Negotiating with Robots

Meshing Plans and Resolving Conflicts in Human-Robot Collaboration

by

AJung Moon

B.A.Sc., Mechatronics Engineering, The University of Waterloo, 2009

M.A.Sc., Mechanical Engineering, The University of British Columbia, 2012

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(Mechanical Engineering)

The University of British Columbia
(Vancouver)

June 2017

Abstract

For both humans and robots, one of the key elements of collaboration is the collaborating agents' ability to communicate and mesh their individual plans with each other. Even after the collaborators have decided on a joint task, the details of the task such as the *how*, *when*, and *where* are often determined as the collaborative activity unfolds. The primary objective of this thesis is to enable fluent communication between human and robotic agents such that they can interactively figure out unspoken details and resolve unforeseen conflicts that arise during a human-robot collaboration.

The author first explores whether robot nonverbal cues inspired by human behaviours can elicit desirable responses from a human user to interweave unspoken – yet essential – spatial-temporal details of an interaction. Results from a series of experiments demonstrate that a robot cue, like gaze, can have a significant influence on *when* a human recipient reaches out to receive an object from a robot.

Subsequently, the author focuses on hesitation gestures – a type of gesture humans naturally use to express uncertainty – to explore whether members of a human-robot dyad can negotiate a desired outcome of an interaction through a nonverbal dialog. The author presents a reactive, real-time trajectory generator, the Negotiative Hesitation Generator (NHG), which has been devised to enable such nonverbal negotiation to take place between a human and a robot. The NHG was implemented on a robot for human-robot interaction experiments where, by design, spontaneous resource conflicts often arose between the two agents. Results from these studies suggest that use of the NHG can enable a type of nonverbal negotiation of resource conflicts to take place. They also demonstrate how such real-time negotiations between a human-robot dyad can lead to a faster resolution of conflicts and a significantly improved outcome of the collaborative task, without jeopardizing the safety of the user.

This thesis advances our understanding of the influence that nonverbal robot behaviours can have on human users. It also demonstrates the feasibility and efficacy of nonverbal negotiations as a mode of interaction for human-robot collaboration.

Lay Summary

Communication is essential to successful collaboration between people. The same is true for human-robot collaboration. The objective of this thesis is to investigate how robots can use nonverbal behaviours to better communicate and collaborate with people.

In this work, a robot was programmed to use human-inspired gaze cues as a nonverbal communication mode while handing over an object to a person. Experimental results suggest that people reach for the object earlier when a robot uses human-inspired gaze behaviours than when it does not. In a following experiment, a robot was programmed to appear hesitant when it and a person reached for the same object simultaneously during a collaborative task. Results suggest that such behaviour can allow a robot to quickly and safely resolve the conflict with a person.

Based on this work, future human-robot collaborations can be designed to be more effective and safer for both parties.

Preface

This thesis is comprised of a number of collaboratively written publications. Chapter 3 contains edited versions of the following:

- Moon, A., Troniak, D. M., Gleeson, B., Pan, M. K. X. J., Zheng, M., Blumer, B. A., MacLean, K., Croft, E. A. (2014). Meet me where I'm gazing: How shared attention gaze affects human-robot handover timing. In International Conference on Human-Robot Interaction (pp. 334-341). Bielefeld, Germany: ACM/IEEE. [105][1]

- Zheng, M., Moon, A., Gleeson, B., Troniak, D. M., Pan, M. K. X. J., Blumer, B. A., Meng, M. Q. H., Croft, E. A. (2014). Human behavioural responses to robot head gaze during robot-to-human handovers. In 2014 International Conference on Robotics and Biomimetics (pp. 362-367). Bali, Indonesia: IEEE. [167]

The author, AJung Moon, contributed to [105] in the design of the experiments in collaboration with Brian Gleeson and Minhua Zheng. The author led the analysis of the results with the assistance of Minhua Zheng. Daniel Troniak spearheaded the technical implementation of the system used in the experiment along with Benjamin Blumer, and Matthew Pan. Minhua Zheng, in collaboration with the author, conducted a qualitative analysis of the same experiment published in [167].

A follow-up study (Study 6) of the material presented in Chapter 3 has been included in Section A.1 as a supplementary research material. The follow-up study has been published in:

- Zheng, M., Moon, A., Croft, E. A., Meng, M. Q.-H. (2015). Impacts of Robot Head Gaze on Robot-to-Human Handovers. International Journal of Social Robotics, 7(5). [168]

Minhua Zheng and the author collaboratively designed the *in situ* experiment for [168]. The experiment and the majority of the analysis were conducted by Minhua Zheng with the supervision of the author and Elizabeth A. Croft.

The studies presented in Chapter 4 are being prepared for submission as a journal publication:

- Moon, A., Billard, A., Van der Loos, H. F. M., Croft, E. A., (n.d.). Development of Negotiative Interaction for Nonverbal Resolution of Human-Robot Conflicts. (*in preparation*)

[1] This paper has won the Best Paper award at IEEE/ACM Conference on Human Robot Interaction 2014

The author led the design and analysis of the experiments with the supervision of Elizabeth Croft and H. F. Machiel Van der Loos. Aude Billard supervised the data exploration of the first of the two human-subjects experiments (Studies 3 and 4) as well as the development of the Negotiative Hesitation Generator (NHG) controller inspired by the results of the experiments.

The final study (Study 5) discussed in Chapter 5 is also being prepared for submission as a journal publication:

- Moon, A., Croft, E. A., Van der Loos, H. F. M., Billard, A. (n.d.). Bidirectional Interweaving of Subplans using Negotiative Interaction in Human-Robot Collaborative Assembly. (*in preparation*)

The author designed the experiment and conducted the analysis of the results. Aude Billard supervised this process along with Elizabeth Croft and H. F. Machiel Van der Loos.

The University of British Columbia (UBC) Behavioural Research Ethics Board approved all user studies conducted as part of this thesis (H10-00503, "HRI Cues"). Study 5 presented in Chapter 5 has been conducted at Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. It has been approved by EPFL Human Research Ethics Committee (HREC: 001-2016 / 12.012016) in addition to the approval from UBC.

Financial support for this research was provided by the Natural Sciences and Engineering Research Council of Canada, the Canada Foundation for Innovation, the UBC Institute for Computing, Information and Cognitive Systems (ICICS), and the European Union project, AlterEgo, under the grant agreement #600010.

Table of Contents

List of Tables

List of Figures

Glossary

AHP Acceleration-based Hesitation Profile, a characteristic trajectory profile commonly observed in a particular type of hesitation gesture as elaborated in [104].

ANOVA Analysis of Variance, a set of statistical techniques to identify sources of variability between groups

DS Dynamical System

HH Human-Human condition

HR Human-Robot condition

HHI Human-Human Interaction

HRI Human-Robot Interaction

HCI Human-Computer Interaction

H2T Human to Target

ICC Intra-class Correlation Coefficient

KD Kickback Distance

LDS Linear Dynamical System

MLM Multi-level Modelling

NHG Negotiative Hesitation Generator

RA Re-attempts

REML Restricted Maximum Likelihood

ROS Robot Operating System

R2T Robot to Target

SSP Social Signal Processing

SVM Support Vector Machine

TS Trigger State

ZVC Zero Velocity Crossing

Acknowledgements

This thesis would not have come to fruition without the help of many writers and academics in my life. These mentors, friends, and family members fought alongside me as I learned to silence my inner sceptics. They include: My mother, Mi-Ja Park, who continues to fight her own battles as she expresses the life of expatriates through fiction; François Garillot, whose love, commitment to provide freshly brewed cups of coffee, and gifts of backup systems proved to be invaluable as I shipped the many chapters in this thesis; Shalaleh Rismani, who not only shared a daily thesis writing tracker to keep both of us on track, but also motivated me with daydreams about what is awaiting us at the end of the tunnel; John Havens, Executive Director of the IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems, who has been there to support my role as a chair of one of the initiative's committees as I intermittently disappeared to focus on my thesis.

I would also like to thank: My father, Young-Gu Moon, who sparked my interests in technology and has never second guessed or doubted my technical abilities as an engineer even in the face of my failures; My sister, Hyo-Jung Moon, who was there for me the day I moved across the country to start my research career in Vancouver as well as the eleventh time I moved between continents, cities and apartments as part of this journey.

I am grateful to the following individuals who provided technical support and assistance throughout the studies outlined in Chapters 3 and 4: Minhua Zheng, Daniel Troniak, Brian Gleeson, Matt Pan, Benjamin Bloomer, Karon MacLean, Maneezhay Hashmi, Bikram Adhikari, Fernando Ramirez, Lars Bollmann, Katelyn Currie, Jake Rose, Justin Hart, Junaed Sattar, Eric Fu, and John Petkau.

I would also like to thank Aude Billard and my labmates at Learning Algorithms and Systems Laboratory (LASA), EPFL, who generously welcomed me into the lab and provided me with critical and insightful discussions that led to the work outlined in Chapters 4 and 5. I would especially like to thank Sina Mirrazavi, Nadia Figueroa, Mahdi Khoramshahi, Nicolas Sommer, Lucia Pais Ureche and Guillaum De Chambrier at LASA who were always there to help diagnose numerous technical issues.

Lastly, I would like to thank my advisors Elizabeth A. Croft and Mike Van der Loos, whose mentorship and support extended much beyond the contents of this thesis and are the very foundation of the researcher I have become today.

Chapter 1

Introduction

In 1994, at the American Association for Artificial Intelligence's National Conference on Artificial Intelligence, Barbara J. Grosz delivered her presidential address titled "Collaborative Systems."[56] Taking inspiration from Lochbaum [90], who stated that "a buddy is better than a slave," that is, a system that works with a user to problem solve is better than one that only takes and follows orders, she encouraged the research community to take interest in human-machine collaboration and shared her vision of what such a system would look like [56]. Since then, the field of robotics has been motivated to develop interactive collaborative robotic systems in part as a means to address the global, impending shortage of labour that is expected as a result of rapidly ageing demographics in the industrialized world.[1]

Today, industrial robotic manufacturers such as KUKA (Augsburg, Germany), ABB (Zürich, Switzerland) and Franka Emika (Munich, Germany) are spearheading the vision of robots working side by side with human workers in manufacturing facilities without barricades. They envision such Human-Robot (HR) collaboration to yield higher throughput per human worker. Rethink Robotics' (Boston, MA) collaborative assembly robotic systems,[2] for example, have already been changing manufacturing environments and processes in ways that computers and single-purpose automation systems could not [77].

In step with the rise of collaborative robotic systems for manufacturing has been an increase in the number of robotic products promising to enter our homes and offices as social companions and assistants. For instance, Pepper is a 20 degree-of-freedom (DOF) humanoid robot from SoftBank Robotics (Tokyo, Japan) that has been marketed in 2016 as a "genuine day-to-day companion whose number one quality is his ability to perceive emotions" [133]. Jibo is a 3-DOF robot advertised as "The World's First Social Robot for the Home," and its founders raised over $3.7M USD through a 2014 crowd-funding

[1] According to a July 2015 report from Statistics Canada [143], more Canadians are over the age of 65 (16%) than under the age of 15 (16%). Canada and the United States have the lowest proportion of the population over the age of 65 compared to other countries in the G7. The population ageing phenomenon is found globally, with Japan having 26% of the population being 65 and older as of 2015. The field of robotics has been especially motivated to provide a technological solution to the projected societal consequences [92].

[2] Rethink Robotics is one of the few companies to develop and deploy interactive manufacturing robots. Called Baxter and Sawyer, they can be operated safely with humans in their workspace.

campaign [21]. Both robots are targeted at small businesses and the consumer market.

The rise in social and collaborative robotic systems is in part enabled by the development of new interfaces. Compared to the 1990s, when the keyboard, mouse, and physical buttons dominated the way people interacted with everyday technologies, two decades of research and development in robotics and computer science have introduced new user interfaces and ubiquitous technologies that allow for more natural and seamless interaction with today's electronic devices. An iconic example is the way people interact with smart phones. With the advancement of natural language processing, people can provide verbal commands to their smart phones and other electronic devices to accomplish simple tasks such as scheduling reminders, navigating the streets, or checking the weather forecast without the need to type commands or press numerous buttons (e.g., Google Now [55], Apple's Siri [7], and voice-command-enabled cars allow for a hands-free experience). However, while our interaction with small or so-called 'smart' electronic devices has been deeply integrated into the daily lives of those in developed economies, interactive robots have yet to achieve the same penetration into the consumer market. The human-machine collaboration Grosz envisioned has yet to be realized in robotics.

What does collaboration comprise and what challenges remain? In Bratman's model, supported and re-iterated by Grosz and Kraus [58], there are four key elements to a collaboration:[3] mutual responsiveness, commitment to the joint activity, commitment to mutual support, and meshing of subplans [20]. All of these elements require collaborating agents to communicate with each other, and imply the dynamic (i.e., constantly changing or evolving) nature of a collaborative process. For example, an agent's state of commitment to the joint activity and attitude toward mutual support can change over time, thereby breaking a collaborative process into separate, individual processes. Moreover, while individual agents may have their own subplans that are framed to serve a joint goal, details (e.g., spatiotemporal information) of the subplans are often not explicitly communicated to the other parties at the onset. These details need to be communicated and interwoven with those of the other collaborating partners who themselves have subplans to help achieve the joint goal. This interweaving process includes negotiating the details of a subplan when an agent's subplan conflicts with those of the others or when an unforeseen problem arises [56]. Given the foundational role communication plays in establishing the four elements of collaboration, the author posits that the development of natural communication mechanisms for meshing and negotiating details of subplans is key to realizing fluid and effective HR collaboration.

Unlike interaction with other electronic devices, robots are embodied in hardware, occupying physical space and manipulating objects in that space. The physical embodiment of robots not only provides new opportunities for what technology can deliver for human society – such as physically assisting an older person rise from a chair or lifting heavy objects to a more ergonomic position for human workers in manufacturing facilities – but also a unique modality for interaction and communication through physical presence and kino-dynamic behaviour that can facilitate an HR collaboration.

The main objective of this thesis is to investigate the role human-inspired robot nonverbal cues can

[3] Bratman used an alternate term, *shared cooperative activity*, to describe the type of collaborative activity pertinent to this discussion. While researchers have used other terminologies to describe similar multi-agent activities, this thesis uses the term *collaboration* for consistency.

have in the interweaving and negotiating of subplans between interacting and collaborating human and robotic agents. In particular, this work is focused on the last of the four elements of collaboration (meshing of subplans) where the intention to collaborate, the task at hand, and the roles assigned to each of the agents are not left to question, yet details of the interaction such as the *how*, *when*, and *where* remain to be determined.

1.1 Examples: Physical Environment and Human Behaviour

To take advantage of the physical embodiment of robots, roboticists are tasked with the problem of how a robot should best use, occupy, and share spaces and objects with human users in the environment. Communication in Human-Human (HH) collaboration includes both verbal and nonverbal means. HRI researchers who focus on nonverbal communication, in particular, examine Human-Human Interaction (HHI) in careful detail, extracting and characterizing interactions that are usually taken for granted in our everyday interactions with other humans. This process helps Human-Robot Interaction (HRI) researchers to discover curious ways in which we humans get along and perform activities with each other that dynamically help shape the outcome of the interactivity, be it psychological, behavioural, performance-oriented and so on.

Take, for example, the case of an able-bodied adult, Jane, walking down a street and another adult, John, walking on the same narrow street from the opposite direction. Jane may notice John and decide to alter her path closer to the right side of the street in order to avoid a head-on collision with John, a stranger. In turn, John may notice this and change his path slightly to his right (i.e., to the opposite side of the street), thereby giving both of them enough room to pass by each other without a collision, or a near-collision. In this scenario, there is a clear resolution on what behaviours result in fluent, safe, and collision-free interaction between the two agents. Jane acts as the main agent that leads what John should do to make the sharing of the space (the street) work between the two.

The effect of how humans influence one another is even more pronounced when considering the case of Jane and John noticing each other in close proximity at the same time. In order to avoid an imminent and unpleasant head-on collision with a stranger, the two pedestrians may unintentionally choose the same edge of the street to yield to the other. Observing this, they further decide to yield again and quickly shuffle to the other edge of the street, thereby finding themselves in the same and unresolved situation of imminent collision. This livelock can continue until their yielding behaviour falls out of synch or one of them decides to explicitly yield the right of way to the other.

These series of events, finding ourselves in an unexpected conflict of a resource (the resource being space, in the above-mentioned cases) and having uncertainties about how it should be resolved,[4] causes us to respond to the situation in interactive and communicative ways. The quick shuffling behaviour in the second scenario is an example of hesitation behaviours that one often observes in human interactions. In HHI, they often result in resolution of such conflicts using only nonverbal cues to negotiate the

[4] For example, a four-way stop intersection in the traffic law of the province of British Columbia specifies that the first car to arrive at the intersection has the right of way. In the case of multiple cars arriving at the same time, it is expected that the right-most car in the intersection has the right of way first. However, much of our everyday interactions do not include prescribed rules about who should yield the right of way.

Figure 1.1: Hesitation behaviours discussed in this thesis are set in the context of two agents reaching for the same resource at the same time. The photo on the left depicts such an instance studied in Chapter 4 where two people are reaching for the same deck of cards at the same time. The photo on the right illustrates an analogous resource conflict scenario in an HRI in which a human and a robot reach for the same object at the same time. The HR scenario is explored in Chapter 5.

outcome.

Ongoing work in the field of HRI related to robot physicality and presence are focused on the design features of a robot that best elicit specific responses from human users [3, 19, 22]. Findings from these studies provide points of reflection on how we affect each other's behaviours and how, increasingly, robots are becoming a source of that influence. As outlined further in Section 2.2, previous studies in HRI have explored ways in which robot nonverbal cues can help establish joint attention, share joint intention, and express internal states. Building on these studies, this thesis focuses on design and implementation of nonverbal communication cues within the context of HR collaboration as a means to interweave and negotiate about subplans of an activity.

Herein, the term interweaving is used as a shorthand to refer to the process in which multiple collaborating agents mesh their independently formed subplans with those of the others in order to accomplish the shared goal. It is used as an umbrella term encompassing processes that are unidirectional – such as the first scenario of Jane and John where Jane elicits a response from John to navigate the street in a mutually agreeable manner – and bidirectional – where Jane and John both influence each other through a nonverbal negotiation[5] of how they want to share the narrow street. While the source of influence in HHI may not be explicitly clear in many of our daily interactions with others, this is made explicit in HRI through the way in which we design interactive robotic systems. For the purpose of this thesis, hesitation is also an interactive behaviour that demonstrates ways in which two agents can negotiate for the resolution of a conflict within a shared, physical environment. Figure 1.1 provides a visual illustration of the context in which hesitation behaviours are studied in this thesis. The author asserts that investigating the feasibility of analogous behaviour to enable HR negotiations helps address the question: "In what ways should a physically embodied robot exert influence on our behaviours and decisions?"

[5] The term negotiation used within the context of this thesis is discussed more in detail in Chapter 2.

1.2 Research Questions

This thesis explores both unidirectional and bidirectional types of interweaving within the context of physical HR collaboration. Robot use of gaze cues and hesitation gestures are investigated as two coordination mechanisms that can help facilitate the interweaving process. In the unidirectional model of interweaving, the author investigates robot use of gaze cues to elicit desired responses from humans. In considering the bidirectional model, nonverbal negotiation between the two members of an HR pair using hesitation gestures is studied as a mechanism that allows interactive coordination of subplans.

The thesis is framed around the following research questions:

1. **Unidirectional Interweaving** Can a robot provide humanlike nonverbal cues to influence people's behavioural responses to an interaction *while* the interaction is taking place?

2. **Bidirectional Interweaving (Negotiation)** Can a robot nonverbally negotiate with a person about what should happen in an interaction? Can an HR negotiation contribute to an improved HR collaboration?

To address the first question, the author considers the activity of robot-to-human handover of an object as an interaction in which the *when* and *where* of the collaborative task is to be implicitly communicated to the human recipient (Chapter 3). While previous studies suggest that robot use of nonverbal cues can affect human behaviours in HR handovers, thereby providing a strong support for the first question, exploration of the question in this thesis explicitly addresses the human behaviours *during* the *interweaving* process of the interaction – that is, coordinating when and where the person should move his/her hand to receive the object the robot is handing over. The author examines effectiveness and efficacy of a robot's use of human-inspired nonverbal cues to elicit a response by the collaborating agent (human), in order to successfully complete the handover task. Like many other HRI studies that have preceded this work, the robot in this interaction does not leave room for the person to initiate and influence the robot and its behaviours.[6] The order of the task to be performed and division of roles in the task are both made clear, and the communication cues used to interweave the *when* and *where* part of the collaboration are unidirectionally provided by the robot to the human recipient.

To address the second question, the author explores an HR collaborative assembly context where, by chance, both human and robot reach for the same resource at the same time and are left to their own devices to resolve the conflict at hand. Hesitation behaviours similar to the second scenario of Jane and John above are investigated as a means for an HR pair to negotiate with each other. In this interaction, the agents must effectively figure out the priority order to access a shared resource that often, spontaneously, comes into conflict for the collaborating HR pair. In order to have a nonverbal dialogue leading to a negotiated conflict resolution the two agents must bidirectionally communicate and respond to each other in real-time.

[6] This statement is not meant as a generalization of the work accomplished in the field of HRI, but only that there are numerous work that share the open-loop interaction model employed in this study. The field has seen some progress in human-to-robot handover interactions, where the user takes the role of initiating the interaction (e.g., [35, 147]).

1.3 Thesis Outline

The remainder of this thesis is organized into five chapters. Chapter 2 provides a general and overarching literature review on HRI with a focus on the influence that interacting agents exert onto each other.

Chapter 3 is dedicated to the first of the two research questions listed above. It describes two studies that have been conducted to understand some of the cues humans use in HH handovers and evaluate whether human-inspired robot behaviours in robot-to-human handovers can elicit desired behavioural responses from human participants. In particular, the studies help address whether a robot's unidirectional, nonverbal communication cue can convey the details of when and where the transfer of an object from robot to human should take place.

Chapter 4 and Chapter 5 focus on the second of the two research questions. Analogous to the process in Chapter 3, two studies are presented in Chapter 4. The first study serves to observe how humans negotiate and resolve resource conflicts with one another using hesitation gestures. With a better understanding of characteristic features of human hesitation gestures, a novel trajectory generator, the Negotiative Hesitation Generator (NHG), was then designed and implemented to control a robot arm. The second study outlined in the chapter evaluates the efficacy of the designed controller and its parameter values. Subsequently, an in-person HRI study is presented in Chapter 5: participants were invited to collaborate with the robot, which exhibited human-inspired negotiative hesitation behaviours based on the NHG. Results from this study demonstrate that human participants do yield to and negotiate with robots in a collaborative activity.

This thesis concludes with Chapter 6 highlighting the main contributions of the thesis, namely: advancing the field's knowledge of nonverbal HRI, especially in robot-to-human handovers and hesitations; a novel hesitation controller that allows an HR pair to negotiate about and interactively and safely resolve resource conflicts in an HR collaborative assembly scenario; and demonstrating the efficacy of nonverbal negotiations as an efficient and fluent mode of interaction.

Chapter 2

Background and Motivating Literature

This chapter presents relevant, motivating literature from fields of study that frame the research questions investigated in this thesis. It first situates the contributions of this thesis on HR communication and collaboration (Section 2.1). Previous studies in HRI are discussed in terms of unidirectional and bidirectional modes of influence previously documented in the field (Section 2.2).

As outlined in Chapter 1, the work presented in this thesis involves investigations of two different types of nonverbal communication cues – gaze cues used in robot-to-human handovers (unidirectional) and hesitations for HR collaboration (bidirectional). Each of them independently contributes to a better understanding of the particular type of communication cues used in HRI. Therefore, previous work that frames the contributions specific to the type of cues investigated in the following chapters are presented in the relevant chapters. This chapter, on the other hand, paints a broader picture of the literature leading to the main contributions of this thesis.

2.1 Human-Robot Collaboration and Communication

Robots have been performing dull, dirty, demanding, and dangerous tasks for decades, particularly in manufacturing facilities.[1] While the industry practice of delegating tasks to robotic systems is not new, most of the traditional robots have been limited to industrial applications and confined to operate within work cells that physically separate the robot's workspace from humans. In comparison, a growing number of robots today are being designed and marketed to work both inside and outside of manufacturing environments and without physical separation from humans. An iconic example that demonstrates the wide-ranging application of today's robotics technology can be found at Yotel [166], a hotel on the densely populated Manhattan Island, New York, NY. At the lobby of the hotel is an IRB 6640 (ABB Group, Zurich, Switzerland), an industrial robot capable of handling 200 kilograms of weight, that has been modified to perform the tasks of storing and retrieving luggage for guests [36] – tasks typically performed by human porters in most hotels. This delegation of task frees human employees to focus on other tasks, and releases them a potentially injurious lifting task. Paralleled by the widening range of

[1] Robots have been deployed to perform tasks in place of humans, such as handling nuclear materials, since 1940s [135]. For more historical coverage of use of robots see [135].

application domains is an increase in the type of tasks that require robots to share spaces and objects with people in a safe and effective manner. For example, medical drug delivery robots, such as HOSPI (Panasonic Asia Pacific, Singapore), perform the traditionally human task of fetching and delivering drugs from one part of a hospital to another [2]. HOSPI is built to autonomously navigate through hospital corridors that are shared by human staff and patients. It also interacts with doctors and nurses to handle the drugs that have been delivered.

These new robotic systems[2] – systems that operate outside the confines of work cells – are creating new technological relationships in our society. While robots can perform certain tasks that are difficult or undesirable for humans to perform in a tireless manner, there are numerous tasks that humans can accomplish with ease that pose complex technical challenges for a robot. For instance, loading a dishwasher may be a trivial task for most able-bodied human adults. However, getting a robot to load the dish washer is a technical challenge that requires the implementation and integration of object recognition, grasping, and manipulation capabilities on a robot. In fact, the dish-loading activity was used as one of the tasks for a robotics competition in 2010 [18]. Therefore, enabling robotic systems to successfully interact and collaborate with users can help complement the human and robotic agent's strengths and weaknesses and create a synergy that maximizes the benefits robotics has to offer.

A number of fields of study ranging from Psychology and Cognitive Science to Entomology document the efficacy of collaboration in human, animal, and insect societies [20, 37, 56–58, 154, 162]. Likewise, findings from studies in HRI suggest that effective collaboration between a human and a robot can reduce task completion time and improve accuracy, quality, enjoyability, and safety of a task in comparison to solo task performance [19, 128]. Previous work also suggests that introducing robots that work alongside humans may be preferred to those that replace humans [150].[3] For example, Wong et al. [163] conducted a study in which a small humanoid robot, Nao (SoftBank Robotics, Tokyo, Japan), recited stories to teenagers either by itself or with a human storyteller. When the robot told the story by itself, it used gestures and gaze to perform the scenes of the story. When the robot was collaborating with the human storyteller, the robot narrated the story while the person performed the scenes. Although the content of the story was the same in both conditions, the participants preferred the second condition involving human-robot collaboration. The researchers attribute this result to the fact the human and the robot complement each other's strengths and weaknesses in interacting with the participants.

As discussed in Chapter 1, communication is an essential component of collaboration. However, the impact robot communication cues have on an HR collaboration can be complex and is the subject of on-going research. For instance, recent studies report that performance improvements of an HR collaboration can be salient in complex collaborative tasks than in simple tasks in which the improvements may

[2] The term New Robotics has indeed been used to refer to this shift from traditional industrial robotic systems to the growing number of robotic systems meant to be deployed and operated outside the confines of work cells [137].

[3] To say that the public's acceptance of robots that work alongside people is higher than those that work in place of people is a crude generalization. A survey-based study conducted by Takayama et al. [150] provides empirical evidence that people prefer robots that work *with* people than robots that *replace* people. The findings from the study were later contested by a subsequent study conducted by the same researchers [72]. The researchers note in discussing the contrasting results from the two studies that much more work is required to understand how this preference of robots' roles in our society varies across different demographics. The author's previous work also demonstrates that public acceptance of robotics technology and its role in our society can vary according to the application domain in question [98, 102, 118–120].

not be present at all. Admoni et al. [3] investigated the impact a robot's nonverbal gestures (pointing and gazing) can have on performing tasks of different difficulty levels. Employing a task that involves memorizing and following a set of instructions given by the robot, the authors varied difficulty levels of the task by adding more steps to the instruction or by introducing interruptions. The results of this study demonstrate that the improvement of the participant's task performance with the robot's use of nonverbal gestures was much more pronounced when the task was difficult than when it was easy. The positive effect nonverbal HR communication seems to have on team performance of complex rather than simple tasks is also echoed in two HRI studies conducted by Gleeson et al. [52]. They investigated whether the nonverbal communication of tapping and pushing to control a robot in a bolt insertion task outperforms the traditional button-based interface to achieve the same objective. The results of their studies demonstrate that for a simple insertion task the physical, tap-and-push interaction with the robot did not improve the team's task performance, nor did participants prefer the traditional button-based interaction. However, it did significantly outperform the button-based interface when the task was less scripted and more complex to execute.

Results from these studies suggest that a robot's behaviours can be designed to enable humans to achieve more when a task is performed with robots. They also suggest that much work remains to be done in understanding the nature of the impact communicative cues can have on an HR collaboration.

2.2 Robots as a Mechanism of Influence

As briefly mentioned in Chapter 1, the physical embodiment of robots makes them a powerful source of influence for human behaviours compared to artificial intelligence software alone. This section provides a short review of the role robot behaviours can play in unidirectionally and bidirectionally influencing human behaviours and decisions.

The distinction the author makes between unidirectional and bidirectional influence is based on how the interactive system is designed. Unlike HHI, in which it is often hard to measure how much one person is affecting the other, the directionality of influence in HRI can be discussed much more explicitly by considering the design of the robotic system. How much or in what ways the system incorporates implicit or explicit input signals from the user is a direct consequence of the designers' decisions. On the other hand, the amount and type influence a user is subjected to during an HRI are not factors that are explicitly controlled by the users. A HRI system can be designed to be reactive to human commands and behaviours, such that human users unidirectionally influence the system. A system can also be designed to proactively elicit desired responses from humans without itself reciprocally being influenced by the person.

This thesis is focused on the influence robots can have on humans in the process of interweaving details of a collaborative task. Therefore, in Section 2.2.1, only the literature in which the robot is designed to unidirectionally influence humans is discussed. In discussing bidirectional influence (Section 2.2.2), the author refers to systems that are designed to influence and be influenced by humans through implicit nonverbal cues.

2.2.1 Unidirectional Influence

The field of HRI has uncovered a variety of ways that robot behaviours can unidirectionally affect humans. For instance, Ju and Sirkin [70] demonstrated that when a physical gesture by a robot is compared against an analogous gesture displayed on a screen, the effect the gesture has on human behaviour differs drastically. They conducted a study where a robotic kiosk placed at a bookstore and a building lobby either physically gestured using its hands or projected the gesture on its screen. The authors found that the robot engaged up to twice as many passersby when it physically gestured versus when it used a screen-based gesture. In fact, robots can significantly affect human behaviours just by its presence. In a recent study, Hoffman et al. [62] found that humans cheat just as much in the presence of a robot as they would with a person. The robot used in the study was built for the specific purpose of conveying social presence rather than demonstrating any monitoring behaviour. Yet, the participants cheated less when a person or the robot was present than when the participants were left alone in a room. This finding echoes the media equation put forth by Reeves and Nass [130], which suggests that humans treat machines as though they are social beings. Moreover, studies have found that a robot's motion can trigger different affective as well as physiological responses to a human observer even when the start and end points of the robot's motions remain the same [83].

A robot's ability to affect humans is a necessary condition for the robot to interweave details of a task with a person in an HR collaboration. Studies in HRI suggest that implementing communication behaviours on a robot, often inspired from human or animal communication behaviours, can positively impact the HR team's performance and the user's perception of the robot [22, 24, 52, 134]. The positive impacts of robot nonverbal cues in HR collaboration are often attributed to the fact that implicit and explicit nonverbal gestures help establish joint attention [65], share joint intention [45, 99, 155], and express the robot's internal states [22, 45, 134] to the users in an intuitive manner. Romat et al. [134] conducted a study where a human participant and a robot were tasked to collaboratively build a tower using Duplo blocks. Due to the kinematic constraints of the robot, the participant needed to help the robot reach the desired blocks placed outside the robot's range of motion by understanding the robot's needs and moving the blocks closer to the robot. The authors found that when the robot provided nonverbal cues to indicate its need to grab blocks outside of its reach the participants were significantly quicker in providing the desired assistance than when the robot did not exhibit the cues. Likewise, Breazeal et al. [22] demonstrated that a robot's use of implicit nonverbal gestures[4] (e.g., shrug and gaze) used in conjunction with explicit nonverbal gestures (e.g., pointing gestures) significantly improves the user's understanding of the robot's internal states and the HR team's overall task performance.

The type of influence that is especially relevant to this thesis is how nonverbal modes of communication affect an HR dyad as the agents interweave spatiotemporal details of a collaborative task. A robot's use of gaze cues, for example, has been demonstrated to provide supplementary information to human users in in-person HRI, thereby facilitating the interweaving process. In an HRI experiment

[4] Here, Breazeal et al. [22] describe implicit nonverbal gestures as nonverbal behaviours that convey information inherent in the agent's behaviour rather being purposefully communicated. This is distinguished from explicit nonverbal gestures in which the agent intends to convey specific information to the recipient. This is analogous to the way implicit interaction is understood in Ju and Takayama [71] and other HRI literature.

involving a table-top manipulation task, Boucher et al. [19] explored whether a robot's use of gaze cues can decrease the reaction time of human collaborators whose role was to manipulate items that were verbally announced by the robot. The results of their study suggest that when the robot used head and eye gaze, subtly conveying which item was being attended to, the participants reacted significantly faster in reaching for the desired items than when the gaze cues were occluded from the participants. The authors attributed this result to the fact that robot gaze cues can serve the function of establishing joint attention and help human observers predict what the next desired item may be. This echoes the finding from a study conducted by Mutlu et al. [113]. In this study, a humanoid robot, either a Geminoid or a Robovie [4], played a guessing game where the participant had to identify which of the items on a table the robot chose for them to guess. The authors found that the participants performed significantly better when the robots provided quick gaze cues as though giving a hint. Both of these studies demonstrate that the robot's use of gaze cues is effective in nonverbally communicating the desired target object to the participant. Based on the results of these studies, gaze cues implemented on a robot can likely convey spatiotemporal information to a human collaborator.

In exploring the first research question (*Can a robot provide humanlike nonverbal cues to influence people's behavioural responses to an interaction while the interaction is taking place?*) in Chapter 3, human-inspired gaze cues are implemented on a robot in a robot-to-human handover context. The studies presented in the chapter help test the hypothesis that nonverbal supplementary cues, such a gaze, that are exhibited during a robot-to-human handover interaction can influence when and where human recipients will reach out to accept objects handed over to them.

Many of the robot nonverbal behaviours that positively impact task performance are designed based on observed human and other animal behaviours [23, 61, 64, 71, 151]. This is not surprising given that the human ability to read someone else's internal states and intentions and engage in collaborative tasks are developed in infancy and through mimicry of other humans.[5] Anthropomorphic communication cues in a robotic system, therefore, have a better chance of being understood by human users. In this thesis, the author also looks to human behaviours to inspire the design of robot nonverbal behaviours for an improved HR collaboration. In particular, this thesis focuses on two different types of nonverbal cues that can be superimposed on the functional motions required to complete a task: in the studies presented in Chapter 3, various gaze behaviours are used to supplement the functional motions of a robot handing over an object to a person; in Chapter 4 and Chapter 5, hesitation gestures are superimposed on a robot's reaching motion as it moves toward a button to be pressed. Results from these studies also support that human-inspired nonverbal cues help improve the performance of the HR collaboration.

In addition, there is evidence that the impact robots have on human behaviours may be attributed to how we neurologically respond to physical motions performed by robots. In neuroscience, it has been established that motion properties of one person can activate the mirror neurons of an observer, and influence the properties of motions exhibited by the observing person. Recently, the field of HRI is gathering evidence that this powerful neurological phenomenon can be present in the context of HRI

[5] Numerous examples in developmental psychology provide evidence that such skills are established in infants as young as 12 months old, who mainly use nonverbal, protosocial behaviours to communicate with adults and other children [31, 154].

– that is, motions of a robot can activate mirror neurons of the humans who observe them and, in turn, influence the properties of motions performed by the observers. In a study conducted by Bisio et al. [17], the human participants were asked to observe either a humanoid robot or a person demonstrate a series of reach motions. The agents performed the reach motion either to a specified location or to transfer an object from one point to another.[6] After the observation, the participants were asked to perform a reach, with which the experimenters measured the similarity of the participants' trajectory profile (velocity) with that of the motions demonstrated to them. The results of this study suggest that regardless of whether a motion was demonstrated by a person or a robot, the participants mimicked the quality of the motion of what was demonstrated to them, except when the demonstrated motion contained a non-biological velocity profile. This demonstrates that the phenomenon of motor contagion, where the motions of one agent are mirrored in the motions of another, is also present when the motion is demonstrated by a robot. While the motion contagion can be bidirectional in HHI, in HRI the direction of contagion is unidirectional (a robot one-sidedly affects the motions of the interacting person) if the motions of the robot are not designed to adapt to that of the person interacting with it.

Therefore, in designing robots that will socially and collaboratively interact with us, it is crucial that we continue to explore the nature of human responses to robot behaviours and how they are similar or distinguished from the way people respond to other technological devices. This new understanding of human responses with respect to robots will not only allow us to better contemplate the societal implications of the technologies, including new social norms that can be formed with the presence of a robot, but also allow HRI practitioners to make informed design choices.

2.2.2 Bidirectional Influence

One of the studies that motivated this thesis is an online survey conducted by the author and her colleagues [30, 106]. In this survey, the authors were interested in finding out what a mail-delivery robot should do when it encounters conflicts with a person in using an elevator. The participants of the survey were given variations of a scenario where a large humanoid robot is carrying mail to be delivered to an important person on another floor. While the robot necessarily needs to use the elevator to travel between floors, it cannot share the elevator with a person due to space and safety reasons. The participants were asked what the robot should do when it is delivering mail and a person is already using or waiting for the elevator. In addition to the location of the person with respect to the elevator, a combination of two other factors (urgency of the mail delivery task, and the status of the person – in a wheelchair, carrying heavy items, or neither) was used to generate multiple versions of the scenario. The participants were given the following choices: do nothing, yield to the person, ask the person to yield, or engage in a dialogue with the person. The sample size of the study was too small to make a general claim. However, when presented with a highly contentious situation (e.g., an urgent mail is to be delivered and a person in a wheelchair is already using the elevator) the participants showed a clear preference that the robot should engage in a dialogue with the person to figure out whether to yield.

Answers to dilemmas vary from one person and context to another. However, as presented in the

[6] The demonstration took place in the same room rather than using a video recording.

example of Jane and John in Chapter 1, people engage in verbal or nonverbal dialogues with the others involved in an uncertain situation to negotiate a way forward. The results from the abovementioned study suggest that an effective way to handle conflicts and contentious situations arising in an HRI may be to enable the collaborating robot and human to negotiate a solution.

In order to engage a robot in a nonverbal HR negotiation, the robotic system must not only be able to express its internal states but also have a means to react to human responses in real-time. Mutual responsiveness is, in fact, one of the key elements of collaboration identified by Bratman [20]. Nikolaidis et al. [115] tested the impact of the bilateral relationship in an HR collaboration context. They developed and tested a mutual adaptation method that allows the HR pair to adapt to each other without an explicit change of turns. In their experiment, the participants were asked to collaboratively lift and transfer a large table with a robot while navigating through a narrow doorway. The authors tested three different conditions to vary the directionality of influence between the human and the robot: fixed, in which the robot used a fixed strategy in maneuvering the table; mutual-adaptation, in which the robot guided the participant toward a strategy when the participant is adaptive, and follows the guidance of the participant when s/he is not adaptive; and a cross-training condition, in which the robot uses one strategy in one phase of the task, and follows the user's strategy in another. Results from this study provide empirical evidence that mutually adaptive systems – in which the system asserts one strategy while allowing the user to negotiate for another – yields significantly improved team performance without sacrificing the user's subjective perception of the robot.

This study provides positive evidence that a robot that influences and is influenced by its users can improve the outcome of the collaborative task. This thesis posits that a robot that nonverbally negotiates with the user in the process of arriving at a joint goal can also improve the outcome of the HR collaboration.

The research presented in Chapter 5 extends this much under-explored line of research by testing artificially generated hesitation trajectories (developed in Chapter 4) as a mechanism for a robot to nonverbally negotiate for a solution to a resource conflict with a human user.

