaster response. *KI-Künstliche Intelligenz*, 29, 193–201.
- Kumar, R. (2024). Unveiling the Dynamics of Human Decision-Making: From Strategies to False Beliefs in Collaborative Human-Robot Co-Learning Tasks.
- Lazaridou, A., & Baroni, M. (2020, July). Emergent Multi-Agent Communication in the Deep Learning Era [arXiv:2006.02419 [cs]]. Retrieved August 29, 2024, from <http://arxiv.org/abs/2006.02419>
- Li, Y., Zhang, Z., & Zhang, F. (2023). Hybrid Co-Learning for Proximate Human-Robot Teaming. *2023 20th International Conference on Ubiquitous Robots (UR)*, 239–244. <https://doi.org/10.1109/UR57808.2023.10202393>

- Madan, C. E., Kucukyilmaz, A., Sezgin, T. M., & Basdogan, C. (2015). Recognition of Haptic Interaction Patterns in Dyadic Joint Object Manipulation [Conference Name: IEEE Transactions on Haptics]. *IEEE Transactions on Haptics*, 8(1), 54–66. <https://doi.org/10.1109/TOH.2014.2384049>
- MATRX Software [original-date: 2020-02-27T12:34:36Z]. (2021, January). Retrieved February 5, 2021, from <https://github.com/matrix-software/matrix>
- Miller, C., Funk, H., Wu, P., Goldman, R., Meisner, J., & Chapman, M. (2005). The play-book™ approach to adaptive automation. *Proceedings of the human factors and ergonomics society annual meeting*, 49(1), 15–19.
- Mioch, T., Sterkenburg, R., Beuker, T., & Neerincx, M. A. (2021). Actionable situation awareness: Supporting team decisions in hazardous situations. *ISCRAM*, 62–70.
- Mörzl, A., Lawitzky, M., Kucukyilmaz, A., Sezgin, M., Basdogan, C., & Hirche, S. (2012). The role of roles: Physical cooperation between humans and robots. *The International Journal of Robotics Research*, 31(13), 1656–1674. <https://doi.org/10.1177/0278364912455366>
- Pepe, K., & Hutchison, N. (2022). AI4SE and SE4AI: Setting the Roadmap toward Human-Machine Co-Learning [eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/inst.12417>]. *INSIGHT*, 25(4), 80–84. <https://doi.org/10.1002/inst.12417>
- Schoonderwoerd, T. A. J., Zoelen, E. M. v., Bosch, K. v. d., & Neerincx, M. A. (2022). Design patterns for human-AI co-learning: A wizard-of-Oz evaluation in an urban-search-and-rescue task. *International Journal of Human-Computer Studies*, 164, 102831. <https://doi.org/10.1016/j.ijhcs.2022.102831>
- Shafti, A., Tjomsland, J., Dudley, W., & Faisal, A. A. (2020). Real-World Human-Robot Collaborative Reinforcement Learning [arXiv: 2003.01156]. *arXiv:2003.01156 [cs]*. Retrieved March 8, 2021, from <http://arxiv.org/abs/2003.01156>
- Taniguchi, T. (2016). Symbol Emergence in Robotics for Long-Term Human-Robot Collaboration*. *IFAC-PapersOnLine*, 49(19), 144–149. <https://doi.org/10.1016/j.ifacol.2016.10.476>
- Teo, Y. M., Luong, B. L., & Szabo, C. (2013). Formalization of emergence in multi-agent systems. *Proceedings of the 1st ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, 231–240. <https://doi.org/10.1145/2486092.2486122>
- Uitdewilligen, S., Waller, M. J., & Pitariu, A. H. (2013). Mental Model Updating and Team Adaptation [Publisher: SAGE Publications Inc]. *Small Group Research*, 44(2), 127–158. <https://doi.org/10.1177/1046496413478205>
- van Diggelen, J., & Johnson, M. (2019). Team Design Patterns. *Proceedings of the 7th International Conference on Human-Agent Interaction*, 118–126. <https://doi.org/10.1145/3349537.3351892>
- van der Waa, J., van Diggelen, J., Cavalcante Siebert, L., Neerincx, M., & Jonker, C. (2020). Allocation of Moral Decision-Making in Human-Agent Teams: A Pattern Approach. In D. Harris & W.-C. Li (Eds.), *Engineering Psychology and Cognitive Ergonomics. Cognition and Design* (pp. 203–220). Springer International Publishing. https://doi.org/10.1007/978-3-030-49183-3_16

A

LIST OF OBSERVED BEHAVIORS IN CHAPTER 2

Table A.1: A list of all observed behaviors in Chapter 2, including their category, a description and the more abstract concepts attributed to them.

Category	Description	Concepts
Stable situation	Let the robot lead but remain in touch (closely watching the robot, proactively following its expected course and keeping the leash tight)	Following; Actively on top of things, actively supervising; Active observation
	Human observes what the robot does for a while	Active observation
	Directing the robot around while letting the robot do its job in small intervals	Alternating control
	Human steadily leading	Leading
	Human dragging the robot along a predefined route (that is not directed to the goal)	Taking over the other's task
	Robot is in the lead but human actively runs around robot taking into account the route	Following - Proactive following
	Robot steadily leading	Following
	Human speeds forward dragging the robot along	Dragging the other along while doing all the work yourself, the other is a burden
	Human trying to remain in touch/control by keeping the leash stretched while moving with the robot	Following; Actively on top of things, actively supervising

	Human trying to lead directly towards the goal but giving the robot space to catch up when it lags behind, causing the robot to lead for short intervals	Focus on yourself, but leave time for the other to catch up
Sudden adaptation	Let the robot do its thing until it changes direction (to something that doesn't match the expectation)	Unexpected action by a team member
	Let the robot do its thing until it has finished a task (picked up an object and stood still)	Waiting for the other to finish their task
	Take the lead once the task is seemingly finished (by letting the robot lead for most of it) to explore if there is more to gain	Trying to finish the other's task when the other is done
	Robot bumps into participant, human starts directing the robot around	Partner-hurting mistake
	Human changes direction, thereby loosening the leash, setting the robot 'free'	Losing contact with the other due to focus on own task
	Robot is in the lead until it deviates from the route expected by the human: human regains control	Unexpected action by a team member
	Human pulls the leash and moves to the goal when getting close to the goal	Being close to finishing the task
	Human picks up the robot to make up for its slow turning	Actively making up for the other's limitations
	Human pulls the leash after an object sound, then waits for the robot to come	Task achievement; Urging the other to be more active, 'come on'
	Human pulls the leash after the robot stands still for a while, then waits for the robot to come	Urging the other to be more active, 'come on'
	While moving towards the goal, the robot suddenly moves away, while the human stands still: leash is stretched due to this	Unexpected action by a team member; Team member stops with what they're doing, waits
	Human lingers around the goal when crossing it, causing the leash to be stretched as robot does move away	Task achievement; Losing contact with the other due to focus on own task
	Human pulls the robot into the direction of where an object was in previous rounds when they get close to that spot	Repeating previous behavior patterns

	Human starts with a stretched leash, but follow the robot when they notice the robot goes somewhere autonomously Robot turns to follow human quickly when human changes direction because the leash was already stretched while the human was following the robot Human pulls leash and moves towards goal after picking up an object Human lets the robot lead when it comes into movement after it was standing still	Recognizing the autonomy of team member Quick response to leadership shifts due to continuous connection Task achievement Team member becomes active after being inactive
Gradual adaptation	Slow overall change to more following over time/number of switches Loosening the leash progressively when learning to better predict the robot's behavior (being able to better follow it) Human follows robot for a bit, trying to regain control, then regains control Human randomly explores different strategies Gradually the robot leads more as the human trusts the robot more	Gradually letting the other do more Learning to predict the other's behavior Trying to regain control in different ways until eventually taking the lead Exploring different styles of leading/following Gradually letting the other do more
Active negotiation	Negotiate control by alternating pulling the leash and loosening it Alternating pulling the robot in a specific direction, waiting for the robot to go, then following the robot	Executing leading in short intervals Executing leading in short intervals

B

FLOWCHARTS OF AGENT BEHAVIOR IN CHAPTER 3




Figure B.1: A flowchart showing the rule-based decision making the agent would go through when using Macro-Action 1.




Figure B.2: A flowchart showing the rule-based decision making the agent would go through when using Macro-Action 2.

C

QUESTIONS ASKED DURING THE EXPERIMENT IN CHAPTER 3

C.1. COLLABORATION FLUENCY QUESTIONNAIRE

1. The human-robot team worked fluently together.
2. The human-robot team's fluency improved over time.
3. The robot contributed to the fluency of the interaction.

These questions were evaluated on a slider ranging from 0 for 'strongly disagree' to 100 for 'strongly agree'. The questions were taken from Hoffman, G. 2019. "Evaluating Fluency in Human-Robot Collaboration." *IEEE Transactions on Human-Machine Systems* 49 (3): 209–18. <https://doi.org/10.1109/THMS.2019.2904558>.

C.2. INTERVIEW QUESTIONS

Confidence interval:

- On a scale from 1 to 10, how confident are you that your strategy is the right strategy?

Interview questions:

- What was the strategy of the team in the end?
- Can you explain who did what in the task?
- Did the strategy change over the course of the experiment? How?

D

LIST OF OBSERVED BEHAVIORS AND INTERACTION PATTERNS IN CHAPTER 3

Table D.1: A list of all observed behaviors in Chapter 3, including their category, a description and the more abstract concepts attributed to them.