2.3 Summary

This chapter presented a broad overview of the motivating literature framing this thesis. Acknowledging the rising trend for and potential usefulness of collaborative robotic systems that share spaces and objects with humans, the challenge remains for the field of HRI to develop communication mechanisms that can help realize that potential. Previous studies demonstrate that nonverbal cues implemented on a robot can help communicate internal states of a robot (e.g., the robot's need for assistance from a person) and positively affect the task outcome of an HR collaboration. While the use of subtle nonverbal cues has previously been employed to convey details of a task, such as which item is relevant for the activity at hand, the impact such cues can have on conveying spatiotemporal details of an interaction – thereby unidirectionally eliciting the person to move to the desired location at the desired time – remains to be investigated. The author hypothesizes that a robot's use of humanlike nonverbal cues can help interweave spatiotemporal details of a task, and tests this in a robot-to-human handover scenario in Chapter 3. There are only a limited number of examples in the HRI literature that demonstrate the usefulness of bidirectional influence in an HR collaboration. However, results from the few studies suggest that bidirectional influence in an HR team can improve the HR collaboration in terms of task outcome. These findings inspire the hypothesis that HR negotiation, which necessitates the agents to influence each other using nonverbal gestures, is a possible and desirable mode of interaction in HR collaboration. Chapter 4 and Chapter 5 of this thesis explore the usefulness of hesitation gestures as a type of communication cue that enables both agents in an HR dyad to influence each other in a resource conflict scenario for an improved HR collaboration.

Chapter 3

Unidirectional Interweaving of Timing and Space in Robot to Human Handovers

3.1 Introduction

This chapter[1] addresses the interaction of a robot handing over an object to a person (robot-to-human handover). This interaction is chosen as a means to explore first of the two research questions explored in this thesis: "Can a robot provide humanlike nonverbal cues to influence people's behavioural responses to an interaction *while* the interaction is taking place?" More specifically, the studies presented in this chapter focus on the extent to which a robot's use of nonverbal cues helps elicit desirable responses from a human user as a means to unidirectionally interweave the spatiotemporal details of the interaction.

Object-handover is an important interaction to study in HRI. Enabling successful and fluent handovers between an HR pair is crucial in order for robots to take on more assistive roles for humans at homes and workplaces. Many application scenarios, including manufacturing and home environments, can involve situations where it is useful for a robot to fetch and hand over an object to a person. Implementing an effective HR handover interaction, however, is a challenge. In HH handovers, a great variety of subtle signals mediates the handover event. Body position, hand and arm pose, gaze, and grip force are used to communicate not only the intent to engage in a handover, but also when and where the handover is to occur. These subtle signals help create a fluent and fast interaction while ensuring that the object is not dropped (e.g., [15, 34, 86, 108, 145, 146]). When a robot does not provide appropriate cues, HR handovers can fail in a variety of ways: people do not recognize that the robot is giving them an object [27], objects can be dropped [34], or people can feel uncomfortable or unsafe during the

[1] ©ACM/Springer. The majority of this chapter has been modified/reproduced, with permission, from the published article Moon et al. [105]. Supplementary research on the same topic has been conducted by the author in collaboration with Ms Minhua Zheng, who took lead authorship. Zheng et al. [167] and Zheng et al. [168] have been published based on the supplementary research. The majority of this chapter's content is on the first publication, Moon et al. [105], for which AJung Moon led the project and authorship.

Findings from Zheng et al. [167], a qualitative analysis of the study conducted in Moon et al. [105], is presented as part of the results in this chapter. A follow-up study conducted afterwards, Zheng et al. [168], is briefly mentioned in this chapter as part of the discussion presented in Section 3.5. A more detailed summary of Zheng et al. [168] is presented in Appendix A.1.

interaction [42, 82].

Robot-to-human handover is an interaction that requires both agents to coordinate their motion through space and time to accomplish their shared objective. While the desired outcome of a handover (an object offered by a giver is successfully handed over to a receiver without dropping the object) is often well understood by a giver-receiver pair at or before the onset of the interaction, precise spatiotemporal details about *when* and *where* the person should reach out to grasp the offered object are usually not explicitly communicated by the interacting agents. In everyday HHI, a giver-receiver pair seems to naturally reach an agreement about these details as the interaction takes place and often without a need for verbal dialogue. In HRI, such natural process needs to be understood and appropriate behaviours implemented onto a robot for the interacting agents to interweave the unspoken details with each other.

Studies presented in this chapter address whether gaze can be used to augment an HR handover event, subtly communicating handover location, handover timing, and providing acceptable social interaction signals to modulate the human recipient's behavioural decision on *when* and to *where* s/he should reach for the offered object. Gaze cues, in either HHI or HRI, have proven to be efficient for communicating attention [81, 110]. During a handover, givers use verbal or nonverbal cues to direct the receiver's attention to an object. Successful handovers typically take place when the two parties achieve shared attention on the same object. Previous studies [86, 145, 146] indicate that gaze can be used by robots to signal handover intent to users before the handover event. However, these studies did not explore the effect of robot gaze during the handover on the timing of the handover event.

It is hypothesized that the use of human-inspired gaze cues during HR handover can influence a person's behavioural decision on handover timing and the subjective experience of the handover by implicitly increasing communication transparency and perception of naturalness for the interaction. This chapter presents two handover studies conducted to investigate the effect of gaze cues on handover timing during an HR handover event.

The first study, Study 1 (Section 3.3), was conducted to observe the type of gaze patterns humans use when one hands over an object to another. Results from Study 1 informed and inspired the gaze patterns implemented for HR studies 2 and 6. In Study 2, a PR2 humanoid robot used two different human-inspired gaze patterns observed from Study 1 to address whether the use of gaze cues affects people's first-time behavioural and self-reported responses to HR handover. Results from Study 2 suggests that the subtle and supplementary gesture of gaze does significantly affect an untrained human recipient's decision on when to reach for the offered object, thereby eliciting a faster, more fluent HR handover interaction. In addition to understanding participants' first-time handover response, it is important to understand whether robot gaze cues can have lasting effect on robot-to-human handovers even after a series of handovers. Following Study 2, the HRI experiment in Study 6 presented in Appendix A.1 reaffirmed the effect of gaze in non first-time responses to handover events and investigated whether the effect of gaze persists in repeated HR handovers. Positive results from these studies serve as a motivation for the studies in the following chapter, Chapter 4, where the outcome of the interaction is uncertain and needs to be resolved by the agents. All of the studies mentioned in this chapter have been approved by the University of British Columbia Behavioural Research Ethics Board.

3.2 Background

This section outlines a body of work specific to HR handovers and the use of gaze in HRI. In particular, to home in on the discussions relevant to the three studies presented, the focus is on studies where a robot is the giver and a human the receiver.

3.2.1 Human-Robot Handover

Previous research in HR handovers can be broadly categorized by the aspect of the handover under investigation: approach for handover, handover trajectory and pose, and the handover event itself.

Studies of approach for handovers consider situations where a mobile robot navigates towards a human to initiate a handover. These studies generally focus on human preference for approach directions and on creating robot behaviours that clearly communicate the intent to initiate a handover. Basili et al. studied how humans hold objects as they approach for a handover [15]. Koay and Sisbot [82] studied human preferences for coordinated arm-base movement in handover approach. Mainprice et al. [91] designed an approach planner that considers the mobility of the receiver. While the studies in this chapter do not involve a robot's handover approach (i.e., the participants approached our robot in our experiment), findings from the above studies guided the experimenters' decision on placement of the robot, as discussed in Section 3.4.1.

Other researchers have investigated handover trajectory and pose, reporting guidelines for how a robot arm should be positioned for handover and how that position should be achieved. In a series of studies, Cakmak et al. [28] and Strabala et al. [145] studied how handover trajectories and final handover poses can best signal the intent to initiate a handover. They found that the final handover pose should feature a nearly fully extended arm in a natural (human achievable) pose with the elbow, wrist, and distal point on the object positioned, respectively, from closest to furthest away from the body in all three dimensions. The object should be held in its canonical orientation (right side up) and positioned to allow easy grasping by the human. A related study emphasized the importance of the physical cues in HR handovers, showing that poorly designed handover poses and trajectories were often unsuccessful in communicating intent and ultimately resulted in handover failure [27]. They found that intent is best communicated by having high contrast between the pose used for holding the object and the pose used for handing over the object. In the studies presented in this chapter, the above guidelines were followed in the design of handover pose and trajectory, as described in Section 3.4.1.

Other studies have investigated the velocity profile of handover motions and have found that trajectories that minimize end-effector jerk make people feel safer in handover interactions [42, 66]. Other studies of handover trajectory include a human-based potential field planner for handover trajectories [74].

Chan et al. [34] studied the actual handover event, measuring grip and load forces in HH handovers and using these data to design a robust robot handover grip controller that imitates human handover behaviour. This controller has been adapted for use in Studies 2 and 3.

3.2.2 Gaze in Human-Robot Interaction

Gaze is an important and useful cue in HHI. People repeatedly look each other in the eye during social interaction and people do not feel that they are fully engaged in communication without eye contact [41]. Studies in psychology have shown various functions of gaze in social interaction, such as seeking and providing information, regulating interaction, expressing intimacy, exercising social control, etc. [41, 81, 124]. Gaze can be named differently in different social situations [110]; for example, mutual gaze or eye contact is defined as two people looking into each other's face or eye region [161], while deictic gaze or shared visual attention is defined as one person following the other's direction of attention to look at a fixed point in space [26].

Previous work has shown the importance of gaze in HRI. For example, Staudte and Crocker [144] demonstrated that humans react to robot gaze in a manner typical of HHI. Since gaze behaviour is closely linked with speech [8], much work has been done on the conversational functions of gaze in HRI [84, 89, 94, 111, 112, 158]. Gaze is particularly effective in regulating turn-taking during HR conversation. Kuno et al. [84] developed gaze cues for a museum guide robot to coordinate conversational turn-taking. Matsusaka et al. [94] used gaze cues to mediate turn-taking between participants in a group conversation.

Another large body of literature focus on using gaze to direct people's attention in HRI [16, 59, 67, 131, 140]. Gaze was combined with pointing gestures in [16, 59, 67] to direct people's attention, which the authors believed would make the interaction more human-like [16] while minimizing misunderstanding [59]. In Rich et al. [131] four types of "connection events" were identified from HHI videos, namely directed gaze, mutual facial gaze, adjacency pairs and backchannels. Implementing them in an HRI game showed a high success rate in forming HR connection or joint attention. In Sidner et al. [140], people directed their attention to the robot more often in interactions where gaze was present, and people found interactions more appropriate when gaze was present.

Introducing gaze cues can also benefit HRI in other ways. In Mutlu et al. [111] and Skantze et al. [141], gaze increased human performance in certain HR tasks. In Kuno et al. [84] and Sidner et al. [140], gaze heightened HR engagement and in Liu et al. [89], gaze cues contributed to the perceived naturalness of a communicating robot.

In the study of HR handovers, other researchers have shown that gaze can be useful in communicating the intent to initiate a handover. Lee et al. [86] studied human motion and gaze cues as people approached each other for handovers. They found that people looked at the object or the receiver as they approached the receiver. Strabala et al. [146] examined the signals that humans use to communicate handover intent before a handover takes place. They initially acknowledged gaze as one of the important features that mark the difference between different phases in handover, but they did not find gaze to be an effective predictor of handover intent. In contrast, Kirchner et al. [80] demonstrated how robot gaze can be effective in targeting an individual recipient out of a group of people for a robot initiated handover. Atienza and Zelinsky [9] augmented handover interactions with gaze cues, demonstrating a system that allowed a human to request an object for handover by looking at it.

While the above studies address gaze in pre-handover cuing and communication of intent to handover, the studies outlined in this chapter examine the use of gaze *during* the handover event. Although the effectiveness of gaze in regulating handover intent remains an open question, gaze may have a positive effect when used during the handover event. Gaze can help establish shared attention to a location in space and could elicit human receivers to move their hands to the desired location. Hence, gaze may be useful in improving the handover itself by establishing shared attention and influencing timing at which the human recipient reaches out to receive the object being handed over.

3.3 Study 1: Observing Gaze Patterns in Human-to-Human Handovers

An HH study, Study 1, was conducted to help understand what kind of gaze people use in HH handovers. In this study, a pair of subjects handed over a water bottle to each other multiple times. The experimenters observed the gaze behaviour of the giver during the handover by collecting and analyzing video recordings of HH handovers (see Figure 3.1). While other researchers have observed gaze in HH handovers (e.g., Strabala et al. [146]) *before* the handover event, the study presented in this section augmented these previous results by focusing on gaze *during* the handover event.

3.3.1 Experimental Procedure

Twelve volunteers (10 male, 2 female) participated in this study. The giver was asked to hand over ten bottles from a side table to the receiver one at a time. The receiver was asked to bring the bottles to a collection box about two meters behind them one at a time, requiring him/her to walk away from the giver between handovers. This process repeated until all ten handovers were completed. Each participant performed the role of the giver, then was paired with another participant and performed the role of the receiver, resulting in twelve giver-receiver pairs (120 handover recordings in total).

In order to collect human gaze patterns that can inform the design of nonverbal HR handovers, the giver and receiver were instructed not to talk during this process. The giver was also instructed to pick up the bottles from the side table only after the receiver returned from the collection box and had put his/her hands on the table. By requiring the receiver to turn and walk away, the common attention between the giver and the receiver was interrupted after each handover and participants needed to re-connect for the next handover.

3.3.2 Results

Annotation of a frame-by-frame video analysis of the givers' gaze patterns indicates that the giver's gaze during a handover can shift between three positions: the object being transferred, the expected handover position, or the receiver's face. Figure 3.2 shows a typical timeline of the five gaze patterns (gaze direction and timing) and corresponding frequencies observed from this study.

The following describes the five gaze patterns:

Shared Attention Gaze (Attn) The most frequent gaze pattern (68% of all handovers observed) consists of the giver gazing at a projected handover location as s/he reaches out to execute the handover. After picking up the bottle, the giver turns to face the receiver, looks down at a midpoint between the giver and the receiver, and keeps the gaze there until the receiver takes control of the

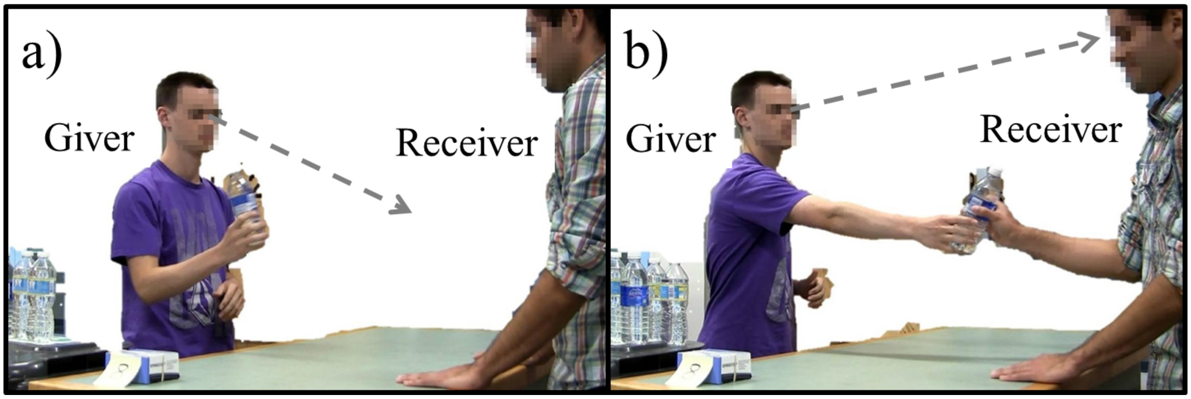

Figure 3.1: Demonstration of two frequently observed gaze behaviours from the HH handover study, Study 1. a) shared attention gaze: the giver looks at the location where the handover will occur, and b) face gaze: the giver looks up at the receiver's face. The five gaze patterns observed in Study 1 consist of different combinations of these two gaze behaviours. (©2014 IEEE/ACM)

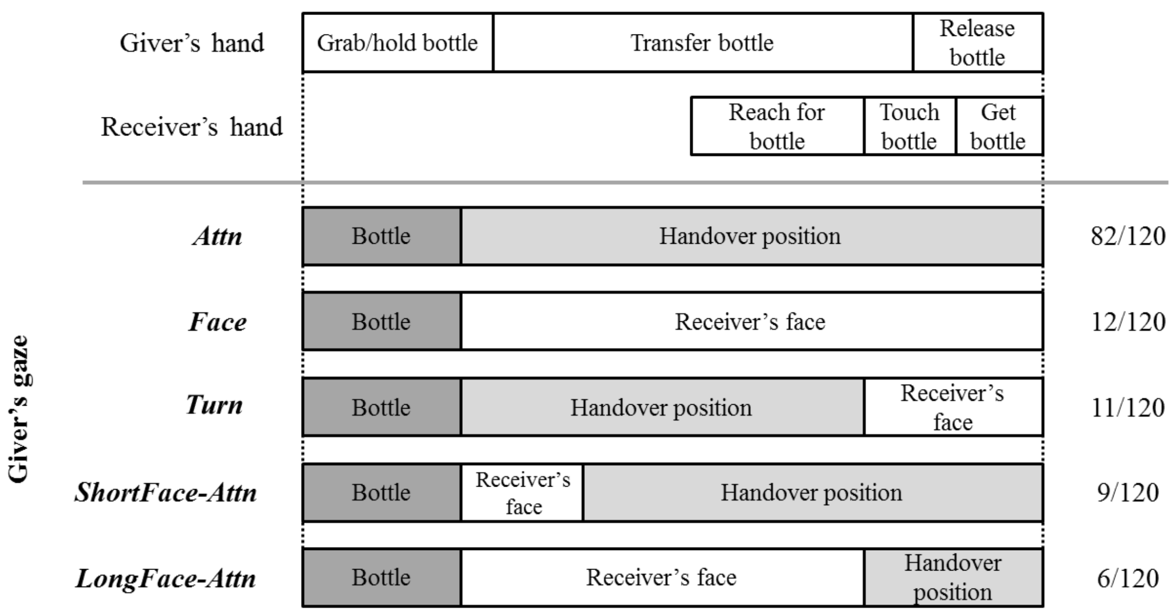

Figure 3.2: Giver's gaze patterns observed from HH handovers. Attn: continual shared attention gaze; Face: continual face gaze; Turn-Taking: long shared attention gaze followed by a short face gaze; ShortFace-Attn: short face gaze followed by a long shared attention gaze; LongFace-Attn: long face gaze followed by a short shared attention gaze. (©2015 Springer)

bottle. This midpoint is approximately where the handover takes place. There is no eye contact between the giver and the receiver throughout this handover gaze pattern (see Figure 3.1a). The gaze towards handover location is labeled shared attention gaze, to indicate the use of gaze to draw the subject's attention to the handover location.

Face Gaze (Face) In some other (10%) handovers, the giver gazes at the receiver's face, perhaps to establish an eye contact, throughout the handover. This gaze behaviour towards the receiver's face is labeled face gaze.

Turn-Taking Gaze (Turn) Some (9%) of the handovers observed showed a slight variation of the shared attention gaze. In addition to gazing at a projected handover location while reaching out, the giver also looked up to make eye contact with the receiver near the end of the handover motion (face gaze), at approximately the time that the receiver made contact with the bottle (see Figure 3.1b).

ShortFace-Attn In 8% of the handovers, the giver looked at the receiver's face and quickly glanced at the bottle when the receiver is about to touch the bottle.

LongFace-Attn In 5% of the cases, the giver glanced at the receiver before but not during handover, and shifted the gaze to the handover location when the receiver is about to touch the bottle.

3.3.3 Discussion

The results of this study suggest that humans use a variety of gaze patterns while handing over an object to another person. In general, the giver tends to shift his/her gaze from the object being handed over to the receiver's face (face gaze), the projected location at which the handover should take place (shared attention gaze), or a combination of the two. The shared attention gaze in the Attn, Turn-Taking, ShortFace-Attn and LongFace-Attn patterns can be interpreted as serving the function of communicating where the physical transfer of the object should happen. The long face gaze in the Face and LongFace-Attn patterns serves a similar function as the face gaze in verbal conversation of regulating a turn [80]; in verbal conversations, the speaker typically ends his/her utterance with a sustained gaze at the listener, signaling willingness to hand over the speaker role, while in this case, to hand over the object. The short face gaze in the Turn-Taking and ShortFace-Attn patterns serves a monitoring function [80], appearing to observe whether the receiver is paying attention to, or is ready for, the transfer of the object.

These results inspire the hypothesis that an implementation of analogous gaze cues for an HR handover could serve similar functions and help the HR dyad interweave the spatiotemporal details of the interaction during the handover. Hence, the gaze patterns observed in this study informed the design of experimental conditions in Study 2. Section 3.5 provides more discussions of the results from the two studies contrasting HHI and HRI.

3.4 Study 2: Impact of Human-Inspired Gaze Cues on First-Time Robot-to-Human Handovers

To examine the impact robot gaze has on first-time human receiver behaviour, the experiments in Study 2 and 6 employed a PR2 humanoid mobile robotic platform (Willow Garage Inc., Menlo Park, CA) with a pan-tilt head and two 7-DOF arms, each with a two-fingered, 1-DOF gripper. The following section (Section 3.4.1) outlines the physical handover cues the PR2 used; Section 3.4.2 describes the three experimental gaze conditions inspired from Study 1 and selected for HR handovers in Study 2; and Section 3.4.3 and 3.4.4 outline the experiment design and technical implementation.

3.4.1 Physical Handover Cues

Based on findings from Basili et al. [15] and Koay and Sisbot [82], the robot in Study 2 was positioned such that it was facing the participant approximately 1 meter away.

The robot executed the handover with its right gripper, as recommended in Koay and Sisbot [82]. At the beginning of each handover, the robot starts its motion at the *grasp position* with its end-effector prepared to grasp a bottle sitting on a table at the robot's right side. When the subject is ready, the end-effector grabs the bottle (marking a *start time*, $t = 0$ of the interaction), then moves the bottle horizontally to a position in front of the robot's centreline (*ready position*). Then the robot moves from the *ready position* forward to the *handover location*. Joint-angle goals of the *grasp position*, *ready position*, and *handover location* are predefined such that when the robot's end-effector is extended, the arm is positioned in a natural pose: the elbow located below the shoulder, and the gripper located below the distal point on the bottle, as shown in Figure 3.3. The *handover location* is designed in accord with the recommendations of previous work [80, 86]. The three locations are constant for all three gaze conditions. While other researchers have proposed handover controllers that adapt to the position of the human's hand, for example Erden et al. [47], a constant *handover location* and gaze cues that vary only during handover events is used in this study.

When the robot's arm reaches the *handover location*, the robot waits for a participant to grasp and pull up on the object. The force the gripper exerts on the bottle is a linear function of the downward gravitational force exerted by the bottle as described by Chan et al. [34]. Thus, as the receiver takes the weight of the bottle, the robot releases its grip (marked as the release time). The PR2's fingertip pressure sensor arrays were used to realize Chan et al.'s handover controller. Finally, after releasing the object, the robot returns to the *grasp position*, ready to grasp and lift the next object.

3.4.2 Experimental Gaze Cues

In this study, the PR2 robot expressed gaze through head orientation. Imai et al. [68] showed that robot head orientation, called head gaze, can be an effective substitute for human-like gaze and that head orientation is interpreted as gaze direction. A single object was used for the handovers to minimize any possible confusion regarding the robot's gaze direction.

This study involved testing of three different gaze patterns in HR handovers, as shown in Figure 3.4.

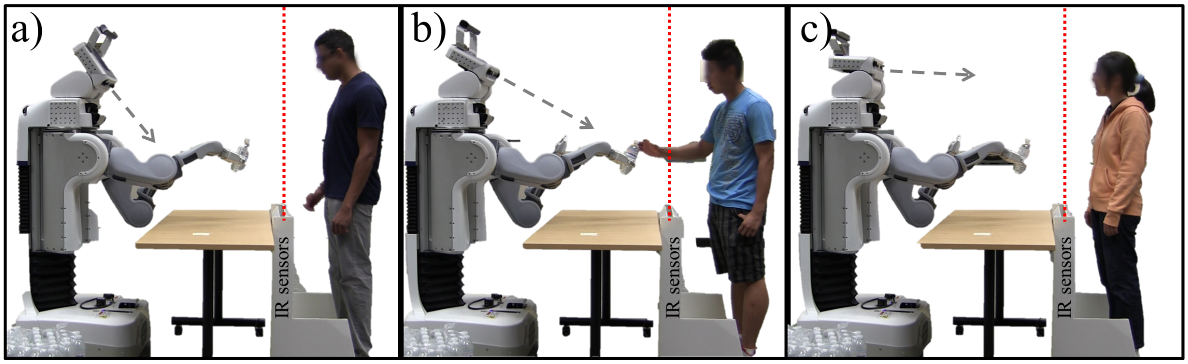

Figure 3.3: Demonstration of the experimental set-up and the three conditions at the *handover location*: a) No Gaze; b) Shared Attention; and c) Turn-Taking. An array of infrared sensors was located at the edge of the table. The red dotted lines represent the location where subject's reach motion is detected. Subjects stood at a specified location marked on the floor. (©2014 IEEE/ACM)

In all conditions, the robot's gaze tracks its end-effector from the *grasp position* to the *ready position* as though the robot is attending to the acquisition of the bottle. When the end-effector arrives at the *ready position*, the robot's head is tilted downwards towards the end-effector. Only when the robot arm transitions between the *ready position* to the *handover location* does the robot transfer its gaze according to the following gaze patterns:

The No Gaze (None) condition is our baseline. The robot head remains looking down towards the ground while the end-effector extends forward for the handover.

The Shared Attention (Attn) gaze condition models the most frequently observed gaze pattern from Study 1. When the robot starts to move from the ready position to the *handover location*, it smoothly transitions its gaze from the bottle to the location in space where the handover will occur, as an implicit cue intended to direct the human's attention towards the projected *handover location*. This condition is necessary to test the hypothesis that shared attention can be established through gaze *during* handovers, and that doing so benefits the handover interaction. Establishing shared gaze at an object or location can serve to direct shared attention (e.g., Imai et al. [67]) and can aid in the successful execution of HR cooperative tasks (e.g., Skantze et al. [141]).

The Turn-Taking (Turn) gaze condition is also derived from the HH handovers in Study 1. When the handover trajectory begins, the robot smoothly transfers its gaze to the *handover location*, as in the Shared Attention condition. Afterwards, as though it is giving a turn to the participant, it shifts its gaze up to the human's face in a quick motion, reaching the final gaze position at approximately the same time that the handover motion completes. Here two hypotheses are tested: that (a) this gaze shift can cue handover timing, and (b) looking at the face can improve the subjective experience of the handover. This type of gaze shift has been shown to be a meaningful HR turn-taking cue [29] and mutual gaze can increase the sense of engagement and naturalness in HRI [89, 140].

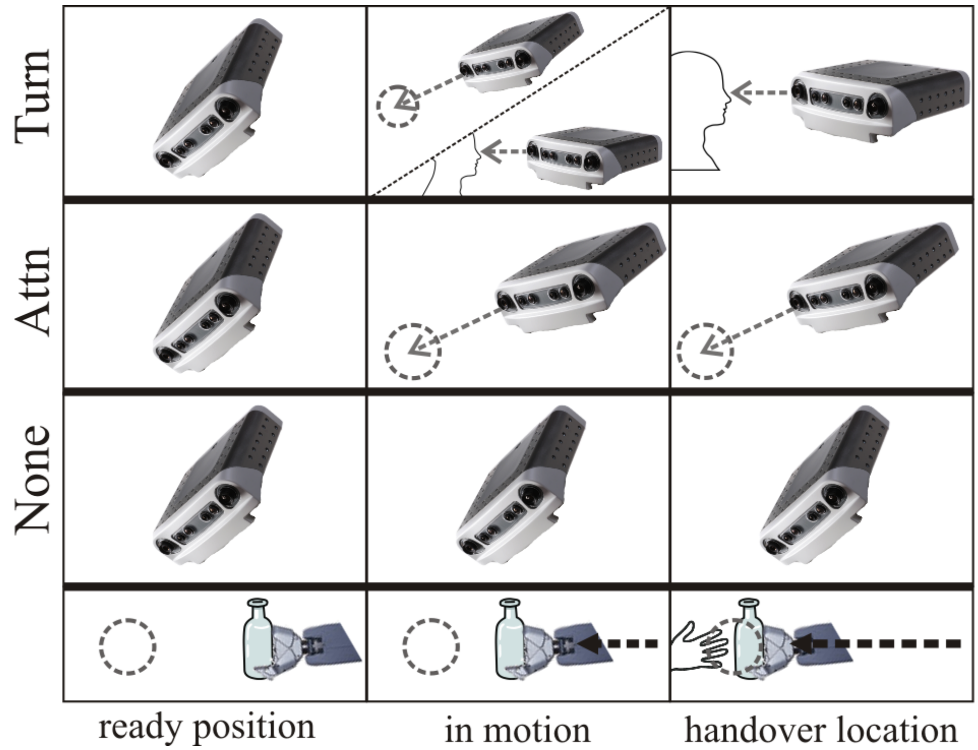

Figure 3.4: Depiction of gaze cues by the head of the PR2 robot: No Gaze (None), Shared Attention Gaze (Attn), and Turn-Taking Gaze (Turn). In the Turn condition, the robot shifts its gaze from the *handover location* to the human's face midway through the handover motion. (©2014 IEEE/ACM)

The following hypotheses are tested using these conditions:

Hypothesis 3.1 *Robot use of gaze during robot-to-human handovers can establish shared attention and improve handover timing.*

Hypothesis 3.2 *Robot shift of gaze from a projected handover location to the recipient's face will cue handover timing.*

Hypothesis 3.3 *Robot gaze towards human recipient's face during a handover can improve the subjective experience of the handover.*

3.4.3 Experimental Procedure

This study involved a paired-comparison HR handover experiment in a controlled room. The study took place on the day of a university orientation event such that many and diverse participants could be rapidly recruited during the public event. The experiment was structured as a balanced incomplete block design ($v = 3$, $b = 96$, $r = 64$, $k = 2$, $\lambda = 32$)[2] to both support rapid trials (maximum 5 minutes)

[2] These variables indicate the structure of a balanced incomplete block design and are necessary for statistical analysis: v =number of treatments (i.e., three conditions – No Gaze, Shared Attention, and Turn-Taking conditions – were tested);

and include only first-time reactions: each participant evaluated one of the three condition pairings. Condition order was randomized and presentation order counterbalanced among trials.

Participants provided informed consent then entered the room where verbal instructions were given (Figure 3.3). They were told to stand at a marked area facing the robot and informed they would participate in a handover interaction. Participants were also told that the robot would pick up the water bottle placed beside it and hand it to them. They were asked to take the bottle from the robot whenever they felt it was the right time to do so. To avoid unintended cueing, during handovers the experimenters sat out of the field of view of participants.

After receiving the first bottle, participants placed the bottle in a box approximately 3 meters behind him/her. This served as a washout between handovers, breaking the participant's focus on the robot and the handover, as was done previously by Cakmak et al. [27]. Participants then returned to the same marked area in front of the robot and participated in a second handover. Participants were permitted to keep the last bottle given to them by the robot.

During each handover, the following events were timestamped: start of robot motion (*start time*), end of robot motion (*end of motion time*), start of release of the robot's gripper (*release time*), and the participant's first reach for the object (*reach time*) as measured by the motion sensor array described in Section 3.4.4.

After the second handover, participants left the room and completed a short questionnaire comparing the two handovers on three subjective metrics: *overall preference*, *naturalness*, and *timing communication*. For each of the following three questions, participants were asked to select either the first or second handover:

1. Which handover did you like better, overall?

2. Which handover seemed more natural?

3. We think that timing is important in human-robot handovers. Which handover made it easier to tell when, exactly, the robot wanted you to take the object?

Participants could also provide additional comments.

3.4.4 Technical Implementation

The control of the PR2 consisted of the Robot Operating System (ROS) [1] with a series of software modules coordinated via the Blackboard architectural design pattern [60] (Figure 3.5). One module controlled the robot's arm and another, its head. The head-control module provided object tracking functionality for bringing the water bottle to the *ready position*, and a smooth, fast gaze transition (average 90 degrees/second) functionality to enable the Shared Attention gaze and Turn-Taking gaze conditions during the handover motion.

b =number of blocks (i.e., observations from a total of 96 participants was analyzed); r =number of replicates (i.e., each condition was tested on a total of 64 participants); k =block size (i.e., each participant saw two conditions); and $\lambda = r(k-1)/(v-1)$.

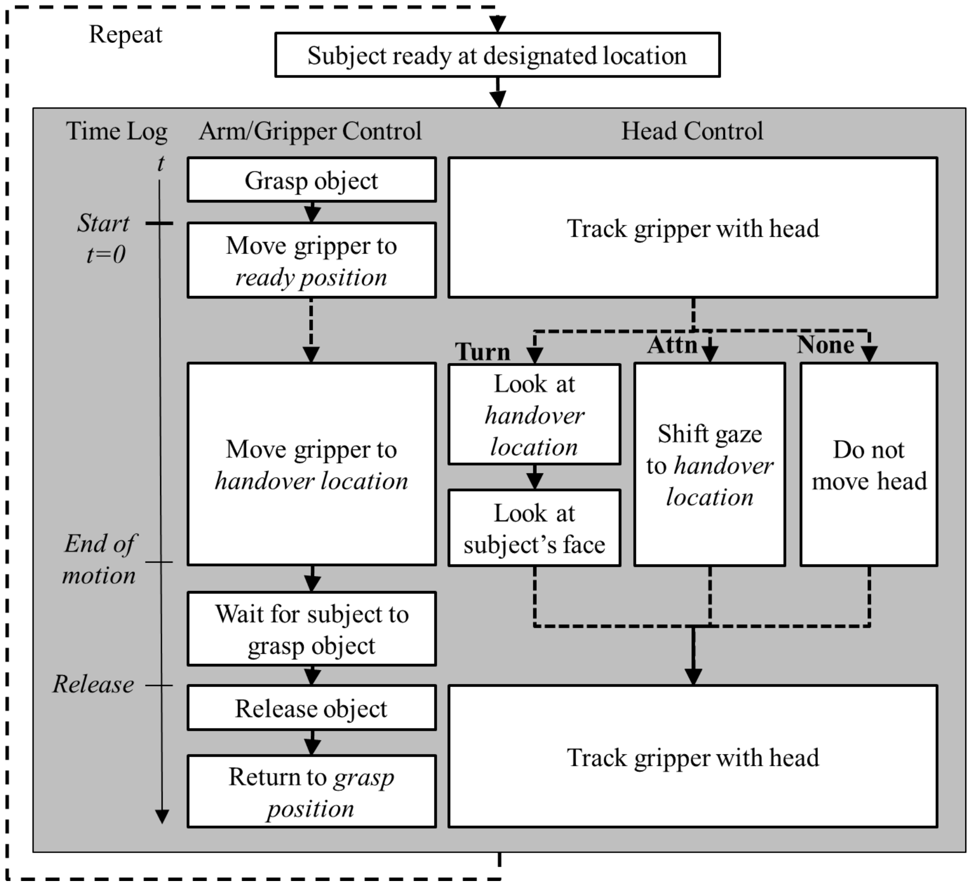

Figure 3.5: Experiment system flow diagram. (©2014 IEEE/ACM)

An independent module logged quantitative measurements of robot's *start time*, *end of motion time*, and *release time*.

An array of three passive infrared motion sensors (SEN-08630, SparkFun Electronics, Boulder, CO) configured as a light curtain was placed at the edge of the table (Figure 3.3), and was used to detect the start of the participant's reach (*reach time*) triggered by the participant's hand crossing the table edge. An Arduino microprocessor relayed the sensor reading to the PC controlling the robot. Sensor readings were logged and time-synchronized with the robot.

3.4.5 Results

A total of 102 volunteers participated in our experiment. Six records were rejected due to the subjects' failure to follow instructions. Therefore, the following analyses include data from 96 participants (63 male, 33 female; age $M = 23, SD = 5.59$). Due to a technical error, *reach time* was not logged in the second handover for five of the participants. This did not affect the analysis of handover timing, since the focus of this study on first-time responses requires *reach time* measures from only the first handovers. No other technical failures occurred, and all handovers were successful (no bottles were dropped).

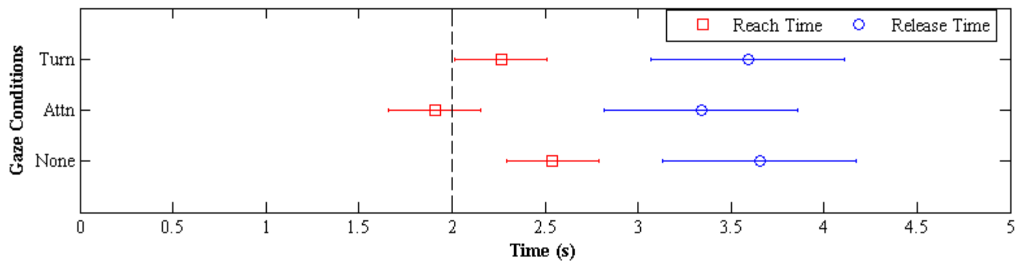

Figure 3.6: HR handover timing results. All times are measured with respect to the robot's start of motion at $t = 0$. The dashed line at 2 seconds indicates the end of robot motion at the *handover location*. *Reach time* indicates the participant's reach toward the proffered object crossing the infrared sensors. Note that in the case of the Attn condition, participants start to reach before the robot has reached *handover location*. The mean *reach time* for the Attn condition is significantly earlier than that of the None condition. The error bars indicate 95% confidence intervals. (©2014 IEEE/ACM)

Handover Timing

Figure 3.6 shows the distribution of three key times: the robot's *end of motion time*, participant's *reach time*, and robot's gripper *release time*. All times are measured relative to *start time*.

The following results are from a one-way Analysis of Variance (ANOVA) conducted on participants' *reach time* across the three conditions. Only the reach time collected during the first of the two handovers performed by each participant is used in this analysis. This is due to a significant learning effect observed between the first and second handover trials ($t(90) = 4.21$, $p < .001$, $d = 0.43$), where *reach time* is earlier in the second handovers. Focusing on only the first handovers also allowed the analysis of truly first-time handover behaviours, which is the focus of this study. The entire robot motion from the *grasp position* to the *handover location* consistently took 2.02 seconds ($SD = 0.01$).

Participants' *reach time* varied across the three gaze conditions ($F(2,93) = 6.49$, $p < .005$) as plotted in Figure 3.6; post-hoc analyses used a Bonferroni correction. Participants reached for the object significantly earlier with Shared Attention ($M = 1.91$, $SD = 0.52$) than with No Gaze ($M = 2.54$, $SD = 0.76$) ($p < .005$). Note that the mean *reach time* for Shared Attention occurs before the robot has stopped moving at the *handover location* (*reach time* $<$ *end of motion time*). These results support Hypothesis 3.1. No significant differences were found between Shared Attention and Turn-Taking ($M = 2.26$, $SD = 0.79$), or between Turn-Taking and No Gaze.

Subjective Experience

Contrasting *overall preference*, *perceived naturalness*, and *timing communication* across the three gaze patterns during handovers involved Durbin's test [46] – analogous to a Friedman test for rank data, but adapted to balanced incomplete block designs – on the aforementioned questionnaire data.

Mann-Whitney U tests revealed no significant gender effects (*overall preference*: $U = 935.0$, $p = .23$, $r = .12$; *naturalness*: $U = 918.5$, $p = .18$, $r = .14$; *timing communication*: $U = 935.5$, $p = .22$, $r = .12$). One-sample Wilcoxon signed rank tests allowed observation of potential bias in selecting

Table 3.1: Ranking of questionnaire results. Each cell represents the number of people who chose the row condition over the column condition. * indicate pairwise comparisons that are significant to $p < .10$ (none were significant to $p < .05$). Note that participants' bias to select the second handover experience was observed regardless of experiment condition.

Overall Preference			
	Turn	Attn	None
Turn-Taking	0	21*	19*
Shared Attention	11	0	17
No Gaze	13	15	0
Naturalness			
	Turn	Attn	None
Turn-Taking	0	20*	19
Shared Attention	12	0	19
No Gaze	13	13	0
Timing Communication			
	Turn	Attn	None
Turn-Taking	0	21*	18
Shared Attention	11	0	19
No Gaze	14	13	0

the first or second handover experience in the questionnaire. Results show a significant bias towards selecting the second handover on the *timing communication* metric ($Z = 2.22$, $p < .05$) and a weak trend to select the second handover on both *overall preference* and *naturalness* metrics ($Z = 1.62$, $p = .11$ and $Z = 1.41$, $p = .16$, respectively). The rank data collected using the questionnaire is insufficient to correct for this bias statistically.

Given this general bias to select the second handover, finding statistical significance to $\alpha = .10$ in questionnaire results is also noteworthy. Hence, observation of trends (results having $p < .10$) is also reported. See Table 3.1 for a summary of the results.

Overall Preference: The results did not show a significant difference in user preference across the three gaze conditions ($T2 = 2.04,\ p = .14$). However, one-tailed pairwise comparisons demonstrate a trend for preference toward Turn-Taking over No Gaze ($p < .10$) and Shared Attention ($p < .10$) conditions.

Naturalness: While the results show no significant difference in perceived naturalness of the handovers across the three gaze conditions ($T2 = 1.82,\ p = .17$), participants tended to choose Turn-Taking as more natural than Shared Attention ($p < .10$) but not over the No Gaze condition.

Timing Communication: No significant differences were found in the perceived communication of timing across the gaze conditions ($T2 = 1.65,\ p = .20$). However, participants tended to choose Turn-Taking over Shared Attention ($p < .10$), but not over No Gaze, as easiest to communicate handover timing.

In total, 59% of all participants provided additional comments (optional) on the questionnaire. Twelve subjects who experienced the Turn-Taking condition explicitly used words such as "head motion", "eye contact" or "looking at me" and expressed the condition in a positive light (e.g., P90 compared No Gaze with Turn-Taking: "During second handover [Turn-Taking], robot made eye contact, which made it easier to tell when the bottle should be taken."; P10 compared Shared Attention and Turn-Taking: "I liked it when robot looked at me. That confirms it's good to take."). However, another twelve subjects expressed that they did not notice any difference between the conditions.

Video Analysis of Participants' Reaching Behaviour

It was unexpected to find that the Turn-Taking condition did not yield a significantly earlier *reach time* than the No Gaze condition, while the Shared Attention condition did. This led to the hypothesis that there are subtle yet important differences in participants' behavioural responses to different robot gaze cues during HR handovers. In order to understand this result, a frame-by-frame video analysis of the recordings of the 97 participants[3] (Turn-Taking: 33, Shared Attention: 32, No Gaze: 32) was carried out using ELAN [95] with a special attention to the participant's reaching (hand) behaviour and gaze behaviour. Two coders annotated the videos with partial overlay. The inter-coder reliability was evaluated through Intra-class Correlation Coefficient (ICC), and showed substantial agreement (all $ICCs > .80,\ p < .01$).

The coders identified a participant's *start motion* (the time when the participant's hand starts to move to the bottle, as a secondary measure of the reach time) and *touch bottle* times (the time when the participant's hand touches the bottle). Figure 3.7 illustrates the distribution of these measures with respect to the robot's behaviours. A one-way ANOVA indicates that *participant start motion* varies across the conditions ($F(2,94) = 4.94,\ p < .01$); post-hoc analysis using Bonferroni correction indicates that *participant start motion* is significantly earlier with Attn ($M = 3.29,\ SD = 0.50$) than with Turn-Taking ($M = 3.70,\ SD = 0.62$) and No Gaze ($M = 3.68,\ SD = 0.65$), but no significant difference was found between Turn-Taking and No Gaze. This is consistent with the findings reported in Section 3.4.5 that relied on reach time measured with IR sensors.

Video Analysis of Participants' Gaze Behaviour

The video analysis suggests that 92% of the participants looked at the robot's head at least once during the handovers. These behaviours are always observed as quick glances, rather than long stares. The distribution of the glances with respect to the robot's behaviours is shown in Figure 3.7.

For all conditions, the glances are clustered around the time when the robot gripper is at the ready position. In the Turn-Taking condition, a cluster of glances occurs after the robot starts shifting its gaze to the participant's face (after $t = 3.6$ s). While this cluster of glances could be due to the robot's gaze, it is also possible that the trigger for these glances is internal (e.g., the participants wanted to make

[3] One of the six participants rejected in the above analysis was still valid for the purposes of the video analysis. Hence, the video analysis includes 97, instead of 96 participants' data.

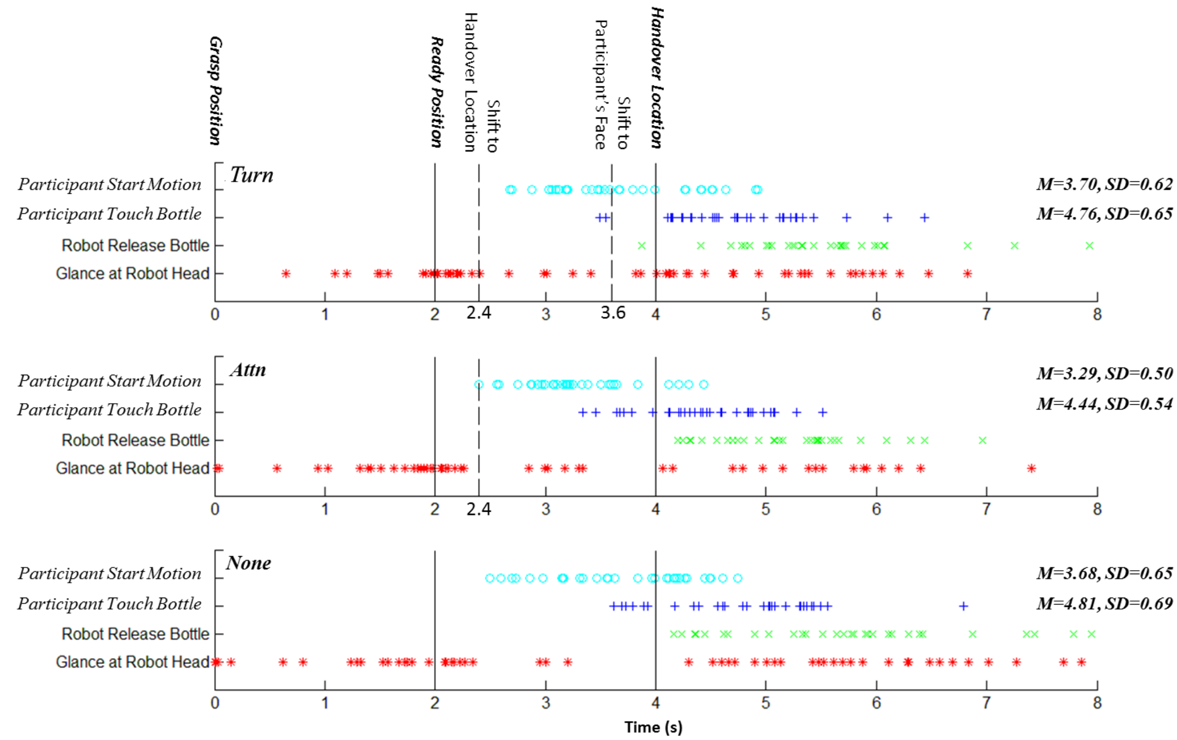

Figure 3.7: Distribution of participants' reaching behaviour (*participant start motion* and *participant touch bottle*) and gaze behaviour (*glance at robot head*) with respect to the robot's gripper behaviour (marked by vertical solid lines) and gaze behaviour (marked by vertical dashed lines). (©2014 IEEE)

sure it was an appropriate time to touch the bottle) since the glances are also close to the *participant touch bottle* time. Result from a Chi-squared test on the number of glances between $t = 3.6$ s and *participant touch bottle* time indicates that participants look at the robot head more in the Turn-Taking condition (12 out of 33) than in the Shared Attention (3 out of 32) and No Gaze (7 out of 32) conditions ($X^2(2) = 6.77,\ p < .05$). This result suggests that the robot's face gaze in the Turn-Taking condition catches some participants' attention and elicits them to make eye contact with the robot.

Could this be why the participants reach for the object later in Turn-Taking than in Shared Attention condition? Figure 3.7 shows that many participants have already started to reach for the bottle before the robot starts shifting its gaze to their face (before $t = 3.6$ s). In the Turn-Taking condition, 58% of the participants (19/33) started to reach before $t = 3.6$ s, so their *participant start motion* time can only be affected by the robot gripper motion and shared attention gaze; 15% of the participants (5/33) started to reach during the transition period (during $t = 3.6$ to 4.0 s), while 27% participants (9/33) started to reach after $t = 4.0$ s, which could also be triggered by the robot's face gaze. Six of the nine participants who started to reach after the face gaze (after $t = 4.0$ s) made eye contact with the robot before starting to reach for the bottle. Hence, *participant start motion time* from these individuals, to a considerable extent, was delayed. Therefore, it is possible that the late appearance of the robot's face gaze contributes to the delayed average *participant start motion time* and average *reach time* in the Turn-Taking condition, compared with the Shared Attention condition. One can test this hypothesis by changing the timing of the robot's gaze toward the participants' face, more specifically, introducing the robot's gaze to the receiver's face earlier in the handover interaction, and measuring whether this modification elicits an earlier *reach time* from participants.

Since a cluster of glances from participants was observed around the time when the robot was at the *ready position* regardless of conditions, it can be suspected that a face gaze immediately after the *ready position* would be more likely to be noticed by the participants. Interestingly, when referring to results from Study 1 (Section 3.3), this proposed gaze pattern is effectively the Face pattern.

3.5 Discussion

Building on previous work that studied communication of intent to handover using gaze, the studies presented in this chapter delved into the use of gaze *during* a handover (i.e., after the intent to handover is already communicated and while the handover is taking place). A handover interaction typically involves well-defined role assignments (a giver and a receiver) and a clear sequence of actions that must take place (the giver grabs and transports the object, and then the receiver receives the object from the giver). However, even in such a well-defined interaction between two agents, the details of *when* and *where* must be interwoven dynamically in order for the interaction to be successful. As an effort to understand what nonverbal cues may help the interweaving process, Study 1 was conducted to identify gaze patterns human givers use when handing over an object to another person. Results of this study not only provide five different gaze patterns human givers use in an HHI handover, but it also suggests that the observed patterns more often than not involve gazing at the projected *handover location*. Following Study 1, Study 2 explored the impact of robot gaze on HR handover timing and user perception of the handover experience. Results show that participants reached for the proffered object significantly earlier when the robot performed a shared attention gaze at the projected *handover location*. In fact, participants reached for the object even before the robot had arrived and stopped at the *handover location* (a mean of 0.11 seconds before the *end of motion time*). This is in contrast to the No Gaze condition where the mean *reach time* is 0.52 seconds after the robot's *end of motion*.

In Study 2, participants were explicitly told that the robot would be handing over objects to them and that they were to take the object from the robot. In addition to this foreknowledge, the robot used highly contrasting poses between the *ready position* and the *handover location* which, according to Cakmak et al. [27], makes the robot's intent to hand over the bottle very clear. Hence, it is unlikely that the observed difference in timing between the gaze conditions is due to uncertainties in understanding the robot's handover intent. Rather, the results suggest that the robot's gaze at the projected *handover location* supplements the communicated intent with implicit information on *where* the handover should take place. This may be helping to establish shared attention on the *handover location* even before the robot arrives there, naturally allowing participants to respond and meet the robot at the location earlier than when such a cue is absent. Thus, the result best supports an increase of fluidity in the execution of the handover as it takes place.

However, the role of mutual gaze used in the Turn-Taking condition required further investigation.

At the beginning of the robot's handover motion, the robot expresses the same locational, shared attention gaze in both the Shared Attention and Turn-Taking conditions. Hence, it was surprising to find that the *reach time* of the Turn-Taking condition is not significantly earlier than that of the No Gaze condition. This finding was supported even when the similar measure, *participant start time* manually collected from the video analysis, was used to verify this result.

The following two hypotheses were tested that involve the Turn-Taking condition: that the Turn-Taking gaze would cue handover timing, and that looking at the participant's face would improve the subjective experience of handover. While there is a trend suggesting that the robot's gaze directed at the face improves the subjective experience of the handover, it remained uncertain whether the shared attention gaze – instead of the face gaze – in the Turn-Taking condition is the main influencer of the handover *timing*. It may be that the face gaze employed in this study served the function of acknowledgment rather than the intended function of giving a turn to the participant. There are also behavioural differences in participants' reaction to Shared Attention and Turn-Taking conditions. According to the video analysis of participants' gaze behaviour, a number of participants waited for the robot to gaze at the subjects' face before starting to reach across the table.

This raised questions about how human *reach time* is affected by the timing of the robot's gaze. How much would varying the robot's gaze timing affect human *reach time*? Is the timing of the robot's gaze a more dominant cue than the location the robot is gazing at? That is, would a robot that shifts its gaze from the object directly to the person's face during handovers have the same effect as the Shared Attention condition?

Also, it is important to note that results from Study 2 alone are representative of first-time participant responses only, where novelty effects may have motivated the participants to observe the robot more carefully than they would if they were more familiar with the robot. Unsurprisingly, there is a significant training effect in the *reach time* data, as well as a bias toward describing the second handover experience more favourably in the questionnaire regardless of the condition experienced. Some of the participants' comments suggest that in certain cases, people did not pay attention to the head of the robot at all. Indeed, it is suspected that in many HH handover scenarios, especially those that are repetitive or trained (e.g., passing a baton in a sprint relay race), people do not use gaze cues at all and still succeed in object handover. Thus, it was hypothesized that robot gaze cues might not have the same effect on trained or familiarized users.

This is effectively the Face condition in Study 6 (Appendix A.1), a follow-up study to Study 2 that contrasts non-naïve participant responses to handover events between the Attn, Face, and LongFace-Attn gaze patterns (Figure 3.2). Results from Study 6 reaffirm findings from Study 2 that gaze significantly affects the timing at which participants reach for the object offered by the robot. It demonstrates that this effect holds true even after the training effect of HR handovers can no longer be observed through repeated handover trials. This implies that the even a subtle gesture of gaze used by a robot during a handover event helps communicate the non-trivial information about the details of the interaction (in particular, the information about *when* the handover can take place) such that the resulting HRI is more fluent and efficient. However, as presented in Appendix A.1, the LongFace-Attn condition, which is designed to be a variant of the Face condition, did not perform as well as the Face condition itself. This result is similar to the findings on *reach time* in Study 2. Given that LongFace-Attn is a variant of Face, just like how Turn-Taking is a variant of Attn, it was expected that the two conditions would elicit similar *reach times*. This indicates that, although gaze does affect the timing involved in a handover, the nature of its effect is not a simple function of the type and timing of gaze implemented on the robot. Further investigation is necessary to understand this relationship fully.

Would there be changes in participants' reach direction if the robot gazed at a different location? Without conducting an additional study, it is difficult to tell, with accuracy, if and when shared attention is established with the participant. A separate experiment with a gaze tracking device would help answer these questions, but this is beyond the scope of this study.

Although the earlier *reach time* of participants in a handover may seem more similar to natural, unscripted HH handovers, this may not necessarily be desired in some HR handover situations. Depending on the handover controller implemented on a robot, handover timing may need to be controlled such that people naturally grab the object only when it is safe to do so. Many of the handover controllers that modulate the release time of the object are built for cases where the robot's gripper is already at the handover location before people grab the object. A situation where the object is grabbed before the robot is ready to release the object could lead people to pull hard on the object, possibly damaging or dropping the object, or resulting in a negative perception of the robot.

3.6 Conclusion

Many previous studies provide evidence that a robot's use of nonverbal cues can affect human behaviours. The two-part investigation outlined in this chapter extends our understanding of the phenomena to cases where the cues are given *during* an atomic interaction. In this chapter, robot-to-human handovers are used as a means to investigate whether subtle nonverbal cues used by a robot can effectively influence details of human action for a successful, fluent HRI. Within the handover context, the detail in question is the timing of the person's reach to receive the offered object.

There were three main contributions to the understanding of HR handovers from the work outlined in this chapter. The results from Study 1 contribute to the understanding of the type of gaze that humans tend to use during HH handovers. Results from Study 1 identified five different types of gaze used during HH handovers. The most frequently observed gaze pattern (Attn) from the giver was from the object to be transferred to the projected location where the handover should take place. Results of Study 2 demonstrate how a robot's use of human-inspired gaze expressions during HR handovers can affect the timing of the handover event in *first-time* participant responses. The study provides empirical evidence that a human-inspired gaze pattern (Attn) implemented on a robot can elicit a human receiver to reach for and retrieve the proffered object earlier than when no gaze (None) cues were provided. Study 6, a follow-up study presented in Appendix A.1, further extends this finding and demonstrates that such an effect persists even after repetitions of HR handovers has taken place. These studies contribute to a better understanding of the effect of gaze on HR handovers.

In a broader context, findings from these studies support the hypothesis that a robot's use of nonverbal cues *during* an interaction can be used to interweave details of the interaction. That is, the studies provide evidence that nonverbal robot cues can not only be used to communicate its "intent" prior to taking an action or to communicate details about its action independent of the user, but also elicit desired behaviours from the user *while* the agents are already engaged in the interaction.

However, the robot's gaze cues used in this study were designed as an open loop interactive behaviour. The robot's use of gaze was not designed into a feedback loop to respond to the person's gaze. Further investigation is required to explore the full extent to which a robot's use of nonverbal cues can help *negotiate* details of an interaction while the interaction is taking place. Investigations outlined in the next chapter delve into this domain by focusing on a proactive and responsive robot nonverbal behaviour designed to interactively negotiate for a solution to unforeseen resource conflicts.

Chapter 4

Development of Negotiative Interaction for Nonverbal Resolution of Human-Robot Conflicts

4.1 Introduction

In the previous chapter, a robot's use of gaze cues during robot-to-human handovers was discussed as a means of affecting a human's behavioural response with regards to *when* to reach to receive the object. In the experimental handover context, there is no contention or ambiguity regarding the outcome of the interaction (e.g., to successfully transfer the object to the person).

However, in many everyday activities, it is common for interacting individuals to experience conflicts regarding priority to access shared resources. When the question of who should get access to the resource first is uncertain, people often nonverbally communicate with each other to negotiate a solution. For example, when two people reach for the same object at the same time, they often manage to resolve the conflict by one yielding the object to the other or by claiming the object before the other obtains it. This chapter considers: What should a robot be programmed to do when such conflicts occur in an HRI? Answers to this question can vary depending on the context of the situation at hand. Investigations presented in this and the next chapter focus on a robot's use of interactive negotiation behaviours as a practical solution to the problem.

Whereas the previous chapter investigated the efficacy of nonverbal cues for interweaving subplans in HRI with predefined roles and outcomes, studies presented in this chapter focus on an interactive scenario when the outcome itself – more specifically, the question of *who* should access the shared resource first – is not predefined and requires negotiation. As Chapter 1 introduced, negotiation refers to the bidirectional interweaving of subplans. As is clear from the dictionary definition, "Discussion aimed at reaching an agreement" [121], *negotiation* is a process that requires bidirectional communication between the parties involved. Within the context of this thesis, only nonverbal communication and dialogue is relevant.

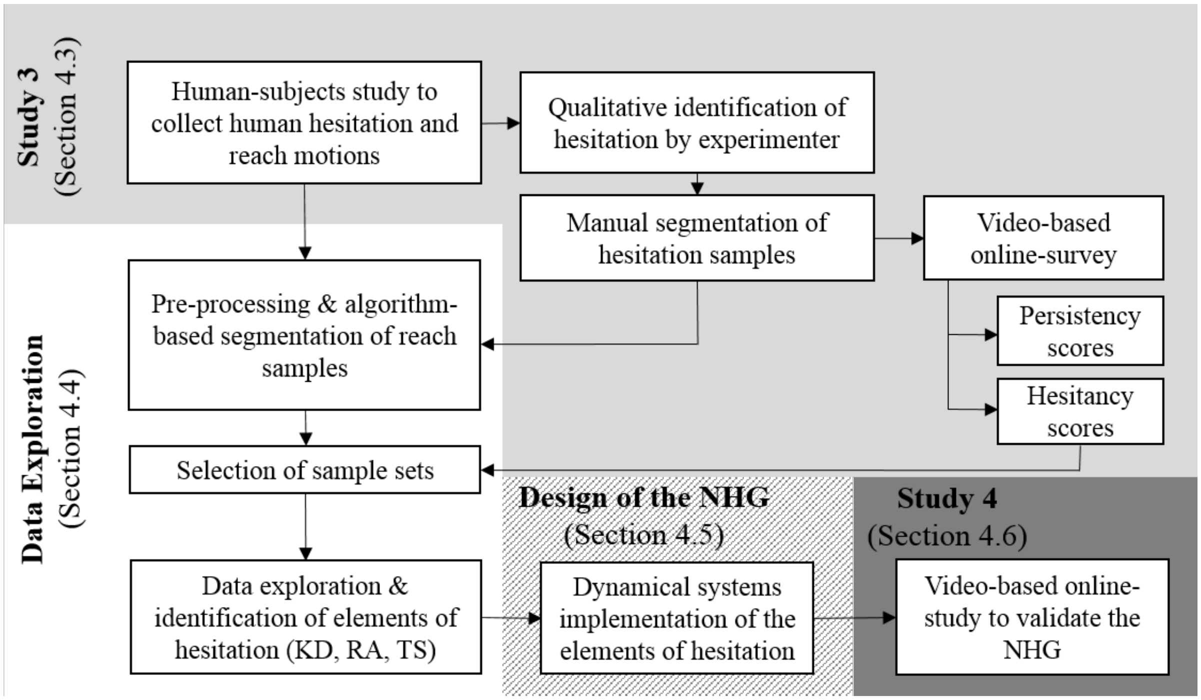

Figure 4.1: Overview of the process taken to devise the Negotiative Hesitation Generator (NHG) that can produce human-inspired, artificial hesitation motions for robotic manipulators. The arrows represent the direction of information from one process that is used to inform the next.

To investigate nonverbal negotiations between a human and a robot, this and the following chapter focus on the nonverbal gesture of hesitation. As Section 4.2 presents in more detail, hesitation behaviours in HHI are used to communicate uncertainties between people in their everyday life. They have been studied as a type of behavioural indicator in psychology and linguistics, and increasingly are being used within the HRI community as a behavioural measure to examine user confusion and uncertainties.

Presented in this chapter are two studies the author conducted (Studies 3 and 4) to investigate the possibility of using hesitations as a means to dynamically negotiate conflicts in HR collaboration. Figure 4.1 provides an overview of the process described in this chapter.

In Study 3 (Section 4.3), the author conducted a two-part experiment with the aim of observing negotiative hesitation gestures humans naturally exhibit in HHI. This study helped confirm and better understand human use of hesitations as an immediate and nonverbal means to negotiate imminent resource conflicts. In the first part of Study 3, an HH dyad experiment was designed to elicit numerous instances of hesitation gestures in a task that involves natural sharing of resources. In the second part of the study, video recordings from the HHI experiment was used to conduct an online survey to measure the level of hesitancy and persistency third party human observers can perceive from the collected human gestures.

With the results and collected data from Study 3, the author subsequently explored the collected and labelled time series data from the study to identify key features that are characteristic of naturally elicited human hesitations (Section 4.4). This data exploration process led to the discovery that the subset of hesitation gestures of interest in this thesis share a common trajectory pattern. This pattern, referred to as hesitation loops (Section 4.4.6), can be reproduced with an implementation of three trajectory-related elements. These elements shaped the design of an artificial negotiative hesitation trajectory generator for a robotic system called the Negotiative Hesitation Generator (NHG) (Section 4.5).

An online, video-based perception study, Study 4, was then conducted to validate the NHG. It was implemented on a 7-DOF articulated robotic platform as a means to generate robotic resource conflict responses that are perceived as humanlike and expressive of hesitation (Section 4.6). Results from this study provide an empirical support that the NHG-generated trajectories are capable of producing robot behaviours that are perceived to be significantly more hesitant, persistent, animate, and anthropomorphic than a conflict response consisting of smoothly stopping and pausing.

Taken together, the work presented in this chapter contributes to a better understanding of kinematic features of hesitation behaviours. It also provides an empirically validated means of generating robot conflict responses that are perceived as hesitations by human observers. Subsequently, human participants were invited to interact with the devised robotic system for an in-person experiment (Study 5) presented in Chapter 5. This helped investigate whether an interacting HR pair can negotiate and resolve conflicts pertaining to shared resources using purely nonverbal means, thereby dynamically determining the outcome of the conflict through interaction.

4.2 Background

Hesitations are social signals often observed as a form of disfluencies and uncertainties in human and animal behaviours. In Psychology, Doob [44] defined hesitation in the temporal domain, as the time that elapses between a stimulus and a corresponding response to the stimulus. Researchers in linguistics refer to filled pauses or disfluencies in speech, such as *uhs* and *ums*, as hesitations [39, 93, 97, 160]. Rober, a clinical psychologist and family therapist, describes a variety of verbal and nonverbal behaviours – such as prolonged silence, a glance, a sigh, and a pause in the flow of the conversation – as hesitations. He suggests that such verbal hesitations are related to the communicator's wavering willingness to speak about a matter due to an internal conflict that has not been verbalized [132]. In Human-Computer Interaction (HCI), hesitation has been used as a measure of a computer user's experienced difficulty with designed interfaces. Reeder and Maxion [129], for example, developed a simple algorithm that detects abnormally long pauses between inputs from keyboard and mouse to detect user hesitations during use of a computer interface. This automated detection of hesitation in user behaviour reduced up to 96% of an experimenter's time required for computer interface analysis.

Given the usefulness of understanding and detecting hesitation behaviours, linguists have used machine learning approaches to detect hesitations in speech such as disfluency and filled pauses [139, 156, 157, 164]. Being able to identify and recognize hesitations is increasingly being recognized as valuable in the study of HRI. Analogous to the HCI studies mentioned above, such natural human behaviours are immediate indicators of interrupted flow and increased uncertainty in an interaction [10, 32, 125]. Researchers have employed the notion of hesitation behaviours as the duration of pause between actions in HRI [12, 13, 76, 96, 142, 142].

For instance, Bartneck et al. [12] conducted a study in which the participants were asked to turn off a robot while the robot begged them not to. The researchers measured the amount of time it took the participant to shut down the robot as an expression of the participant's hesitation in following through with the task. In another HRI study, human participants and a humanoid robot engaged in a verbal interaction as part of an experiment on verbal turn-taking [142]. The robot used either a passive or active mode of interrupting a conversation with the participant in an attempt to give or take speaking turns between the agents. In the active condition, the robot actively interrupted the participants to take over the turn to speak, and it ignored the participants' analogous interruptions to do the same. In the passive condition, the robot hesitated (paused its actions) as a means to yield the turn to the participant when s/he interrupted the robot. In its yielding of the turn to speak, the robot was programmed to hesitate by pausing its actions for a given period until the participant either seized the turn from the robot or yielded to the robot. Results from this study demonstrate that people speak more when the robot is passive and yielding its turns than when the robot is active and interrupting the user.

In addition to the hesitations measured and expressed as pauses, literature in sports, zoology, and HRI provides evidence that hesitation also includes certain kinematic behaviours in humans and animals [43, 104, 114, 116, 165]. However, only a few scholars have studied kinematic hesitation gestures within and outside the context of HRI [43, 103, 104, 114, 116, 148, 165]. This is in part due to the added complexity of processing kinematic hesitation signals. Kinematic hesitations have high dimensionality and provide a noisier signal than auditory pauses or the durations measured between keyboard/mouse inputs used in HCI. One of the few studies that focused on kinematic hesitation gestures in HRI is the author's previous work [103, 104].

Within the framework of HR collaboration involving reaching motions, the author asserts that hesitations can be exhibited as a response to interruptions of the motions. The authors of Moon et al. [104] analyzed acceleration trajectories of the wrist in a subset of human hesitation gestures, called R-type hesitations, and formulated characteristic features from the trajectories as a series of cubic splines in the acceleration domain called the Acceleration-based Hesitation Profile (AHP). R-type hesitations involve the hesitating agent to immediately yield to the other agent upon encountering a resource conflict. Trajectories of these hesitations are characterized as having three key points that represent the temporal location and amplitude of acceleration at launch, braking, and yielding phases of a hesitation motion. Figure 4.2 illustrates sample R-type trajectories. Using the AHP, the authors generated human-inspired R-type hesitation trajectories for a robot. They used an analytic approach to extract characteristic features of the most basic type of human hesitation trajectories that, when used to generate artificial hesitations for HR collaboration, are perceived as hesitations and clearly distinctive from other similar robot motions. Findings from Moon et al. [104] demonstrate that human observers correctly recognize the artificially generated robot-hesitation having the AHP as hesitations and can distinguish them from other similar robot motions. However, the type of hesitation behaviour (R-type) investigated in Moon et al. [104] was limited to the simplest form of hesitation behaviour, and one that is not meant to elicit nonverbal dialogue among the interacting HR pair.[1]

[1] By definition, R-type hesitations are comprised of reacting to the conflict with a jerky halting motion followed by an

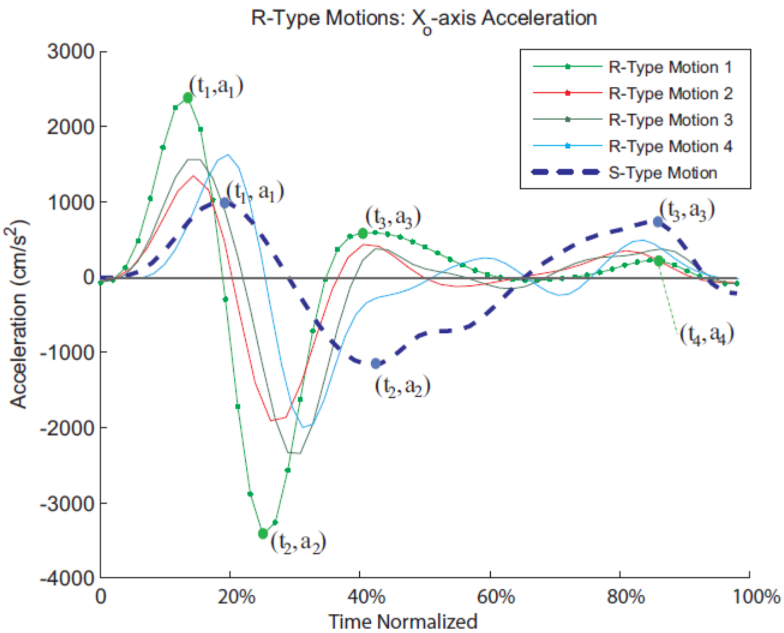

Figure 4.2: Demonstration of the AHP that characterizes R-type hesitations. Human R-type hesitation motions shown here are captured at the wrist and characterized in acceleration space. R-type hesitations involve the hesitating agent to immediately yield to the other agent upon encountering a resource conflict. The reference S-type motion (a Successful reach-retract motion) is characterized by a relatively symmetrical accelerate-decelerate, and then accelerate back to the starting point, motion. R-type hesitation trajectories (samples 1-4 shown) were characterized as four splines that connect the three key points – t_1 and a_1, t_2 and a_2, and t_3 and a_3 – and the start and end points of the motion. The three key points represent the temporal location and amplitude of acceleration at launch, braking, and yielding phases of a hesitation motion. The ratios between these variables in acceleration space have been determined from observed R-type human hesitation trajectories.

In developing a robotic system that can use hesitation gestures to resolve resource conflicts with humans interactively, it is important to equip the robot with a vocabulary of hesitation gestures that proactively engages the user in a nonverbal dialogue. One method of enriching the vocabulary of hesitation gestures for a robot is to understand and mimic how humans negotiate a solution to resource conflicts using hesitations (*negotiative hesitations*). Once a set of trajectory features is identified from negotiative human hesitations, these features can be used to generate analogous robot trajectories for HR negotiation of resource conflicts. The studies presented in this chapter contribute to this end by collecting negotiative hesitation motion samples from a human-subjects experiment and examining trajectory features that can be implemented in an HR collaboration context.

immediate yielding of the resource before any dialogue can take place.

4.2.1 Persistency

One of the elements of negotiative hesitations that distinguish them from R-type hesitations is the expressed level of persistent interest and prolonged need for a resolution of the conflict at hand. By persistency, the author refers to the quality that is expressed in the persistent state of an agent or the act of persisting in something (i.e., how persistent is an agent?). Here, the author refers to the dictionary definition of persistence as "the fact of continuing in an opinion or course of action in spite of difficulty or opposition," and "the continued or prolonged existence of something" [122]. While *persistence* is synonymous to nouns such as perseverance, tenacity, endurance, and tirelessness, the concept of persistency discussed here pertains to the spectrum of persistence conveyed in a communicative signal – in particular, in hesitations.

Persistency expressed as part of human behaviour has been studied in a variety of contexts. In child development literature, in particular, persistence is accepted as a means to determine, measure, and analyze intentional communication, especially in infant preverbal communication signals [51, 53, 85]. In studies involving human infants, persistency is measured as repeating of or augmenting a version of a communicative gesture when the infant's initial attempt to communicate with the recipient fails [51, 53]. Persistence is also one of two criteria for an infant's gesture to be considered an intentional communication [85].[2] Scholars in child development suggest that intentional communication – evidenced by persistence, among other behaviours – is a means for the communicator to manipulate the interlocutor rather than a mere attempt to engage in a conversation [54]. This suggests that negotiative hesitations by a robot that – in contrast to R-type hesitations – express a higher level of persistency are likely to be perceived by humans as intentional communication, possibly leading to the emergence of nonverbal dialogues.

Persistence in communicative gestures is also closely related to negotiation behaviours. Golinkoff [53] states that persistence is seen as a proactive behaviour that is a necessary component of ongoing negotiation between mother and infant. Golinkoff [53] also suggests that even at infancy, humans negotiate with their mothers. An infant's persistent communication with the mother elicits this type of negotiation, which allows the infant to use communicative cues to manipulate its interlocutor/mother. In this context of infant-adult negotiation, what is being communicated by an infant is the infant's desires and intents. The negotiation between them takes place for the purpose of the infant influencing the decision-making of the interlocutor (e.g., a child requesting for an adult to bring a toy that is out of reach).

In this work, the author investigates persistency as it pertains to negotiative hesitations. A level of persistency is assumed to be expressed in negotiative hesitations since negotiative hesitations require the hesitating agent to have a persistent interest in resolving a conflict. Therefore, the studies presented in this chapter explore persistency as a concept related to, but not trivially correlated with, hesitancy expressed in negotiative hesitations.

[2] Elaboration of the original communication signal upon failure to communicate is the other criteria.

4.2.2 Social Signal Processing

Scholars who attempt to generate readily recognizable social robot behaviours often employ empirical methods that involve human subjects. For example, Lim et al. selected four parameters (speed, intensity, regularity, and extent), and varied them across three modes of expressing emotion (gesture, verbal expression, and music) [88]. Their aim was to determine whether these four parameters could be used to generate artificial expression of emotion across the three modes. The studies presented in this chapter are similar in terms of studying the qualitative content of communicative, time-series signals. To validate the usefulness of parameters in emotion expression, Lim et al. [88] measured the percentage of human subjects who recognized the designed, artificial emotional expression correctly. However, such percentage measures for validating the efficacy of parameters studied are not sufficiently informative to allow one to create new signals of the same quality.

Studies in Social Signal Processing (SSP), or socially aware computing, share the similar goal of analyzing time-series signals from human subjects using algorithms. In SSP, sensed signals from an individual, such as audio or visually observable behaviours, are used to infer otherwise hard-to-measure internal states of the individual computationally [136, 160]. Some of the popular work in SSP include inferring individuals' states and actions during group discussions (e.g., dominance or interest level) with dynamic Bayesian networks [109] or layered Hidden Markov Models [117] on the individuals' speaking energy and body language.

While the primary goal of SSP is to perceive measurable social signals to infer internal states from humans, the design of interactive, social robots has the added challenge of needing to respond to the inferred human states with behaviours that express appropriate social signals easily understood by humans. Hence, a number of well-performing approaches in SSP are not directly transferable to HRI.[3] It is not always possible for a robot to exhibit the mix of signals needed to produce a social signal understood by human observers (e.g., the robot may not have a speaker to verbalize a message, or a finger to point at an object) [160]. Therefore, in contrast to employing algorithms that may require the use of numerous features from a social signal, it is advantageous to identify the minimum number of features that a robot can generate to convey desired social signals to human observers. In exploring the human motion data collected from Study 3 (Section 4.4.3), the author attempted to find such a minimum set of features that can be used to generate artificial, negotiative hesitation gestures for a robot.

4.3 Study 3: Observing Hesitations in Human-Human Dyads

This section describes a human-subjects experiment conducted to collect a large set of naturally exhibited, negotiative human hesitations. Section 4.3.1 describes the experiment involving a motion capture system (VICON, Vicon Motion Systems Ltd., Oxford, UK, [159]) and the experimental task that was designed to trigger multiple instances of conflicts about shared resources between two human

[3] For example, it is often suggested that combining different classifiers, each of which classifies different aspect of the same problem, performs well and is recommended in SSP. While this recommendation is useful in classifying even complex social signals, generating the mix of signals for the robot to exhibit in order to send the same social signal to the human observers remains a challenge in social robotics.

participants. As presented in Section 4.3.1, the collected set of human behaviours was then labelled through a video-based online study using the Amazon Mechanical Turk platform [5]. This online study was designed to collect the varying degrees of perceived *persistency* and *hesitancy* people perceive from the recorded instances of hesitations. These samples of hesitation gestures were then used to explore the common features that are present in the trajectories. The exploration of the trajectory data is presented in the following section (Section 4.4).

4.3.1 Experimental Procedure

The author conducted an HH paired experiment in a controlled environment. Paired participants were brought into a laboratory and introduced to their partners. Eight pairs (N=16, 3 male-male, 2 female-female, 3 female-male pairs) of volunteers participated in this study.

The experiment employed nine VICON cameras (six T40, two T160, and a synchronized digital camera) to capture participants' motion at 100 Hz (50 Hz for the digital camera). Both participants in each pair wore seven reflective markers on the joints of their dominant arm and hand as shown in Figure 4.3. This allowed the collection of Cartesian coordinates for each marker with respect to a predetermined reference point in space. Knowing the significant role wrist trajectories played in the previous work, Moon et al. [104], participants wore two markers on either side of the wrist for more accurate measurement of this joint.

After being instrumented with the VICON markers, participants stood facing each other across a table and played a modified version of Solitaire, a card sorting game. On the table was a deck of randomly ordered cards organized into two piles. In front of each participant were two aces of the same colour (either red or black). The participants' task was to order two full sets of cards hierarchically (e.g., ace, 2, 3, and so on) starting from the two aces given to them. They had to order the cards according to an alternating colour pattern, following the rules of the traditional Solitaire. The participants were told to finish the task as fast as they could. They were also instructed to access the two piles of cards in an alternating (left - right or right - left) order and to grab or return only one card at a time. This requirement for alternating access to the shared resource (two piles of cards) added cognitive load on the participants and prevented them from accurately keeping track of which pile of cards the other person would reach for next. By the design of the game, both participants needed to frequently access the shared piles of cards to finish the task, resulting in multiple natural occurrences of hesitations along with numerous reaching motions exhibited by each. Each pair of participants played the game twice, swapping the colour of the starting aces in the second round.

Labelling Human Hesitation Gestures

The experimenter watched all subject motions from the video recordings of the experiment and manually segmented and labelled participant motions that were deemed hesitations. In total, this qualitative process yielded 302 trajectories of hesitation gesture samples (171 from subjects on the left and 131 from the right side of the digital camera).