Category	Description	Concepts
Stable situation	Building a wall to protect victim	Damage control: prevent damage caused by a team member
	Pick up rocks to avoid them from falling	Damage control: prevent damage caused by a team member
	Avoiding the robot to avoid breaking	Actively adapting to the behavior of a team member and synchronizing actions
	Alternating picking up and waiting	Alternating actively working on the task and waiting for a team member
	Participant finishes task themselves	Focusing on own task
	Being careful and waiting	Being generally passive and letting a team member do most of the work
	Making active use of O2 to direct robot	Actively adapting to the behavior of a team member and synchronizing actions

Sudden adaptations

Making active use of O3 to direct robot	Actively adapting to the behavior of a team member and synchronizing actions
Leaving rocks on victim to protect them	Damage control: prevent damage caused by a team member
Waiting for robot action	Waiting for a team member to start acting
Trying to communicate by hovering	Trying to communicate by signalling task actions
Attempting to direct the robot to small rocks	Trying to communicate by signalling task actions
Trying to communicate by being around robot	Trying to communicate by interacting with a team partner
Robot does repeated pickup and dropping	Team member performs an action that makes no sense
Getting into action when robot does	Coming into action when a team member comes into action
Noticing robot strategy change to O3 due to following	Team member changes strategy, which is visible by a behavioral cue
Trying to urge the robot to empty hands	Trying to communicate by interacting with a team partner
Trying to communicate by walking around rock	Trying to communicate by signalling task actions
Clear/drop rocks in a specific place to direct robot	Trying to communicate by signalling task actions
Trying to communicate by attempting pickup	Trying to communicate by signalling task actions
Dropping rocks where robot drops	Following a team member's action
Walking around	Moving around different task components
Make mistake because unsure what to do	Doing useless or harmful actions because there is nothing else to do
Running away	Avoiding communication with a team member
Feeling alone, not helped	Feeling alone, as if team member does not help
Confused by non-human-like behavior	Being confused by non-human-like behavior
Happy that robot does as expected	Being happy that a team member does as expected
Surprised by robot autonomy	Being surprised by unexpected behavior (positive)
Confused that the robot does break not pickup	Being confused by unexpected behavior (negative)

Confused by robot interfering	Being confused by unexpected behavior (negative)
Confused by robot inaction	Being confused by unexpected behavior (negative)
Frustrated/confused/distracted by non-effective robot behavior	Being confused by unexpected behavior (negative)
Learning about directing pick up	Learning about own capabilities
Learning that the robot breaks rocks	Learning about team member's capabilities or strategy
Learning about directing breaking	Learning about own capabilities
Learning that the robot follows	Learning about behavioral cues
Learning that the robot works top to bottom	Learning about team member's capabilities or strategy

Table D.2: A list of all identified Interaction Patterns in Chapter 3 that fall in the category 'Stable situations'

Stable situations	Description
Actively synchronizing actions with a team member	Human understands the capabilities of another team member and actively uses their own actions to make optimal use of the combined capabilities.
Alternating actively working on the task and waiting for a team member	Human switches between performing their own task for a while, then waiting for a team member to perform their task, and so on.
Being generally passive and letting a team member do most of the work	Human is overall passive and lets the other team member do the work.
Damage control: prevent damage caused by a team member	Human performs actions that prevent their team member from causing intentional or unintentional harm or damage.
Focusing on own task	Human performs their own task without paying much attention to their team member.

Table D.3: A list of all identified Interaction Patterns in Chapter 3 that fall in the category 'Sudden adaptations', including the type of trigger

Sudden adaptations	Description	Type
--------------------	-------------	------

Avoiding communication with a team member	One of the team member actively avoids the other team member to avoid unwanted communication interpretations.	Internal trigger and outcome
Being confused by non-human-like behavior	A human team member is confused by non-human-like behavior performed by a team member.	External trigger and internal trigger that follows (emotion)
Being confused by unexpected behavior (negative)	One of the team members is confused or frustrated by behavior performed by their team member that they did not expect.	External trigger and internal trigger that follows (emotion)
Being happy that a team member does as expected	One of the team members is happy that their team member performs the kind of behavior that they expect and hoped for.	External trigger and internal trigger that follows (emotion)
Being surprised by unexpected behavior (positive)	One of the team members is positively surprised by behavior performed by their team member that they did not expect.	External trigger and internal trigger that follows (emotion)
Coming into action when a team member comes into action	A team member starts to actively perform their task after a period of inaction, when their team member also starts to actively perform their task after a period of inaction.	External trigger and outcome
Doing useless or harmful actions because there is nothing else to do	A team member is unable to perform useful actions, therefore starts performing useless or harmful actions.	Internal trigger and outcome
Feeling alone, as if team member does not help	A human team member feels left alone.	External trigger and internal trigger that follows (emotion)
Following a team member's action	A team member follows or copies the action performed by another team member.	External trigger and outcome
Learning about behavioral cues	A team member gains insight into specific behavior performed by another team member.	External trigger and internal trigger that follows (learning)
Learning about own capabilities	A team member gains insight into their own capabilities.	External trigger and internal trigger that follows (learning)

Learning about team member's capabilities or strategy	A team member gains insight into the capabilities or strategy of another team member.	External trigger and internal trigger that follows (learning)
Moving around different task components	A team member moves around different task components without actually performing any task.	Outcome, preceded by any trigger
Team member changes strategy, which is visible by a behavioral cue	A team member observes that another team member changes strategy by a behavioral cue.	External trigger
Team member performs an action that makes no sense	A team member performs a useless action.	External trigger and internal trigger (emotion)
Trying to communicate by interacting with a team partner	A team member attempts to communicate with another team member by directly interacting with them, for example by coming close to them.	In-between-situation, preceded by any trigger, succeeded by another sudden adaptation or stable situation
Trying to communicate by signalling task actions	A team member attempts to communicate with another team member by trying out different actions that they want their team member to perform.	In-between-situation, preceded by any trigger, succeeded by another sudden adaptation or stable situation
Waiting for a team member to start acting	A team member waits for another team member to start performing their task.	Outcome, preceded by internal trigger

E

CODING OF EMERGENT CPs IN CHAPTER 6

Table E.1: All observed Collaboration Patterns per participant per round in Chapter 6

Participant	Condition	Round	CPs
4006	C1	1	Focus on own task, Damage control, Active synchrony
		2	Focus on own task, Alternating, Active synchrony, Passive
		3	Focus on own task, Active synchrony, Alternating
		4	Active synchrony, Focus on own task
		5	Active synchrony, Focus on own task
		6	Active synchrony
		7	Active synchrony
		8	Active synchrony
4007	C1	1	Focus on own task
		2	Focus on own task, Active synchrony, Alternating
		3	Focus on own task, Active synchrony, Alternating, Passive
		4	Focus on own task, Alternating
		5	Focus on own task, Alternating
		6	Focus on own task, Alternating
		7	Focus on own task, Alternating
		8	Focus on own task, Active synchrony, Alternating
4008	C1	1	Focus on own task, Alternating

		2	Focus on own task, Active synchrony, Alternating, Passive
		3	Focus on own task, Alternating
		4	Focus on own task, Alternating, Passive
		5	Focus on own task, Alternating, Passive
		6	Focus on own task, Alternating
		7	Focus on own task, Alternating
		8	Focus on own task, Alternating
4016	C1	1	Focus on own task, Alternating, Passive, Active synchrony
		2	Focus on own task, Active synchrony, Alternating, Passive
		3	Focus on own task, Passive, Alternating
		4	Focus on own task, Active synchrony, Alternating, Passive
		5	Focus on own task, Alternating, Active synchrony
		6	Focus on own task, Alternating
		7	Focus on own task, Alternating
		8	Focus on own task, Alternating, Passive
4051	C2	1	Focus on own task, Alternating
		2	Focus on own task, Alternating
		3	Focus on own task, Alternating
		4	Focus on own task, Alternating
		5	Focus on own task, Alternating
		6	Focus on own task, Alternating, Active synchrony
		7	Focus on own task, Active synchrony, Alternating, Damage control
		8	Focus on own task, Active synchrony, Alternating
4055	C2	1	Focus on own task, Passive, Active synchrony, (jitter)
		2	Focus on own task, Passive, Active synchrony, (jitter), Damage control
		3	(jitter), Passive, Focus on own task
		4	Focus on own task, (jitter), Active synchrony
		5	Focus on own task, Damage control, (jitter), Alternating, Active synchrony
		6	Focus on own task, Passive, (jitter), Active synchrony
		7	Focus on own task, Passive, (jitter), Active synchrony
		8	Passive, (jitter), Alternating
4056	C2	1	Focus on own task, Active synchrony

		2	Focus on own task, Active synchrony
		3	Focus on own task, Active synchrony
		4	Focus on own task, Active synchrony
		5	Focus on own task, Active synchrony
		6	Focus on own task, Active synchrony
		7	Focus on own task, Active synchrony
		8	Focus on own task, Active synchrony
4057	C2	1	Active synchrony, Alternating
		2	Focus on own task, Alternating, Active synchrony
		3	Focus on own task, Active synchrony
		4	Active synchrony, Focus on own task, Alternating
		5	Focus on own task, Alternating, Active synchrony
		6	Focus on own task, Alternating, Active synchrony
		7	Focus on own task, Alternating, Active synchrony
		8	Focus on own task, Alternating
4071	C3	1	Focus on own task, Alternating, (CB) stand still break big rock
		2	(CB) move break big rock bottom, Focus on own task, (CB) stand still break big rock
		3	(CB) move break big rock bottom, (CB) remove small rocks, (CB) stand still break big rock
		4	(CB) remove small rocks, (CB) stand still break big rock
		5	(CB) move break big rock bottom, Passive, (CB) remove small rocks
		6	(CB) move break big rock bottom, (CB) remove small rocks
		7	(CB) move break big rock bottom, Passive, (CB) stand still break big rock, Focus on own task
		8	Passive, (CB) move break big rock bottom, (CB) stand still break big rock, Focus on own task
4073	C3	1	Passive, Focus on own task
		2	(CB) Small rocks on top
		3	(CB) Small rocks on top
		4	(CB) Large rock on top V2, Passive, Focus on own task
		5	Passive, (CB) Small rocks on top V2, Focus on own task
		6	(CB) Small rocks on top V2, (CB) brown rock and small rocks on op, Focus on own task