Subsequently, a video-based online survey was conducted using the segments of videos labelled

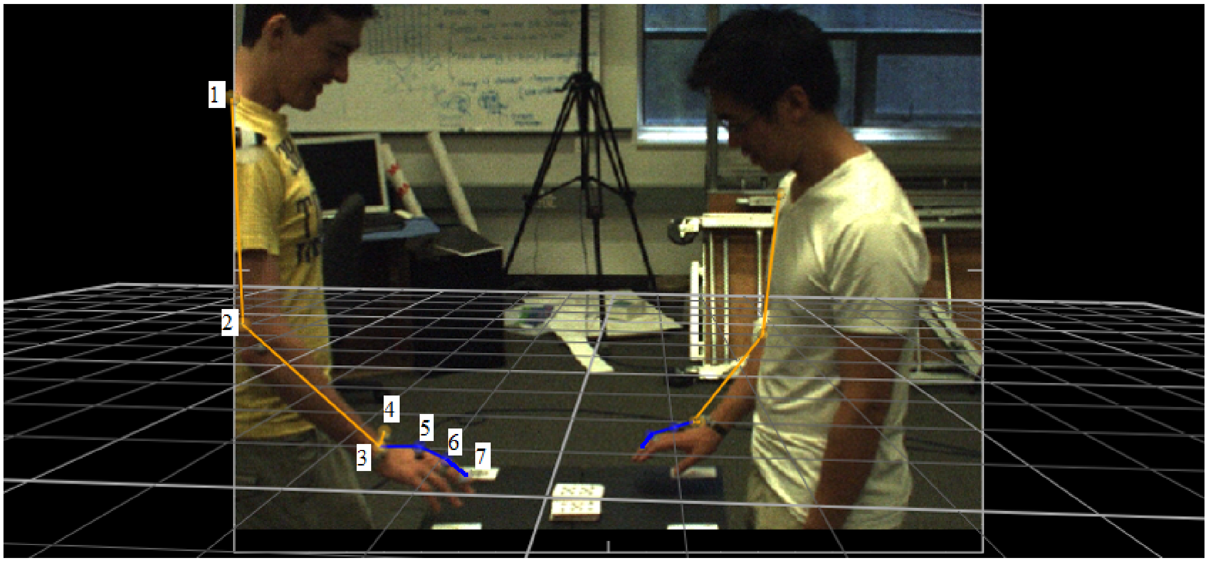

Figure 4.3: A side view of the HHI experiment captured from VICON's synchronized digital camera. Each pair of participants stood facing each other with a table between them and were asked to play a modified game of Solitaire. The numbers indicate the seven markers, used to track the participants' motion, placed on the following locations: 1 – shoulder (scapular acromion), 2 – elbow (humeral lateral epicondyle), 3 – right wrist (ulnar styloid), 4 – left wrist (scaphoid), 5 – knuckle (metacarpophalangeal joint of the second finger), 6 – upper finger (interphalangeal joint, between first and second phalanges, of the second finger), and 7 – lower finger (interphalangeal joint, between second and third phalanges, of the second finger).

as hesitations in order to identify which samples of motion were perceived to express lower or higher degrees of persistency and hesitancy. Participants for this online survey were recruited using the Amazon Mechanical Turk platform [5]. They were required to understand basic English to follow written instructions, and have good vision to be able to watch video recordings of human gestures.

Each video contained a segment identified as a hesitation plus a second before and after the motion segment used to provide the context of the motion. The survey participant watched and scored a random selection of twenty hesitation videos on the following two seven-point scales:

- **Persistency** "How persistent do you think the person on the LEFT is?" (1- Not persistent at all, 7 - Extremely persistent)
- **Hesitancy** "How hesitant do you think the person on the LEFT is?" (1 - Not hesitant at all, 7 - Extremely hesitant)

We changed the word LEFT to RIGHT, as appropriate, to highlight the specific individual of interest in the video.

4.3.2 Results and Discussion

Both the VICON study and the online survey were approved by the University of British Columbia (UBC) Behavioural Research Ethics Board (H10-00503).

A total of 300 participants completed the online survey. For each segment of motion (video), approximately 20 individuals ($M = 19.8$, $SD = 3.25$) provided both *persistency* and *hesitancy* scores. The average *persistency* and *hesitancy* scores were computed for each, and the score averages range from 1.68 to 6.26 ($M = 4.32$) for *hesitancy* and 2.10 to 6.58 ($M = 3.94$) for *persistency*. The average standard deviation of the *hesitancy* and *persistency* scores range from 0.83 to 2.39 ($M = 1.58$) and 0.67 to 2.37 ($M = 1.52$), respectively. These are a reasonable range of scores and standard deviation considering the 7-point Likert scale data. The data collected from this study is used to analyze human motion data as discussed at length in the following sections.

4.4 Exploring Human Hesitation Trajectories

To generate artificial hesitation behaviours for a robot, it is necessary to understand what trajectory features distinguish hesitation gestures from uninterrupted reach motions. The author's previous research [103, 104] demonstrated that there is a characteristic acceleration profile in R-type hesitations that is different from uninterrupted reach motions. R-type hesitations are comprised of an individual's reach trajectory toward a target, and interruption during the reach that leads to an immediate halting of the reach. This is followed by a retraction (yielding) motion that returns the agent back to the starting point of the reach. Based on this characteristic, human-recognizable robot hesitations can be successfully generated even for a 7-DOF robot[4] with a non-anthropomorphic morphology [104].

However, the same research showed that the same trajectory features are absent in negotiative hesitations (formerly introduced as P-type in Moon et al. [103]), and that the acceleration profiles of negotiative hesitations collected in the study are similar to that of uninterrupted reaches. Figure 4.4 demonstrates this finding. Although the negotiative hesitations may seem indistinguishable from reach motions in the acceleration domain, the results of Study 3 suggest that humans can perceive and recognize the subtle differences between the two types of motions. This result suggests that negotiative hesitation trajectories must be characterized in a different domain or using a different set of trajectory features.

This section presents the author's two-part exploration of the negotiative hesitation trajectory dataset as an attempt to identify characteristic features from the trajectories. The first process employs inferential statistics to understand the nature of negotiative human hesitations better. The second process tries to determine the metric of the collected human trajectories that can be used to inspire the design of hesitation behaviours for a robot. As is common in exploratory work such this, much of the analyses conducted on the collected human trajectories did not produce positive results. Some of these failed investigations led to the ultimate findings outlined in this section. Appendix A.2 documents these failed analyses. Appendix A.2 also presents detailed results of the final analyses described in this section so as not to distract the reader from the main findings.

[4] WAM™, Barrett Technologies, Cambridge, MA, USA.

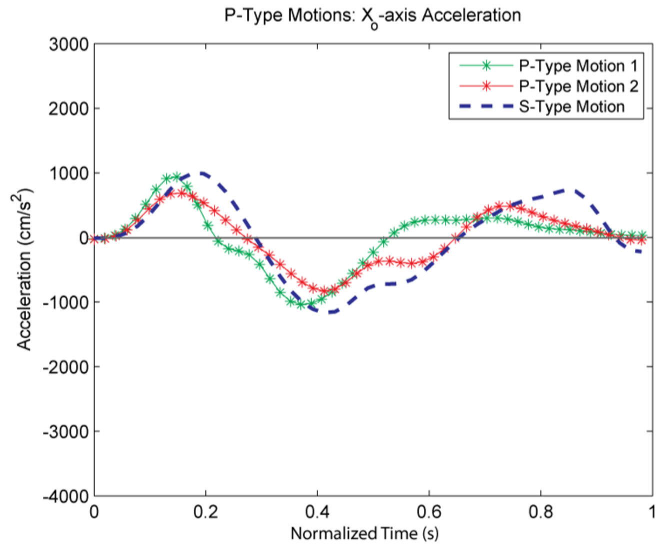

Figure 4.4: Acceleration profiles of sample negotiative hesitations, formerly referred to as P-type hesitations in [103, 104]. Two samples of the hesitations are shown here along with a sample of an uninterrupted reach motion (S-type). These samples demonstrate the negotiative hesitation gestures' similarity to an uninterrupted reach motion in this domain. (©2014 IEEE)

In the following section, Section 4.4.1, the author provides a description of the segmentation and pre-processing conducted on the stream of trajectory data collected in Study 3. The processed samples of the negotiative hesitations and uninterrupted reach motions were selected and divided into four sample sets. This division of sample sets allowed the author to conduct analyses that require a balanced number of samples as well as those with robustness to unbalanced number of samples (Section 4.4.2). Subsequently, 75 trajectory features were computed and analyzed from the motion samples. Section 4.4.3 describes this process.

Results from this exploration point to two metrics with which the hesitation trajectories can be distinguished from reach motions. This result, discussed in Section 4.4.4, informs the subsequent investigations on negotiative hesitation motions (Section 4.4.5) and ultimately, the design of the NHG.

4.4.1 Pre-Processing and Segmentation

Before analyzing the collected trajectories from Study 3, the raw recordings of human motions were processed and prepared as follows. First, the recorded Cartesian coordinates of the joints were transformed with respect to the coordinates of the shoulder. The average of the two markers placed on either side of each participant's wrist was computed for a more consistent treatment of the wrist joint measurement. The trajectories were filtered using a Butterworth low-pass filter at the cutoff frequency of 10 Hz, the highest frequency at which voluntary human motion can occur. After a number of failed analyses using trajectories of all joints, the author narrowed the focus of the analyses to the wrist joint. This decision was based on the findings from the author and colleagues' earlier work, Moon et al. [103], demonstrating that robot motions generated by tracing the wrist trajectories of human hesitations are also perceived as expressive of hesitation.

Herein, $M(t)$ and $P(t)$ refer to the Cartesian coordinates of the main and partner participants' trajectories, respectively. The orientation of the wrist was ignored to simplify the analysis. The main participant is the one whose motion segment is labelled as a reach or a hesitation. The partner participant's trajectory refers to whatever motion the partner was exhibiting during the main participant's reach or hesitation. Also, dimensionality-reduced expressions of the relative wrist motions from the 3D Cartesian coordinates ($\alpha_1(t)$ for the main and $\alpha_2(t)$ for the partner participant) were produced by selecting the dominant component from the Principal Components Analysis.

One of the key challenges when working with human motion trajectories in an unconstrained experimental task, such as the one described in the previous section, is the need to divide a stream of motion data into segments that can be labelled and studied as a supervised learning problem. Since the participants' motions during the experiment were not constrained spatially (they could move about the task space in whatever direction they chose), it was imperative to segment the vast amount of the collected natural reach motions in a systematic way such that they could be compared to segments of hesitation motions.

Different methods were employed to tackle this problem for each motion type. To segment reach motions, coordinates of *practical* target locations for each participant were computed. Clusters of each participant's trajectory data within 150 mm (approximately half of the distance participants travelled in a reach) from the known centre locations of the two card decks (targets) were identified. These represent the locations (T_{m1} and T_{m2} for the main, and T_{p1} and T_{p2} for the partner participant) where the participants' wrists were located when they grabbed a card from or placed a card onto the deck. Compared to the known centre location of the target objects, T_1 and T_2, these locations correspond to the Cartesian points in space the person tried to reach for, based on his/her preferred way of taking cards from the deck. These locations were typically near one of the four corners of the card closest to the participant's right hand. Figure 4.5 illustrates an example output of the segmentation algorithm used. Since it was not clear which patterns exist in marking the start and end of a hesitation, these samples were manually segmented by watching the video recordings.

4.4.2 Sample Selection

The author excluded outliers from the 302 hesitation samples collected in Study 3. These outliers had ≥ 3 SD in velocity maximum and range. With the above-mentioned procedure, a total of 1706 reach and 298 hesitation motion segments ($N_{(hes \geq 0)} = 298$) were collected and selected for analysis. That is, only 15% of the total samples, $N_{ub1} = 2004$, are hesitation samples. To account for the large unequal sample size, the author employed a random number generator to randomly select the same number of reach samples from a participant as hesitation samples from the same participant (balanced total sample size, $N_{b1} = 596$).

Since one individual manually segmented the hesitation motions, there is a possibility of bias from the individual. To account for this possibility, the author used the hesitation scores collected from Mechanical Turk for each of the segments to filter out samples with hesitancy scores below the median

Figure 4.5: Segmentation of a stream of a representative participant's motion shown in Euclidean distance trajectories, $d_1(t)$, with respect to the two target locations, T_{m1} (solid line) and T_{m2} (dashed line). Unlike the known centre location of the two card decks, T_1 and T_2, T_{m1} and T_{m2} represents the participant's distance to the Cartesian points in space the person tried to reach for (e.g., bottom right corner of the left deck). Points indicated by o and * are the zero velocity crossings of the Euclidean trajectory that coincide with the start and end of a participant's reach motion. The shaded portion indicates an area manually labelled as a hesitation.

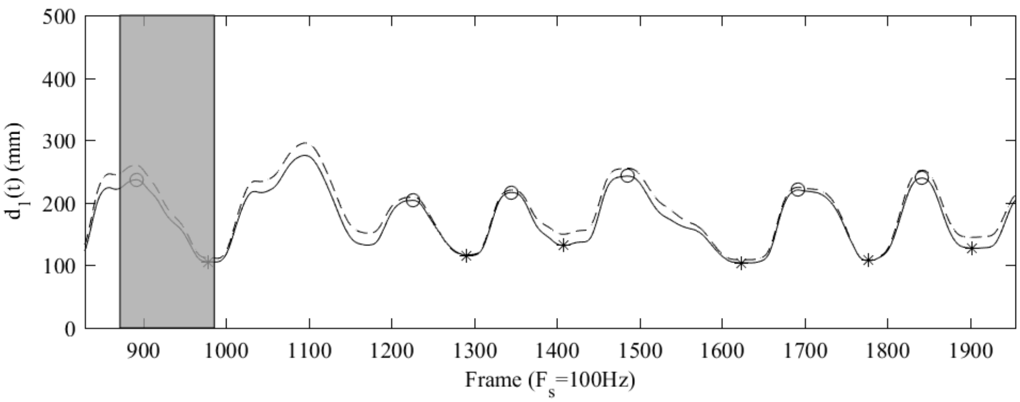

($hesitancy < 4$). This yielded a total of 192 hesitation samples ($N_{(hes \geq 4)} = 192$), resulting in $N_{b2} = 384$ and $N_{ub2} = 1898$.

In summary, this process created four sample sets (N_{b1}, N_{ub1}, N_{b2}, N_{ub2}) with which one can investigate salient features that set hesitation samples apart from reach motions. The following section outlines this data exploration process.

4.4.3 Data Exploration

To explore trajectory features that are characteristic of hesitation gestures, the author obtained and investigated a set of 75 features from the segmented sample trajectories of both hesitation and reach motions. This initial set includes the five types of features – maximum (max), minimum (min), mean (μ), amplitude (A), and number of zero crossings (ρ) – computed for the following fifteen metrics:

- $d_1(t) = \|M(t) - T_m\|$, the main participant's Euclidean distance to target, and its first ($\dot{d}_1(t)$) and second derivative ($\ddot{d}_1(t)$),
- $\alpha_1(t)$, the main participant's principal component of the segment along the direction of travel, and its first ($\dot{\alpha}_1(t)$) and second derivative ($\ddot{\alpha}_1(t)$),
- $d_2(t) = \|P(t) - T_p\|$, the partner participant's Euclidean distance to the same target, and its first ($\dot{d}_2(t)$) and second derivative ($\ddot{d}_2(t)$),
- $\alpha_2(t)$, the partner participant's principal component of the motion segment along the direction of

travel, and its first ($\dot{\alpha}_2(t)$) and second derivative ($\ddot{\alpha}_2(t)$),

- $\delta(t)= \|M(t) - T_m\| - \|P(t) - T_p\|$, the difference between the main and partner participants' Euclidean distances to their respective target locations, and its first ($\dot{\delta}(t)$) and second derivative ($\ddot{\delta}(t)$).

Here, T_m and T_p refer to the static target location relevant for the segment of motion computed for each participant. It is implied that the variables T_m and T_p refer to the appropriate one of the two custom target locations[5] relevant to the particular motion segment. The purpose of examining $d_1(t)$, $d_2(t)$, $\delta(t)$, and their derivatives is to find out which of the metrics deserve further analysis as an attempt to characterize and produce hesitation motions.

With the 75 features considered, the Shooting Algorithm was used as a regularization method [50] for building a logistic regression model of the two types of motions. This process allowed the author to eliminate a large number of the 75 features that do not distinguish one motion type from the other. The optimum λ value, a tuning parameter for the shooting algorithm, used was 16 for all sample sets considered. The regularization path used to obtain the optimum value is presented in Appendix A.2. As is a standard practice, the features used were normalized to have a zero-mean with unit variance across all samples considered. A total of twenty-one features were found to have a non-zero weight (> 0.0001). See Table A.2 in Appendix A.2 for the full list of the features and their weights computed for the four sample sets.

Here, it is important to note that, due to the definition of $\delta(t)$, some of the 75 features considered are linearly correlated with each other. Hence, one can expect that some of the results from the inferential statistics and shooting algorithm will be redundant. That is, for a feature computed across metrics $d_1(t)$, $d_2(t)$, and $\delta(t)$, the shooting algorithm would return non-zero weights in two out of three metrics. However, rather than eliminating one of the metrics entirely from the analysis, the author chose to explore all of these features. This is because these features represent distances, velocities, and accelerations in physical space. Therefore, investigating which two of the three metrics return with a larger weight from the shooting algorithm analysis can be useful, because it helps identify the features that are most numerically sensitive to change and likely most meaningful for implementation in real-life.

The features from the shooting algorithm that yielded non-zero weights at the optimum λ were selected as features worthy of further analysis. Using combinations of these features, the author employed a Support Vector Machine (SVM) having a linear kernel to build a logistic regression model. This process helped examine the significance of the features in classifying one motion type from the other.

In addition, the author conducted t-tests on the balanced sets of samples (N_{b1} and N_{b2}) and Welch tests[6] on unbalanced set of samples (N_{ub1} and N_{ub2}). The inferential statistical analyses were conducted on computed max, min, μ, A, and ρ of $d_1(t)$, $d_2(t)$, and $\delta(t)$ metrics. This analysis helped better test some of the assumptions about the nature of negotiative hesitation behaviours. Although only minor

[5] One target location belonging to the left card deck, and the other belonging to the right card deck, as explained in the Section 4.4.1.

[6] Welch test serves the same function as that of t-test, while computationally accounting for the differences in sample size in the distributions being compared.

differences were expected between the results of the t- and Welch tests, both were conducted to gauge possible biases in the random sampling of the reach motions in generating the balanced sample set. The full results of the t-tests and the list of features with significant results are presented in Table A.3.

4.4.4 Feature Differences in Reach and Hesitation Motion Samples

This section presents a summary of key findings from the t- and Welch tests followed by SVM classification models that were built to confirm saliency of selected features further. The two metrics of interest for further analysis were selected based on the results of the SVM and informed the discussion of trajectory characteristics of hesitation samples in the next section.

Inferential Statistics for Understanding Significant Trajectory Differences

This section presents a brief summary of the inferential statistics conducted on the four sample sets. While all 75 features were tested, only 47 features showed significant differences[7] between hesitation and uninterrupted reach motion samples, and not all of the significant results are non-trivial findings worthy of discussion. Hence, the full results of the statistical tests are presented in Table A.3, and only the discussion of the inferential statistics on the metrics related to $d_1(t)$, $d_2(t)$, and $\delta(t)$ are presented here.

$d_1(t)$– Main participant's Euclidean distance to target: Feature differences observed from the main participant's reach and hesitation trajectories provide empirical confirmation of some of the assumptions about hesitations. First, hesitation motions have a significantly larger minimum distance to the target, $\min(d_1(t))$, and smaller range of motion travelled, $A_{d_1(t)}$, on average than uninterrupted reach motions. This indicates that, as expected, the participants travelled a smaller distance and did not get as close to the target when hesitating as compared to when they successfully reached the target. Observations of zero crossings, $\rho_{d_1(t)}$, suggest that there are multiple zero velocity crossings for a majority of hesitation motion segments in the sample sets. A total of 179 hesitation motion segments have two or more zero velocity crossings, of which 112 have more than three, 56 more than four, and so on. Table A.4 presents this distribution.

$d_2(t)$– Partner participant's Euclidean distance to target: There is a significant difference in the distribution of the partner participant's minimum and average distance to the target ($\min(d_2(t))$ and $\mu_{d_2(t)}$, respectively) in hesitation and reach motions. In the segments of motion where the main participant is hesitating, the partner participants tend to be closer to the target than in segments where the main participant is fully reaching. This finding supports the hypothesis that the spatial state of an interacting agent sharing the same space influences behavioural responses – in this case, hesitations – of another. For the purposes of this thesis, this relationship indicates that the partner participant's location with respect to the target is a key influencer of the main agent's hesitation motion. Moreover, a significant portion of

[7] All inferential statistics reported here were conducted with $\alpha < 0.05$.

$A_{d_2(t)}$, the range of motion covered by the partner's motion, for the main participant's reach segments are found to be zero or near-zero. This indicates that the partner participants were often stationary or not moving along the axis toward or away from the target in these samples. This is in contrast to the larger $A_{d_2(t)}$ for cases where the main participant is hesitating.

$\delta(t)$– Difference between $d_1(t)$ and $d_2(t)$: $\delta(t)$ is a measure of the difference between the main and the partner participant's distance to their respective target locations. Based on the experimental setup of Study 3, it is also a measure of who is closer to the target and by how much. Results of the statistical tests suggest that the hesitation motion segments have a significantly larger and positive $\max(\delta(t))$, the maximum difference between the main and the partner participant's distance to the target, on average than reach motion segments. This indicates that the main participant is farther away from the target than the partner when the participant is hesitating. For reach, the average of $\max(\delta(t))$ hovers around zero, indicating an unbiased mix of cases where one participant is farther away from or closer to the target than the other. Complementing this spatial dynamics is the finding that $\min(\delta(t))$ has a larger negative value in reach motions than in hesitations. This indicates that, for reach segments, the partner is typically farther away from the target than the main participant. This is consistent with the result of $A_{d_2(t)}$ described above in which the participants were interpreted to be stationary or not moving toward or away from the target when the main participants fully reached the target.

From the abovementioned results, it is clear that the main participant's hesitation motions are affected by the trajectories and spatial locations of the partner participant with respect to the shared resource. This suggests that in characterizing negotiative hesitations – unlike the R-type hesitations that can be characterized from a single agent's wrist trajectories as a standalone motion, irrespective of the state or presence of another agent – one must take into account the spatial context of the workspace shared by the dyad.

Verification of Salient Features with SVM

A number of different combinations of features were used to classify one motion type from the other using the SVM approach. For each logistic regression model built with an SVM, a 4-fold cross validation with a 50% train/test ratio was conducted. To obtain a fair measure of the performance of the models, only the two balanced sets of samples, N_{b1} and N_{b2}, were used in this process. The ν parameter[8] was tuned to have approximately 15% of the total samples designated as support vectors. Appendix A.2 presents the full results of the SVM models.

In this process, two features, $\max(\dot{d}_1(t))$ and $\mu_{\dot{\delta}(t)}$, stood out from the rest as strong contributors for accurately classifying hesitations from reach motions – $\max(\dot{d}_1(t))$ represents the maximum speed the main participant was moving toward/away from the target, and $\mu_{\dot{\delta}(t)}$ represents the average of how the main participant's speed toward/away from the target compared against the speed of the partner with respect to the target. In particular, using only these two features for the N_{b2} sample set yields a

[8] The parameter $\nu \in (0,1]$ in the Support Vector method controls for the number of samples used as support vectors in the regression. See Chalimourda et al. [33] for more detail.

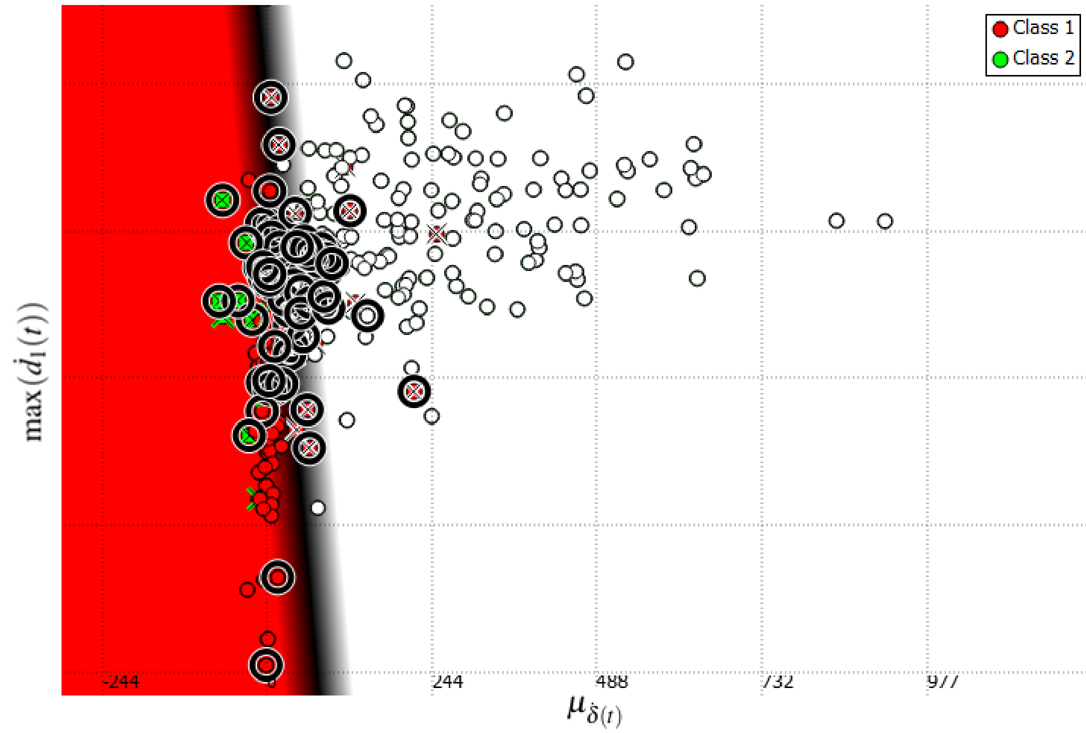

Figure 4.6: Illustration of an SVM model with $N_{b2} = 384$, $\nu = 0.3$, and $SV = 66$. Features $\max(\dot{d}_1(t))$ and $\mu_{\dot{\delta}(t)}$ are used. The highlighted samples (thick circles) represent support vectors. Class 1 refers to reach and Class 2 to hesitation samples. White circles represent the hesitation samples that have been properly classified using the algorithm.

test accuracy of 86% with a 50% test/train ratio and 82% with 75% test/train ratio ($CI : 69.4, 87.0$ and $CI : 68.8, 92.3$ for the 50 and 75% ratios respectively).[9] Figure 4.6 presents this SVM model.

This finding alone does not suffice for producing artificial hesitation behaviours for a robot. However, the saliency of these features suggests that the trajectory metrics $\dot{d}_1(t)$ and $\dot{\delta}(t)$ could provide promising results toward this end upon further analysis. These results motivate the follow-up investigation outlined in the next section that ultimately leads to the design of a human-inspired hesitation trajectory generator.

4.4.5 Understanding Hesitation Loops

Visual investigation of hesitation and reach samples in $\dot{d}_1(t)$ and $\dot{\delta}(t)$ presents trajectory patterns that are common to the majority of hesitation samples but are absent from reach gestures. That is, when hesitation samples are plotted in the state space of $\dot{\delta}(t)$ and $\delta(t)$, the trajectories tend to have the shape of a loop (see Figure 4.7). This is in contrast to reach trajectories visualized in the same state space as shown in Figure 4.8.

To verify this feature as characteristic of hesitation motions in contrast to reach motions, the author counted the number of samples that have a looped shaped in the state space. Results show that a total of

[9] SVM parameters $\nu = 0.3$ and $SV = 67(17\%)$ were used for this model.

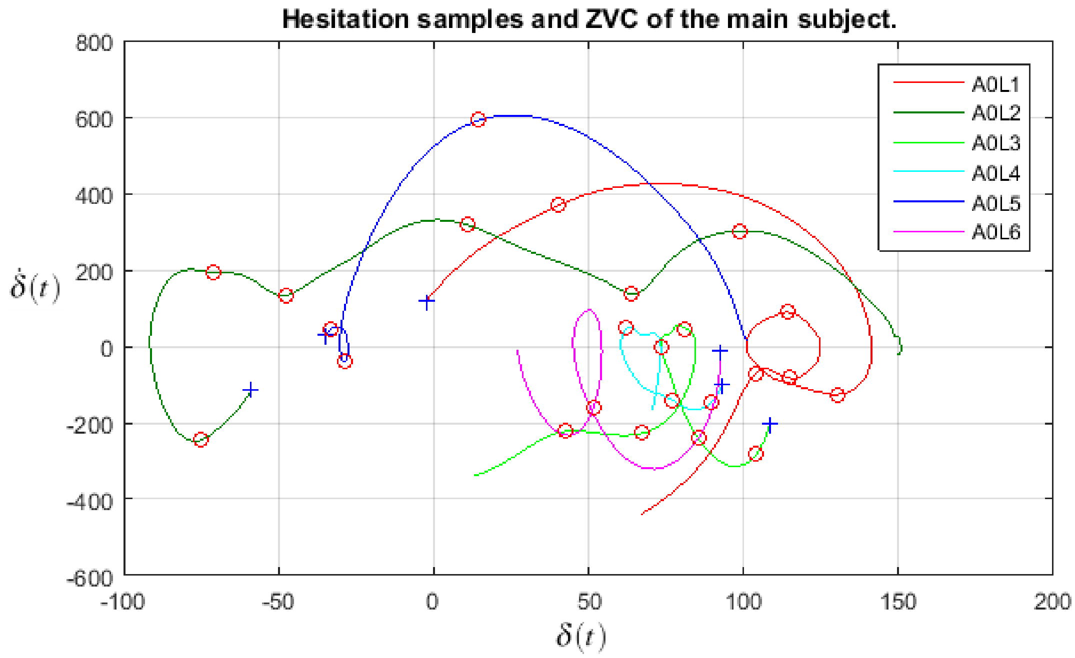

Figure 4.7: Overlay of hesitation samples demonstrating the nature of the $\dot{\delta}(t)$ vs. $\delta(t)$ plots to have circular loops around the zero $\dot{\delta}(t)$ axis. Zero Velocity Crossing (ZVC) of the main participant is marked in red circles. Blue crosses represent the start of the motion segment. As can be observed in sample A0L2, not all but most motions labelled as hesitation loop around $\dot{\delta}(t)$=0. A0L2 in particular consists of a hesitation that has larger sideways than back-and-forth motion with respect to the target.

134 hesitation samples out of $N_{hes \geq 4} = 192$ share this looped trajectory that is absent in reach gestures. Further analysis of the location and direction of the loops suggests that all of the loops in hesitation samples encircle $\dot{\delta}(t) = 0$. Due to the linear relationship between $\dot{\delta}(t)$ and $\dot{d}_1(t)$, these loops can also be observed in $\dot{d}_1(t)$ vs. $\delta(t)$ space.

With the sample hesitations that have these loops, the author computed the Euclidean distance travelled by the main participant between the two zero crossings of the encirclement. This distance represents how much the main participant retracted, if any, while the negotiative hesitation took place. Based on the computation, the distribution of this distance, Kickback Distance (KD), has a mean of 19.6 mm ($SD = 11.8$). However, as shown in Figure 4.9, KD of human hesitations are not normally distributed. This raises the question whether implementing a generator that exhibits a KD of 19.6 mm would produce communicative negotiative hesitation gestures for a robot.

It is important to note that, since behaviours humans perceive as hesitations can take on many different forms, including lack of motion, the characteristic loops referred to as *hesitation loops* in this section are not proposed as a feature that is present in all hesitations. Instead, they are presented as characteristic features of a subset of hesitation behaviours studied in this work, and put forth as components that can be implemented onto a robot to generate human recognizable artificial hesitation behaviours.

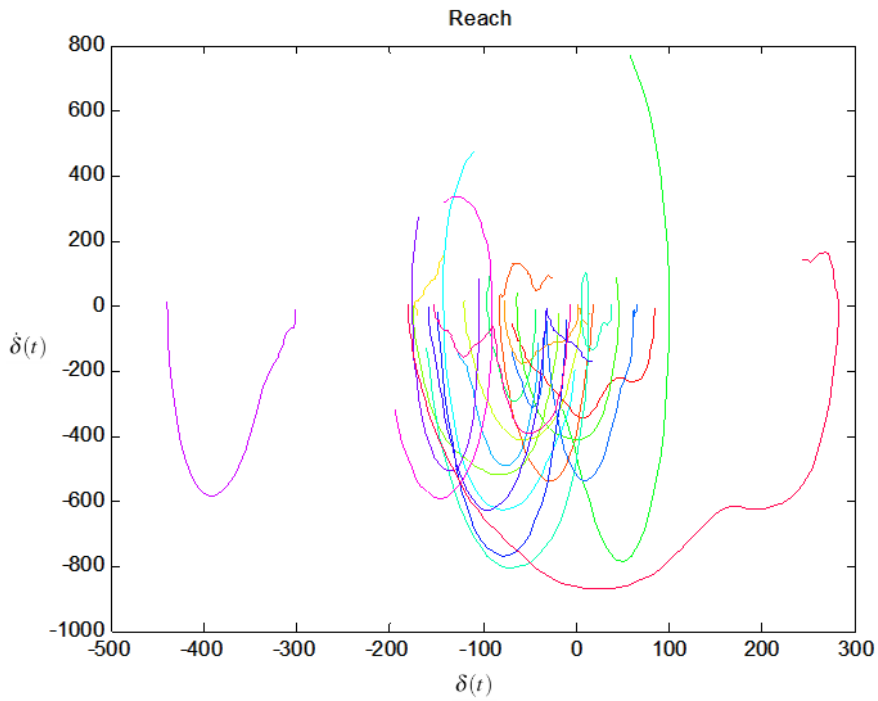

Figure 4.8: Overlay of reach samples in $\delta(t)$ vs. $\dot{\delta}(t)$ state space. As shown in the figure, reach samples demonstrate a fewer number of circular loops around $\dot{\delta}(t)$=0 than negotiative hesitation samples.

The ultimate goal of having such dynamic and reactive behaviours in a robot is to be able to generate negotiative behaviours in HRI that result in an interactive resolution of conflicts. To this end, hesitation loops are discussed in more detail in the following section and are used to frame the design of the NHG.

4.4.6 The Four Cases of Hesitation Loops

Given the discovery of hesitation loops, one can formulate an intuitive understanding of the context embedded in the trajectory in state space.

First, trajectories that move in the positive direction of state space, or crossing $\delta(t) = 0$ from the left to the right quadrants, can indicate that the main participant was moving toward the target or the partner was moving away from the target. In the former case, the trajectory must travel in the upper quadrants ($\dot{\delta}(t) > 0$);[10] in the latter case, the trajectory must travel in the lower quadrants ($\dot{\delta}(t) < 0$).[11] Since hesitation motions of the main participant are of focus, only the former case is interesting for our discussion.

Trajectories that move in the negative direction in the state space represent either that the partner participants were getting closer to the target than the main participant, or that the main participant was moving away from the target while the partner remained stationary. Likewise, only the trajectories that

[10] $\dot{d}_1(t) > \dot{d}_2(t)$
[11] $0 - \dot{d}_2(t) < 0$

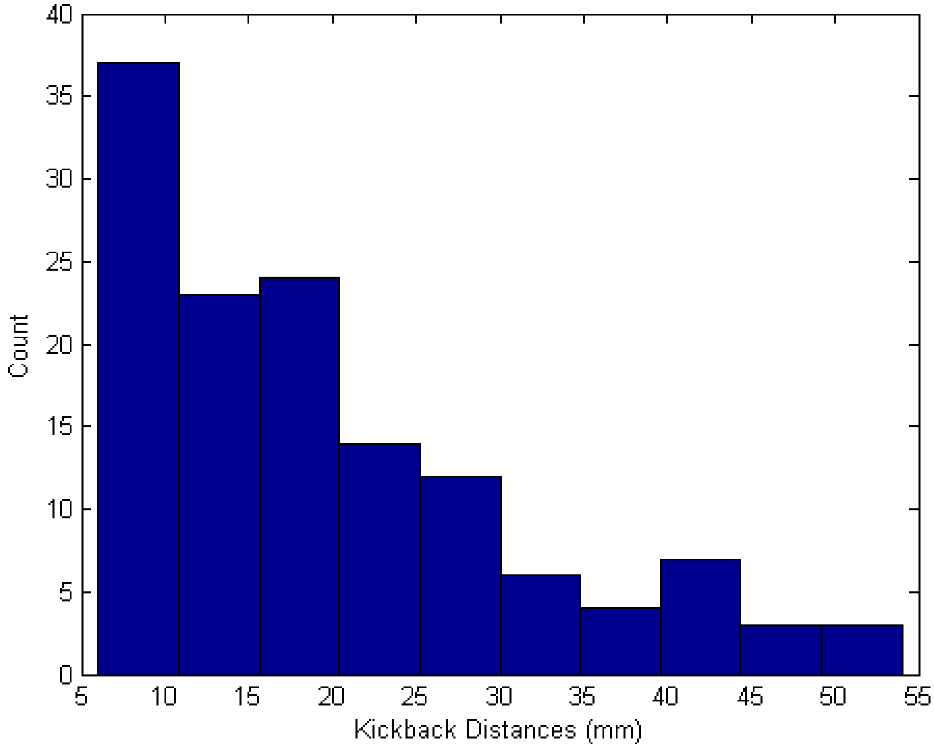

Figure 4.9: Distribution of KD collected from the 134 samples of hesitations in ($N_{hes\geq 4} = 192$) that encircle $\dot{\delta}(t)$=0. The distribution is non-normal with most of the values occurring between 5 to 30 mm.

move from right to left while the partner is moving faster than the main participant ($\dot{\delta}(t) < 0$)[12] are relevant to the discussion of hesitations.

There are four cases of observed hesitations that describe the spatial contexts in which hesitations occur. Figure 4.10 provides a visual demonstration of the cases. This section outlines the HH dyad's interaction dynamics that are represented in these cases and what the entering and exiting directions of a hesitation loop mean in these contexts.

Case 1 This case represents hesitation samples starting with a state of ($\delta(t) > 0$, $\dot{\delta}(t) < 0$). In this case, $\delta(t) > 0$ indicates that the main participant is farther away from the target than the partner at the onset of a hesitation ($d_1(t) > d_2(t)$). $\dot{\delta}(t) < 0$ suggests that the main participant is moving slower toward the target than the partner. Indeed, it is also possible that the same starting state in state space can represent the main participant moving faster away from the target than the partner who may be moving toward or away from the target him/herself. However, these non-contentious scenarios are not considered here.[13]

[12] $\dot{d}_1(t) < \dot{d}_2(t)$

[13] All of the hesitation samples considered in this investigation have a resource conflict to resolve. This quality is ensured from the manual segmentation method employed for the hesitation samples, which involved qualitatively observing recordings of human motions during occurences of resource conflicts.

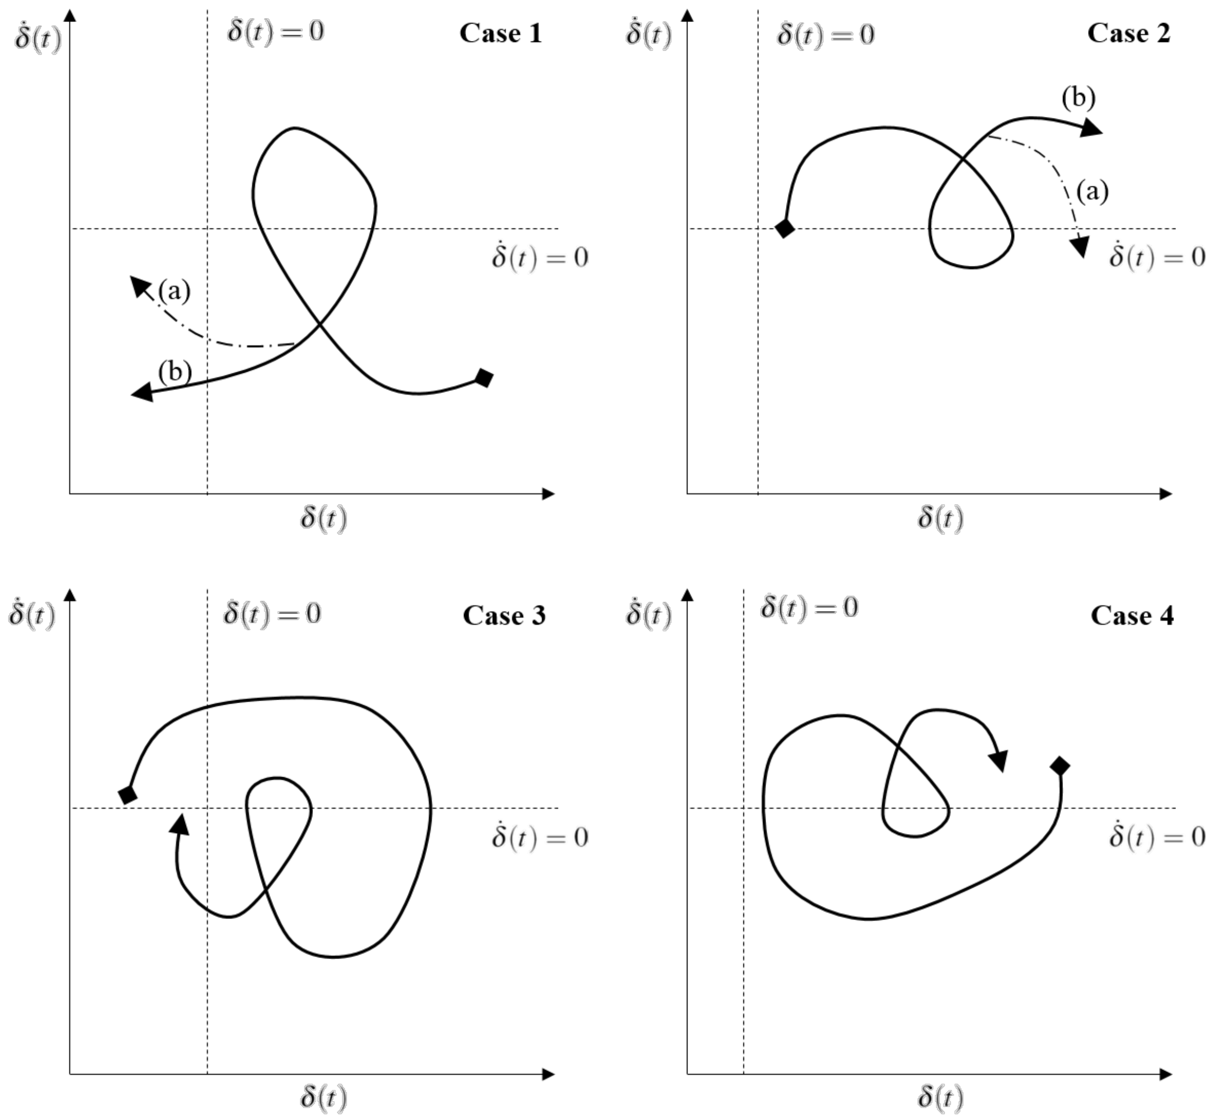

Figure 4.10: The four cases of hesitation loops demonstrated in the $\dot{\delta}(t)$ vs. $\delta(t)$ state space. Each of the cases represent a spatial context between the two interacting agents' motions as outlined in Section 4.4.6. All hesitations loop around $\dot{\delta}(t) = 0$.