4098	C3	7	Passive, Alternating, (CB) brown rock on top
		8	(CB) brown rock on top, Alternating, (CB) large rocks on top V4, Focus on own task
		1	Focus on own task, Passive, Alternating, (CB) Handle large rocks before smash
		2	(CB) Wait for human before break, (CB) Handle large rocks before smash, Focus on own task
		3	(CB) Handle large rocks before smash, (CB) Wait for human before break, Focus on own task
		4	(CB) Move small rocks off large rocks, (CB) Wait for human before break, (CB) Unstack large rocks
		5	(CB) Move small rocks off large rocks, (CB) Wait for human before break, (CB) Unstack large rocks
		6	(CB) Move small rocks off large rocks, (CB) Unstack large rocks
4099	C3	7	(CB) Move small rocks off large rocks, (CB) Prevent brown guillotine
		8	(CB) Move small rocks off large rocks, (CB) Unstack large rocks, (CB) Wait for human before break, Focus on own task, Active synchrony
		1	Passive, Focus on own task, Active synchrony
		2	(CB) BreakBigTop, Focus on own task, (CB) PickBigTop, (CB) PickUpMultipleSmall
		3	(CB) PickUpMultipleSmall, (CB) PickBigTop
		4	(CB) PickBigTop, (CB) PickUpMultipleSmall
		5	(CB) PickUpMultipleSmall, Damage control, (CB) RemoveBigLeftSide
		6	(CB) PickUpMultipleSmall, (CB) PickBigTop, Passive, Focus on own task, (CB) PickTop
7	(CB) PickUpSmall, Focus on own task, (CB) PickUpLarge, Passive		
8	(CB) PickUpSmall, Focus on own task, Passive, (CB) GuidedPickUp		

F

DESCRIPTION OF AWARENESS OF CPs IN CHAPTER 6

Table F1: Descriptions of the awareness of emergent Collaboration Patterns for each selected participant in Chapter 6, based on the coding provided in Appendix E and their answers to open questions about human-machine collaborative strategies.

Condition	Participant	Awareness of CPs
Baseline	4006	The participant starts out by mentioning that they hope the robot would pick up large rocks. In round 3, participant mentions that they can direct the robot ('by holding my hand at a big rock, the robot will pick it up'). In later rounds, they express that they are unsure if this was really the case, although they do keep using CPs in which they direct the robot all throughout the experiment. This is again emphasized at the end: 'I thought for a time I could steer the robot towards rocks I wanted him to grab. But eventually that did not really work...'

Baseline	4007	The participant mentions trying to remove small rocks as fast as possible in the first rounds. In round 4, participant mentions that they ‘waited for the robot to break the big ones’, thereby acknowledging that they had to wait for the robot. In further rounds, the participant alludes to this in an unsure manner by saying ‘I was waiting for the robot to understand what I was doing [...] but we failed’ and ‘the robot seemed to understand what I was doing but then it seems it stopped working’. After the experiment, the participant seemed to think that the robot autonomously decided to focus on large rocks (‘It got faster and decided to focus on the big rocks’), although the CP in which the human directs the robot towards large rocks was used consistently throughout the experiment.
Baseline	4008	Throughout the experiment, the participant does not mention the robot behavior in detail. Participant does mention that they need to free up a path for the robot after the experiment (‘[I removed] the smaller blocks, so that the robot has the only option to go for the larger blocks’), saying that their understanding increased over the course of the experiment (‘[...] later it went more smoothly, when I understood better how the robot works and could anticipate [...]’). They used the same CP in which they alternate picking up rocks and waiting for the robot consistently throughout the experiment.
Baseline	4016	Participant mentions that they waited for the robot to pick up large rocks in round 1 (‘I moved the rocks, then waited for the robot to move or break the larger rocks’), but calls robot behavior random in round 2. In further rounds the participant mostly describes clearing away small rocks to avoid the robot picking those up, and to prevent mistakes by the robot. The dyad does converge to a stable CP in which the participant alternates picking up rocks and waiting for the robot in the second half of the experiment.

Basic communication	4051	The participant emphasizes that the robot does not understand them: 'if the robot could understand more what I was working on, [...] then it could perform much better', explaining how they removed small rocks to leave large rocks for the robot, although according to them this did not work: 'unfortunately the robot would also focus on smaller rocks'. They did mention that sometimes 'it felt like the robot saw what was trying to be achieved and went in the right direction'. The dyad used a CP in which the participant alternated between picking up rocks and waiting for the robot throughout the experiment, and additionally also a CP in which the participant directed the robot to certain large rocks in later rounds.
Basic communication	4055	The participant mentions several times that it looked like the robot understood what they wanted ('there was one instance where it looked like it did what I wanted to do'), but later attributed this to chance ('I believe it was by chance, but the robot picked up the right rock'). They do mention doing damage control in round 6 by placing rocks in a specific position. As this participant did not develop other consistent CPs, this fits with their behavior.
Basic communication	4056	The participant mentions waiting for the robot and signaling to the robot several times (round 1, 4 and 8; '[...] waited for the robot to start doing something', 'trying to hover over the big rocks needed to be picked up'). In the beginning they phrase it as 'trying' to direct the robot, later on they state doing it more firmly ('the signaling worked better suddenly'). This fits with the consistent use of CPs in which they directed the robot throughout the experiment.

Basic communication	4057	Participant mentions that team members need to wait for each other (round 2, 'team members must wait and analyze the other team member's actions') and pay attention to each other (round 7, 'trying to understand what the other team member is focusing on'), but they do not concretely describe why. They describe the order in which actions need to be done ('break and remove big rocks on the top first', round 8). Throughout the experiment, the dyad used CPs in which the participant alternated picking up rocks and waiting for the robot as well as CPs in which they directed the robot to certain rocks.
Collaboration Book	4071	In round 1, participant formalizes a CP that matches an emergent CP ('a way for me to order the robot to break a big rock'). In round 2, participant starts creating CPs that do not match emergent behavior, but the initial CP is kept and tweaked to work as desired. Adjustments are made in later rounds that, as the participant mentions, 'allow for more control over the actions of the robot'. After the experiment they mention that their strategy was to 'for the human to sign when what needed to be done', and that 'ideally the human signed when a big rock needed to be moved and the robot then removed that rock'. Additionally, they mention adjusting their speed: 'I had to be much slower in my removing of the rocks so that the collaboration could go more smoothly'.
Collaboration Book	4073	Participant designed simple CPs with little collaborative behavior from the start, that did not resemble earlier emergent behavior. They specifically mention formalizing the CPs 'because it seems useful'.
Collaboration Book	4098	In round 1, participant formalizes a CP that matches an emergent CP ('instruct the robot to break up large rocks'). From round 2 onwards, the participant designed additional CPs that did not match emergent behavior, but the original CP was used until the end of the experiment although improvements were made several times.

Collaboration 4099
Book

Participant started designing the robot behavior early in the experiment ('I wanted to break it instead of picking it up so I can also have some control over it'), however, this was not something that matched any emergent behavior. In round 7 and 8 more elaborate and collaborative CPs were formalized in which the human directed the robot, but again these did not originate from emergent behavior.

G

DESCRIPTION OF REUSE OF CPs IN CHAPTER 6

Table G.1: Descriptions of the reuse of emergent Collaboration Patterns for each selected participant in Chapter 6, based on the coding provided in Appendix E.

Condition	Participant	Awareness of CPs
Baseline	4006	CPs emerged in the first round, but in rounds 1 – 3 there was still quite a lot of exploration. In rounds 4 – 5 participant capitalized on the most successful CP, using only this CP for rounds 6 – 8.
Baseline	4007	In round 1 the participant only focused on themselves, but started exploring different ways of collaborating in round 2 – 3. In round 4 – 7 the used CPs were stable, while the participant deviated in round 8 after discovering a more successful CP.
Baseline	4008	A CP emerged already in the first round. In round 2 and round 5 some exploration was done, but in all other rounds the same CPs were used.
Baseline	4016	This participant explored a lot in round 1 – 4. In round 5, this started converging to a stable CP, which was used in round 6 – 8.
Basic communication	4051	Stable CPs emerged quickly, which were used consistently in round 1 – 5. New, more successful CPs emerged and were explored in round 6 – 8.
Basic communication	4055	All throughout the experiment, different CPs were used momentarily. The participant moved around a lot to try to communicate with the robot but did not stick with a strategy.

Basic communication	4056	Stable CPs emerged quickly, that were used consistently all throughout the experiment.
Basic communication	4057	Different CPs were used momentarily throughout round 1 – 4. In round 5 – 8, these CPs became more stable and were used consistently.
Collaboration Book	4071	In round 1, CPs started emerging; this was formalized and then used in round 2 – 4, although additional CPs were specified. In round 5 – 8, CPs specified in an earlier stage were used in different ways.
Collaboration Book	4073	Human and robot behaved relatively separate from the start. This was formalized in CPs outlining the robot behavior. This behavior was kept in different variations in round 2 - 5. In round 6 – 8, new variations were made, while new, more collaborative CPs started emerging outside of the formalized CPs.
Collaboration Book	4098	In round 1 a CP started emerging, which was formalized and used together with some additional formalized (but designed) CPs in round 2 – 5. In round 6 and 7 the behavior was reduced to directing the robot to pick up certain rocks via CPs, to return to the original emergent CP in round 8.
Collaboration Book	4099	A CP emerged in round 1, but it was not formalized. Throughout the rest of the rounds, different designed CPs directing the robot to pick up specific rocks were used while the human acted separately from the robot.