From this starting location, the main participant remains farther away from the target than the partner $\delta(t) < 0$, but closes the distance gap as the main participant gains speed and catches up to the speed of the partner (crosses $\dot{\delta}(t) = 0$ in the upward direction). The main participant is briefly faster than the partner, but quickly slows down (crosses $\dot{\delta}(t) = 0$ in the downward direction). In some of the cases, the main participant immediately speeds up again with respect to the partner, suggestive of the main participant's second attempt at accessing the target (see Case 1 (a) in Figure 4.10). In other cases, the main participant slows down to $\dot{d}_1(t) = 0$, with the hesitation motion segment ending with $\dot{\delta}(t) = 0 - \dot{d}_2(t)$. This is shown as Case 1 (b) in Figure 4.10.

Case 2 Hesitation samples in Case 2 have a starting state of ($\delta(t)$< 0, $\dot{\delta}(t)$≤ 0). Often the value of $\dot{\delta}(t)$ is near zero, indicating that the partner participant may be engaged in some other activity (e.g., shuffling cards sideways on his/her side of the table) or remaining stationary when the main participant starts to reach for the target (i.e., $\dot{d}_1(t)$∼ 0). Depicted in Figure 4.10 is the main participant located almost equidistant to the target as the partner at the onset of the motion. Then, the main participant moves faster and gets closer to the target than the partner when the partner starts his/her motion or moves faster towards the shared target. Either by a brief speeding up of the partner's motion or slowing down of the main participant's motion (hesitation), the partner moves faster toward the target before the main participant re-attempts to continue his/her reach. As illustrated in Figure 4.10 Case 2(a), the participants may go through multiple cycles in which one of the participants moves slower than the other, crossing $\dot{\delta}(t)$= 0 downward. Alternatively, the partner may yield to the main participant as depicted in Figure 4.10 Case 2(b). Going through multiple cycles of the trajectory pattern is indicative of multiple re-attempts by the main participant to gain access to the target resource. The number of times this re-attempt takes place is heretofore referred to as Re-attempts (RA).

Case 3 Cases 3 is a variation of Cases 1 described above. In Case 3, the main participant reaches for the target, but the speed toward, and subsequently, the distance to the target is overtaken by the partner who moves faster – perhaps despite having started his/her reach later than the main participant. This dynamics brings the pair's trajectory to cross the $\dot{\delta}(t)$= 0 state downwards once before crossing it again in the manner described in Case 1.

Case 4 Likewise, Case 4 is a variation of Case 2, where the dyad crosses $\dot{\delta}(t)$= 0 upward before going through the same downward hesitation loop in Case 2. Case 4 illustrated in Figure 4.10 starts at a state where the main participant is closer to the target than the partner. However, the partner's reach toward the target overtakes the dynamics, until the main participant's motion toward the target overtakes the speed of the partner. Their speed and respective distances to the target lead them to eventually go through the same hesitation loop as described in Case 2.

Understandably, because $\delta(t)$ and $\dot{\delta}(t)$ combine the effects of both the main and the partner participants, the zero crossings of the system depend on both agents' actions. Hence, generation of hesitation loops in HRI requires a translation of this trajectory pattern into a system in which only one of the agents' (robot's) motions can be precisely controlled – needless to say, motions of the other agent (human's) remain outside of our control.

For the purpose of developing a trajectory generator that can recreate such hesitation dynamics in HRI, one can detect when such encirclement of $\dot{\delta}(t)$= 0, such as detection of the zero crossing, occurs and manipulate the robot's reaching motion to produce the kickback motion (move backwards by a KD) that completes the loop. Herein, a $\dot{\delta}(t)$= 0 occurrence is referred to as the Trigger State (TS) to designate the state at which a hesitation behaviour of a robot is to be triggered.

4.5 Design of the Negotiative Hesitation Generator

To investigate the efficacy of using negotiative hesitations in HRI as a mode of resolving resource conflicts, interactive behaviours of a robot must be able to generate the elements of human hesitation behaviours mentioned in the previous section. This artificial hesitation behaviour must also be implemented on a robot in such a way that the system can respond to human hesitations in real-time to allow negotiative nonverbal dialogues to emerge.

Among other alternatives, the author chose to use the Linear Dynamical System (LDS) approach to generate reference trajectories of artificial hesitations. A LDS is defined as the following:

$$\dot{x} = Ax + b \tag{4.1}$$

In this differential equation, A is a matrix and b is a vector describing a vector field in the designated state space. These parameters represent the linear dynamics that exist between the state variable, $x \in \mathbb{R}^d$, and the rate of change expressed in $\dot{x}$. Hence, iterative computation of the differential equation results in the determination of the trajectory of the linear system through the state space. In robotics, x typically represents unambiguous states of a robot, such as joint angles. Therefore, computed states and state velocities can be used as reference trajectories with which to control a robot.

One of the reasons for choosing to implement artificial hesitation behaviours with an LDS is that the Dynamical System (DS) approach offers reference trajectory generators that are reactive, real-time, and robust to disturbances or changes to the environment. In contrast, the author and colleagues' previous work used trajectories generated using splines [104]. The splines function as an interpolation between two boundary states defined early in the motion trajectory. Hence, unexpected disturbances that push the robot out of its planned trajectory, for example, typically result in an undesirable, high-jerk movement of the robotic platform. This happens as the system tries to catch up to larger than planned difference between the state that has been planned for its next time step and the current, disturbed state of the robot. In comparison, trajectories that are generated using a DS iteratively uses the robot's current state for the computation of the desired next state. With a DS that converges to a target location, any sudden disturbances to the system only changes the value of the input state used to compute the desired next state of the system. Since the reference state for the next time step is computed based on the vector field, sudden changes to its input states are seamlessly handled. This highly reactive nature of DS-based systems has also been showcased in complex real-time obstacle avoidance tasks, as well as tasks that involve catching flying target objects [79, 101].

With the development of algorithms that can guarantee the stability of a DS during the training phase of the system, such as Stable Estimator of Dynamical Systems [78], the DS has also been used to encode a variety of tasks for a robot as a teach-by-demonstration method [48, 63, 123]. Given its potential to develop a large vocabulary of activities for robots to perform, implementing hesitation behaviours on a DS opens up the possibility for artificial hesitation behaviours to be triggered during tasks not envisioned or tested in this thesis. With this in mind, the author designed a DS-based trajectory generator that takes into consideration point-to-point motions involved in reaching for a shared resource. Such motions are not only foundational to many tasks, but they are also the main type of motion used in the type of resource conflicts discussed in this thesis.

Coincidentally, implementation of the TS – the state at which the controller should trigger a hesitation behaviour – and KD – the distance to retract as part of the hesitation behaviour – in a LDS yields a simple control architecture. As discussed in the previous section, for hesitation loops to be generated in an HRI, one must remember that only one of the agents, the robot, is within the scope of our control. Moreover, $\delta(t)$ and $\dot{\delta}(t)$ represent ambiguous states that can be translated into an infinite combination of states in the physical world.[14] Therefore, rather than devising a control signal in $\delta(t)$ and $\dot{\delta}(t)$ space, one can monitor the state space to detect the TS in which an encirclement of the $\dot{\delta}(t)= 0$ is desired.

To do this, the author employed a simple LDS in $d_1(t)$ vs. $\dot{d}_1(t)$ space ($\dot{d}_1(t)= Ad_1(t)$) that can be used to reach for a target, and implemented an upper layer module that monitors for the TS in $\delta(t)$ vs. $\dot{\delta}(t)$ state space. Knowing the location of the shared target resource, T, the values of $d_1(t)$ and $\dot{d}_1(t)$ can be translated to Cartesian positions and velocities along a desired path of motion. This transforms the ambiguous one-dimensional states, $d_1(t)$ and $\dot{d}_1(t)$, to the unambiguous position and velocity $M(t)$ and $\nabla M(t)$.

Once a TS is detected, the upper layer algorithm moves the equilibrium point, or the point of convergence, of the LDS by a KD. The devised system that produces reference trajectories using this control regime is the Negotiative Hesitation Generator (NHG). To demonstrate that such a layered approach produces the desired encirclement of $\dot{\delta}(t)= 0$, this regime was simulated in MATLAB (The MathWorks Inc., Natick, MA, USA) with quintic splines that represent two agents reaching for the same target. As shown in Figure 4.11, the generated reference trajectory from the simulation encircles the $\dot{\delta}(t)= 0$ observed from HH negotiative hesitations. The upper layer module was simulated using splines, since this was practical due to the off-line nature of the simulation. A spline-based implementation of the TS and KD for real-time, in-person HRI will require added complexities in the control architecture to ensure responsiveness of the system without generating undesirable, residual motion. The system architecture to implement the NHG using a LDS is described in detail in Section 4.6.2.

In the next section, Section 4.6, an online study is used to demonstrate that this particular reactive system can communicate hesitation to third party observers. In the following chapter, Chapter 5, the author determines the system's efficacy in triggering negotiative responses from humans in in-person HRI.

4.6 Study 4: Validating the Negotiative Hesitation Generator

Although the design of the NHG is based on the observations of human motions from Study 3, the efficacy of the trajectories generated by the NHG has yet to be tested. Does the NHG produce humanlike hesitation responses? Does the LDS implementation of the elements of hesitation convey a level of hesitation a human can observe? These are some of the grounding questions that need to be addressed to justify the use of the NHG for bidirectional, negotiative interaction between an HR dyad.

[14] $\delta(t)$ and $\dot{\delta}(t)$ are functions of relative Euclidean distances to the target. Hence, $\delta(t)$=2 cm can mean one agent being anywhere in a 2 cm radius from the target and the other being 4 cm away.

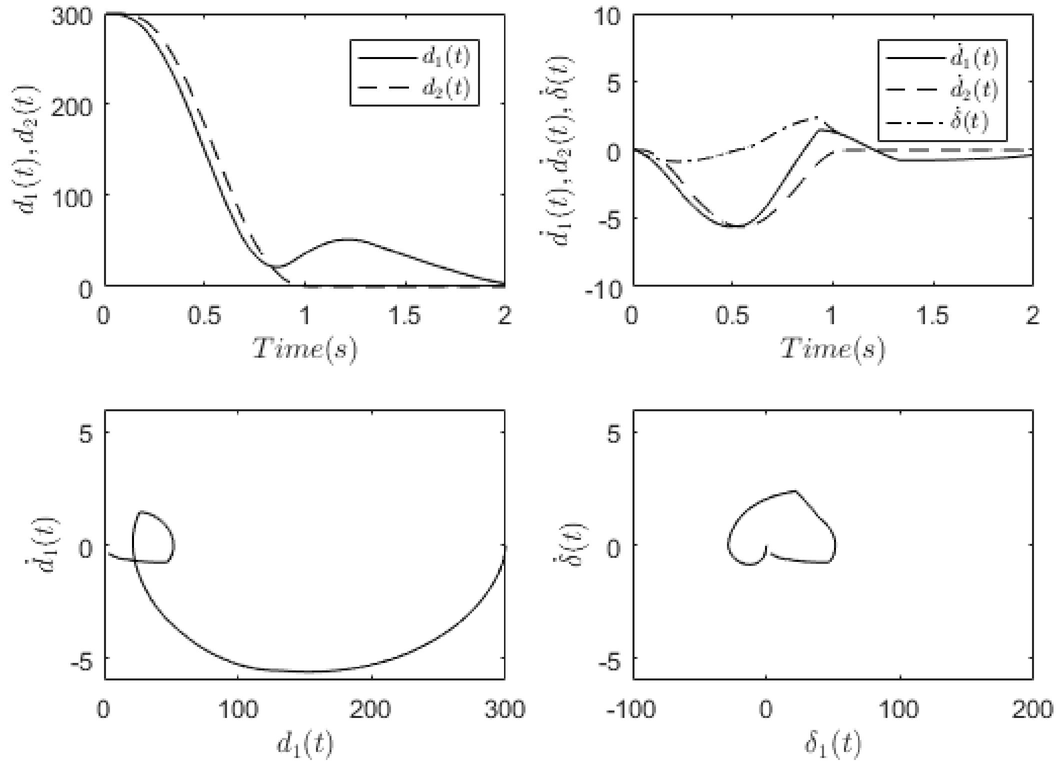

Figure 4.11: Simulation output demonstrating the NHG implementation using quintic splines. Two quintic splines are generated to represent reach motions of two agents reaching for the shared target a fraction of a second (0.05 s) apart from each other. The main agent's quintic spline is interrupted when a TS is detected. This triggers the main agent to move back by KD before reaching for the target again. This results in a circular motion in the $\delta(t)$ state space (representing difference between the main and the partner participants distance to their respective target locations) around (0, 0) of $\dot{\delta}(t) = 0$.

As outlined in Section 4.5, a premise of the NHG is in the presence of hesitation loops found in the $\dot{\delta}(t)$ vs. $\delta(t)$ state space. When the state space is interpreted with the context of an in-person HRI, it allows one to understand the dynamic, spatial interaction between the two agents. In this state space, the NHG calls for the implementation of the following three elements: TS, a state at which human agents tend to hesitate in response to an imminent resource conflict; KD, the amount of distance the hesitating agent retracts as part of the hesitation behaviour before re-attempting to access the resource in conflict; and re-attempts (RA), the maximum number of times the robot should re-attempt to access the conflicted resource before yielding. While TS refers to a state at which hesitation behaviour should be triggered, implementation of an NHG involves assigning values of KD and RA, both of which are adjustable parameters. In the HH hesitation samples collected in Study 3 (Section 4.3), a range of values of both KD and RA were observed. Hence, within the range of observed parameter values, it is useful to investigate what pairs of values for these parameters are acceptable for the NHG to generate artificial robot motions that convey a state of hesitation.

In Study 4, presented in this section, the efficacy of a set of KD and RA parameter values are evaluated against each other, in addition to a smooth stopping motion typically implemented in robotic collision avoidance behaviours. The following hypotheses are tested:

Hypothesis 4.1 *Robot hesitation responses generated by the* NHG *are perceived to be more hesitant and persistent than a smooth stopping behaviour.*

Hypothesis 4.2 *Robot hesitation responses generated by the* NHG *are perceived to be more animate and anthropomorphic than a smooth stopping behaviour.*

Hypothesis 4.3 *Robot hesitation responses generated by the* NHG *are perceived to be more dominant and useful than a smooth stopping behaviour.*

Hypothesis 4.4 *Robot hesitation responses generated with a human-inspired kickback distance (*$0 \leq KD < 19$ *mm) are perceived to be more hesitant than those with* KD *outside this range.*

Hypothesis 4.5 *Robot hesitation responses showing a larger number of* RA *are perceived to be more persistent, dominant and useful than those with a larger* RA.

Obtaining empirical support for these hypotheses would validate the use of the trajectory generator as a mechanism that is adequate for creating artificial robot hesitation responses for in-person interactions with humans. In addition, it would provide a measure of the dominance, animacy, and anthropomorphic qualities perceived by human observers in the generated trajectories. These perception measures provide valuable information to consider in implementing robot behaviours for an improved HRI experience. The remainder of this section outlines the experimental procedure used to conduct the online survey (Section 4.6.1), the technical system used to implement the NHG (Section 4.6.2), and empirical and qualitative findings from the analysis (Section 4.6.3). Section 4.6.4 discusses implications of the findings in detail.

4.6.1 Experimental Procedure

This online survey-based study was structured to be similar to Study 3. The participants were recruited via Amazon's Mechanical Turk system to watch and report on their impression of a series of videos. All survey participants gave consent online by explicitly selecting the option "I consent to participate." Upon giving consent, participants provided their demographical information (age and gender). Afterward, participants were asked to watch a short introductory video before proceeding with the survey. This introductory video was made to provide the context of the interaction that the participants need to know to understand the videos that followed. The introductory video showed the experimenter and the robot, facing each other and collaborating on an assembly task. In front of the robot were two liquid pumps, and between the two agents was a dispenser operated with a button on top. Near the experimenter was a bin containing uncooked lentil beans of two different colours. The following text appeared at the bottom of the screen:

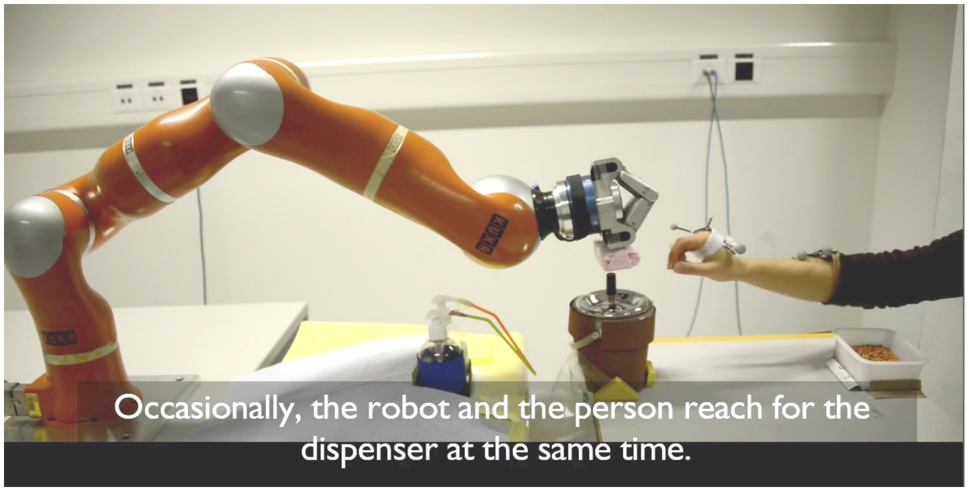

Figure 4.12: A screen capture of the introductory video shown to participants in Study 4. All other non-introductory videos shown to the participants were captured from the same camera angle, showing only the experimenter's arm, the robot, and the experimental setup.

> You will be watching a person and a robot working on a collaborative task. In this task, the person sorts and places items onto the dispenser in the centre. Meanwhile, the robot helps by handling the liquid pumps, and flushing the dispenser using the black button. Occasionally, the person and the robot reach for the dispenser at the same time.

Near the end of the video, the video showed the robot and the experimenter reaching for the dispenser at the same time. See Figure 4.12 for a screen capture of this segment of the introductory video. The video faded to black before revealing who got the right of way.

In the following pages of the survey, participants watched a total of fifteen videos. All videos started with the motion of the robot returning from the dispenser after successfully having reached and pressed it. Afterward, each of the videos showed the robot pressing on one of the liquid pumps, then reaching for the dispenser again. But this time, the experimenter reached for the dispenser as well, creating a conflict of resource. The videos showed the robot responding to the conflict with either a smooth stopping behaviour, or a motion generated by the NHG with one of the fourteen parameter pairs (see Figure 4.13 for the full list of conditions). For the purpose of testing Hypothesis 4.4, values 0, 10, 19, and 40 mm are used to represent the range of KD observed from Study 4. The KD value of 19 mm represents an approximate average of the distribution of KD observed in human hesitation behaviours. The KD values 0 and 40 represent the two extremes of the distribution.

Ideally, the parameter pair would include five recordings for each value of KD, spanning RA values from 0 to 4. However, videos for the parameter set ($KD = 10$, $RA = 4$) and ($KD = 40$, $RA = 0$) were lost from our recording due to a technical failure. They could not be re-recorded due to logistical constraints of the international collaborative research.[15] In addition, the KD value of 0 inherently does not provide

[15] All of the videos were recorded at Ecole Polytechnique Fédérale de Lausanne (EPFL) at a location and with an experimental setup that are no longer accessible to the author.

Smooth Stop (Control Condition)
X

The NHG Conditions						
		RA				
		0	1	2	3	4
KD	0	X				
	10	X	X	X	X	
	19	X	X	X	X	X
	40		X	X	X	X

Figure 4.13: Outline of the conditions tested for Study 4. The structure of this study is a non-standard two-factor factorial design that has a control condition outside of the two factors of interest. The conditions tests are marked with an 'X.'

Table 4.1: Internal reliabilities of the self-reported measures used for Study 3 are presented here. Hesitancy and Persistency measures only used a single item and their Cronbach's alpha value cannot be computed. All except the *Usefulness* measures used in this study have a Cronbach's alpha $\geq .70$.

Measures	Cronbach's alpha	Items
Hesitancy	N/A	Hesitant
Persistency	N/A	Persistent
Dominance	.92	Dominant, Aggressive
Usefulness	.66	Useful, Efficient
Animacy	.78	Artificial, Mechanical
Anthropomorphism	.88	Artificial, Machinelike

any back and forth motion, rendering the value of RA meaningless. Given the lack of the re-attempt behaviour, the robot to quickly retreat to its home position as soon as an imminent conflict is detected. The parameter set ($KD = 0$, $RA = 0$) serves as a reference point to study the main effects of KD and RA as is explained in more detail in the results section (Section 4.6.3). The smooth stop response is used as a control condition with which to contrast all NHG-based trajectories.

Presented below each video were a set of questions that echo the questionnaire used in Study 3. The first nine questions were designed to collect six self-report measures: three questions for measuring *Anthropomorphism* and *Animacy*, derived from the Godspeed Questionnaire [14]; four questions for measuring *Dominance* and *Usefulness*, two perceived team measures from Moon and Nass [107]; and *Hesitancy* and *Persistency* measures.[16] Table 4.1 presents the questionnaire items used for this study.

Findings of Study 3 suggest that responses to *Persistency* and *Hesitancy* may not be enough to describe the quality of motion the participants pay attention to. Therefore, a textbox was included to collect optional, open-ended qualitative responses to the question "What adjectives would you use to

[16] The *Hesitancy* and *Persistency* measures were used in Study 3 with word pairs, (Hesitant – Not hesitant), and (Persistent – Not persistent). The same word pairs were used in this study.

describe the motion of the robot observed in the video? (Optional)".

It was also made explicit that participants could replay the video as many times as they liked to complete the survey. Each video was 16 to 22 seconds in length. All participants watched all fifteen videos in random order. Based on a pilot study, the total time to complete the survey was estimated to be less than 30 minutes. Participants were offered a financial reward of $1.60 USD for completing the survey.

4.6.2 Technical Implementation

The experimental system consisted of a 7-DOF KUKA LWR 4+ (KUKA Robot Group, Augsburg, Germany) manipulator controlled using the Fast Research Interface (FRI) at a 1 kHz frequency, and an OptiTrack (NaturalPoint Inc. DBA OptiTrack, Corvallis, OR) motion capture system operating at a 120 Hz sampling rate. The motion capture system was used to track markers on the experimenter's hand and forearm.

In all of the conditions, the robot's task involved reaching and retraction motions to and from the two liquid pumps and the dispenser button. Depending on the experimental condition, the robot used either a Hermite quintic spline or the NHG to generate the reaching motion to and from the dispenser as further described below. For all conditions, once the robot reached the dispenser button or one of the pumps, it was programmed to press on the button/pump using a Hermite quintic spline. This ensured that all aspects of the generated robot motions remain equal for all conditions except for the reaching motion and the robot's response to conflicts.

In the fourteen conditions that involve values of KD and RA, the LDS implementation of the NHG was used to generate reference Cartesian position trajectories of reaching motions. The robot was programmed to re-attempt (i.e., reach for the resource again immediately after returning to the starting position) a maximum of four times before giving up its access to the resource and resuming its task of tending to the liquid bottles.

In the Smooth Stop (control) condition, the robot's reaching motion was generated using a Hermite quintic spline. This technique takes the start and end states of the robot and the desired travel completion time to interpolate between the two points in a minimum-jerk manner. The resulting trajectory allows for the robot to trace smooth reference trajectories considered humanlike [49]. When the robot detects a conflict near the onset of its reach motion, it exhibited the R-type hesitation using the AHP. When the conflict was detected after the time window t_1 (see Figure 4.2), it was programmed to remain at its current location, leading it to come to a full stop and hold its position for 1 second before retreating. Since this condition does not employ the notion of the TS used in the NHG to determine the state at which the robot should trigger a conflict response behaviour, the system was programmed to use proximity and direction of motion between the robot and the experimenter to detect possible HR resource conflict states. When the robot and the experimenter moved toward each other and were closer than 5 cm from each other, the robot was programmed to remain at the current location such that it comes to a full stop. After 1 second of pause in the stopped state, the robot retracted back to its home position.

System Architecture

The system architecture employed for this study was implemented in Robot Operating System (ROS) [126]. This section provides a high-level overview of the architecture. Three main components were custom developed for this experiment:

1. a communication node dedicated to receiving, transforming, and publishing sensed information from OptiTrack,
2. an action server that registers and maps requested actions from action clients onto a set of predefined action policies that ultimately generate desired reference trajectories, and
3. a client action node that acts as the master script to trigger different robot actions at appropriate times.

First, the data stream from OptiTrack was broadcast through a local area network and converted into a data format appropriate for further use of the data in ROS. The ROS action client-server protocol was employed to interface the trajectory generator algorithms with the FRI interface controlling the robot. The experiment's main action client acted, in part, as a server for smaller action clients. The main client is responsible for the continuous monitoring of human motion captured via OptiTrack. When initialized, it launches a dedicated node to receive and convert the stream of OptiTrack data into a local coordinate system as well as Euclidean distance to the shared resource (dispenser). The final ROSTOPIC containing human motion information is ultimately communicated to the action server to trigger different trajectory responses for the robot. The smaller action clients trigger various behaviours, such as reach or retract, to the main client, which in turn requests actions to be serviced by the action server. The action server, upon receiving a request, generated and communicated the reference trajectories for the FRI control interface to actuate the robot.

4.6.3 Results

A total of 50 people participated in the study (16 females, 33 males, 1 preferred not to disclose). On average, it took approximately 18 minutes for the participants to complete the survey (M =18:27, SD =9:03). The age of the participants ranged from 21 to 52 ($M = 32$, $SD = 7.0$).

As mentioned in Section 4.6.1, this experiment was structured as a non-standard, two-factor, factorial design where variables KD and RA serve as the two crossed factors. The five levels (0, 1, 2, 3, 4) of RA are factored with the four levels of KD (0, 10, 19, 40). Two elements of this experiment structure make this factorial design non-standard. First, none of the values of KD and RA apply to the control condition, requiring the analysis to include a control group outside the KD-RA factorial. Second, there are some combinations of KD-RA that are missing from the experiment due to the reasons described earlier. While this is not ideal, subjecting the participants to a smaller number of conditions shortened the length of the within-subjects study, thereby avoiding potential effects of pencil-whipping.[17]

[17] Pencil-whipping within the context of survey-based studies refers to the tendency of participants to become careless or provide answers to questions without giving much thought, often due to boredom, and produces undesirable noise in the dataset.

To test the aforementioned hypotheses, the responses to the six measures from the survey are analyzed using a regression analysis with a linear mixed effects model. A Multi-level Modelling (MLM) using the lme4 package in R with Restricted Maximum Likelihood (REML) was used for model fitting. Given that the structure of this study is a non-standard two-factor factorial design (see Figure 4.13 for the structure of the conditions), a more familiar statistical modeling such as ANOVA cannot be used. On the other hand, MLM allows for data from such non-standard study designs to be captured such that inferential statistics can still be conducted. Keeping in mind the repeated measures aspect of this study (all participants saw all 15 videos), *Participant* was used as a random factor and KD and RA as fixed factors. In this analysis, the mean responses to the (KD=0, RA=0) pair, μ_0, is used as the reference point with which the responses to the remaining parameter pairs are modelled:

$$(KD = 0, RA = 0) \sim \mu_0 \tag{4.2}$$

$$Smooth\_Stop \sim \mu_0 + \tau_{Smooth\_Stop} \tag{4.3}$$

$$(KD = \alpha, RA = 0) \sim \mu_0 + KD_\alpha \tag{4.4}$$

$$(KD = \alpha, RA = \beta) \sim \mu_0 + KD_\alpha + KD_\alpha : RA_\beta \tag{4.5}$$

Here, τ_{Smooth_Stop} refers to the difference in survey response between the Smooth Stop and the (KD=0, RA=0) condition; KD_α refers to the main effect of $KD = \alpha \in [10,\ 19,\ 40]$ compared to KD=0 where α is one of the non-zero KD levels tested; RA_β refers to the main effect of $RA = \beta \in [1,\ 2,\ 3,\ 4]$ compared to RA=0, where β is one of the non-zero RA levels tested; $KD_\alpha : RA_\beta$ is the interaction effect between KD_α and RA_β. This approach allows us to identify the effects of the NHG conditions against the control condition while taking into account the main effects of KD and RA.

The internal reliability of the measures was computed for all of the standardized questionnaires used. All except the *Usefulness* measure have a Chronbach's $\alpha > .70$ (see Table 4.1). Therefore, the *Usefulness* measure is excluded from the analysis.

All of the statistical analyses presented below use a significance level of $\alpha = .05$. The Tukey method was used to conduct post-hoc analyses. The presentation of the results below is organized according to the questionnaire measures. Table 4.2 provides a summary of the findings.

Hesitancy

Results from the MLM analysis on *Hesitancy* demonstrate that the NHG-generated robot motions are perceived to be significantly more hesitant than that of the control condition ($F(1,\ 692) = 30.4,\ p < .001$), providing support for Hypothesis 4.1. It also shows that there are significant differences in *Hesitancy* across the different values of KD ($F(3,\ 692) = 3.98,\ p < .01$) and RA ($F(4,\ 692) = 7.80,\ p < .001$). These main effects of KD and RA can be interpreted independently since the interaction effect between the two is negligible ($X^2(6) = 10.5,\ p = .11$).[18]

[18] This Chi-square result reflects the difference between the two regression models built to predict the same response variable – in this case, *Hesitancy*. For all of the response variables discussed in this section, a regression model that includes an interaction effect between KD and RA (Equation 4.5) is compared against a model that does not. A significant finding in the Chi-square result from this analysis indicates that the interaction effect of KD and RA plays a significant role in predicting the

Table 4.2: Findings from Study 4 are summarized here in the order of the hypotheses tested.

Measure	Result	Conclusion
Hypothesis 4.1 Robot hesitation responses generated by the NHG are perceived to be more hesitant and persistent than a smooth stopping behaviour.		
Hesitancy	$(F(1,692) = 30.1,\ p < .001)$	Supported
Persistency	$(F(1,686) = 16.3,\ p < .001)$	Supported
Hypothesis 4.2 Robot hesitation responses generated by the NHG are perceived to be more animate and anthropomorphic than a smooth stopping behaviour.		
Animacy	$(F(1,692) = 11.9,\ p < .001)$	Supported
Anthropomorphism	$(F(1,692) = 10.4,\ p < .01)$	Supported
Hypothesis 4.3 Robot hesitation responses generated by the NHG are perceived to be more dominant and useful than a smooth stopping behaviour.		
Dominance	Significant effect in the opposite direction	Not supported
Usefulness	Internal reliability not sufficient	Inconclusive
Hypothesis 4.4 Robot hesitation responses generated with a human-inspired kickback distance ($0 \leq KD < 19$ mm) are perceived to be more hesitant than those with KD outside this range.		
Hesitancy		Not supported
Hypothesis 4.5 Robot hesitation responses showing a larger number of RA are perceived to be more persistent, dominant and useful than those with a larger RA.		
Persistency	Significant KD & RA interaction effect	Inconclusive
Dominance	Significant KD & RA interaction effect	Inconclusive
Usefulness	Internal reliability not sufficient	Inconclusive

Post-hoc analyses across the levels of KD indicate that robot motions generated with KD=0 are perceived to convey a significantly higher level of *Hesitancy* than those with KD=10 and KD=19 (see Figure 4.14). This is interesting given that the robot does not perform any re-attempts in the (KD=0, RA=0) condition – the only condition with a value of KD= 0 – and, therefore, is the least negotiative of all NHG conditions. This result does not support the Hypothesis 4.4 that robot hesitation responses generated with a human-inspired kickback distance ($0 \leq KD < 19$ mm) are perceived to be more hesitant than those with higher values of KD. Rather, this result demonstrates that there are significant differences on *Hesitancy* among the human-inspired values of KD. In addition, the results suggest that *Gender* significantly affects the perception of *Hesitancy* ($F(1,\ 692) = 6.50,\ p < .01$). Across the conditions tested, female participants rated robots to be more hesitant than males ($Z = -2.89,\ p < .01$).

response variable.

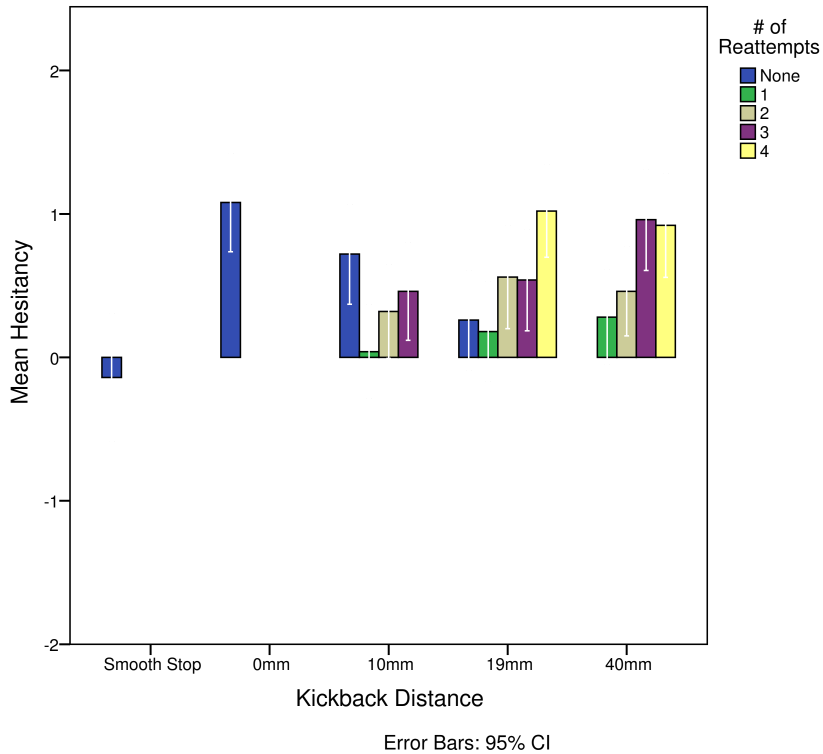

Figure 4.14: The perceived *Hesitancy* scores (a 5-point Likert scale) collected for the different levels of KD and RA. The scores are higher for the case of KD$=0$ than most of the other NHG-generated motions. Analyses of the scores across all NHG-generated motions suggest that there are significant differences in *Hesitancy* among the human-inspired values of KD.

Persistency

The analysis of *Persistency* scores suggests that there is a significant interaction effect between KD and RA ($X^2(6) = 12.65,\ p < .05$). This indicates that the effects of factors KD and RA cannot be interpreted independently from each other. Hence, the results do not offer a straight-forward support for Hypothesis 4.5 that robot hesitation responses showing a higher number of RA are perceived to be more persistent than those with lower RA values.

Nonetheless, the analysis provides an empirical support that the NHG-generated robot motions are perceived to be significantly more persistent than smooth stop motions ($F(1,\ 686) = 16.3,\ p < .001$). As shown in Figure 4.15, the *Persistency* score is higher for KD$= [10,\ 19]$ for most non-zero RA values than that of the control condition or the reference (KD=0, RA=0) condition, although this is not true for all values of RA.

Gender is not found to be a significant factor of *Persistency* ($X^2(2) = 4.87,\ p = .09$).

Animacy and Anthropomorphism

Results of the analysis on *Animacy* and *Anthropomorphism* suggest that the NHG-generated motions are perceived to be more animate and anthropomorphic than the smooth stopping motions in the control

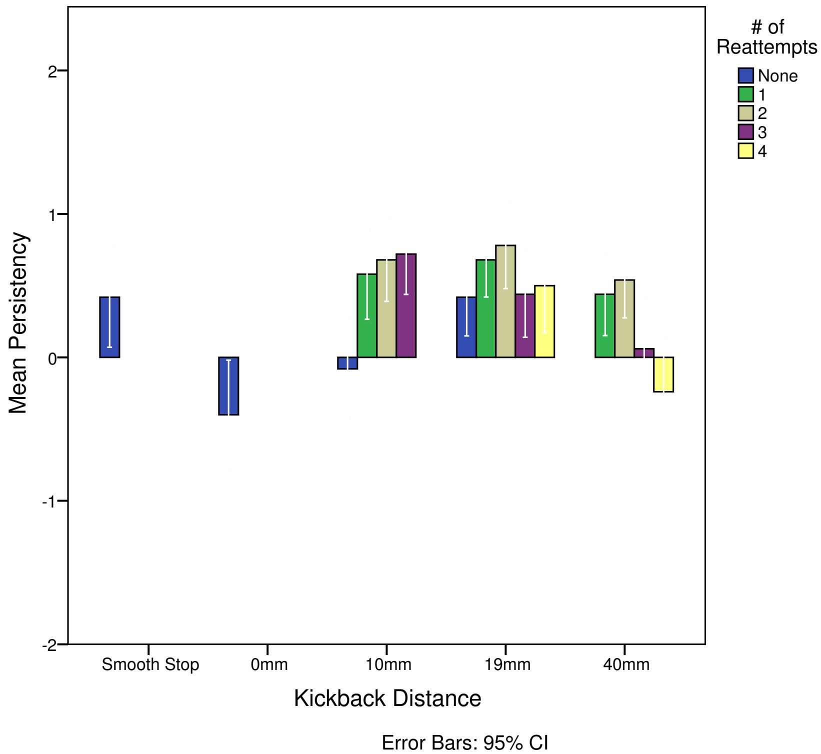

Figure 4.15: Perceived *Persistency* of different KD by RA values. There is a significant interaction effect between KD and RA. Nonetheless, NHG-generated motions are perceived to be significantly more persistent than the motion implemented in the Smooth Stop condition.

condition (*Animacy*: $F(1,\ 692) = 11.9,\ p < .001$; *Anthropomorphism*: $F(1,\ 692) = 10.4,\ p < .01$). This supports Hypothesis 4.2. No significant interaction effect between KD and RA is found for either of the measures (*Animacy*: $X^2(6) = 6.24,\ p < .40$; *Anthropomorphism*: $X^2(6) = 5.21,\ p = .52$). No significant *Gender* effect is found in either of the measures (*Animacy*: $X^2(2) = 2.54,\ p = .28$; *Anthropomorphism*: $X^2(2) = 4.28,\ p = .12$).

The results of the analysis on *Animacy* and *Anthropomorphism*, along with *Hesitancy* mentioned above, provide strong support for Hypotheses 4.1 and 4.2 that robot conflict responses generated by the NHG are perceived to be more humanlike hesitations than the smooth stopping behaviour typically used in industries. Figure 4.16 presents perceived *Animacy*, and Figure 4.17 presents perceived *Anthropomorphism* scores.

Dominance

Results of the regression analysis on *Dominance* provide evidence rejecting Hypothesis 4.3. Contrary to the hypothesis, the results suggest that the NHG-generated motions are perceived to be significantly *less* dominant than the smooth stopping motions of the control condition ($F(1,\ 686) = 35.2,\ p < .001$). The results also demonstrate that the male participants rated the robot's motions to be more dominant than the females did ($X^2(2) = 6.55,\ p < .05$). In exploring the effects of KD and RA on perceived

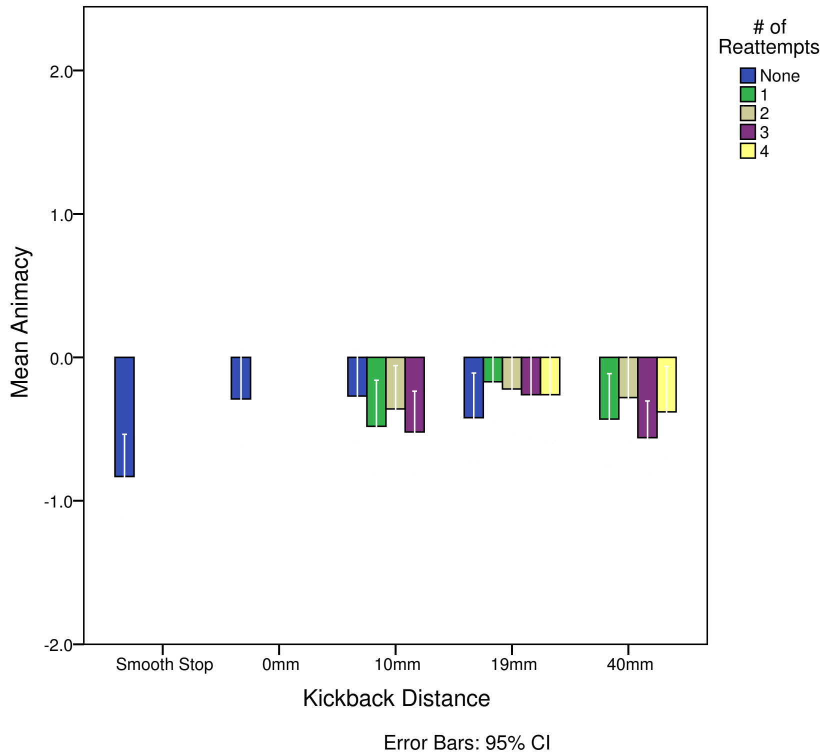

Figure 4.16: Perceived *Animacy* of different KD by RA values. The NHG-generated behaviours are, in general, perceived to be more animate than that smooth stopping behaviour demonstrated in the Smooth Stop condition.

dominance of NHG-generated motions, the results find that there is a significant interaction effect between the two variables ($X^2(6) = 31.6,\ p < .001$). Hence, the effects of factors KD and RA cannot be interpreted independently from each other, and support for Hypothesis 4.5 remains inconclusive. See Figure 4.18 for the distribution of the *Dominance* measure.

Qualitative Feedback

The adjectives the survey participants optionally submitted were collected for qualitative analysis. Keeping in mind the principles of grounded theory, two experimenters coded the adjective entries and extracted emerging themes with respect to the above-mentioned hypotheses. The responses captured using an open text field were mostly adjectives, as requested in the questionnaire. However, some participants included full sentences. The author included these responses in the analysis. Since these entries were not a mandatory part of the survey, only 50% of the participants provided at least one adjective throughout the survey. Most of these participants submitted responses to four or fewer conditions. Given the small number of responses collected, this qualitative analysis serves the purpose of complementing the quantitative findings discussed above.

The types of adjectives that participants used fell into approximately three categories: adjectives describing the perceived expression or conveyed personality of the robot (e.g., hesitant, confident, scared),

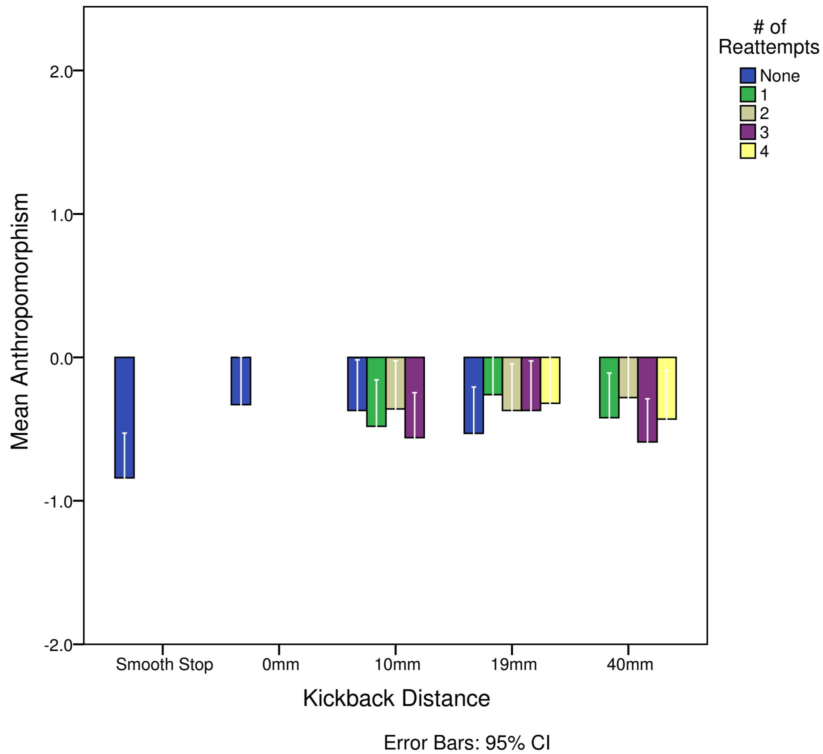

Figure 4.17: Perceived *Anthropomorphism* of different KD by RA values. Echoing the results of *Animacy*, NHG-generated motions are perceived to be more anthropomorphic than the smooth stopping behaviours of the robot in the Smooth Stop condition.

motion quality of the robot (e.g., smooth, jerky), perceived anthropomorphism (e.g., humanlike, calculated, artificial), and perceived performance of the robot (e.g., efficient, smart). In all of the conditions, one of the words 'hesitant' and 'persistent' was mentioned by at least one participant in all but condition (KD=40, RA=4).

When the participant responses to Smooth Stop and (KD=0, RA=0) are compared, both of the robot motions are described as hesitant. However, whereas the Smooth Stop condition is described with words hinting at low anthropomorphism (e.g., awkward, cautious, mechanical and robotic), motion quality of (KD=0, RA=0) is described with positive adjectives such as safe and smart. The Smooth Stop condition is described as expressing a sense of dominance (e.g., aggressive and assertive), whereas the (KD=0, RA=0) condition is described to be apprehensive and wary, despite being responsive (e.g., pushover, scared, submissive). Based on the quantitative results mentioned above, the (KD=0, RA=0) condition received the lowest mean *Persistency* and *Dominance* scores, while scoring highly on the *Hesitancy* measure. The qualitative description of the robot's motion in this condition supports and complements these results. The Smooth Stop condition also received the lowest *Animacy* and *Anthropomorphism* scores, which is corroborated in the qualitative finding.

In comparison, the participants perceived the NHG-generated motions with KD=10 and 19 to be efficient, calculated, and cautious. The robot motions with KD=19, in particular, were described as more

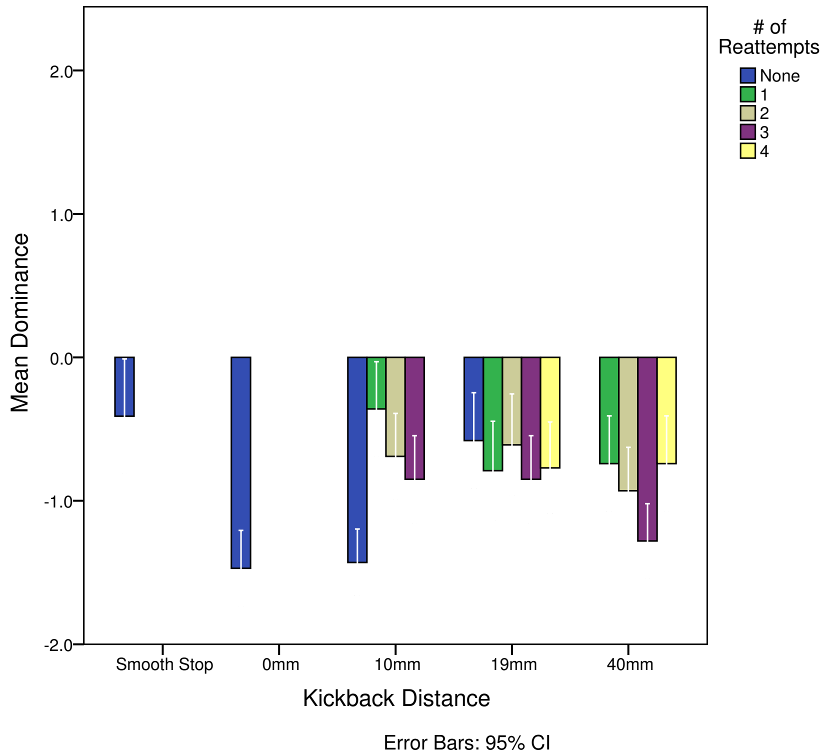

Figure 4.18: Perceived *Dominance* of different KD by RA values. Contrary to Hypothesis 4.3, NHG-generated motions were perceived to be less dominant than the smooth stopping alternative. There is an interaction effect between KD and RA as demonstrated by the inconsistent distribution of *Dominance* across the two factors.

life-like and aware across the RA values, except for RA=4, which was described to be submissive and unsure. The NHG-generated motions with KD=40 were described as timid and shy. This affirms the low *Dominance* scores collected for these conditions. In addition, KD=40 motions with lower RA values were perceived as compromising and slow, while those with higher RA values were seen as scared.

4.6.4 Discussion

The purpose of this study was to validate the efficacy of the NHG as a mechanism to generate human recognizable robot hesitation behaviours, and to establish the effects of KD and RA on human perception of the generated motions. The main finding of this study is that the NHG-generated motions are indeed distinguished from smooth stopping behaviours of a robot. The motions produced by the NHG were perceived to be more hesitant and persistent. They were also observed to be more animate and anthropomorphic. These results provide empirical support for using the NHG for in-person HRI where human behavioural responses to the artificial hesitations can be observed.

However, there were also surprising findings. First, counterintuitively, the (KD=0, RA=0) condition received the highest *Hesitancy* score. The qualitative understanding of the participants' perception of the motion offers some explanation for this finding (i.e., the robot motion with this parameter setting was considered to constitute a highly apprehensive and hesitant behaviour). Moreover, this condition received the lowest *Persistency* score. This is consistent with the concept of persistency established in Section 4.2, in which a higher value of RA would be considered to exhibit a higher *Persistency*. This interpretation does not generalize across all levels of KD, as demonstrated by the significant interaction effect of KD and RA observed in the *Persistency* scores. Given that this condition did not include any re-attempt behaviours and, therefore, was non-negotiative in nature, it does not make sense to use this set of parameters to investigate in-person HR non-verbal negotiations. On the other hand, the human-inspired KD values, 10 and 19, received high *Persistency* scores across the various RA values. Moreover, taking into account variabilities by RA values, KD=19 received a relatively high level of *Hesitancy* in comparison to other values of KD in general. This also coincides with the average KD value discovered in HH hesitations (see Section 4.4.5), and confirms the efficacy of using this value for generating artificial hesitations with the NHG.

Second, it was unexpected to find that, contrary to Hypothesis 4.3, the motions produced by the NHG are perceived to be less dominant than that of the smooth stopping motions. This is surprising because the smooth stopping motion employed in this study only had a short pausing behaviour before the robot retracted back. This is in contrast to motions generated by the NHG with non-zero RA values that exhibited multiple back-and-forth motions that could be perceived as more dominant than a pause. In observing the distribution of *Dominance* scores, in general, motions with a high RA value are seen to be less dominant than those with a lower RA – although this trend is not significant. This result suggests the possibility that the resulting negotiative hesitation behaviours designed to exhibit multiple RA are perceived to be timid and apprehensive.

4.7 Conclusion

The work presented in this chapter establishes a framework with which one can investigate the role robot negotiative hesitations can play in an interactive HR resolution of resource conflicts.

To this end, naturally occurring negotiative hesitations in humans were collected from an HHI experiment (Study 3). A video-based online study involving 300 participants helped rate each sample of hesitations on its expressed level of *Hesitancy* and *Persistency*. Exploration of the numerous features of the human hesitation trajectories revealed that, unlike R-type hesitation behaviours investigated in the author's previous work, negotiative hesitations take into consideration motions of both agents. This led to the discovery of hesitation loops, which are trajectory patterns found in the majority of the negotiative hesitation samples collected in Study 3. The design of the NHG was inspired by the hesitation loop trajectory patterns. It was designed such that the change in the relative distance between the two interacting agents to the shared target is used to determine when the robot should trigger its hesitation behaviour.

In implementing the NHG for in-person HRI, it is important that the trajectory output of the NHG be established as communicative of hesitation. Study 4 helped validate the efficacy of the NHG-based hesitation behaviours. It provides empirical evidence that robot motions generated by the NHG are perceived to be more hesitant, animate, and anthropomorphic than those of a smooth stopping behaviour of the same robot. In addition, it helped better understand the relationship various parameter values of the NHG have on human perception of the robot motions generated. For example, the results suggest that a KD value of 19 mm, an approximate average of the KD values found in human negotiative hesitations, is acceptable for producing communicative hesitation gestures.

Taken together, this chapter contributes to a better understanding of negotiative hesitations naturally found in HHI. It also provides a validated trajectory generator with which one can explore the dynamics that could emerge as HR nonverbal negotiation of resource conflicts. In the following chapter, the author takes the first step in investigating such in-person HRI using NHG-generated robot hesitations.

Chapter 5

Study 5: Bidirectional Interweaving of Subplans using Negotiative Interaction in Human-Robot Collaborative Assembly

5.1 Introduction

In Chapter 3 the author addressed the unidirectional perspective of how robots exert influence on our behaviours in a robot-to-human handover context. Chapter 4 presented a novel mechanism, the NHG, that permits exploration of the dynamics of bidirectional interweaving (negotiation) of subplans between an HR collaborating dyad. The results of a video-based online survey, Study 4, suggested that the artificial robot hesitation behaviours generated using the NHG are observed as expressing a state of hesitation. This provides crucial building blocks for investigating the second research question of this thesis: “Can a robot nonverbally negotiate with a person about what should happen in an interaction?” Study 5 presented in this chapter builds on the previous chapter’s investigation of hesitation to address this question. The main focus of this in-person study is to investigate whether the generated hesitation behaviour of a robot can be used in an HR collaboration to interactively negotiate for access to a shared resource with a human user in real-time.

Given the numerous studies on nonverbal HRI – including the results of Study 2 (Chapter 3) – that suggest that robot motions can influence human behaviours, there are reasons to believe that the answer to the research question is trivial and affirmative. However, previous work in HRI also provides evidence of interactions where human users dominate and jeopardize the efficacy of the interaction, or even obstruct the robot from performing its task.[1] For example, in a study involving the use of robots in a mall environment in Japan, Brscić et al. [25] found that the robots were subject to both verbal and physical harassment by children. In order to prevent physical damage to the robots, the researchers devised a novel path-planning algorithm that preemptively prevents robots being abused by children.

[1] Brscić et al. [25] defined the term *robot abuse* to refer to the increasingly documented “persistent offensive action, either verbal or nonverbal, or physical violence that violates the robot’s role or its human-like (or animal-like) nature.”

A previous in-person experiment conducted by the author and her colleagues [105] also provides a relevant example. In the experiment, an HR pair engaged in a collaborative assembly task in which they often coincidentally reached for the same resource at the same time. The robot was programmed to respond to the resource conflict in one of three ways: ignore and risk the possibility of colliding with the user; exhibit an R-type hesitation behaviour and retract; or trigger an emergency-like stop response before retracting. Results from the study showed that when the robot immediately yielded to the user (as was the case in the conditions involving hesitation or emergency-like stopping behaviours), the HR pair took significantly longer to complete the assembly task than when the robot ignored and sometimes even collided with the user to accomplish the task.[2] These studies suggest that timid behaviours in a robot may not be useful or desirable in HRI, and that such interactions with a robot, over time, could cause the robot to be ignored by users.

There are also reports about robots deployed outside of laboratory environments that echo this trend. For example, Dietmar Exler, the current Chief Executive Officer of Mercedes-Benz USA, noted that bullying of self-driving vehicles by humans is one of the hindrances when deploying such vehicles onto the roads [100]. He remarked that human drivers may take advantage of autonomous, safety-prioritizing driving patterns. This concern is shared by others in the automotive industry. Volvo, for example, stated that the self-driving vehicles to be pilot tested on the streets of London in 2018 would not be flagged as self-driving because of this concern [38]. This decision by Volvo came after the release of a report by the London School of Economics and City University of London on public perception of, and attitudes toward, self-driving vehicles [152]. The report indicates that drivers who have a combative style of driving also feel that they could take advantage of self-driving vehicles. This potential for humans to suboptimally share the road with self-driving vehicles results from the human expectation that these vehicles will strictly adhere to safe driving practices at all times, and will always yield to drivers who cut them off.

Hence, if we wish to see a future in which robots share our physical world in friendly and efficient ways, it is important to investigate novel behaviours that robots can use to proactively address conflicts about spaces and resources shared with humans. One possibility is to design communicative behaviours that robots can use to negotiate for their right of way or access to a shared resource.

As demonstrated in Chapter 4, humans use hesitation behaviours to interactively negotiate for solutions to spontaneously occurring resource conflicts. The description of the NHG captures aspects of these human hesitation characteristics. The results of Study 4 helped validate the NHG's effectiveness in generating artificial hesitation motions for a robotic arm. Building on the findings from Chapter 4, this chapter presents a within-subjects study, Study 5, that contributes to a better understanding of hesitation-based HR negotiations. Study 5 consists of an in-person HRI experiment in which human participants interacted with a robot in an HR collaborative assembly task. Each participant performed the same task four times, twice with a robot responding to the resource conflicts with artificial hesitations (the *Negotiate* condition), and twice with the robot smoothly stopping to avoid collision with the participant (the *Stop* condition).

[2] Safety measures were taken for this experiment such that collisions with the robot were harmless.

The following hypotheses were tested:

Hypothesis 5.1 *Robot conflict responses generated by the* NHG *are perceived more favorably than the smooth stopping behaviour often used in industries.*

Hypothesis 5.1a *Smooth stopping behaviours are perceived as less animate, anthropomorphic, and likeable than hesitation responses generated by the* NHG.

Hypothesis 5.1b *Responses generated by the* NHG *are perceived to be less dominant, more emotionally satisfactory, and more useful than a smooth stopping behaviour.*

Hypothesis 5.2 *Hesitation responses generated by the* NHG *enable faster completion of the task than a smooth stopping behaviour.*

Hypothesis 5.3 *Hesitation responses generated by the* NHG *do not threaten perceived safety, nor jeopardize the actual safety, of the interaction in comparison to a smooth stopping behaviour.*

Hypothesis 5.4 *Hesitation responses generated using the* NHG *will lead to a faster resolution of resource conflicts with humans than can be achieved with a smooth stopping behaviour.*

Continuing the discussion of the literature on hesitation, Section 5.1.1 provides a literature review on hesitations observed as behavioural responses exhibited by humans and animals. Subsequently, Section 5.2.1 outlines the experimental task in more detail, including a pilot study that was conducted to inform the experiment design (Section 5.2.3). Following the experimental procedure outlined in Section 5.2, technical details of the robotic system are presented in Section 5.3. Findings from this study are discussed in Section 5.4. These findings not only support the hypothesis that nonverbal HR negotiation is possible using this paradigm, but also provide evidence that such negotiative interaction can improve the performance of the HR team without jeopardizing the perceived or actual safety of the human. The implications of the results are discussed in Section 5.5, followed by a conclusion in Section 5.6.

5.1.1 Background

The discussion of hesitations in Section 4.2 establishes that verbal and nonverbal expression of hesitation behaviours have been studied in psychology as behaviours closely linked with uncertainty or indecision. It also suggests that observing human hesitations and expressing analogous robot hesitations are increasingly becoming useful in the study of HRI. This section extends the discussion on hesitation by focusing on the nature of nonverbal hesitation behaviours and their relationship to the stimuli that trigger them.

Hesitations in the form of a kinetic dialogue are found in primitive animals as well as humans. For example, Levin [87] studied the approach-withdrawal behaviour of fish to investigate how fish decide to flee, pursue, or ignore something when the difference between prey, predator, and the natural environment is difficult to perceive. While a fish should pursue its prey, it must flee from its predators in order to survive. Levin [87] reports that fish stop swimming in uncertain cases. Analogous hesitations in humans have also been studied as a kinetic response to situations that require either action or cessation of action. Netick and Klapp [114] observed hesitation behaviours in human participants engaged in a continuous tracking task. In investigating what behavioural processes are involved in hesitation behaviours (which they define as "a reaction in which an ongoing action is halted"), Netick and Klapp posit that hesitations are a response to a stimulus, rather than a part of the strategic processing of the stimulus. They theorize that rather than a strategic and controlled response to a stimulus, kinetic hesitations are reflexive behaviours that occur when attention is diverted from an on-going task, thereby interrupting control over the task, which results in the freezing of an action.[3]

This theory of the reflexive nature of human hesitation is crucial to the understanding of the role that hesitation-based kinetic negotiations can play in HRI. First, it presents hesitations as a universal behaviour that can be used to kinetically resolve conflicts between an HR pair over a shared resource. Second, having a hesitation-based kinetic dialogue may be intuitive even for a naïve user, given that hesitation is a universal response that can be observed and understood by humans even from the motions (or lack thereof) of primitive animals such as fish.

Moreover, the reflexive nature of hesitation responses makes hesitation a promising technique to enable a robot to kinetically engage in a dialogue with humans without its actions being undesirably dominated – or bullied as described in the introduction. For example, in Austin, Texas, a cyclist and an autonomous car unintentionally engaged in a nonverbal dialogue as they negotiated each other's right of way at a four-way stop [40]. The car had stopped slightly earlier than the cyclist and had the right of way. The cyclist recounts:

> "The car remained motionless for several seconds and I continued to maintain my balance without moving. As the car finally inched forward, I was forced to rock the handlebars to hold my position. Apparently, this motion was detected by one of the sensors and the car stopped abruptly. I held the bike in balance and waited for another several seconds and the cycle repeated itself ... the car inched forward, I shifted my weight and re-positioned the bars and the car stopped. We did this little dance three separate times and the car was not even halfway through the intersection."

In this incident, the system was confused about the state of the cyclist who balanced his bicycle by rocking the handlebars rather than putting his foot down on the ground. It is not an unfamiliar experience for most drivers to have nonverbal dialogues such as this one at an intersection that somehow result in a resolution of the conflict using purely nonverbal means. This example suggests that equipping a robot with an ability to negotiate its right of way using hesitation behaviours could lead to the development of a natural and interactive mode of conflict resolution for HRI.

[3] To support this claim, Netick and Klapp [114] provide evidence that electromyographic recordings of human hesitations from their study resemble the patterns of electromyographic signals that are expected in active cessation of movement.

5.2 Experimental Procedure

This section outlines the procedures of a within-subjects HRI experiment conducted in a controlled laboratory environment to examine the impact that negotiative robot behaviours have on real-time HR collaboration. This project was conducted as a collaboration between the Learning Algorithms and Systems Laboratory (LASA) at EPFL and the Collaborative Advanced Robotics and Intelligent Systems (CARIS) laboratory at UBC. The experiment itself was conducted at LASA, EPFL, Lausanne, Switzerland.

The devised experimental context consisted of a collaborative task requiring both the participant and a robotic manipulator to frequently access a shared resource. The author conducted a pilot study, consisting of the *Yield* and *Negotiate* conditions, to fine-tune the HR collaborative task devised for this study. The pilot study also served to inform the design of a control condition, the *Stop* condition, such that the negotiative robot behaviour, the *Negotiate* condition described below, is compared against a non-trivial alternative. This selection process also helped minimize the number of tested conditions to two in the full study, thereby limiting the duration of the experiment to an hour and minimizing possible effects of fatigue.

Participants were recruited from within and outside of the EPFL campus and were offered 20 CHF compensation for their participation. At the beginning of the experiment, participants provided informed consent and were asked to fill out a short demographics questionnaire on age, gender, dominant hand, and their familiarity in working with robotic arms. They were then introduced to the 7-DOF robotic arm, KUKA LWR 4+ (KUKA Robot Group, Augsburg, Germany), and the experimental task to be performed in collaboration with the robot. Afterwards, participants sat facing the robot arm with a small table between them. Figure 5.1 illustrates the HR workspace setup. The experimenter placed clusters of reflective OptiTrack markers on the participants' dominant forearms and hands to track the participants' motions in real-time. The entire experiment was video recorded. It was also time-synchronized with human and robot trajectory recordings via integration of the OptiTrack system and the robot's joint trajectory readings.

5.2.1 Experimental Task

The participants were told that the goal of the experimental task was for the participant and the robot to collaborate with each other to create a small concoction. The concoction was to consist of two different liquids[4] and orange-coloured dry lentil beans. The participant's job was to sort orange-coloured lentils from a mix of orange and green lentil beans in a bin and place them onto a dispenser in front of them. The robot's job was to pump the two liquids in a specific order into a transparent cup, and also to press the knob in the centre of the dispenser so as to flush the orange lentils into the cup.

The participants were told that the robot had a fixed number of pumps it needed to perform before the concoction would be considered complete, and that the team's performance would be measured based on how many orange lentils were in the concoction at the end of the trial, as well as on how quickly the team finished the task. They were also told that incorrect-coloured lentils in the concoction would be

[4] Coffee and water were used, although not specified to the participant.

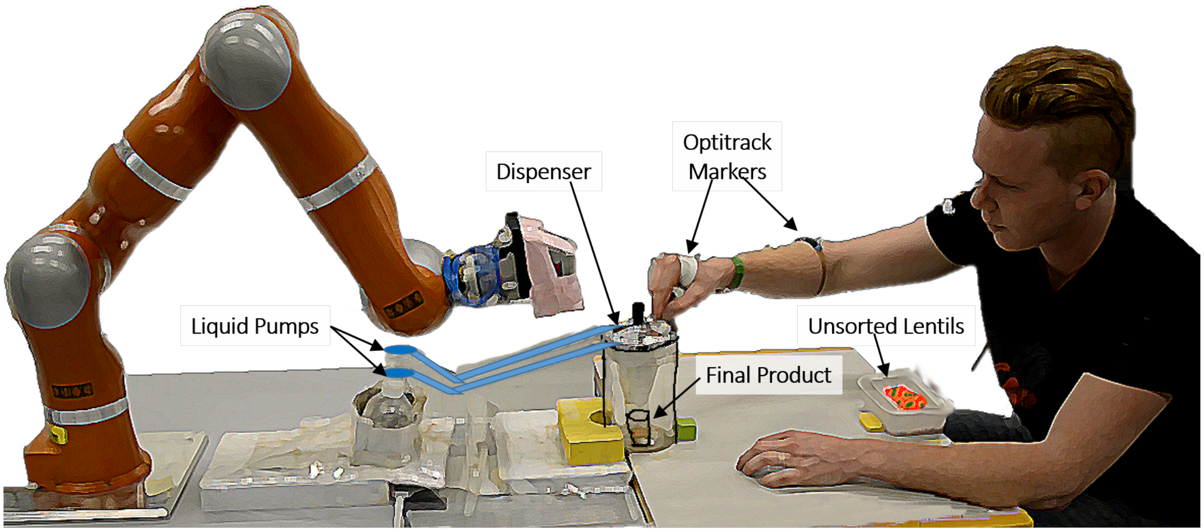

Figure 5.1: Experiment set-up of Study 5. The robot and a participant sat facing each other with a dispenser mechanism located on the table between them. The shared use of the dispenser allowed spontaneous HR resource conflicts to occur throughout the experiment.

counted as a penalty for the team. The dispenser was intentionally made of transparent material so that participants could monitor the progress of the task based on the amount of liquids and solids deposited into the cup. This also served to provide a visual confirmation that the lentil sorting by the participant, and the liquid pumped by the robot, came together as one final product.

Each trial started with the robot reaching for the dispenser and pressing on the dispenser knob to demonstrate the robot's reaching motion. The participants were asked to start immediately after this demonstration. Meanwhile, the experimenter sat diagonally behind the participants, out of their immediate field of view to ensure that the experimenter would not cause unintentional visual bias. This arrangement allowed the experimenter to observe the workspace continuously, so as to be able to trigger an emergency stop if necessary.

The two experimental conditions, described in detail in Section 5.2.4, were randomized in such a way that each participant encountered the two conditions in the first two trials of the experiment in random order. They also encountered the two conditions again in another random order in the last two trials of the experiment. In total, each participant completed four trials of the experimental task. Figure 5.2 provides a visual overview of this procedure.

5.2.2 Questionnaire and Interview

At the end of each trial, the participants were accompanied away from the robot to fill out a questionnaire. In order to avoid participant bias toward the robot in their reports (as reported as possible in Reeves and Nass [130]), the room was configured in such a way that the participants faced away from the robot while filling out the questionnaire. This also served as a washout period during which the participant was distracted while the experimenter prepared the robot and the experimental workspace

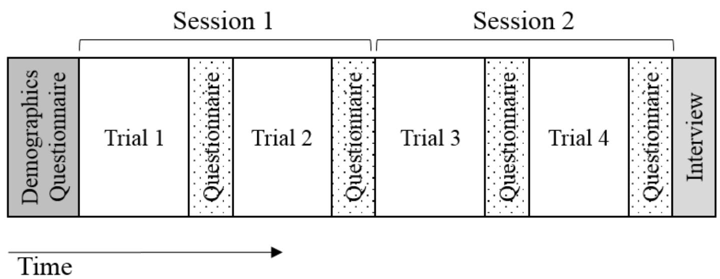

Figure 5.2: Experiment procedure overview of Study 5. The two conditions (*Stop* and *Negotiate*) were randomly assigned to the trials for each session, such that the participants encountered both conditions at least once in Session 1 and once again in Session 2. The questionnaires between the trials served as a wash-out period. The transition between the two sessions were not announced to the participants.

for subsequent trials.

The questions in the questionnaire were selected from two widely used questionnaires: 1) Godspeed questionnaire [14] used to measure *Animacy*, *Anthropomorphism*, *Perceived Safety*, *Likeability*, and *Perceived Intelligence* and 2) Moon (no relation to the author) and Nass questionnaire [107] used to evaluate human-machine teamwork including *Dominance*, *Usefulness*, and *Emotional Satisfaction*. Internal reliability values from a similar previous in-person experiment, Moon et al. [104], informed the selection of the most promising set of questions for the eight measures. These standardized user perception and teamwork measures helped collate and contrast the participants' impressions of the robot across the two experimental conditions.

At the end of the fourth and last trial, the experimenter conducted a semi-structured interview to collect qualitative feedback on the participants' perception of the robot, and preference for and interpretation of different robot behaviours in the two experimental conditions. The interview was video recorded with the camera angled in such a way that it captured the participants' hand gestures without recording their faces.

5.2.3 Pilot Study

A pilot study was conducted in order to design conflict response behaviours for the robot that would be comparable to the devised hesitation behaviours. Two conditions were tested in the pilot study: *Yield* and *Negotiate*.

Yield The *Yield* condition is comprised of the same components of the control (Smooth Stop) condition implemented and tested in Study 4. In this condition, the robot was equipped to respond using a human-inspired, R-type hesitation controller called the Acceleration-based Hesitation Profile (AHP) controller [104]. By design, the controller requires the robot to automatically yield to the person upon detection of an imminent conflict with no continued nonverbal dialogue between the robot and the person. In addition, this controller is limited in terms of the relative time-frame within which the human-inspired hesitation response can be triggered. If the robot detects that a conflict persists past the time window immediately after the robot has peaked in its forward acceleration toward the person (practically immediately after the robot has started its reaching motion from its starting position), the robot was programmed to come to a stop (a smooth stop) and hold its position for one second before retreating. The robot was programmed to re-attempt (i.e., reach for the resource again immediately after returning to the starting position) a maximum of three times before giving up on access to the resource and resuming its task of tending to the liquid bottles.

Negotiate The *Negotiate* condition was designed using the negotiative hesitation controller described in the previous chapters and further discussed in Section 5.2.4. In contrast to the *Yield* condition, the robot does not yield to the person immediately upon detection of a conflict, and actively reattempts to access the shared resource until a resolution of the conflict is reached. Similarly to the *Yield* condition, the robot was programmed to re-attempt access to the resource three times before giving up and returning to the starting position and resuming its other tasks.

Using the procedure described above, the experimenter conducted the pilot study with seven biased participants (two females, five males).[5] The average age of the participants was 31.1 years old, with a relatively high familiarity-with-robots score of 3.14, indicative of the fact that the participants had robotics background. Due to the small number of samples used for this pilot study, no inferential statistical analyses were conducted. Also, the robot was programmed to complete a total of fourteen dispensing motions instead of the forty used in the full experiment. Performance measures suggested that there could be significant differences between the *Yield* and *Negotiate* conditions in how fast the HR team completed the task. The average duration spent per trial for the *Negotiate* condition is much shorter ($M = 149.18$, $SD = 8.39$) than that of the *Yield* condition ($M = 184.41$, $SD = 21.12$). A qualitative review of video-recorded interviews with the pilot study participants suggests that the *Negotiate* condition is preferred over the *Yield* condition.

Participants also reported that the first two trials of the experiment were exciting because of the novelty effect of working with the robot, and perceived the last two trials to be more representative experiences of working with the robot in the two conditions. This supported the experiment design decision to conduct a within-subjects, repeated-measures experiment in which participant responses to the same experimental conditions are collected twice.

Although the results from the qualitative study provide optimistic results for the *Negotiate* condition, the participants' qualitative feedback hinted that the immediate retraction behaviour of the *Yield* condition is perceived as inefficient. This retraction behaviour also unfairly disadvantages the respective trials' measured completion time for the *Yield* condition, because the extra distance the robot must travel for the retraction is absent from the *Negotiate* condition. Hence, the control condition was redesigned as the *Stop* condition outlined in Section 5.2.4. The *Stop* condition, used in the main experiment, does not have the extra travel time required by the robot to return to the starting position.

[5] These participants were colleagues and acquaintances of the experimenter, thereby having a chance of providing biased responses to the experiment.

5.2.4 Experimental Conditions

Informed by the pilot study, the following near-collision responses were implemented to form the two experimental conditions: *Stop* and *Negotiate*.

Stop The *Stop* condition was designed as a control against which to compare the more dynamic *Negotiate* condition. Instead of the NHG, which was built using a linear dynamical system (refer to Section 4.5), this condition was a modified version of an already tested hesitation controller from Moon et al. [104]. In the original controller, the author had employed a series of quintic splines with parameters derived from human hesitations. However, the original controller only generated hesitations that led to automatic yielding behaviour on the part of the robot. Also, due to the parameter constraints that are necessary to generate this yielding hesitation, the robot must resort to a different type of collision avoidance behaviour after it has travelled a short distance. Hence, in the *Stop* condition, the original controller was modified such that the robot triggers the yield-type hesitation if it can do so. Otherwise, the robot stopped as soon as a near collision situation was detected, and paused at that position for 1.0 second before reaching again. This generated a behaviour where the distance travelled by the robot in each condition is the same, except that the generated robot motion while negotiating for access to the shared resource was different.

Negotiate The *Negotiate* condition uses an implementation of the NHG controller developed in Chapter 4. This controller allows the robot to exhibit human-inspired hesitation behaviour in which the robot's persistent interest in the shared resource is expressed. In this condition, the robot reached for the dispenser, and when a near collision situation was detected the robot moved away, in a direction opposite to its original direction, by a *kickback-distance* empirically determined from Chapter 4, before attempting to reach for the shared resource again. To avoid a possible livelock[6] by perpetual hesitation on the part of the robot and the participant, for each instance of hesitation, the robot repeated this kickback motion for a maximum of 15 times before retracting to its starting position and yielding to the participant.

5.3 Technical Implementation

The same robotic system described in Section 4.6 was used in this study. Instead of the *Smooth Stop* condition, equivalent to the *Yield* condition tested in the pilot study, the *Stop* condition was implemented. Details of the trajectories generated for each condition are outlined in Section 5.3.1. The safety features of the system are presented in Section 5.3.2.

5.3.1 Trajectory Generation

The *Stop* condition employs the spline-based trajectory generator used in Moon et al. [104]. The controller uses quintic Hermite splines derived from a cubic acceleration profile. When hesitation is

[6] The author uses the definition of livelock provided by Zöbel [169] who defines it as "all forms of imprevisible delay, often called starvation, permanent blocking and indefinite delay."

triggered very early in the robot's reach, four different quintic splines are stitched together with continuous boundary conditions. This generates a smooth quintic-based trajectory that a robot can exhibit to immediately and anthropomorphically yield the resource to the participant. However, this particular system requires resource conflict to be detected at the onset of the robot's reach, limiting the robot's capability to respond to conflicts that are detected later on in the interaction. Hence, for conflicts detected afterwards, the spline-based trajectory generator was designed to smoothly stop and pause the robot for 1.0 second, wherever the robot might be at the time. Unlike the pilot study, in which the robot returned to its starting position upon detection of a conflict, the robot in the *Stop* condition was designed to continue to reach for the target (the dispenser) after the pause, thereby generating a smooth and pause-based re-attempt behaviour.

In order to generate a negotiative conflict response for the *Negotiate* condition, the NHG, a LDS (defined by $\dot{x} = mx + b$) implementation of the characteristics observed in human hesitations in Section 4.4 was employed. The LDS was tuned such that one round of reach, press, and retract motion took 3.2 seconds, matching the time it takes the robot to complete the same action in the *Stop* condition. Figure 5.3 demonstrates an NHG-generated and a spline-generated trajectories without encountering any conflict.

The linear slope, shown in the Robot to Target (R2T) plot of the figure, demonstrates the use of the LDS in reaching toward the target object. The lower right-hand corner also demonstrates a sharp change in velocity. This is because the robot moves faster the farther it is from the target due to the nature of LDS. Hence the robot quickly reached its desired velocity at the beginning of its reach motions. The curved trajectory shown on the upper left-hand side of the plot demonstrates the robot's pressing of the dispenser button. This particular motion used the same quintic trajectory generator as in the *Stop* condition. This was to ensure that the only difference between the two conditions was in the robot's conflict response, rather than other extraneous task-related trajectories such as pressing of the liquid pumps.

Upon detection of a conflict, an upper layer module modulating the LDS moved the target location of the DS to 19 mm (i.e., Kickback Distance (KD)= 19 mm) behind the current location of the robot. This choice of the NHG kickback parameter, KD, was based as the average KD distance observed in human hesitations. This effectively moved the robot backwards by KD before the upper layer module moved the target location of the DS again to the dispenser button for a re-attempt to access the resource. Figure 5.6 shows conflict response trajectories generated for the *Stop* and *Negotiate* conditions. In order to prevent the robot from reaching a livelock condition in which the robot is in motion but is never able to successfully reach the resource – presumably because of a participant refusing to yield – the experimenter counted the number of re-attempts (RA) made by the robot and triggered the robot to yield/retract after 15 re-attempts. This value of RA, a value over and above what is observed in HHI, was chosen as a large value that will help contrast the maximum observed RAs in HRI against that of HHI. It was assumed that an HR pair would not reach a situation where the robot would use all 15 RAs. See Figure 5.7 for a sample of a trajectory with multiple RAs generated during the experiment.