H

CP CLUSTERS IN CHAPTER 6

Situation clusters:

- Small rock + a location (mostly Top of rock pile, sometimes something more specific such as Bottom or On top of Victim)
- Large rock + Top of rock pile
- Large rock + specific location (e.g. Bottom, On top of something)
- Rock + On top of another rock
- Brown rock + a location
- Human + Left side of field
- Robot + Left side of rock pile
- Large rock + Top of rock pile + Human + Right side of field
- Large rock + On top of Victim
- Large rock + Robot + location

Human actions clusters:

- Stand still + location (mostly Left side of field)
- Empty
- Move back and forth + location
- Move to + object/actor/location
- Pick up + rock (mostly Small) (and sometimes + Drop + location)

Robot actions clusters:

- Pick up + Large rock + location + Drop + location (sometimes with a 'Move to' added)
- Pick up + Large rock + location
- Pick up + Small rock + location + Drop + location
- Break + Large rock + location (sometimes with a 'Move to' added)
- Brown rock + action (Pick up or Break) + Drop
- Stand still in + location

Full CP clusters:

- Brown rock present + Robot does something with it
- Large rock present + Robot picks it up
- Large rock present + Human stands still + Robot picks up/breaks rock (sometimes after moving to human)
- Small rock present + Human picks up Small rock + Robot picks up Small rock

I

RESULTS OF LINEAR MIXED MODELS IN CHAPTER 6

Table I.1: Results of the Linear Mixed Model for time in ticks for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	3101.26	117.18	419.21	26.465	< 2e-16 ***
Roundnr	-184.30	41.23	358.99	-4.470	1.05e-05 ***
C1	-165.14	167.78	419.21	-0.984	0.32556
C2	-125.35	163.83	419.21	-0.765	0.44461
B. Brown rock	-350.93	336.93	371.69	-1.042	0.29829
Round:C1	165.61	59.03	358.99	2.805	0.00530 **
Round:C2	136.26	57.64	358.99	2.364	0.01861 *
Round:B. Brown rock	184.11	63.75	368.02	2.888	0.00411 **
C1:B. Brown rock	-212.96	512.38	371.63	-0.416	0.67792
C2:B. Brown rock	141.88	477.21	372.00	0.297	0.76639
Round:C1:B. Brown rock	-95.06	94.30	366.63	-1.008	0.31411
Round:C2:B. Brown rock	-154.56	89.11	367.71	-1.734	0.08367 .

Table I.2: Results of the Linear Mixed Model for time in ticks for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	3101.26	108.08	250.41	28.695	< 2e-16 ***
Roundnr	-184.30	36.57	189.00	-5.040	1.08e-06 ***
C1	-165.14	154.74	250.41	-1.067	0.28692
C2	-125.35	151.10	250.41	-0.830	0.40754
Round:C1	165.61	52.36	189.00	3.163	0.00182 **

Round:C2	136.26	51.12	189.00	2.665	0.00836 **
-----------------	--------	-------	--------	-------	------------

Table I.3: Results of the Linear Mixed Model for time in ticks for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	2.735e+03	3.576e+02	1.357e+02	7.648	3.37e-12 ***
Roundnr	8.168e-03	5.468e+01	1.299e+02	0.000	1.000
C1	-3.487e+02	5.476e+02	1.346e+02	-0.637	0.525
C2	-1.570e+01	5.072e+02	1.363e+02	-0.031	0.975
Round:C1	6.951e+01	8.276e+01	1.266e+02	0.840	0.403
Round:C2	-1.119e+01	7.642e+01	1.296e+02	-0.146	0.884

Table I.4: Results of the Linear Mixed Model for victim harm for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	752.38	131.83	498.24	5.707	1.97e-08 ***
Roundnr	-140.95	46.06	441.00	-3.060	0.00235 **
C1	-104.88	188.75	498.24	-0.556	0.57869
C2	-204.65	184.30	498.24	-1.110	0.26736
B. Brown rock	-352.26	328.97	441.00	-1.071	0.28485
Round:C1	152.95	65.95	441.00	2.319	0.02085 *
Round:C2	147.32	64.40	441.00	2.287	0.02264 *
Round:B. Brown rock	147.79	65.15	441.00	2.269	0.02377 *
C1:B. Brown rock	837.76	471.01	441.00	1.779	0.07599 .
C2:B. Brown rock	728.62	459.91	441.00	1.584	0.11385
Round:C1:B. Brown rock	-221.79	93.27	441.00	-2.378	0.01784 *
Round:C2:B. Brown rock	-223.25	91.08	441.00	-2.451	0.01462 *

Table I.5: Results of the Linear Mixed Model for victim harm for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	752.38	137.38	250.27	5.477	1.05e-07 ***
Roundnr	-140.95	48.72	189.00	-2.893	0.00426 **
C1	-104.88	196.70	250.27	-0.533	0.59436
C2	-204.65	192.07	250.27	-1.066	0.28766
Round:C1	152.95	69.75	189.00	2.193	0.02954 *
Round:C2	147.32	68.11	189.00	2.163	0.03180 *

Table I.6: Results of the Linear Mixed Model for victim harm for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	400.124	283.783	209.629	1.410	0.1600
Roundnr	6.841	42.434	189.000	0.161	0.8721
C1	732.876	406.316	209.629	1.804	0.0727
C2	523.967	396.743	209.629	1.321	0.1881
Round:C1	-68.841	60.756	189.000	-1.133	0.2586
Round:C2	-75.932	59.325	189.000	-1.280	0.2021

Table I.7: Results of the Linear Mixed Model for remaining rocks for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	19.825381	1.865787	386.409014	10.626	<2e-16 ***
Roundnr	-0.452588	0.725277	340.317528	-0.624	0.533
C1	-2.566123	2.714238	386.524353	-0.945	0.345
C2	-2.256953	2.645624	385.808071	-0.853	0.394
B. Brown rock	3.147004	4.901651	339.997873	0.642	0.521
Round:C1	-0.734715	0.999513	338.109040	-0.735	0.463
Round:C2	-0.945848	1.007195	344.969860	-0.939	0.348
Round:B. Brown rock	-0.006667	1.006372	339.941164	-0.007	0.995
C1:B. Brown rock	-6.786384	6.820376	336.189709	-0.995	0.320
C2:B. Brown rock	-6.385168	6.860465	339.173708	-0.931	0.353
Round:C1:B. Brown rock	1.659066	1.381647	336.566549	1.201	0.231
Round:C2:B. Brown rock	1.894693	1.402778	342.160542	1.351	0.178

Table I.8: Results of the Linear Mixed Model for remaining rocks for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	20.0333	1.8588	185.4186	10.778	<2e-16 ***
Roundnr	-0.6496	0.7166	143.1334	-0.907	0.366
C1	-2.8875	2.7078	185.5375	-1.066	0.288
C2	-2.5912	2.6419	185.8511	-0.981	0.328
Round:C1	-0.4868	0.9862	140.1983	-0.494	0.622
Round:C2	-0.6481	0.9968	147.0753	-0.650	0.517

Table I.9: Results of the Linear Mixed Model for remaining rocks for Round 5 - 8.

	Estimate	Std. Error	df	t	p
--	-----------------	-------------------	-----------	----------	----------

(Intercept)	22.2260	4.0813	172.6756	5.446	1.75e-07 ***
Roundnr	-0.3435	0.6096	153.7296	-0.563	0.574
C1	-8.2381	5.6077	170.6133	-1.469	0.144
C2	-9.2127	5.6777	171.1842	-1.623	0.107
Round:C1	0.7684	0.8263	149.2728	0.930	0.354
Round:C2	1.0372	0.8495	152.2877	1.221	0.224

Table I.10: Results of the Linear Mixed Model for human idle time for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	0.922827	0.010778	264.776820	85.622	< 2e-16 ***
Roundnr	-0.009468	0.003078	440.999998	-3.076	0.00223 **
C1	-0.036367	0.015432	264.776819	-2.357	0.01917 *
C2	-0.002161	0.015068	264.776819	-0.143	0.88606
B. Brown rock	-0.013152	0.021980	441.000002	-0.598	0.54990
Round:C1	0.004088	0.004407	440.999999	0.928	0.35414
Round:C2	-0.003818	0.004303	440.999999	-0.887	0.37544
Round:B. Brown rock	0.008555	0.004353	441.000001	1.965	0.04998 *
C1:B. Brown rock	-0.042144	0.031470	441.000002	-1.339	0.18121
C2:B. Brown rock	-0.047257	0.030729	441.000002	-1.538	0.12480
Round:C1:B. Brown rock	0.001998	0.006232	441.000001	0.321	0.74870
Round:C2:B. Brown rock	0.007059	0.006085	441.000001	1.160	0.24665

Table I.11: Results of the Linear Mixed Model for human idle time for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	0.922827	0.009888	200.201865	93.331	<2e-16 ***
Roundnr	-0.009468	0.002879	189.000000	-3.289	0.0012 **
C1	-0.036367	0.014157	200.201865	-2.569	0.0109 *
C2	-0.002161	0.013823	200.201865	-0.156	0.8759
Round:C1	0.004088	0.004122	189.000000	0.992	0.3226
Round:C2	-0.003818	0.004024	189.000000	-0.949	0.3440