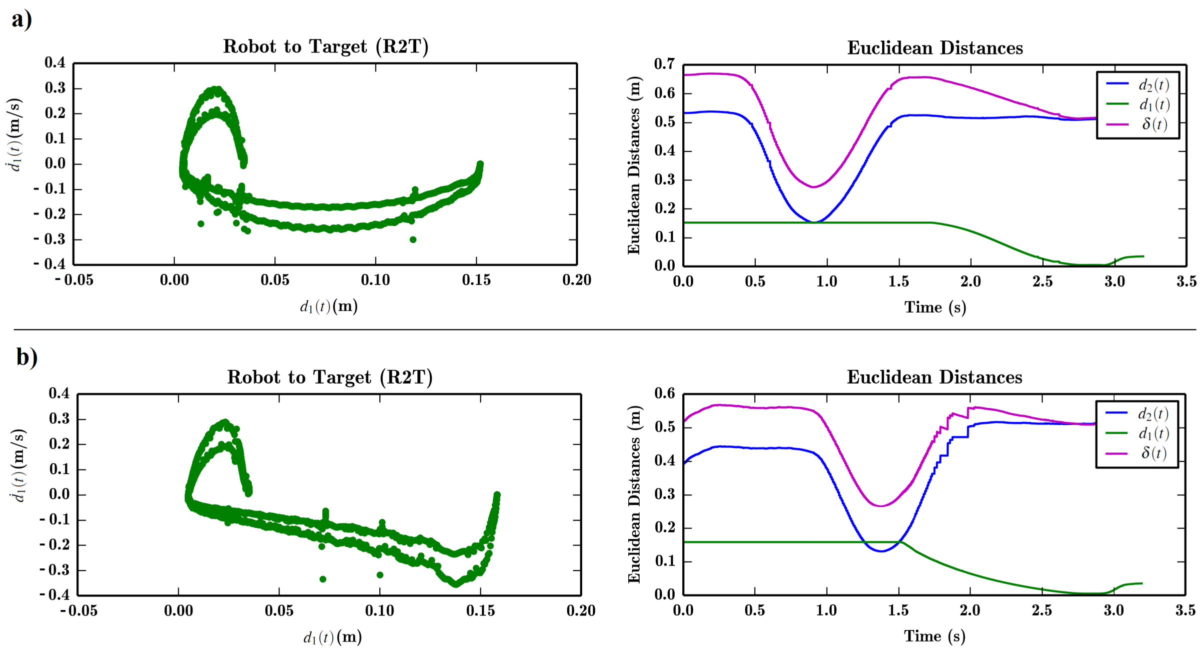

Figure 5.3: A spline-generated and NHG-generated robot motions ($d_1(t)$, $\dot{d}_1(t)$) without encountering any conflict in the *Stop* and *Negotiate* conditions. As defined in Chapter 4, $d_1(t) = \|M(t) - T\|$, in which $M(t)$ is the main agent's (robot's) Euclidean distance to the static target T, and $\dot{d}_1(t)$ represents its first derivative. Shown in both R2T plots are the robot's motion toward the target from 0.15 m away, reaching the target, and pressing down on the button (shown as the curved upside-down J shape) before returning to the original position. a) illustrates a spline-generated motion used for the *Stop* condition, and b) shows an NHG-generated motion used for the *Negotiate* condition. The plots on the right-hand column represent the respective Euclidean distances over time ($d_1(t)$, $d_2(t)$, and $\delta(t)$). $d_2(t)$ represents the Euclidean distance of the participant to the target, and $\delta(t)$ represents that difference between $d_1(t)$ and $d_2(t)$.

5.3.2 Safety Features and Motion Tracking

The system also included technical features to ensure the safety of the participants interacting with the robots. For example, the control panel of the KUKA LWR robot is equipped with an emergency button. The button must be pressed in order for the robot to operate. Hence, a release of the button brings the robot to a total stop immediately. To use this safety constraint, the experimenter sat near the workspace and continuously held the button depressed, ready at all times to release it.

In addition, the motion capture system utilized in Section 4.6, OptiTrack, was also used for this study. The motion capturing system provided a stream of the participant's hand and arm locations at a sampling frequency of 120 Hz with a known latency of 4.2 ms to sense the participant movement. The OptiTrack system yields easily identifiable values (series of zeroes) when the system fails to detect the designated markers accurately. The experimental system was designed to detect this anomaly such that the robot would come to a full and immediate stop upon losing connection with the sensor, or being unable to detect the participant's hand/arm location.

Previous studies report that the average maximum speeds at which humans feel safe while working in close proximity to a robot's working envelope is 41 to 64 cm/s depending on the size of the robot [75]. Since it is not desirable to threaten the participants, the robot was set with a velocity limit of 41 cm/s, such that the robot would stop its motion immediately if it is commanded to move at a higher speed. Note that this maximum speed is higher than what international standards – ANSI/RIA R15.06-2012 [6] and ISO 10218-1:2011 [69] – consider as a safe enough speed for a person to withdraw from hazardous motion (25 cm/s).

Since the purpose of this study is to test the collision avoidance mechanism, the aforementioned safety features were implemented with the aim of preventing any collision between human and robot. However, in the case where the robot comes to a stop and the person's hand continues to travel toward the fully stopped robot, the end effector of the robot was padded with a soft cushion-like material, such that the participant's contact with the robot was soft and harmless. This approach and velocity level have been used in a previous study Moon et al. [104] where, in certain experimental conditions, a padded moving robot collided with the participants without causing any physical harm.

In addition, the range of participant's hand motion during the manipulation task had a maximum of 6 cm overlap with the robot's end-effector at their fully extended positions. This overlapping region of the workspaces is also the end of a robot's reaching motion toward the shared object. Hence, by design, the robot always moved below its peak velocity within the workspace region it shared with the participants.

5.4 Results

Out of the 39 participants recruited, data from only 33 participants (14 females and 19 males) were analyzed. Data from the remaining six participants were rejected due to technical failures. The participants were 20 to 40 years of age ($M = 26$, $SD = 5.6$). In contrast to the biased participants recruited for the pilot study, most of the participants indicated that they were not at all familiar working with robots ($M = 0.73$, $SD = 1.2$).

This section is structured based on the type of collected measures: Analyses of participants' self-reported subjective experiences from questionnaire responses and interviews are presented in Section 5.4.1; quantitative measures of the HR teams' performance, such as task completion time and number of conflicts triggered, are given in Section 5.4.2; differences observed in human behaviours across the two conditions are in Section 5.4.3. References to the hypotheses stated in Section 5.1 are made throughout these sections as appropriate. Possible type I errors for multiple pairwise comparisons were corrected using the Bonferroni method.

5.4.1 Subjective Experience

In this section, Hypothesis 5.1 is tested using the self-reported subjective experience of the participants, which was analyzed using the questionnaire and interview responses.

Self-Reports on Perception of a Robot as a Partner

As presented in Table 5.1, all eight of the self-report measures used in this study have internal reliability greater than 0.7, and are all deemed reliable for analysis. A repeated-measures ANOVA with Condition and Encounter as within-subjects fixed effects was employed to analyze the measures and is summarized in Table 5.2. Overall, the participants viewed the *Stop* condition to be more favourable.

Regarding Hypothesis 5.1b, participants perceived the *Negotiate* condition responses to be significantly more dominant than the *Stop* condition. However, there are significant interaction effects between Condition and Order on *Usefulness* ($F(1,\ 32) = 9.66,\ p < .01$) and *Emotional Satisfaction* ($F(1,\ 32) = 6.32,\ p < .05$) measures. This indicates that the perceived *Usefulness* and *Emotional Satisfaction* of the conditions change from the participants' first encounter of the condition to the second. While this significant interaction effect neither fully supports nor rejects Hypothesis 5.1b on the *Usefulness* and *Emotional Satisfaction* measures, the results do suggest that both the *Negotiate* and the *Stop* conditions received an above-median *Emotional Satisfaction* score (3.52 and 3.86 respectively on a 5-point scale).

In contrast to Hypothesis 5.1a, the participants perceived the *Stop* condition to be more animate, anthropomorphic, and likeable than the *Negotiate* condition. The robot was also perceived to be more intelligent in the *Stop* than in the *Negotiate* condition. However, there was no significant difference as to how safe the participants perceived the robot to be in the two conditions.

There are significant order effects on some but not all of the eight measures. The participants perceived the robot to be more useful, animate, and anthropomorphic when the conditions were encountered for the second time (Session 2). The robot was also perceived to be significantly safer in the second encounters, irrespective of the condition.

Table 5.1: Internal reliabilities of the eight self-reported measures are presented here. All measures used in this study have Cronbach's alpha $\geq .70$. These values are comparable to the values reported by the original authors of the standardized questionnaires Moon and Nass [107] and Bartneck et al. [14].

Measures	Cronbach's alpha	Items
Dominance	.76	Aggressive, Dominant, Forceful
Usefulness	.79	Efficient, Helpful, Useful
Emotional Satisfaction	.83	How much did you like this robot?, How much did you like working with this robot?, Enjoyable
Perceived Safety	.83	Anxious, Agitated
Likeability	.86	Like, Kind, Pleasant, Friendly
Animacy	.71	Apathetic, Artificial, Mechanical
Anthropomorphism	.87	Artificial, Fake, Machinelike
Perceived Intelligence	.70	Foolish, Intelligent (reverse scale), Incompetent, Ignorant

Table 5.2: Repeated-measures ANOVA results are presented for all eight perception measures with Condition and Encounter as fixed factors. Since both factors only have two levels, Mauchly's Test of sphericity does not apply. $^{+}$*Likeability* score is a reverse measure, in which a lower value indicates a more favourable score.

Condition			
Measure	ANOVA	*Negotiate*, Mean (SE)	*Stop*, Mean (SE)
Dominance*	$F(1,\ 32) = 6.06,\ p < .05$	2.52 (0.16)	2.11 (0.13)
Usefulness	$F(1,\ 32) = 0.89,\ p = .35$	3.68 (0.13)	3.76 (0.10)
Emotional Satisfaction*	$F(1,\ 32) = 6.63,\ p < .05$	3.52 (0.13)	3.86 (0.10)
Animacy*	$F(1,\ 32) = 4.35,\ p < .05$	2.70 (0.12)	2.93 (0.12)
Anthropomorphism*	$F(1,\ 32) = 4.87,\ p < .05$	2.31 (0.14)	2.60 (0.14)
Perceived Safety	$F(1,\ 32) = 2.59,\ p = .12$	3.96 (0.14)	4.21 (0.14)
Perceived Intelligence**	$F(1,\ 32) = 8.16,\ p < .01$	3.33 (0.10)	3.56 (0.11)
Likeability^{*+}	$F(1,\ 32) = 6.23,\ p < .05$	2.87 (0.12)	2.52 (0.12)
Encounter			
Measure	ANOVA	1st, Mean (SE)	2nd, Mean (SE)
Dominance	$F(1,\ 32) = 0.00,\ p = 1.00$	2.32 (0.14)	2.32 (0.12)
Usefulness*	$F(1,\ 32) = 5.91,\ p < .05$	3.62 (0.11)	3.82 (0.12)
Emotional Satisfaction	$F(1,\ 32) = 0.02,\ p = .89$	3.70 (0.11)	3.69 (0.11)
Animacy*	$F(1,\ 32) = 6.91,\ p < .05$	2.66 (0.10)	2.98 (0.14)
Anthropomorphism*	$F(1,\ 32) = 5.42,\ p < .05$	2.32 (0.12)	2.60 (0.15)
Perceived Safety*	$F(1,\ 32) = 4.96,\ p < .05$	3.94 (0.15)	4.22 (0.12)
Perceived Intelligence	$F(1,\ 32) = 0.02,\ p = .90$	3.44 (0.10)	3.45 (0.11)
Likeability^{+}	$F(1,\ 32) = 0.64,\ p = .43$	2.74 (0.10)	2.65 (0.11)

Qualitative Analysis of User Interviews

To supplement some aspects of the subjective experience that may not be captured in, and to validate participants' responses to, the questionnaire, the experimenter conducted a semi-structured interview framed around the following questions:

1. Out of the four trials, which trial was your favourite?
2. Did you feel that you and the robot were collaborating together on a task, or did you feel that you were both working on a task independently from each other?
3. Was there any point in time that you felt the robot was being too aggressive?

Only 25% of the participants chose trials with the *Negotiate* condition as their favourite trial, compared to 75% who chose a *Stop* trial. This supports the quantitative finding on *Likeability*, rejecting Hypothesis 5.1a. The participants were also biased in choosing later trials as a favourite, with 78% of the participants choosing one of the last two trials over one of the first two.

In addressing the issue of the robot being too aggressive, 60% of the participants said that there was no point at which they felt the robot to be too aggressive. The remaining 40% of the participants either mentioned the way the robot pressed the dispenser, or the high jerks observed at the start of reaching motions in the *Negotiate* condition trials. These high jerks are inevitable due to the nature of LDS, where the highest forward velocity is commanded from a resting state at the onset of the robot's reach toward the target.

On the question of perceived collaboration, 57% of the participants saw the experimental task to be collaborative, whereas the remaining participants said that they worked independently from the robot, despite the fact that the final product required the contribution of both the robot and the participant.

5.4.2 Collaborative Task Performance

To better understand the full effect of the two conditions in an HR collaborative task, it is important to observe the HR pairs' overall collaborative task performance. In this section, the number of conflicts triggered within a trial, the task completion time and the teams' throughput are discussed as team performance measures.

Number of Conflicts Triggered

While the number of dispenser presses by the robot was fixed at 40 per trial, the number of conflicts the robot encountered varied due to the spontaneous nature of the conflicts triggered. Chi-square tests were conducted on this measure across Gender, Condition, and Encounter as factors.

Across the two conflict resolution responses by the robot (Condition), there is no significant difference in the number of conflicts triggered. In both conditions, approximately 37% of all robot motions in a trial encountered a conflict with a participant ($X^2(1) = 0.796,\ p = .82$). However, surprisingly, gender played a role in how many conflicts the robot had with the participants ($X^2(1) = 5.00,\ p < .05$). Female participants triggered more than the expected number of conflicts, whereas male participants triggered less than the expected number. Looking at a more detailed analysis of whether there were significant differences in the number of conflicts triggered across conditions within the same gender, there is no significant effect ($p = .21$ and $p = .14$ for female and male participant groups, respectively).

However, when looking at gender differences in the number of conflicts triggered, holding conditions constant, there are significant differences in conflicts triggered in the *Stop* condition only. In the *Negotiate* condition, 38% and 37% of all robot motions encountered conflicts with female and male participants, respectively ($X^2(1) = 0.054,\ p = .82$). In the *Stop* condition, female participants triggered a significantly larger number of conflicts (40%) than male participants (35%) ($X^2(1) = 8.59,\ p < .01$).

The robot encountered more conflicts in the second encounters than in the first ($X^2(1) = 15.8,\ p < .001$). While in the first two trials, the robot came across conflicts 35% of the time, it saw conflicts 40% of the time in the second two trials. A Pearson's chi-square analysis on each set of encounters with respect to the response conditions reveals that this training effect is common in both conditions ($X^2(1) = 0.043,\ p = .84$ for the first two trials, and $X^2(1) = 0.06,\ p = .81$ for the second two trials across the two conditions).

Task Completion Time

A repeated-measures ANOVA was conducted on task completion times, a measure representing how long it took for the HR pair to complete a trial. The start and end of a trial are defined as the robot's start of the first and end of the last reach (the 40^{th}). Factors including Condition, Gender, and Encounter were considered.

The ANOVA results indicate that there is a significant difference in the task completion time between the two conditions ($F(1,\ 31) = 100.3,\ p < .0001$) and across the two encounters ($F(1,\ 31) = 7.62,\ p < .01$). The HR team completed each trial in the *Negotiate* condition significantly faster ($M = 298.3,\ SE = 0.92$) than the trials in the *Stop* condition ($M = 320.0,\ SE = 2.43$), supporting Hypothesis 5.1b. This supports Hypothesis 5.2, that is that hesitation responses generated using the NHG enable faster completion of the task than smooth stopping behaviour. Perhaps surprisingly, the HR teams took longer in the last two trials (the second encounter, $M = 311.0,\ SE = 1.69$) than in the first two (first encounter, $M = 307.3,\ SE = 1.58$), demonstrating a reverse of the direction of the order effect one would typically expect. However, as seen in Figure 5.4, this order effect is much more pronounced in the *Stop* condition in comparison to the *Negotiate* condition. There is no significant effect of Gender ($F(1,\ 31) = 0.663,\ p = .422$), nor interaction effect between Condition, Gender, and Encounter.

Team Throughput

In this experiment, measure of throughput for an HR collaboration consists of the following measures: the number of correctly processed and incorrectly processed lentils, and two customized throughput scores as a function of the lentil count.

As per previous analyses, a repeated-measures ANOVA was conducted considering Condition, Gender, and Encounter. Table 5.3 provides a summary of the results. In both correctly and incorrectly processed products, there are no interaction effects across Conditions, Encounters, and Gender. Results indicate that female participants ($M = 89.4,\ SE = 4.26$) have a significantly higher number of correctly processed lentils ($p < .01$) than male participants ($M = 72.9,\ SE = 3.65$). Despite this significant difference in the number of lentils correctly processed, there were no significant gender differences in the number of incorrectly processed lentils ($p = .92$). This suggests that female participants were able to do the sorting task more quickly, without impacting the quality of the sorting process.

The robot's response to conflict (the Condition factor) also significantly affected the number of lentils processed. While the number of correctly processed lentils was higher ($F(1,\ 31) = 21.17,\ p < .001$) for the *Stop* condition ($M = 85.1$, $SE = 3.25$) compared to the *Negotiate* condition ($M = 77.2$, $SE = 2.58$), the number of mistakes made by the participants was also higher ($F(1,\ 31) = 6.60,\ p < .05$) in the *Stop* condition ($M = 0.33,\ SE = 0.08$) than in the *Negotiate* condition ($M = 0.64,\ SE = 0.13$).

Since it is highly likely that a larger number of processed lentils is correlated to a higher probability of making mistakes, a *Penalized Output Score*, computed as the number of correctly processed minus incorrectly processed lentils, was employed in the analysis. In addition, the *Rate of Output* (Penalized Output Score / Total Duration of Trial), the rate at which the participant processed the output correctly, was computed.

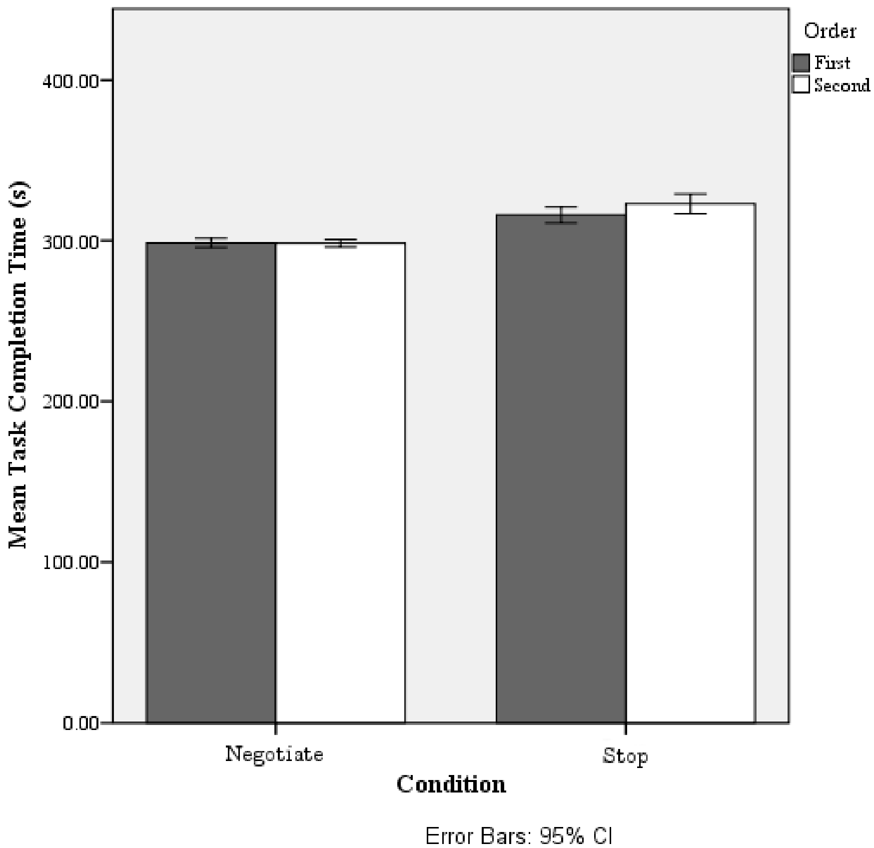

Figure 5.4: Task completion time for the *Stop* vs. *Negotiate* conditions. The HR team completed each trial in the *Negotiate* condition significantly faster than in the *Stop* condition. The HR teams took longer in the second encounter than the first encounter. However, this order effect is much more pronounced in the *Stop* condition in comparison to the *Negotiate* condition.

The repeated-measures ANOVA results on these customized throughput scores are presented in Table 5.4. Analysis of the penalized score indicates that there is a significant difference between *Negotiate* ($M = 76.9,\ SE = 2.58$) and *Stop* ($M = 84.8,\ SE = 3.25$). This indicates that even though the number of incorrectly processed lentils is higher for the *Stop* condition, the number of correctly processed lentils overall remains significantly higher for the *Stop* over the *Negotiate* condition. On the other hand, there is no significant difference in the Rate of Output of *Negotiate* ($M = 0.258,\ SE = 0.009$) versus *Stop* ($M = 0.262,\ SE = 0.009$). The fact that the Rate of Output of *Stop*, with a significantly larger number of mistakes observed, is not noticeably higher than *Negotiate* suggests that the mistakes observed in the *Stop* condition cannot be attributed to the larger number of lentils processed. Rather, it is highly likely that the two conditions contributed to this result.

One should note that there are no significant differences in the Rate of Output, while there is a significant difference in the task completion times between the two conditions. This demonstrates that a significant amount of time is saved in HR resolution of conflicts when the robot exhibits negotiative hesitation behaviours instead of the stopping behaviours used in the *Stop* condition.

Table 5.3: Repeated-measures ANOVA results with Condition and Encounter as fixed, within-subjects factors, and Gender as a between-subjects factor are presented for the correctly and incorrectly (mistake) processed products (lentils). Results from Levene's test suggest that the collected data-set does not violate the homogeneity of variance assumption.

Measure	ANOVA
Correct	
Condition***	$F(1,\ 31) = 21.2,\ p < .001$
Order***	$F(1,\ 31) = 25.2,\ p < .001$
Gender**	$F(1,\ 31) = 8.64,\ p < .01$
Condition*Gender	$F(1,\ 31) = 0.343,\ p = .56$
Order*Gender	$F(1,\ 31) = 0.090,\ p = .77$
Condition*Order	$F(1,\ 31) = 0.408,\ p = .53$
Condition*Order*Gender	$F(1,\ 31) = 0.138,\ p = .71$
Mistakes	
Condition*	$F(1,\ 31) = 6.60,\ p < .05$
Order	$F(1,\ 31) = 1.42,\ p = .24$
Gender	$F(1,\ 31) = 0.010,\ p = .92$
Condition*Gender	$F(1,\ 31) = 0.685,\ p = .41$
Order*Gender	$F(1,\ 31) = 0.129,\ p = .72$
Condition*Order	$F(1,\ 31) = 0.523,\ p = .48$
Condition*Order*Gender	$F(1,\ 31) = 0.055,\ p = .82$

As expected, the ANOVA results show a significant order effect on the counts of lentils correctly processed. Fewer correctly processed lentils were found in the first encounters ($M = 76.9,\ SE = 2.65$) compared to the second encounters ($M = 85.4,\ SE = 3.18$). Fewer incorrectly processed lentils were observed in the second encounters ($M = 0.41,\ SE = 0.08$) than in the first ($M = 0.56,\ SE = 0.13$); however, this is not significant ($p = .24$).

5.4.3 Trajectory and Behaviours

In addition to the subjective experience of the participants and the performance of the HR team considered above, the author also analyzed the participants' trajectory characteristics as indicators of their behaviour patterns. This section outlines trajectory-related findings from segments of motion that triggered conflict responses in the robot. Since resource conflicts were randomly triggered during the HR interaction, the number of conflicts that occurred in a trial for each subject consequently varied. To account for the unequal number of conflicts within and between subjects, the author conducted a linear mixed-model analysis using REML (Multi-level Modelling (MLM), lme4 package, using the statistics program, R (R, The R Foundation) with Condition, Encounter, and Gender as fixed factors, and Participants as a random factor. Possible interaction effects between Condition and Encounter were accounted for in the model, resulting in the following mixed model:

$$measure \sim Condition + Encounter + Gender + Condition * Encounter + (1|Subject). \qquad (5.1)$$

Table 5.4: Repeated-measures ANOVA results with Condition and Encounter as fixed, within-subjects factors, and Gender as a between-subjects factor are presented for the Penalized Output Score and the Rate of Output. Levene's test on these newly created measures suggests that they do not violate the homogeneity of variance assumption.

Measure	ANOVA
Penalised Output Score	
Condition***	$F(1,\ 31) = 18.8,\ p < .001$
Order***	$F(1,\ 31) = 26.2,\ p < .001$
Gender**	$F(1,\ 31) = 8.66,\ p < .01$
Condition*Gender	$F(1,\ 31) = 0.268,\ p = .61$
Order*Gender	$F(1,\ 31) = 0.105,\ p = .75$
Condition*Order	$F(1,\ 31) = 0.470,\ p = .50$
Condition*Order*Gender	$F(1,\ 31) = 0.125,\ p = .73$
Rate of Output	
Condition	$F(1,\ 31) = 0.995,\ p = .33$
Order***	$F(1,\ 31) = 20.4,\ p < .001$
Gender**	$F(1,\ 31) = 9.25,\ p < .01$
Condition*Gender	$F(1,\ 31) = 0.201,\ p = .66$
Order*Gender	$F(1,\ 31) = 0.006,\ p = .94$
Condition*Order	$F(1,\ 31) = 0.001,\ p = .98$
Condition*Order*Gender	$F(1,\ 31) = 0.262,\ p = .61$

Discussed below are the factors that were found to be significant predictors of the measurement of interest.

Does it take longer to resolve conflicts when a robot uses negotiative hesitations?

One of the hypotheses (Hypothesis 5.4) of this study is that the amount of time it takes for a resource conflict to be resolved between an HR pair would be smaller for the *Negotiate* condition in comparison to the *Stop* condition. As stated before, the amount of time it takes for the robot to reach for and return from the dispenser is the same in both conditions. Hence, any residual time it takes for the robot to perform a reach-and-return motion when a conflict is detected is a measure of how long the HR pair took to resolve the conflict in the segment of motion.

The MLM analysis of this duration measure indicates that Condition is indeed a significant factor in determining how long a robot's motion segment in conflict took place ($X^2(2) = 78.2,\ p < .001$).[7] Motion segments in the *Stop* condition are predicted to take an average of 1.38 seconds ($SE = 0.068$) longer than that of an equivalent motion segment in the *Negotiate* condition, and supports Hypothesis 5.4. This finding is despite the fact that the robot re-attempted a greater number of times during a conflict in the *Negotiate* condition than in the *Stop* condition. As shown in Figure 5.5, only in the *Negotiate* condition did the robot trigger its maximum programmed RA value of 15 before yielding to

[7] This means that the regression model, Equation 5.1, provides a significantly better prediction of duration than a regression model without Condition as a factor.

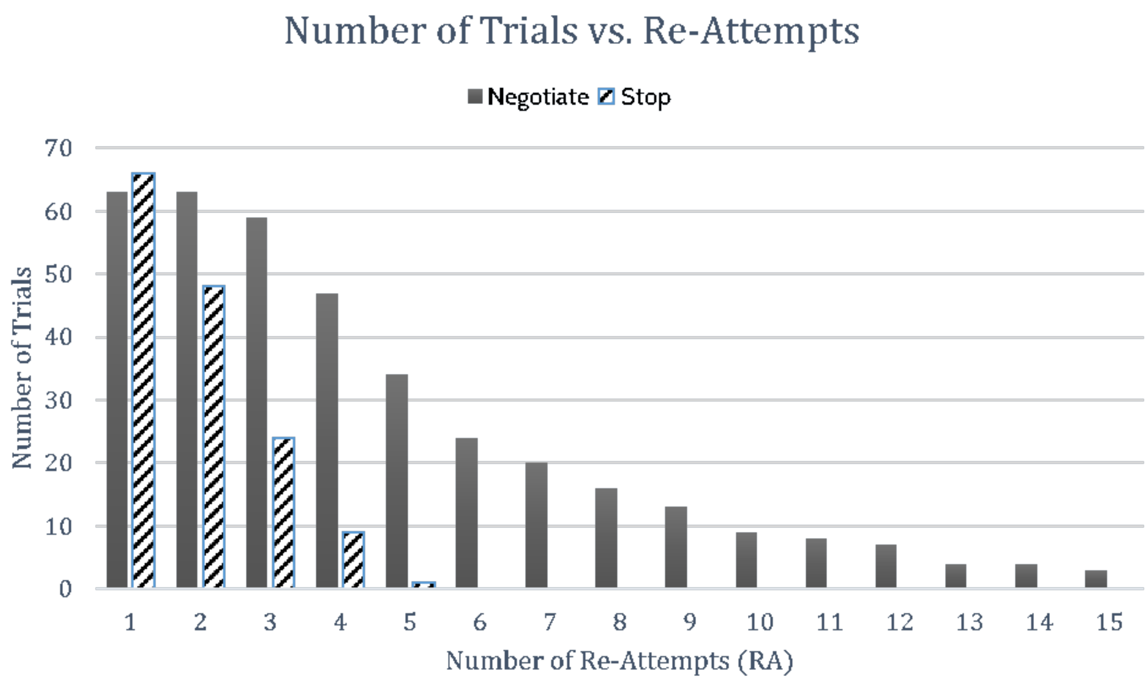

Figure 5.5: The number of trials that triggered RA values greater than zero are presented here in a cumulative manner. For example, Participant 301 in Trial 1 interacted with the robot in the *Negotiate* condition and had a conflict situation that resulted in triggering of all 15 RAs. This incident is reflected as a trial in all fifteen bar charts for the *Negotiate* condition. Since the robot's maximum RA value was set to be 15, the robot yielded to the participant afterwards. A higher number of RAs were triggered under the *Negotiate* condition than the *Stop* condition, although the conflicts were resolved much more quickly in the *Negotiate* condition than the *Stop* condition.

the participant. These extreme instances reflect cases where curious participants intentionally kept their hand out near the dispenser to see what the robot would do next.[8] Other than these extreme cases, all of the HR conflicts were resolved before the robot fully yielded to the participant.

There was no significant interaction effect between Condition and Encounter, and no significant main effects of Gender and Encounter in the MLM analysis.

Do participants stay farther away from the shared resource in negotiative HRI?

One possible indicator of how safe participants perceived the robot to be is the location of their hands with respect to the workspace. If the participants kept their hands farther away from the shared resource, it would provide a behavioural indication of their perceived safety or comfort level with respect to the HRI.

Results from the MLM analysis indicate that Condition is a significant predictor of mean Human to

[8] In fact, in two out of the three trials represented in Figure 5.5, the same subject triggered the maximum RA value in two separate trials.

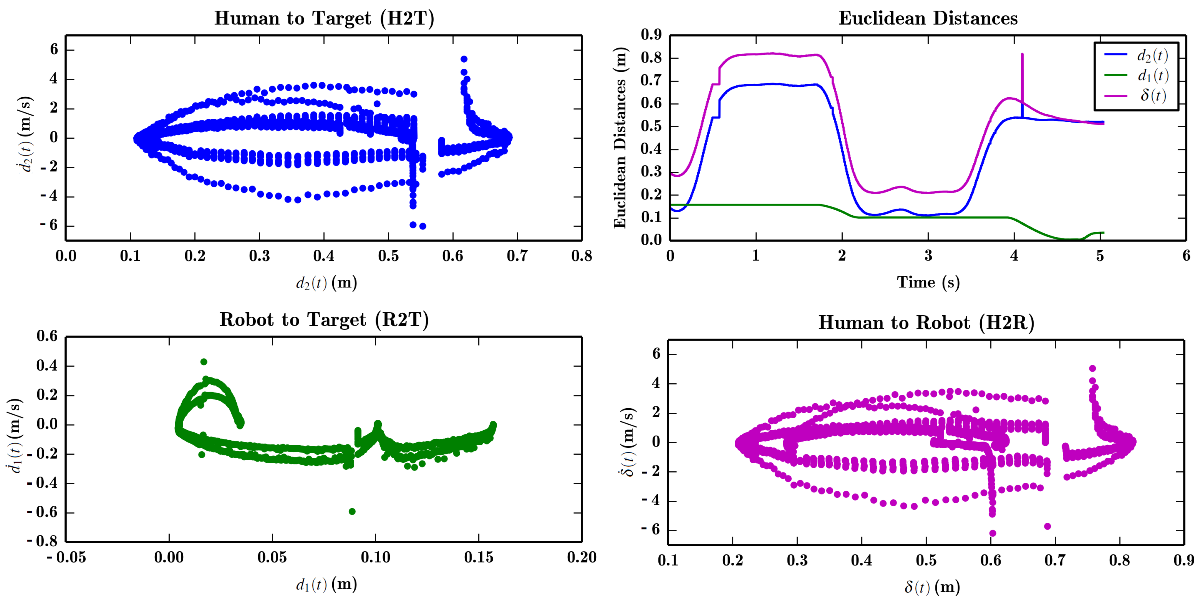

Figure 5.6: One representative participant's behavioural response to a resource conflict during a *Stop* condition trial (Subject: 1, Trial: 2, Motion: 38). As illustrated in the Euclidean Distances plot, the participant ($d_2(t)$) hesitated upon the robot's start of pause and yielded the right of way. The peak shown at (0.10, 0.0) of the R2T plot demonstrates the point at which the robot stopped to pause before proceeding to the target location.

Target (H2T) distance ($X^2(2) = 53.2,\ p < .001$). However, Encounter and Gender are not significant predictors of average H2T distance. The participants' mean Euclidean distance (measured from the wrist of the dominant hand to target:H2T) during conflict segments is smaller in the *Negotiate* condition than in the *Stop* condition by 1.3 cm ($t(1870) = 6.22,\ p < .001$). Participants, on average, had their hands closer to the dispenser in *Negotiate* than in *Stop*.

A more telling indicator of a person's comfort level with a robot is the minimum Euclidean distance to the target (minimum H2T). However, analysis of this measure shows that there is a significant interaction effect between Encounter and Condition ($t(1870) = -2.23,\ p < .05$). The use of both Condition ($X^2(2) = 6.50,\ p < .05$) and Encounter ($X^2(2) = 6.13,\ p < .05$) significantly improves prediction of minimum H2T; however, these factors cannot be interpreted independently from each other. Hence, this analysis provides only a partial support for Hypothesis 5.3, that robot hesitation responses generated using the NHG do not threaten perceived safety of the HRI in comparison to a smooth stopping behaviour.

5.5 Discussion

With the in-person HRI study presented in this chapter, the author attempted to address the second of two research questions stated at the beginning of this thesis: "Can a robot nonverbally negotiate with a person about what should happen in an interaction? Can an HR negotiation contribute to an improved HR collaboration?" The previous chapter described how the NHG was devised and validated as a mechanism

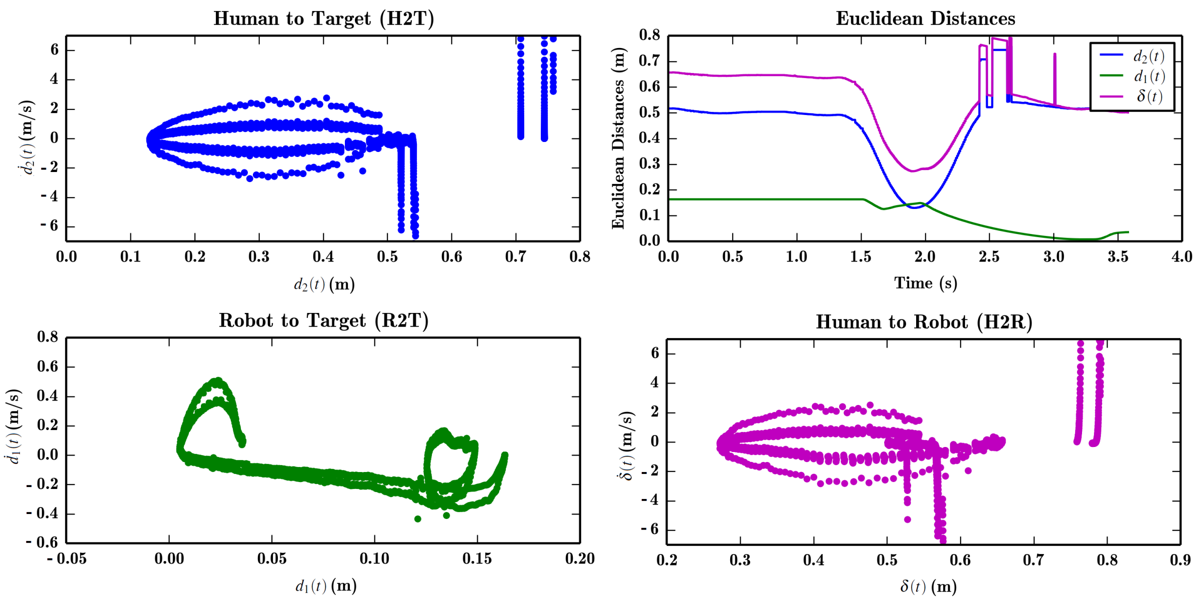

Figure 5.7: One representative participant's behavioural response to a resource conflict during a *Negotiate* condition trial (Subject: 1, Trial: 1, Motion: 21). As illustrated by the Euclidean Distances plot, the participant ($d_2(t)$) reached for the target and retracted right away as the robot triggered its negotiative hesitation behaviour before proceeding to access the target. The loop shown in the R2T plot demonstrates one re-attempt (RA) by the robot in this segment of interaction. This is also shown as the cluster of points encircling $\dot{\delta}(t)$=0.0 in the $\delta(t)$ vs. $\dot{\delta}(t)$ plot, which is absent in the H2T plot.

that enables such negotiation-based conflict resolution to take place in HRI. This chapter presents how an implementation of the NHG to generate artificial, hesitation-based conflict responses from a robot was used in in-person HR collaborative assembly.

HR negotiations lead to faster conflict resolution without jeopardizing the perceived and actual safety of the user

The results of the experiment provide empirical evidence that a robot's resource conflict response generated using the NHG contributes to an efficient resolution of a conflict and supports an affirmative answer to the research question. For instance, the analysis of the collaborative task performance supports Hypothesis 5.2 – that participants completed the task significantly faster in the *Negotiate* than in the *Stop* trials. This is despite the fact that the number of conflicts triggered in the two conditions is not significantly different.

It is also important to note that there is a training effect (in the reverse direction) in the number of conflicts triggered per trial, and in the task completion time. In both conditions, more conflicts were triggered in the second session than in the first. This suggests the possibility that repeated exposure to both types of conflict responses may result in the participants ignoring the robot and dominating the workspace completely, as was the case observed in Moon et al. [105]. This brings us back to the reports of human users abusing self-driving vehicles' systematic, safety-prioritizing responses discussed in the

Introduction [40, 100]. While hesitation-based kinetic dialogue in HRI has been proposed as a solution to a human's potential unwillingness to yield to a robot, thereby undesirably dominating the interaction, the results of this study suggest that repetition of such dialogue may be ineffective. With the increase in the number of conflicts triggered in the second session (40% of robot's reach motions were interrupted, in comparison to 35% in the first session), trials in the second session for both conditions took longer to complete than those in the first.

Upon a closer look at this trend, there is a case to be made for employing a persistent or pushy robot in resolution of HR resource conflicts. As seen in Figure 5.4, there is a much larger increase in the task completion time of the *Stop* condition from the first to the second session in comparison to that of the *Negotiate* condition. This suggests that, even if human users tend to inevitably take advantage of conservative, safety-driven robotic systems after repeated encounters, a negotiative behaviour of the robot can help it to reach resolution of a conflict in a more efficacious manner. In support of this, the results show that the resource conflicts were resolved more quickly in the *Negotiate* condition than the *Stop* condition, which supports Hypothesis 5.4. Hence, while the participants may have started to take advantage of the robot's conflict response in both conditions in the second session, this did not hinder the overall performance of the team in the *Negotiate* condition due to better handling of the resource conflict using the negotiative mode of interaction.

At this point, it is important to note that a robot behaviour that is perceived as more persistent and dominant does not always translate to an HRI that is threatening, harmful, or disliked. From the questionnaire responses reporting on the perceived dominance of the robot, the *Negotiate* condition is considered more dominant than the *Stop* condition (supporting Hypothesis **??**). The mean *Dominance* score for the *Negotiate* condition was 2.52 on a 5-point scale – approximately the middle of the scale. However, behavioural measures on the participants' average distance to the shared resource show that the participants kept their wrist closer to the resource in the *Negotiate* condition than they did in the *Stop* condition. This demonstrates that the participants were, in general, comfortable interacting with the robot in the *Negotiate* condition. If they felt threatened by the robot's motions in the *Negotiate* condition, one would expect the average distance to the shared resource to be larger in comparison to the *Stop* condition.

Qualitative analysis of the interview further supports that while the robot may be perceived as more aggressive in the *Negotiate* condition, it was never seen as too aggressive or threatening. The interview results also suggest that the source of the perceived dominance is not necessarily the robot's conflict response behaviour. For example, the participants referred to the way the robot was pressing the dispenser button while discussing dominance of the robot, although the same pressing behaviour was used in both conditions. This suggests that the perceived dominance of the robot may stem from another source, or that the jerky behaviour of the robot at the onset of reach motions in the *Negotiate* condition caused participants to project a perception of dominance to other parts of the robot's motion, such as pressing of the dispenser by the robot. Moreover, the *Perceived Safety* measure from the questionnaire responses also demonstrates that the robot behaviours in both conditions score highly (averages of 3.96 and 4.21 on a 5-point scale for *Negotiate* and *Stop* conditions, respectively), and that there is no significant difference in this measure across the two conditions ($p = .12$). Hence, the results provide a multitude of evidence for Hypothesis 5.3, namely that hesitation responses generated using the NHG do not threaten the perceived safety nor jeopardize the actual safety of the interaction[9] in comparison to a smooth stopping behaviour.