Table I.12: Results of the Linear Mixed Model for human idle time for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	9.097e-01	2.096e-02	2.441e+02	43.406	< 2e-16 ***
Roundnr	-9.126e-04	2.933e-03	1.890e+02	-0.311	0.75607
C1	-7.851e-02	3.001e-02	2.441e+02	-2.616	0.00944 **

C2	-4.942e-02	2.930e-02	2.441e+02	-1.687	0.09295 .
Round:C1	6.085e-03	4.200e-03	1.890e+02	1.449	0.14904
Round:C2	3.242e-03	4.101e-03	1.890e+02	0.790	0.43027

Table I.13: Results of the Linear Mixed Model for robot idle time for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	0.311777	0.048859	482.901692	6.381	4.12e-10 ***
Roundnr	0.043737	0.016728	440.999998	2.615	0.00924 **
C1	-0.102120	0.069956	482.901691	-1.460	0.14500
C2	-0.092499	0.068307	482.901691	-1.354	0.17632
B. Brown rock	0.208094	0.119461	441.000002	1.742	0.08222 .
Round:C1	-0.002798	0.023951	441.000000	-0.117	0.90705
Round:C2	-0.003158	0.023386	441.000000	-0.135	0.89264
Round:B. Brown rock	-0.038324	0.023657	441.000001	-1.620	0.10595
C1:B. Brown rock	-0.112576	0.171043	441.000001	-0.658	0.51077
C2:B. Brown rock	-0.023795	0.167013	441.000001	-0.142	0.88677
Round:C1:B. Brown rock	0.014846	0.033871	441.000001	0.438	0.66137
Round:C2:B. Brown rock	0.005529	0.033073	441.000001	0.167	0.86730

Table I.14: Results of the Linear Mixed Model for robot idle time for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	0.311777	0.042166	241.983997	7.394	2.32e-12 ***
Roundnr	0.043737	0.013722	189.000001	3.187	0.00168 **
C1	-0.102120	0.060372	241.983997	-1.692	0.09203 .
C2	-0.092499	0.058950	241.983997	-1.569	0.11793
Round:C1	-0.002798	0.019647	189.000001	-0.142	0.88690
Round:C2	-0.003158	0.019184	189.000001	-0.165	0.86942

Table I.15: Results of the Linear Mixed Model for robot idle time for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	0.519871	0.112823	218.651770	4.608	6.91e-06 ***
Roundnr	0.005413	0.016641	189.000003	0.325	0.745
C1	-0.214696	0.161538	218.651769	-1.329	0.185
C2	-0.116294	0.157732	218.651769	-0.737	0.462
Round:C1	0.012048	0.023826	189.000002	0.506	0.614
Round:C2	0.002371	0.023264	189.000002	0.102	0.919

Table I.16: Results of the Linear Mixed Model for human move actions for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	7.451e-02	9.743e-03	2.641e+02	7.648	3.81e-13 ***
Roundnr	8.928e-03	2.780e-03	4.410e+02	3.211	0.00142 **
C1	2.915e-02	1.395e-02	2.641e+02	2.090	0.03758 *
C2	-1.027e-03	1.362e-02	2.641e+02	-0.075	0.93995
B. Brown rock	1.203e-02	1.985e-02	4.410e+02	0.606	0.54495
Round:C1	-4.327e-03	3.981e-03	4.410e+02	-1.087	0.27762
Round:C2	2.325e-03	3.887e-03	4.410e+02	0.598	0.55012
Round:B. Brown rock	-7.999e-03	3.932e-03	4.410e+02	-2.034	0.04252 *
C1:B. Brown rock	2.943e-02	2.843e-02	4.410e+02	1.035	0.30118
C2:B. Brown rock	4.080e-02	2.776e-02	4.410e+02	1.470	0.14231
Round:C1:B. Brown rock	-2.361e-04	5.629e-03	4.410e+02	-0.042	0.96656
Round:C2:B. Brown rock	-5.870e-03	5.497e-03	4.410e+02	-1.068	0.28612

Table I.17: Results of the Linear Mixed Model for human move actions for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	0.074515	0.009156	200.849970	8.138	4.13e-14 ***
Roundnr	0.008928	0.002670	189.000000	3.344	0.000997 ***
C1	0.029154	0.013109	200.849970	2.224	0.027267 *
C2	-0.001027	0.012800	200.849970	-0.080	0.936126
Round:C1	-0.004327	0.003823	189.000000	-1.132	0.259099
Round:C2	0.002325	0.003733	189.000000	0.623	0.534203

Table I.18: Results of the Linear Mixed Model for human move actions for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	8.654e-02	1.876e-02	2.433e+02	4.614	6.38e-06 ***
Roundnr	9.291e-04	2.631e-03	1.890e+02	0.353	0.7244
C1	5.858e-02	2.685e-02	2.433e+02	2.181	0.0301 *
C2	3.977e-02	2.622e-02	2.433e+02	1.517	0.1306
Round:C1	-4.563e-03	3.768e-03	1.890e+02	-1.211	0.2273
Round:C2	-3.546e-03	3.679e-03	1.890e+02	-0.964	0.3363

Table I.19: Results of the Linear Mixed Model for robot move actions for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	0.53847	0.03547	490.25915	15.183	< 2e-16 ***
Roundnr	-0.05713	0.01225	441.00000	-4.665	4.1e-06 ***
C1	0.05877	0.05078	490.25915	1.157	0.24771
C2	0.07107	0.04958	490.25915	1.433	0.15241
B. Brown rock	-0.23109	0.08745	441.00000	-2.642	0.00852 **
Round:C1	0.01991	0.01753	441.00000	1.136	0.25665
Round:C2	0.01388	0.01712	441.00000	0.811	0.41790
Round:B. Brown rock	0.05497	0.01732	441.00000	3.174	0.00161 **
C1:B. Brown rock	0.06862	0.12521	441.00000	0.548	0.58395
C2:B. Brown rock	0.06543	0.12226	441.00000	0.535	0.59280
Round:C1:B. Brown rock	-0.02208	0.02480	441.00000	-0.891	0.37363
Round:C2:B. Brown rock	-0.02249	0.02421	441.00000	-0.929	0.35343

Table I.20: Results of the Linear Mixed Model for robot move actions for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	0.53847	0.03086	249.61693	17.447	< 2e-16 ***
Roundnr	-0.05713	0.01039	189.00000	-5.500	1.22e-07 ***
C1	0.05877	0.04419	249.61693	1.330	0.185
C2	0.07107	0.04315	249.61693	1.647	0.101
Round:C1	0.01991	0.01487	189.00000	1.339	0.182
Round:C2	0.01388	0.01452	189.00000	0.956	0.340

Table I.21: Results of the Linear Mixed Model for robot move actions for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	0.307378	0.083316	216.204286	3.689	0.000285 ***
Roundnr	-0.002158	0.012336	188.999997	-0.175	0.861316
C1	0.127390	0.119290	216.204287	1.068	0.286758
C2	0.136501	0.116480	216.204287	1.172	0.242533
Round:C1	-0.002168	0.017662	188.999999	-0.123	0.902414
Round:C2	-0.008609	0.017246	188.999999	-0.499	0.618209

Table I.22: Results of the Linear Mixed Model for human productive actions for Round 1 - 8.

	Estimate	Std. Error	df	t	p
--	-----------------	-------------------	-----------	----------	----------

(Intercept)	8.858e-03	3.447e-03	2.387e+02	2.570	0.0108 *
Roundnr	1.376e-03	9.538e-04	4.410e+02	1.443	0.1497
C1	7.815e-03	4.936e-03	2.387e+02	1.583	0.1146
C2	3.725e-03	4.819e-03	2.387e+02	0.773	0.4404
B. Brown rock	5.155e-03	6.811e-03	4.410e+02	0.757	0.4495
Round:C1	-4.858e-04	1.366e-03	4.410e+02	-0.356	0.7222
Round:C2	8.027e-04	1.333e-03	4.410e+02	0.602	0.5475
Round:B. Brown rock	-1.644e-03	1.349e-03	4.410e+02	-1.219	0.2236
C1:B. Brown rock	1.126e-02	9.752e-03	4.410e+02	1.155	0.2487
C2:B. Brown rock	4.805e-03	9.522e-03	4.410e+02	0.505	0.6141
Round:C1:B. Brown rock	-9.724e-04	1.931e-03	4.410e+02	-0.503	0.6149
Round:C2:B. Brown rock	-3.249e-04	1.886e-03	4.410e+02	-0.172	0.8633

Table I.23: Results of the Linear Mixed Model for human productive actions for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	8.858e-03	2.383e-03	1.833e+02	3.717	0.000268 ***
Roundnr	1.376e-03	6.639e-04	1.890e+02	2.073	0.039502 *
C1	7.815e-03	3.412e-03	1.833e+02	2.290	0.023136 *
C2	3.725e-03	3.332e-03	1.833e+02	1.118	0.265079
Round:C1	-4.858e-04	9.505e-04	1.890e+02	-0.511	0.609917
Round:C2	8.027e-04	9.281e-04	1.890e+02	0.865	0.388228

Table I.24: Results of the Linear Mixed Model for human productive actions for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	1.401e-02	6.643e-03	2.516e+02	2.109	0.0359 *
Roundnr	-2.674e-04	8.753e-04	1.890e+02	-0.305	0.7604
C1	1.908e-02	9.512e-03	2.516e+02	2.006	0.0459 *
C2	8.530e-03	9.288e-03	2.516e+02	0.918	0.3593
Round:C1	-1.458e-03	1.253e-03	1.890e+02	-1.164	0.2461
Round:C2	4.778e-04	1.224e-03	1.890e+02	0.390	0.6966

Table I.25: Results of the Linear Mixed Model for robot productive actions for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	5.842e-02	8.114e-03	4.403e+02	7.200	2.62e-12 ***
Roundnr	-2.811e-03	2.678e-03	4.410e+02	-1.050	0.2945
C1	-8.380e-03	1.162e-02	4.403e+02	-0.721	0.4711

C2	-7.052e-03	1.134e-02	4.403e+02	-0.622	0.5344
B. Brown rock	1.844e-03	1.913e-02	4.410e+02	0.096	0.9232
Round:C1	7.433e-03	3.835e-03	4.410e+02	1.938	0.0532 .
Round:C2	6.358e-03	3.745e-03	4.410e+02	1.698	0.0902 .
Round:B. Brown rock	8.557e-04	3.788e-03	4.410e+02	0.226	0.8214
C1:B. Brown rock	3.653e-02	2.739e-02	4.410e+02	1.334	0.1829
C2:B. Brown rock	-5.959e-03	2.674e-02	4.410e+02	-0.223	0.8238
Round:C1:B. Brown rock	-8.219e-03	5.423e-03	4.410e+02	-1.516	0.1303
Round:C2:B. Brown rock	-2.384e-03	5.296e-03	4.410e+02	-0.450	0.6528

Table I.26: Results of the Linear Mixed Model for robot productive actions for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	0.058421	0.007458	229.091683	7.833	1.77e-13 ***
Roundnr	-0.002811	0.002336	189.000000	-1.203	0.2303
C1	-0.008380	0.010678	229.091683	-0.785	0.4334
C2	-0.007053	0.010427	229.091683	-0.676	0.4995
Round:C1	0.007433	0.003344	189.000000	2.223	0.0274 *
Round:C2	0.006359	0.003266	189.000000	1.947	0.0530 .

Table I.27: Results of the Linear Mixed Model for robot productive actions for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	6.027e-02	1.879e-02	2.148e+02	3.207	0.00154 **
Roundnr	-1.956e-03	2.788e-03	1.890e+02	-0.701	0.48392
C1	2.815e-02	2.690e-02	2.148e+02	1.046	0.29658
C2	-1.301e-02	2.627e-02	2.148e+02	-0.495	0.62089
Round:C1	-7.863e-04	3.992e-03	1.890e+02	-0.197	0.84405
Round:C2	3.974e-03	3.898e-03	1.890e+02	1.020	0.30921

Table I.28: Results of the Linear Mixed Model for subjective collaboration fluency for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	29.167	5.804	379.548	5.025	7.76e-07 ***
Roundnr	6.915	1.833	441.000	3.773	0.000183 ***
C1	-6.617	8.310	379.548	-0.796	0.426405
C2	-12.148	8.114	379.548	-1.497	0.135184
B. Brown rock	17.715	13.090	441.000	1.353	0.176624
Round:C1	-3.721	2.624	441.000	-1.418	0.156901
Round:C2	-2.369	2.562	441.000	-0.924	0.355761

Round:B. Brown rock	-7.262	2.592	441.000	-2.802	0.005310 **
C1:B. Brown rock	-13.759	18.741	441.000	-0.734	0.463238
C2:B. Brown rock	-10.358	18.300	441.000	-0.566	0.571683
Round:C1:B. Brown rock	4.849	3.711	441.000	1.307	0.192060
Round:C2:B. Brown rock	4.347	3.624	441.000	1.199	0.231006

Table I.29: Results of the Linear Mixed Model for subjective collaboration fluency for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	29.167	5.521	232.514	5.282	2.93e-07 ***
Roundnr	6.915	1.746	189.000	3.962	0.000105 ***
C1	-6.617	7.905	232.514	-0.837	0.403467
C2	-12.148	7.719	232.514	-1.574	0.116894
Round:C1	-3.721	2.499	189.000	-1.489	0.138165
Round:C2	-2.369	2.440	189.000	-0.971	0.332931

Table I.30: Results of the Linear Mixed Model for subjective collaboration fluency for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	46.8819	12.4338	224.7671	3.771	0.000208 ***
Roundnr	-0.3467	1.8155	189.0000	-0.191	0.848768
C1	-20.3759	17.8025	224.7671	-1.145	0.253613
C2	-22.5061	17.3831	224.7671	-1.295	0.196748
Round:C1	1.1277	2.5993	189.0000	0.434	0.664909
Round:C2	1.9777	2.5381	189.0000	0.779	0.436829

Table I.31: Results of the Linear Mixed Model for self-efficacy for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	57.4762	6.7347	190.4312	8.534	4.38e-15 ***
Roundnr	1.2048	1.7242	441.0000	0.699	0.4851
C1	-12.9012	9.6426	190.4312	-1.338	0.1825
C2	-22.7489	9.4154	190.4312	-2.416	0.0166 *
B. Brown rock	17.9714	12.3131	441.0000	1.460	0.1451
Round:C1	-2.7898	2.4687	441.0000	-1.130	0.2591
Round:C2	-0.8457	2.4105	441.0000	-0.351	0.7259
Round:B. Brown rock	-4.4714	2.4384	441.0000	-1.834	0.0674 .
C1:B. Brown rock	-6.3764	17.6297	441.0000	-0.362	0.7178
C2:B. Brown rock	-25.0442	17.2144	441.0000	-1.455	0.1464
Round:C1:B. Brown rock	2.9014	3.4912	441.0000	0.831	0.4064

Round:C2:B. Brown rock 4.5851 3.4090 441.0000 1.345 0.1793

Table I.32: Results of the Linear Mixed Model for self-efficacy for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	57.4762	6.9793	142.9934	8.235	1.03e-13 ***
Roundnr	1.2048	1.7026	189.0000	0.708	0.4801
C1	-12.9012	9.9928	142.9934	-1.291	0.1988
C2	-22.7489	9.7574	142.9934	-2.331	0.0211 *
Round:C1	-2.7898	2.4377	189.0000	-1.144	0.2539
Round:C2	-0.8457	2.3803	189.0000	-0.355	0.7228

Table I.33: Results of the Linear Mixed Model for self-efficacy for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	75.4476	11.2370	248.8674	6.714	1.27e-10 ***
Roundnr	-3.2667	1.5449	189.0000	-2.114	0.0358 *
C1	-19.2776	16.0890	248.8674	-1.198	0.2320
C2	-47.7931	15.7099	248.8674	-3.042	0.0026 **
Round:C1	0.1117	2.2119	189.0000	0.050	0.9598
Round:C2	3.7394	2.1598	189.0000	1.731	0.0850 .

Table I.34: Results of the Linear Mixed Model for subjective performance for Round 1 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	20.6429	6.2489	427.3682	3.303	0.00104 **
Roundnr	8.5238	2.0431	441.0000	4.172	3.64e-05 ***
C1	-7.4179	8.9471	427.3682	-0.829	0.40752
C2	-7.5519	8.7363	427.3682	-0.864	0.38783
B. Brown rock	11.6143	14.5909	441.0000	0.796	0.42646
Round:C1	-3.2988	2.9253	441.0000	-1.128	0.26007
Round:C2	-3.0056	2.8564	441.0000	-1.052	0.29327
Round:B. Brown rock	-7.5286	2.8894	441.0000	-2.606	0.00948 **
C1:B. Brown rock	-0.1493	20.8911	441.0000	-0.007	0.99430
C2:B. Brown rock	8.8948	20.3989	441.0000	0.436	0.66302
Round:C1:B. Brown rock	2.9186	4.1370	441.0000	0.705	0.48089
Round:C2:B. Brown rock	1.4513	4.0396	441.0000	0.359	0.71957

Table I.35: Results of the Linear Mixed Model for subjective performance for Round 1 - 4.

	Estimate	Std. Error	df	t	p
(Intercept)	20.643	5.916	237.222	3.490	0.000577 ***
Roundnr	8.524	1.896	189.000	4.496	1.21e-05 ***
C1	-7.418	8.470	237.222	-0.876	0.382020
C2	-7.552	8.270	237.222	-0.913	0.362089
Round:C1	-3.299	2.714	189.000	-1.215	0.225764
Round:C2	-3.006	2.650	189.000	-1.134	0.258224

Table I.36: Results of the Linear Mixed Model for subjective performance for Round 5 - 8.

	Estimate	Std. Error	df	t	p
(Intercept)	32.2571	13.5601	221.3627	2.379	0.0182 *
Roundnr	0.9952	1.9913	189.0000	0.500	0.6178
C1	-7.5671	19.4151	221.3627	-0.390	0.6971
C2	1.3429	18.9577	221.3627	0.071	0.9436
Round:C1	-0.3802	2.8511	189.0000	-0.133	0.8940
Round:C2	-1.5543	2.7839	189.0000	-0.558	0.5773

J

DISCOVERED CPs IN CHAPTER 6

Participant 4071 used the following CPs in the task:

1. Stand still break big rock: if there is a large rock at the top of the rock pile, the human stands still on the side of the field and the robot breaks a large rock at the top of the rock pile.
2. Move break big rock bottom: if there is a large rock at the bottom of the rock pile, the human moves back and forth in that location and the robot breaks a large rock at the bottom of the rock pile.
3. Remove small rocks: if there is a small rock at the top of the rock pile, both human and robot pick up small rocks and drop them.

Participant 4073 used the following CPs in the task:

1. Small rocks on top: if there is a small rock at the top of the rock pile, human picks up small rocks and drops them while the robot stands still.
2. Large rock on top V2: if there is a large rock at the top of the rock pile, human stands still on the side of the field and robot picks up a large rock at the top of the rock pile.
3. Small rocks on top V2: if there is a small rock at the top of the rock pile, human picks up small rocks and drops them while the robot stands still.
4. Brown rock on top: if there is a brown rock at the top of the rock pile, human stands still on the side of the field and robot picks up a brown rock at the top of the rock pile.
5. Large rocks on top V4: if there is a large rock at the top of the rock pile, human stands still on the side of the field and robot breaks a large rock at the top of the rock pile.

Participant 4098 used the following CPs in the task:

1. Wait for human before break: if there is a large rock at the right side of the rock pile, the human moves to the right side of the rock pile, the robot moves to the human and breaks a large rock at the right side of the rock pile.
2. Handle large rocks before smash: if there is a large rock at the top of the rock pile, the robot breaks a large rock at the top of the rock pile, moves to a small rock, then picks up a small rock at the top of the rock pile.
3. Move small rocks off large rocks: if there is a small rock on top of a large rock, the robot picks up a small rock on top of a large rock.
4. Unstack large rocks: if there is a large rock on top of a large rock, the robot moves to a large rock and picks up a large rock on top of a large rock. Then the robot drops the large rock on the left side of the field.
5. Prevent brown guillotine: if there is a brown rock on top of a large rock, the robot breaks the brown rock on top of the large rock.

Participant 4099 used the following CPs in the task:

1. BreakBigTop: if there is a large rock on top of the rock pile, the robot breaks a large rock on top of the rock pile.
2. PickBigTop: if there is a large rock on top of the rock pile, the robot picks up a large rock on top of the rock pile.
3. PickUpMultipleSmall: if there is a small rock on top of the rock pile, the robot picks up a small rock on top of the rock pile; this action is repeated three times.
4. RemoveBigLeftSide: if there is a brown rock on the left side, the robot picks up a large rock on the left side.
5. PickUpLarge: if there is a large rock on top of the rock pile, the robot picks up a large rock on top of the rock pile.
6. GuidedPickup: if there is a brown rock on top of the rock pile, the human stands still on top of a large rock, and the robot picks up a large rock.

K

QUESTIONNAIRE USED IN CHAPTER 6

Questions asked about Locus of Control before the experiment:

- I'm my own boss.
- If I work hard, I will succeed.
- Whether at work or in my private life: What I do is mainly determined by others.
- Fate often gets in the way of my plans.

Questions about Collaboration Fluency, Self-efficacy, Subjective performance and Collaboration Patterns after each round of the experiment:

- I believe we can succeed at this task if I set my mind to it.
- We, the human-robot team, worked fluently together.
- Our human-robot team's fluency improved over time.
- The robot contributed to the fluency of the interaction.
- Our human-robot team's performance improved over time.
- The robot's performance improved over time.
- C1 and C2: Did you have a strategy for solving the team task? If so, please describe the strategy. C3: For each stored Collaboration Pattern, please explain why this sequence of collaborative actions is useful for the team.
- How would you rate the performance of the team?

Questions asked after the experiment:

- What was the strategy of the team in the end?
- Can you explain who did what in the task?
- Did the strategy change over the course of the experiment? How?
- How fluently did you and the robot collaborate? Did this change over the course of the experiment?

ACKNOWLEDGEMENTS

The PhD journey is one that I enjoyed and that I felt fit me like a glove, although it came with some unexpected challenges, and a lot of life happened in between. I always tried to see the PhD as a job first and foremost, but ironically, being at home for different reasons (lockdowns, maternity leave) made me realize how important that job was for me. From the mundane things like having lunch with colleagues to the somewhat larger things like developing a vision on a topic and brainstorming about new research directions: they are part of what makes me feel alive, capable and empowered. I'm still as convinced as I was at the beginning of this journey that research is what I should be doing. In the following sections I will switch between Dutch and English to thank the different people that were (and for some still are) part of my journey through research and academia.

Ik wil graag beginnen met mijn begeleiders, omdat dit proefschrift er zonder jullie niet geweest was. Mark en Karel, vanaf het eerste gesprek over dit promotietraject was het duidelijk dat wij met elkaar op één lijn zaten over de invulling. Jullie hadden 'toevallig' een vacature geschreven waar mijn ideeën perfect in pasten, dat vond ik, maar volgens mij vonden jullie dat ook. Karel, jij bent wat mij betreft vanaf dag één een dagelijks begeleider geweest zoals het hoort. Ik heb het ontzettend gewaardeerd dat je altijd oog had voor mij als persoon en voor hoe ik me voelde, wat zeker in die eerste periode waarin we elkaar alleen online spraken heel erg fijn was. Dankjewel voor het minutieus becommentariëren van ieder woord wat ik in dit proefschrift heb geschreven, voor alle verhalen over je wandelvakanties, voor je statistiek-inbreng en voor het accepteren van mijn emoties op kwetsbare momenten. Mark, ik zeg altijd tegen iedereen dat ik mazzel had dat ik jou zo vaak en zo veel sprak, die luxe hebben niet veel promovendi. Regelmatig gaf je me met een paar woorden precies de hint die ik nodig had om zelf verder te kunnen denken en het theoretisch kader van mijn werk sterker te maken. Dankjewel voor het altijd toevoegen van waardevolle referenties, voor abstract commentaar wat me aan het denken zette en voor de gesprekken over een academische carrière bij TNO. David, hoewel je later bent aangesloten bij dit traject, voelde het voor mij vanaf het eerste moment als een natuurlijke en nuttige samenwerking. Juist door jouw frisse blik en net andere invalshoek is mijn werk zowel breder als grondiger geworden. Hoe jij verschillende mensen bij elkaar aan tafel krijgt is voor mij een voorbeeld. Dankjewel voor je enthousiasme over kwalitatief onderzoek en ontwerpers, voor het laten zien dat je een succesvol academicus kunt zijn in een vierdaagse werkweek en voor het doorgeven van mijn telefoonnummer aan de ambtenaar van SZW. Bedankt dat jullie vertrouwen hadden in mijn visie en ideeën, en dat jullie me altijd het gevoel hebben gegeven dat ik de ruimte had om mijn PhD in te richten op mijn manier. Dat bedoel ik zowel inhoudelijk als in de combinatie met andere projecten en mijn persoonlijke leven.

I would like to continue with the members of my doctoral committee, Joost de Winter, Annette ten Teije, Alessandro Bozzon and Frank Flemisch. Thank you for taking the time to read and assess my work. The four of you make up a great combination of dis-

ciplines, something I really wanted for my committee. I would like to especially thank Joost and Frank for this full-circle moment, thinking back to the conversations we had in the summer of 2019.

Before starting my PhD journey, there were several people who helped me discover that I wanted to go into research, and who supported me in developing a vision on AI, Design, and Human-Centered Technology. Emilia Barakova, thank you for believing in me early on and for teaching me the first few things about AI. Frank Delbressine, thank you for coaching me towards the right choice of Master programs. Matthias Rauterberg, thank you for always pushing me further. Chris Janssen, Mehdi Dastani, Frank Dignum and the other teachers: thank you for admitting me into the AI Master program and for making this designer an AI scientist. Jurriaan, Marieke en Anita, bedankt dat jullie me leerden dat mens-machine teaming een bestaand onderzoeksonderwerp was, en dat TNO zulk tof onderzoek doet!

Over TNO gesproken: het onderzoek wat TNO doet, vooral in Soesterberg, sluit zo goed aan bij mijn onderzoeksvisie en interesses dat ik er niet weg lijkt te komen. Voorlopig blijf ik zeker! Tijdens mijn promotieonderzoek heb ik meegewerkt aan een aantal projecten, dit zorgde ervoor dat mijn werk afwisselend bleef. Tjeerd, Romy, Birgit, Anthia en Ivana, bedankt voor de samenwerking op het onderwerp co-learning, de gezelligheid en de mooie publicaties die daaruit voort zijn gekomen. Ook bedankt aan alle FATE collega's, die proefkonijnen waren voor de methode die ik samen met Tjeerd heb uitgewerkt om AI-onderzoek meer human-centered te maken. Chris en Willeke, bedankt dat jullie ook wilden dat ik bleef. Alle andere collega's, bedankt voor de fijne start als werknemer bij TNO ondanks dat ik mijn proefschrift nog af moest maken.

To move on to my other research colleagues: the Interactive Intelligence section has felt like the perfect academic home for me from the start. A big thank you to all the staff for being such a great combination of disciplines and for being good lunch company. Thank you to Bernd, Ilir, Elena, Miguel, and Eli for welcoming me into the group as the first new PhD student in a long time when I just started out. Thank you to Enrico, Sid, Carolina, Nele, Ruben, Pei-Yu, Mani, and Masha, for becoming the group with which we first returned to the office, for the many great conversations during lunch, tea, or at random moments during the working day. Thank you also for our very serious Overcooked study evening and for the other fun activities we did over the years. Thank you to all the other PhDs and Postdocs, all of you are the reason I feel sad about officially leaving the group now.

I would also like to thank the people within the HRI group at ME who were eager to collaborate with me. Niek, thank you for your enthusiasm, and Luka, for putting up with my two maternity leaves and still wanting to supervise another student and write another paper with me. Talking about those students: Hugo, thank you very much for being the first Master student that I supervised. I couldn't have asked for a better first student and I thoroughly enjoyed how you took the complex topic that we gave you and made it your own. Jesse, thank you for extending our work with your ideas and for making me feel like human-machine co-learning was a well-enough established topic for a myriad of thesis projects.

An enormous amount of gratitude goes out to all the people who participated in my weird and frustrating (pilot) experiments. I hope that, in the end, it was kind of fun to

participate even though I confused many of you with unexplainable robot behavior. The fact that you made it to the end of an experiment should be seen as an achievement!

Now on to my paranymphs. I have known for a while that I wanted to ask both of you to be a paranymph for me, as both of you have been a very important part of my PhD journey, so I'm very happy that you said yes. Carolina, I think we were among the few people who loyally joined the online meetings, and for me it quickly felt like we clicked. I know that I have sometimes been absent, maybe due to my Dutchness or life happening, but throughout working together during lockdowns in Haarlem, visiting the HRI event in Barcelona and the many, many conversations in the office, I definitely consider you a friend. Thank you for being my paranymph. Marina, wat moet ik nou nog tegen jou zeggen, ik beschouw jou al zo lang als mijn beste vriendin. Ik vind het heel mooi hoe wij allebei het onderzoek in zijn gegaan, misschien was dat wel voorbestemd na al die filosofische gesprekken die wij altijd hadden op de fiets. Dankjewel dat je er altijd voor me bent en bent geweest, en dat je nu onderdeel bent van mijn promotie.

Lieve Wreedbeesten, dankjulliewel dat ik altijd bij jullie terecht kan voor fijne weekenden en avonden, ook al zien we elkaar niet altijd even vaak. Iris, ik weet dat ik altijd op je kan rekenen, dankjewel daarvoor. Paula, dankjewel dat je bent uitgegroeid tot een onmisbaar onderdeel van ons groepje. Jullie maken dat er nóg meer leven is naast een PhD en een gezin.

Pap en mam, ik weet dat jullie altijd trots op me zijn, als ik maar mijn best doe. Dankjulliewel voor het hebben van precies de verwachtingen waaraan ik kan voldoen, en voor het ondersteunen van mijn reis door de universiteit ondanks dat jullie er zelf geen ervaring mee hebben.

Joris, jij helpt me om het beste uit mezelf te halen, en op mijn beurt laat ik jou zien dat het soms goed genoeg is. Dankjewel dat je me overeind heb gehouden tijdens maanden van thuiszitten, en dankjewel voor hoe trots je altijd over me praat. Jij helpt me om beter voor mezelf op te komen, juist als ik zelf doe alsof het niet nodig is.

Len en Kai, jullie zijn mijn lieve moppies. Terwijl ik dit schrijf met een draagdoek om weet ik dat ik regelmatig gefrustreerd ben geweest omdat ik niet alles tegelijk kon, maar het was het meer dan waard.

CURRICULUM VITÆ

Emma Maxime VAN ZOELEN

12-07-1995 Born in Wester-Koggenland, The Netherlands

EDUCATION

- 2020–2025 **Ph.D. in Computer Science**
Delft University of Technology, Delft, The Netherlands
- 2016–2020 **Master of Science in Industrial Design**
Eindhoven University of Technology, Eindhoven, The Netherlands
- 2016–2019 **Master of Science in Artificial Intelligence**
Utrecht University, Utrecht, The Netherlands
- 2013–2016 **Bachelor of Science in Industrial Design**
Eindhoven University of Technology, Eindhoven, The Netherlands

EXPERIENCE

- 2024–... **TNO**, Soesterberg, The Netherlands
Scientist Human-Machine Teaming
- 2022–2024 **Staatscommissie Demographic Developments 2050**,
The Hague, The Netherlands
Member
- 2018–2019 **TNO**, Soesterberg, The Netherlands
Graduate intern
- 2015 **STBY**, London, United Kingdom
Design research intern

LIST OF PUBLICATIONS

2025

1. Karel van den Bosch, **Emma M. van Zoelen**, Tjeerd Schoonderwoerd, Anthia A. Solati, Birgit van der Stigchel, Ivana Akrum. 2025. Design and Effects of Co-Learning in Human-AI Teams. In *Journal of Artificial Intelligence Research*, 82, 1445-1493.

2024

1. **Emma M. van Zoelen**, Hugo Veldman-Loopik, Karel van den Bosch, Mark Neerincx, David A. Abbink, Luka Peternel. 2024. Enabling Embodied Human-Robot Co-Learning: Requirements, Method, and Test With Handover Task. In *IEEE Robotics and Automation Letters*, 1425 - 1432.
2. Carolina C. Jorge, **Emma M. van Zoelen**, Ruben Verhagen, Siddharth Mehrotra, Catholijn M. Jonker, Myrthe L. Tielman. 2024. Appropriate context-dependent artificial trust in human-machine teamwork. In *Putting AI in the Critical Loop*, 41-60. Academic Press.

2023

1. **Emma M. van Zoelen**, Tina Mioch, Mani Tajaddini, Christian Fleiner, Stefani Tsaneva, Pietro Camin, Thiago S. Gouvea, Kim Baraka, Maaïke H. de Boer, Mark A. Neerincx. 2023. Developing team design patterns for hybrid intelligence systems. In *HAI 2023: Augmenting Human Intellect*, 3-16. IOS Press.

2022

1. **Emma M. van Zoelen**, Karel van den Bosch, David A. Abbink, Mark A. Neerincx. 2022. Ontology-Based Reflective Communication for Shared Human-AI Recognition of Emergent Collaboration Patterns. In *International Conference on Principles and Practice of Multi-Agent Systems*, 621-629. Springer International Publishing.
2. Tjeerd A. Schoonderwoerd, **Emma M. van Zoelen**, Karel van den Bosch, Mark A. Neerincx. 2022. Design patterns for human-AI co-learning: A wizard-of-Oz evaluation in an urban-search-and-rescue task. In *International Journal of Human-Computer Studies*, 164, 102831.
3. Maaïke H. de Boer, Steven Vethman, Roos M. Bakker, Ajaya Adhikari, Michiel Marcus, Joachim de Greeff, Jasper van der Waa, Tjeerd A. Schoonderwoerd, Ioannis Tolios, **Emma M. van Zoelen**, Fieke Hillerström, Bart Kamphorst. 2022. The FATE System Iterated: Fair, Transparent and Explainable Decision Making in a Juridical

Case. In *Proceedings of the AAAI 2022 Spring Symposium on Machine Learning and Knowledge Engineering for Hybrid Intelligence (AAAI-MAKE)*.

2021

- 1. **Emma M. van Zoelen**, Karel van den Bosch, Matthias Rauterberg, Emilia I. Barakova, Mark Neerinx. 2021. Identifying interaction patterns of tangible co-adaptations in human-robot team behaviors. In *Frontiers in Psychology*, 12, 645545.
- 2. **Emma M. van Zoelen**, Karel van Den Bosch, Mark Neerinx. 2021. Becoming team members: Identifying interaction patterns of mutual adaptation for human-robot co-learning. In *Frontiers in Robotics and AI*, 8, 692811.
- 3. **Emma M. van Zoelen**, Karel van Den Bosch, Mark Neerinx. 2021. Human-robot co-learning for fluent collaborations. In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 574-576.

2020

- 1. **Emma M. van Zoelen**, Emilia I. Barakova, Matthias Rauterberg. 2020. Adaptive leader-follower behavior in human-robot collaboration. In **2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)**, 1259-1265. IEEE.
- 2. Tessa Aarts, Linas K. Gabrielaitis, Lianne C. de Jong, Renee Noortman, **Emma M. van Zoelen**, Sophia Kotea, Silvia Cazacu, Lesley L. Lock, Panos Markopoulos. 2020. Design card sets: Systematic literature survey and card sorting study. In *Proceedings of the 2020 ACM designing interactive systems conference*, 419-428.
- 3. **Emma M. van Zoelen**, Anita Cremers, Frank P. Dignum, Jurriaan van Diggelen, Marieke M. Peeters. 2020. Learning to communicate proactively in human-agent teaming. In *Highlights in Practical Applications of Agents, Multi-Agent Systems, and Trust-worthiness. The PAAMS Collection: International Workshops of PAAMS 2020, L'Aquila, Italy, October 7-9, 2020, Proceedings 18*, 238-249. Springer International Publishing.

2019

- 1. **Emma M. van Zoelen**, Laure Peeters, Sander J. Bos, Feng Ye,. 2019. Shared control and the democratization of driving in autonomous vehicles. In *Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications: Adjunct Proceedings*, 120-124.

 Included in this thesis

LIST OF RELEVANT DATASETS

1. **Emma M. van Zoelen**, Karel van den Bosch, Mark Neerincx, David A. Abbink. Evaluating Shared Representations of Emergent Collaboration Patterns in Human-Machine Co-Learning, data underlying PhD thesis chapter: Shared Representations of Emergent Collaboration Patterns.
<https://doi.org/10.4121/af52cc1f-4fb6-481d-81c1-1058e22f713d>, to be published.
2. **Emma M. van Zoelen**, Karel van den Bosch, Mark Neerincx, David A. Abbink. Identification and Formalization of Human-Machine Collaboration Patterns, data underlying the publication: Ontology-based Reflective Communication for Shared Human-AI Recognition of Emergent Collaboration Patterns.
<https://doi.org/10.4121/2d0f80df-1ae8-407e-bc5f-f2e570bbb306>, 2025.
3. **Emma M. van Zoelen**, Karel van den Bosch, Matthias Rauterberg, Emilia Barakova, Mark Neerincx. Leader-Follower Behavior in Human-Robot Co-Adaptation, data underlying the publication: Identifying Interaction Patterns of Tangible Co-Adaptations in Human-Robot Team Behaviors.
<https://doi.org/10.4121/8fbb4c07-4ba2-4892-8753-2c909e0a4609>, 2025.
4. **Emma M. van Zoelen**, Karel van den Bosch, Mark Neerincx. Mutual Adaptation for Human-Robot Co-Learning - USAR task, data underlying the publication: Becoming Team Members.
<https://doi.org/10.17632/r2y8z6bzbz8.1>, 2021.



Emma believes that in order for humans and artificially intelligent agents to live together symbiotically, they need to collaborate as team partners. Her PhD research centers around the patterns of collaborative behavior that emerge when both human and machine adapt and learn continuously, and how these patterns can be used to improve mutual understanding and team performance.