Perception of the interaction with negotiative robot behaviours is not superior to that of smooth stopping behaviours, but is still viewed in a positive light

Self-reports of participants' experience with the two conditions were measured using the questionnaire response and the interview. Results of these measures provide some mixed findings. First, *Emotional Satisfaction* and *Usefulness* (two out of the three team perception measures collected) for both the *Stop* and the *Negotiate* conditions were above the median (above 3.5 on 5-point scales). Both conditions provided a positive experience in terms of emotional satisfaction and perceived usefulness of the system. However, participants reported a significantly higher *Emotional Satisfaction* for the *Stop* condition than the *Negotiate* condition. In addition, there was no significant difference in the *Usefulness* score of the two conditions. These results fail to support Hypothesis 5.1b that the hesitation responses generated using the NHG are perceived to be more emotionally satisfactory than a smooth stopping behaviour. However, the fact that the *Negotiate* condition scores highly on both measures is a positive indicator that interacting with a robot that uses the NHG to negotiate for a resource conflict with them is not perceived as threatening, but rather as emotionally satisfying and useful.

Second, the author hypothesized (Hypothesis 5.1a) that the robot in the *Negotiate* condition will be perceived in a more favourable light than the *Stop* condition. However, the results indicate that the robot in the *Stop* condition was perceived to be more animate, anthropomorphic, likeable, and intelligent than in the *Negotiate* condition. This fails to support Hypothesis 5.1a, and is in contrast to the results of Study 4, in which the NHG-generated robot responses were perceived to be more anthropomorphic and animate than a non-NHG alternative. This finding suggests that a perception gap exists between third party observation of and in-person interaction with a robot that uses the NHG.

The higher *Animacy* and *Anthropomorphism* scores of the *Stop* condition may be owing to the fact that the reach trajectories of the condition were designed using quintic splines. Quintic splines generate inherently smooth and humanlike trajectories compared to the LDS-generated trajectories of the *Negotiate* condition, which provide a jerkier motion at the onset of a robot's reach. The higher *Likeability* score of the *Stop* condition in comparison to the *Negotiate* condition may be related to the higher *Anthropomorphism* and *Animacy* scores of the *Stop* condition. The collected *Likeability* score is not subject to a significant training effect. This suggests that the participants' preference for the *Stop* condition was retained throughout the experiment. This is also supported by the results of the post-experiment interviews.

It is also interesting to note that the participants were divided about their perception of the task as either collaborative or independent. Results from the interview show that 57% of the participants per-

[9] With the description of the safety features of the system presented in this chapter, no physical collision or safety issue occurred in this experiment for any of the participants.

ceived the task as collaborative, versus 43% perceiving the experimental tasks as independent tasks that involve sharing of a resource. Reflecting on this feedback, it is likely that the team metrics (*Dominance*, *Usefulness*, *Emotional Satisfaction*) are inadequate to capture the full nature of the HR collaboration that was established. In describing the experiment, the participants' perception of the collaborative task did not involve terms such as *Dominance* and *Usefulness*. Rather, the interaction was described in terms that are related more closely to the ideas of flow and fluency.

5.6 Conclusion

In this chapter, an in-person HRI experiment was presented as a means to demonstrate the efficacy of HR nonverbal negotiations using hesitation gestures. While designed hesitation behaviours using the NHG are not preferred over a more traditional conflict response, the results of this study strongly support the effectiveness of this interaction mode on the basis of improved HR collaborative performance. In addition, while the NHG-based negotiative behaviour of the robot is perceived to be more dominant than the smooth stopping alternative, this perception does not equate to a negative experience or perception of the robot. Therefore, the results from this experiment successfully present the NHG as a proof of concept of a new mode of interaction that enables interactive, shared decision-making about an outcome of a conflict in an HR collaboration. This mode of interaction also does not jeopardize the safety of the user nor the performance of the collaboration.

The journey to conducting this experiment also involved a few secondary contributions. The artificial hesitation behaviours generated using the NHG were successfully tested in an in-person interaction with human participants, and validates its efficacy in in-person HRI. The devised system provides a proof-of-concept of real-time, negotiative conflict resolution between a collaborating human and a robot using tracking of the user's wrist motion. It also demonstrates that the hesitation-based conflict responses generated using the NHG deliver an interaction experience distinguished from smooth stopping behaviours used in the *Stop* condition.

In the next chapter, the findings from the studies presented in this thesis – including the supplementary Study 6, presented in Appendix A.1 – are summarized with respect to the two research questions framing this thesis.

Chapter 6

Conclusion

An interacting agent's ability to interweave its plans with another agent is a key element of collaboration [20]. These plans include details of a task such as the *who*, *when*, and *where* that must be communicated with the other agents to successfully collaborate on the task together. Noting that these details are often unspoken at the onset of a collaboration and organically determined as the task unfolds, the author asserted that fluid communication between the collaborating agents is imperative for the agents to interweave these plans. To contribute to the development of robotic systems that collaborate with people in an efficient, safe, and friendly manner, this thesis focused on communication cues a robot can be equipped with to communicate and interweave these details with human users.

In particular, the physical actuation capability of a robot is an essential feature that sets apart the promise of robotics from that of other technologies. This thesis leveraged this capability by exploring the ways in which kino-dynamic behaviours of a robot can be designed to facilitate the interweaving process in an HR collaboration in two exemplar situations, object handover and resource sharing. The often unspoken, yet essential details of the specific tasks explored in this thesis were: *when* and *where* a human recipient should reach out to receive an object from the robot in a robot-to-human handover, and *who* should get access to the shared resource first when both a human and a robot are reaching for it at the same time. A series of six human-subjects studies collectively explored the following two research questions:

1. **Unidirectional Interweaving** Can a robot provide humanlike nonverbal cues to influence people's behavioural responses to an interaction *while* the interaction is taking place?
2. **Bidirectional Interweaving (Negotiation)** Can a robot nonverbally negotiate with a person about what should happen in an interaction? Can an HR negotiation contribute to an improved HR collaboration?

Results from the six studies support that human-inspired nonverbal cues are an effective means to unidirectionally and bidirectionally enable the process of interweaving in an HR team. The following sections present a summary of the findings, research contributions, and future work that pertain to each of the questions. In addition, this chapter concludes with two points of discussion on the implications of the findings from this thesis.

6.1 Unidirectional Interweaving of Subplans

Can a robot's use of nonverbal cues help interweave spatiotemporal details of a task? As discussed in Chapter 2, even simple motions of an automated door or presence of a robot in a room can affect human decisions and behaviours. Nonverbal cues of a robot have been used in various HRI contexts to establish joint attention with and communicate a robot's intent and internal states to a user. Building on the findings from previous studies, this thesis focused on the role nonverbal robot behaviours can play in helping to interweave plans with a person.

The investigations presented in Chapter 3 used the context of robot-to-human handover interaction. In this context, the order of the tasks to be performed and the division of roles was made explicit from the beginning. The desired outcome was also clear from the onset, while the spatiotemporal details of the interaction needed to be determined by the interacting agents while the interaction took place. By design, only the participants were influenced by the interaction. That is, the robot led the interaction with communication cues that unidirectionally influenced the participant, while the robot itself was not influenced by the communication signals from the participants.

Study 1 ($n = 12$) was conducted to investigate the pattern of gaze cues humans use when handing over an object to each other. Qualitative analysis of the collected data resulted in the identification of five types of gaze patterns. This result was used to inform the design of human-inspired gaze patterns a robot can use when handing over an object to a person.

The most frequently observed gaze patterns from Study 1 – the shared attention gaze – was implemented on a humanoid robot for an in-person HRI study, Study 2 (n =102). Study 2 investigated whether the robot's use of human-inspired gaze behaviour can help interweave unspoken spatiotemporal details of a robot-to-human handover interaction. In contrast to previous HRI studies [86, 146] that focused on establishing the joint intent to engage in a handover *before* a handover event takes place, the investigations in Study 2 explicitly addressed the impact a robot's gaze has on human behaviours *during* a handover. The results of Study 2 demonstrate that the robot's use of gaze can impact when and where the participants decide to move their hands to receive the object from the robot. The participants reached out for the object significantly earlier when the robot exhibited the shared attention gaze than when it did not provide a human-inspired gaze pattern. This suggests that the implementation of nonverbal cues on a robot, such as gaze, can influence people's behavioural responses to an interaction *while* the interaction is taking place. In addition, with the shared attention gaze, the participants reached to the projected handover location before the robot had fully reached the location. This provides a glimpse of the power nonverbal cues can have in the interweaving process of an HR collaboration. By cueing the participants to reach out earlier to meet the offered object before the robot has finished moving, the interaction is not only more efficient but also more fluid. Results of a follow-up study, Study 6 (n =30) led by Minhua Zheng and presented in Appendix A.1, further demonstrate that this effect of gaze is sustained even after a period of repeated HR handovers.

These results supplement the previous work in HRI discussed in Section 2.2.1. In previous studies, gaze cues were demonstrated as effective in communicating a robot's target object and internal states to an observing person. Findings from this thesis add to the previous literature and demonstrate that robot gaze cues, even when exhibited during a handover interaction, can be used to subtly communicate spatiotemporal details of a handover to the human receiver and elicit a desired kinodynamic behaviour from the person. This work not only contributes to increasing the fluidity of robot-to-human handovers, but also inspired the investigations on how a robot's use of nonverbal gestures can impact the bidirectional interweaving process of an HR collaboration.

6.1.1 Limitations and Future Work

The role of gaze cues explored in Studies 1, 2, and 6 only focused on the context of robot-to-human handovers. This limits the level at which one can generalize the findings of the studies' results with respect to the impact human-inspired robots cues. Much more extensive studies involving various contexts and nonverbal cues would be necessary before being able to draw general and detailed conclusions about the role a robot's human-inspired nonverbal cues can have on unidirectionally eliciting desirable human behaviours for collaboration.

In addition, although the author and her colleagues proposed and led the qualitative analyses to identify the gaze behaviour of the participants during the handover event, the quality of information that can be gathered on a participants' gaze seen from recordings of a digital camera is limited. In the future, a gaze detector would be necessary to gather an accurate insight about when the participant understands the robot's communication of the desired spatiotemporal details.

While limited in scope, however, the results of this work provided empirical support that even the subtle, and supplementary nonverbal cues such as gaze can play a sigificant role in improving the fluency of an interaction and informing the interweaving process.

6.2 Bidirectional Interweaving (Negotiation) of Subplans

Would people negotiate with a robot? Would an HR negotiation lead to an improved outcome of an interaction? In Chapters 4 and 5, the author focused on negotiative hesitation behaviours as a means to investigate whether a human and a robot can affect each other to negotiate for a desired outcome of an interaction. A series of experiments was conducted to study the dynamics that emerge during in-person HR negotiations of resource conflicts. The experiments considered a collaborative assembly scenario where, while necessarily sharing the same spaces and resources, the collaborating agents reached for the same object at the same time, creating a conflict that they must resolve before they can proceed. In contrast to the unidirectional interweaving discussed in Chapter 3, the outcome of the conflict (i.e., the question of who should access the object first) is not determined at the onset of the interaction.

As presented in Chapter 4, a human-subjects study, Study 3 (n =16) was conducted to collect samples of naturally occurring human hesitation behaviours from an HH collaboration context. Samples of hesitations were observed by online participants (n =300) who rated them according to the expressed *hesitancy* and *persistency* of the motions captured. The results of this study demonstrate that hesitations are indeed used as a mechanism to nonverbally and dynamically resolve the conflict of resources that spontaneously arise in HHI.

Upon analysis of the sample human hesitations and uninterrupted reach gestures collected from

Study 3, the author discovered that, when viewed in state space, a trajectory pattern exists in negotiative hesitations, called *hesitation loops*. These loops represent the changing speeds and distances of the two collaborating agents with respect to the shared target object. While the trajectory pattern can be observed post hoc, generating the same trajectory pattern in HRI is non-trivial. Only one of the two agents (the robot) that shape the hesitation loops is within our control.

Subsequently, the author identified a set of parameters that can be used to convert the hesitation loops into a trajectory generation scheme, called the NHG, which can be used to produce negotiative hesitation trajectories for a robot. The NHG was implemented using the LDS approach to develop a system that is highly responsive to the changes in the workspace. In Study 4 (n =50), trajectories generated by the NHG were used to validate the efficacy of the motions produced by the trajectory generator. The results of this video-based online study demonstrate that the motions produced by the NHG are not only perceived to be more hesitant than an industrial alternative, but also more animate and anthropomorphic. In addition, the feature space used to discover the hesitation loops presented a domain in which the interplay of motions in a nonverbal HR negotiation can be evaluated. Findings from the study also confirmed that the parameter values derived from observed human motions are adequate for generating effective negotiative hesitation gestures for a robot.

The investigations presented in Chapter 4 contribute to a better understanding of hesitation behaviours that are increasingly becoming useful in the field of HRI. The design of the NHG also contributes to the field by providing a mechanism with which one can generate human-inspired negotiative hesitation behaviours. This allows one to study the negotiative dynamics of an HR dyad using reactive, hesitation-based conflict responses with parameter values that are empirically validated by human observers.

In Chapter 5, the author used the NHG as a mechanism to study the negotiation dynamics that emerge in an HR collaboration. A final in-person human-subjects study, Study 5 (n =39), was conducted with a collaborative assembly context in which the HR pair naturally and often reached for the same object at the same time.

Results from this study demonstrate that the NHG can be used to enable nonverbal HR negotiation of resource conflicts to take place. It also illustrates how an HR pair can determine the outcome of a conflict interactively and nonverbally. Not only did the participants yield access to the shared resource to the robot when it exhibited NHG-based hesitation behaviours, the HR team finished the task significantly faster and with fewer mistakes than a non-negotiative alternative, in which the robot paused its motion upon detection of a conflict. Moreover, the real-time negotiations between a human and a robot led to a faster resolution of conflicts without jeopardizing the actual safety of the user or the user's perceived safety of the interaction.

The findings from Study 5 empirically address the second research question. They support that an HR pair can dynamically figure out a desired outcome through nonverbal negotiation. The findings also suggest that, similar to how people typically behave toward another person in such conflicts, people do yield to the robot rather than dominating the shared space. This contributes to the field of HRI by demonstrating the efficacy of nonverbal negotiations as an efficient and fluent mode of interaction.

Furthermore, such negotiative interactions can shift the point of decision-making about an outcome of an interaction from the designers, who are often not directly affected by the outcome, to the users, who are engaged in the interaction in-person.

6.2.1 Limitations and Future Work

Results of Studies 4 and 5 provide a positive outlook on the use of negotiations as a means for users to take part in determining the outcome of an interaction. This shifts the decision making process from the designers and engineers who do not take part in the interaction to the users themselves. However, the investigations presented in this thesis did not evaluate the quality of the negotiated outcome other than the overall task performance. Given this work as a foundation, it would be imperative to develop an experiment where an HR nonverbal negotiation can involve an obvious good or bad outcome. Such experiments would allow us to understand the role robot motions can play in helping people to make the right, split-second decisions, or whether a person's decisions are less likely to be affected by the motions of the robot in contexts where the person is presented with obvious good decisions.

Note that, as further elaborated in the discussions below, the negotiative mode of interaction in Study 5 was not favoured over the alternative. It was also perceived to be less animate and anthropomorphic. This contradicts the findings from Study 4, where the NHG-based motions were perceived more favourably against a non-negotiative alternative. This contradicting finding may be indicative of the fact that an observer's preference or perception of a robot motion can understandably be different if they are third party observers of the interaction versus an in-person participant in the interaction. It may also suggest that the implementation of the NHG using the LDS approach requires further improvements in order to increase its perceived animacy and anthropomorphism.

While the results of Study 5 illustrated that the NHG-based conflict responses can contribute to an improvement of team performance indicators, the results also illustrated a trend that the participants triggered more conflicts as they got used to the experiment trials. If such negotiative behaviours are to be implemented in repetitive collaborative task scenarios, the training effect on the number of conflicts triggered must be further investigated to evaluate if the performance improvements are maintained after a prolonged duration of interaction with the robot.

Going forward, it is important to note that the NHG provides a mechanism with which one can study the dynamics of an HR nonverbal negotiation. For example, one of the next steps that utilize the NHG would be to modify the parameters of the NHG to express higher or lower levels of dominance. The impact of the robot's expressed dominance in a nonverbal negotiation could be examined in such investigations along with an examination of the resulting quality of the HR negotiated outcome. Participants' personality types could also be used as a factor to better understand the type of negotiation dynamics that can be expected in an HR conflict resolution. Such investigations would provide an important insight in how human motions at a subconscious level can be affected by motions of the robot in a nonverbal negotiation setting.

6.3 Efficient HRI vs. Preferred HRI

The results of Studies 2 (Chapter 3) and 5 (Chapter 5) illustrate that the two human-inspired nonverbal cues designed and tested in the studies improve the efficiency of the HR collaboration. However, the author also found that the participants do not prefer these more efficient, human-inspired cues over the less efficient alternatives. In Study 2, the robot's use of the human-inspired, Shared Attention gaze in a robot-to-human handover improved the efficiency of the HRI by encouraging the participants to reach for and receive the object earlier. However, the participants did not show a significant preferrence for this gaze over the Turn-Taking gaze, a less efficient human-inspired alternative, or not having a human-inspired gaze at all – the No Gaze condition where the robot simply looked downwards. Likewise, in Study 5, the Smooth Stop condition, which was not designed based on human behaviours, was preferred over the NHG-generated behaviours. This was despite the fact that the NHG-generated motions were human-inspired and resulted in a superior team task performance.

This follows a puzzling observation within the field of HRI: a robot behaviour that improves the task performance does not necessarily result in a more preferred subjective experience. For example, in a study by Ragni et al. [127], a robot that provides imperfect answers in a memory game with a person was perceived in a more positive light, while it negatively impacted the participant's performance of the task. In addition, Baraglia et al. [11] conducted a study where a robot helped a person on a collaborative manipulation task in one of the following ways: the user controlled when the robot should help; the robot reactively helped the user when the need for help was detected; and the robot proactively provided help whenever it was able to. The researchers found that the HR team worked more fluently when the robot was proactive in helping the user. However, the users preferred it when they had the explicit control of when the robot should help, despite the fact that such interaction resulted in a suboptimal task performance.

In these studies, a user's preference for an HRI does not seem to have a positive correlation with the performance improvements the interaction can provide. Implementation of a robot's behaviour, locus of control, and the nature of the collaborative tasks investigated in the abovementioned studies are all different from one another. They are also factors that are likely to affect user preferences. Therefore, it is hard to make any definitive claims about this intriguing trend observed in these studies.

In order to further investigate this trend, it is necessary to devise a standardized and comprehensive method for measuring the impact an interactive behaviour has on an HR collaboration. The two standardised questionnaires used within the field of HRI – the Godspeed questionnaire [14] and the Negative Attitudes towards Robots Scale [149] – both focus only on self-reported user perception of robots. Utilizing such questionnaires may be sufficient to evaluate the quality of interactive robot behaviours that do not emphasize functional efficiency of the resulting interaction. However, evaluating an interactive behaviour for an HR collaboration cannot rely only on questionnaire results, and must also consider various measures of task performance. For collaborative systems that involve repetitive tasks, studies with much longer duration of interaction may also be necessary before one can identify the eventual preference and efficacy of a collaborative system. In the future, a more comprehensive metric that combines self-reports of user experiences along with task outcomes may be necessary in order to better contrast one collaborative system against another. This will also help demystify the observed tradeoff between efficiency and preference of interactive robot behaviours in HR collaboration.

6.4 Broader Implications and Future Directions: Inclusivity and Bidirectionality

Unlike HHI, HRI is often a process in which one of the agent's precise behaviours and motions are determined by a third party (designers and engineers) who does not partake in the interactivity and are not directly affected by the consequence or the outcome of the interaction. Nonetheless, HRI practitioners are the ones tasked with the design decisions that must be made about future interactions that users will have with the systems. As the author established in Chapter 2, affecting human behaviours and decision-making with a moving object is inevitable. Designing behaviours of an interactive robot is, therefore, an activity that is directly related to determining what type of systematic influence the users will encounter. With the rise of interactive robotic systems that can manipulate our physical environment, it is imperative that HRI practitioners closely examine this relationship we have with the system's users. It may be harmless to implement gaze cues on a robot that can consistently elicit a desired kinodynamic response from a person in a handover interaction. A system that unidirectionally influences users to make certain purchasing decisions, on the other hand, may be considered unethical [138]. This relates to the uneasy feelings depicted in many works of science fiction where a dystopian future with robots involves humans who are helplessly subjected to unidirectional influence of the so-called robot overlords. In reality, the robotic systems that are being developed today lie on a spectrum-of-influence model with unidirectional influence on one extreme and complete bidirectional influence on the other. Where a system lies on this spectrum is determined by design decisions that map the system's ability to sense its environment, including users around it, and its ability to respond to the sensed signals. As robots can be viewed as a new 'designed' species deployed into the world, it is up to the HRI practitioners to be mindful of this influence directionality spectrum and its implications to the outcome of an interaction [73].

A number of researchers are concerned about this social distance between designers and users, and advocate for the principle of transparency as a means to instill the right level of trust in a user of a robotic system [153]. One way to increase such transparency between the user and the system is to enable more fluent and intuitive modes of communication for the HRI. In fact, systems on any point on the influence directionality spectrum can be designed to be transparent to the user as long as they are equipped with appropriate abilities to convey their internal states to the users. However, one-sided communication does not allow room for continuous changes that may be desired in HRI. The fact that communication is imperative to every element of collaboration is founded upon the notion that collaborative processes are typically dynamic in nature, allowing a team of agents to respond to changes in the situation, environment, and even internal states or desires. To realize the type of human-machine collaboration Barbara Grosz envisioned, the dynamic nature of our everyday tasks must be something that robots can accommodate.

The value further supported by designing systems to be more bidirectional in its interaction with users is that it allows room for the user to have a say in what should happen in the interaction, which can be different from one context and person to another. When presented with an uncertain or a contentious situation that involves a person and a robot, the author's previous work suggests that people prefer the robot to engage in dialogues with the stakeholder than determine the interaction outcome by itself (see Section 2.2.2). The results of Study 5 (Chapter 5) illustrate that, at the most basic level of interaction that involves subtle movements of the hands, a robot is able to engage in a dialogue with a person toward a resolution of a conflict, even if it is not favoured over a robotic alternative. Regardless of the rapid advancements in artificial intelligence and machine learning technologies, which could help determine an optimal solution to a problem, a system that is opaque in its decision-making may be seen as a hindrance to the user who may feel the need to exercise control and have a sense of autonomy about a situation at hand. The negotiative behaviours explored in this thesis attempt to put forth inclusivity – which the author refers to as a system's ability to include the stakeholders (including users) into the process of making decisions that affect the collective – as a principle that can be fostered in interaction designs as we deal with more systems that will come across unintended conflicts and unpredicted situations with people.

Bibliography

[1] ROS.org. URL http://wiki.ros.org. Accessed on 2013-09-08. → pages 25

[2] E. Ackerman. Panasonic revives hospital delivery robot, May 2014. URL http://spectrum.ieee.org/automaton/robotics/medical-robots/panasonic-hospital-delivery-robot. Accessed on 2015-01-03. → pages 8

[3] H. Admoni, T. Weng, B. Hayes, and B. Scassellati. Robot nonverbal behavior improves task performance in difficult collaborations. In *ACM/IEEE International Conference on Human-Robot Interaction*, volume 2016-April, pages 51–58. ACM/IEEE, 2016. ISBN 9781467383707. doi:10.1109/HRI.2016.7451733. → pages 4, 9

[4] A. Albu-Schäffer, S. Haddadin, C. Ott, A. Stemmer, T. Wimböck, and G. Hirzinger. Robovie: an interactive humanoid robot. *Industrial Robot: An International Journal*, 28(6):498 – 504, Aug 2001. ISSN 0143-991X. doi:10.1108/01439910710774386. → pages 11

[5] Amazon.com. Amazon Mechanical Turk: artificial artificial intelligence, 2016. URL https://www.mturk.com/mturk/welcome. Accessed on 2016-12-10. → pages 42, 43

[6] ANSI/RIA. *ANSI/RIA R15.06-2012 American National Standard for Industrial Robots and Robot Systems - Safety Requirements (revision of ANSI/RIA R15.06-1999)*. American National Standards Institute, 2012. → pages 85

[7] Apple Inc. iOS - Siri - Apple (CA), 2011. URL http://www.apple.com/ca/ios/siri/. Accessed on 2016-12-12. → pages 2

[8] M. Argyle and M. Cook. *Gaze and Mutual Gaze*. Cambridge University Press, Cambridge, UK, 1976. → pages 18

[9] R. Atienza and A. Zelinsky. *Intuitive Human-Robot Interaction Through Active 3D Gaze Tracking*, pages 172–181. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005. ISBN 978-3-540-31508-7. doi:10.1007/11008941_19. URL http://dx.doi.org/10.1007/11008941_19. → pages 18

[10] W. a. Bainbridge, J. W. Hart, E. S. Kim, and B. Scassellati. The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics*, 3:41–52, 2011. ISSN 18754791. doi:10.1007/s12369-010-0082-7. → pages 37

[11] J. Baraglia, M. Cakmak, Y. Nagai, R. Rao, and M. Asada. Initiative in robot assistance during collaborative task execution. *ACM/IEEE International Conference on Human-Robot Interaction*, 2016-April:67–74, 2016. ISSN 21672148. doi:10.1109/HRI.2016.7451735. → pages 104

[12] C. Bartneck, M. van der Hoek, O. Mubin, and A. Al Mahmud. “Daisy, Daisy, give me your answer do!”. *Proceedings of the ACM/IEEE international conference on Human-robot interaction - HRI ’07*, (2007):217, 2007. doi:10.1145/1228716.1228746. → pages 37, 38

[13] C. Bartneck, T. Kanda, O. Mubin, and A. Al Mahmud. Does the design of a robot influence its animacy and perceived intelligence? *International Journal of Social Robotics*, 1(2):195–204, Feb 2009. ISSN 1875-4791. doi:10.1007/s12369-009-0013-7. → pages 37

[14] C. Bartneck, D. Kulić, E. Croft, and S. Zoghbi. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, 1(1):71–81, Nov 2009. ISSN 1875-4791. doi:10.1007/s12369-008-0001-3. → pages 62, 80, 86, 104, 124

[15] P. Basili, M. Huber, and T. Brandt. Investigating human-human approach and hand-over. *Human Centered Robot Systems*, pages 151–160, 2009. doi:DOI=10.1007/978-3-642-10403-9_16. → pages 15, 17, 22

[16] M. Bennewitz, F. Faber, D. Joho, M. Schreiber, and S. Behnke. Towards a humanoid museum guide robot that interacts with multiple persons. In *5th IEEE-RAS International Conference on Humanoid Robots, 2005.*, pages 418–423, Dec 2005. doi:10.1109/ICHR.2005.1573603. → pages 18

[17] A. Bisio, A. Sciutti, F. Nori, G. Metta, L. Fadiga, G. Sandini, and T. Pozzo. Motor contagion during human-human and human-robot interaction. *PLoS ONE*, 9(8), 2014. ISSN 19326203. doi:10.1371/journal.pone.0106172. → pages 12

[18] S. Bouchard. Robot manipulation challenge: Clean up dining table, load dishwasher, Nov 2009. URL http://spectrum.ieee.org/automaton/robotics/robotics-software/mobile-manipulation-challenge. Accessed on 2016-12-12. → pages 8

[19] J. Boucher, U. Pattacini, A. Lelong, G. Bailly, F. Elisei, S. Fagel, P. F. Dominey, and J. Ventre-Dominey. I reach faster when I see you look: gaze effects in humanhuman and humanrobot face-to-face cooperation. *Frontiers in Neurorobotics*, 6(3), 2012. ISSN 1662-5218. doi:10.3389/fnbot.2012.00003. → pages 4, 8, 11

[20] M. E. Bratman. Shared cooperative activity. *The Philosophical Review*, 101(2):327–341, 1992. → pages 2, 8, 13, 99

[21] C. Breazeal. Jibo, the world’s first social robot for the home, 2014. URL https://www.indiegogo.com/projects/jibo-the-world-s-first-social-robot-for-the-home. Accessed on 2016-11-14. → pages 2

[22] C. Breazeal, C. Kidd, A. Thomaz, G. Hoffman, and M. Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 383–388, Edmonton, Canada, 2005. IEEE. ISBN 0-7803-8912-3. doi:10.1109/IROS.2005.1545011. → pages 4, 10

[23] C. Breazeal, A. Takanishi, and T. Kobayashi. Social robots that interact with people. *Springer Handbook of Robotics*, pages 1349–1369, 2008. doi:10.1007/978-3-540-30301-5_59. → pages 11

[24] K. O. Brien, J. Sutherland, C. Rich, and C. L. Sidner. Collaboration with an autonomous humanoid robot: a little gesture goes a long way. In *Proceedings of the 6th international conference on Human-robot interaction - HRI '11*, pages 215–216, Lausanne, Switzerland, 2011. ACM. → pages 10

[25] D. Brscić, H. Kidokoro, Y. Suehiro, and T. Kanda. Escaping from children's abuse of social robots. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '15, pages 59–66, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-2883-8. doi:10.1145/2696454.2696468. → pages 74

[26] G. Butterworth. The ontogeny and phylogeny of joint visual attention. In A. Whiten, editor, *Natural theories of mind: Evolution, development, and simulation of everyday mindreading*, pages 223–232. Basil Blackwell, 1991. → pages 18

[27] M. Cakmak, S. S. Srinivasa, M. K. Lee, S. Kiesler, and J. Forlizzi. Using spatial and temporal contrast for fluent robot-human hand-overs. In *Proceedings of the 6th International Conference on Human-robot Interaction*, HRI '11, pages 489–496, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0561-7. doi:10.1145/1957656.1957823. → pages 15, 17, 25, 31

[28] M. Cakmak, S. S. Srinivasa, Min Kyung Lee, J. Forlizzi, and S. Kiesler. Human preferences for robot-human hand-over configurations. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1986–1993. IEEE, Sep 2011. ISBN 978-1-61284-456-5. doi:10.1109/IROS.2011.6094735. → pages 17

[29] E. Calisgan, A. Haddadi, H. F. M. Van der Loos, J. A. Alcazar, and E. A. Croft. Identifying nonverbal cues for automated human-robot turn-taking. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, pages 418–423, Sept 2012. doi:10.1109/ROMAN.2012.6343788. → pages 23

[30] E. Calisgan, A. Moon, C. Bassani, F. Ferreira, F. Operto, G. Veruggio, E. A. Croft, and H. F. M. Van der Loos. Open Roboethics pilot: accelerated policy design, implementation and demonstration of socially acceptable behaviors. In *We Robot2*, Stanford, CA, 2013. → pages 12

[31] M. Carpenter, K. Nagell, and M. Tomasello. Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development*, 63(4):i–vi, 1–143, 1998. ISSN 0037-976X. → pages 11

[32] J. Cassell and K. R. Thorisson. The power of a nod and a glance: envelope vs. emotional feedback in animated conversational agents. *Applied Artificial Intelligence*, 13(4-5):519–538, May 1999. ISSN 0883-9514. doi:10.1080/088395199117360. → pages 37

[33] A. Chalimourda, B. Schölkopf, and A. J. Smola. Experimentally optimal ν in support vector regression for different noise models and parameter settings. *Neural Networks*, 17(1):127–141, 2004. ISSN 08936080. doi:10.1016/S0893-6080(03)00209-0. → pages 50

[34] W. P. Chan, C. A. Parker, H. F. M. Van der Loos, and E. A. Croft. A human-inspired object handover controller. *The International Journal of Robotics Research*, 32(8):971–983, Jul 2013. ISSN 0278-3649. doi:10.1177/0278364913488806. → pages 15, 17, 22

[35] W. P. Chan, K. Nagahama, H. Yaguchi, Y. Kakiuchi, K. Okada, and M. Inaba. Implementation of a framework for learning handover grasp configurations through observation during human-robot object handovers. In *15th IEEE-RAS International Conference on Humanoid*

Robots, Humanoids 2015, Seoul, South Korea, November 3-5, 2015, pages 1115–1120, 2015. doi:10.1109/HUMANOIDS.2015.7363492. → pages 5

[36] A. Chang. NY midtown robots allow for conversation-free hotel service, May 2015. URL http://cnb.cx/2mLSg5b. Accessed on 2016-12-12. → pages 7

[37] P. R. Cohen and H. J. Levesque. Teamwork. *Noûs*, 25(4):487–512, 1991. → pages 8

[38] S. Connor. First self-driving cars will be unmarked so that other drivers don't try to bully them: Volvo fears that other road users could drive erratically in order to take advantage, 2016. URL https://www.theguardian.com/technology/2016/oct/30/volvo-self-driving-car-autonomous. Accessed on 2016-12-12. → pages 75

[39] M. Corley and O. W. Stewart. Hesitation disfluencies in spontaneous speech: the meaning of um. *Linguistics and Language Compass*, 2(4):589–602, 2008. ISSN 1749818X. doi:10.1111/j.1749-818X.2008.00068.x. → pages 37

[40] S. Crowe. A cyclist's encounter with an indecisive google self-driving car, Aug 2015. URL http://www.roboticstrends.com/article/a_cyclists_encounter_with_an_indecisive_google_self_driving_car/P2. Accessed on 2016-10-03. → pages 77, 96

[41] D. C. L. David J. Knight, Daniel Langmeyer. Eye-contact, distance, and affiliation: the role of observer bias. *Sociometry*, 36(3):390–401, 1973. ISSN 00380431. URL http://www.jstor.org/stable/2786340. → pages 18

[42] F. Dehais, E. A. Sisbot, R. Alami, and M. Causse. Physiological and subjective evaluation of a humanrobot object hand-over task. *Applied Ergonomics*, 42(6):785–791, Nov 2011. ISSN 00036870. doi:10.1016/j.apergo.2010.12.005. → pages 16, 17

[43] C. Dondrup, N. Bellotto, and M. Hanheide. Social distance augmented qualitative trajectory calculus for human-robot spatial interaction. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pages 519–524, 2014. ISBN 978-1-4799-6765-0. doi:10.1109/ROMAN.2014.6926305. → pages 38

[44] L. W. Doob. *Hesitation: Impulsivity and Reflection*. Greenwood Press, Westport, CT, 1990. ISBN 0313274460. → pages 37

[45] A. D. Dragan, K. C. Lee, and S. S. Srinivasa. Legibility and predictability of robot motion. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 301–308. IEEE, Mar 2013. ISBN 978-1-4673-3101-2. doi:10.1109/HRI.2013.6483603. → pages 10

[46] J. DURBIN. Incomplete blocks in ranking experiments. *British Journal of Statistical Psychology*, 4(2):85–90, 1951. ISSN 2044-8317. doi:10.1111/j.2044-8317.1951.tb00310.x. → pages 27

[47] M. S. Erden, K. Leblebicioğlu, and U. Halici. Multi-agent system-based fuzzy controller design with genetic tuning for a mobile manipulator robot in the hand over task. *J. Intell. Robotics Syst.*, 39(3):287–306, Mar 2004. ISSN 0921-0296. doi:10.1023/B:JINT.0000021039.56110.c8. → pages 22

[48] N. Figueroa, A. L. P. Ureche, and A. Billard. Learning complex sequential tasks from demonstration: A pizza dough rolling case study. In *ACM/IEEE International Conference on Human-Robot Interaction*, volume 2016-April, pages 611–612, 2016. ISBN 9781467383707. doi:10.1109/HRI.2016.7451881. → pages 57

[49] T. Flash and N. Hogan. The coordination of arm movements: Mathematical model. *Journal of Neuroscience*, 5(7):1688–1703, 1985. → pages 63

[50] W. J. Fu. Penalized regressions: the bridge versus the lasso. *Journal of Computational and Graphical Statistics*, 7(3):397–416, 1998. ISSN 10618600. doi:10.2307/1390712. → pages 48

[51] K. Gillespie-Lynch, P. M. Greenfield, Y. Feng, S. Savage-Rumbaugh, and H. Lyn. A cross-species study of gesture and its role in symbolic development: implications for the gestural theory of language evolution. *Frontiers in psychology*, 4(June):160, Jan. 2013. ISSN 1664-1078. doi:10.3389/fpsyg.2013.00160. → pages 40

[52] B. Gleeson, K. Currie, K. MacLean, and E. Croft. Tap and push: Assessing the value of direct physical control in human-robot collaborative tasks. *Journal of Human-Robot Interaction*, 4(1): 95, 2015. ISSN 2163-0364. doi:10.5898/JHRI.4.1.Gleeson. → pages 9, 10

[53] R. M. Golinkoff. i beg your pardon?: the preverbal negotiation of failed messages. *Journal of Child Language*, 13(03):455–476, Sept. 1986. ISSN 0305-0009. doi:10.1017/S0305000900006826. → pages 40

[54] R. M. Golinkoff. When is communication a meeting of minds? *Journal of child language*, 20 (1):199–207, Feb. 1993. ISSN 0305-0009. → pages 40

[55] Google. The Google app. URL https://www.google.com/search/about/. Accessed on 2016-12-12. → pages 2

[56] B. J. Grosz. Collaborative systems. *AI Magazine*, 17(2):67–85, 1996. → pages 1, 2, 8

[57] B. J. Grosz and L. Hunsberger. The dynamics of intention in collaborative activity. *Cognitive Systems Research*, 7(2-3):259–272, 2006. ISSN 13890417. doi:10.1016/j.cogsys.2005.11.006. → pages

[58] B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996. ISSN 00043702. doi:10.1016/0004-3702(95)00103-4. → pages 2, 8

[59] M. Häring, J. Eichberg, and E. André. Studies on grounding with gaze and pointing gestures in human-robot-interaction. In *Proceedings of the 4th International Conference on Social Robotics*, ICSR'12, pages 378–387, Berlin, Heidelberg, 2012. Springer-Verlag. ISBN 978-3-642-34102-1. doi:10.1007/978-3-642-34103-8_38. → pages 18

[60] B. Hayes-Roth. A blackboard architecture for control. *Artif. Intell.*, 26(3):251–321, Aug 1985. ISSN 0004-3702. doi:10.1016/0004-3702(85)90063-3. → pages 25

[61] G. Hoffman and W. Ju. Designing robots with movement in mind. *Journal of Human-Robot Interaction*, 3(1):89, 2014. ISSN 2163-0364. doi:10.5898/JHRI.3.1.Hoffman. → pages 11

[62] G. Hoffman, J. Forlizzi, S. Ayal, A. Steinfeld, J. Antanitis, G. Hochman, E. Hochendoner, and J. Finkenaur. Robot presence and human honesty. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction - HRI '15*, pages 181–188,

New York, New York, USA, 2015. ACM Press. ISBN 9781450328838. doi:10.1145/2696454.2696487. → pages 10

[63] B. Huang, M. Li, R. L. De Souza, J. J. Bryson, and A. Billard. A modular approach to learning manipulation strategies from human demonstration. *Autonomous Robots*, 40(5):903–927, 2016. ISSN 15737527. doi:10.1007/s10514-015-9501-9. → pages 57

[64] C.-M. Huang and B. Mutlu. The repertoire of robot behavior: Designing social behaviors to support human-robot joint activity. *Journal of Human-Robot Interaction*, 2(2):80–102, 2013. ISSN 2163-0364. doi:10.5898/JHRI.2.2.Huang. → pages 11

[65] C.-M. Huang, S. Andrist, A. Sauppé, and B. Mutlu. Using gaze patterns to predict task intent in collaboration. *Frontiers in psychology*, 6(July):1049, 2015. ISSN 1664-1078. doi:10.3389/fpsyg.2015.01049. → pages 10

[66] M. Huber, M. Rickert, A. Knoll, T. Brandt, and S. Glasauer. Human-robot interaction in handing-over tasks. In *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, pages 107–112, Aug 2008. doi:10.1109/ROMAN.2008.4600651. → pages 17

[67] M. Imai, T. Ono, and H. Ishiguro. Physical relation and expression: joint attention for human-robot interaction. In *Robot and Human Interactive Communication, 2001. Proceedings. 10th IEEE International Workshop on*, pages 512–517, 2001. doi:10.1109/ROMAN.2001.981955. → pages 18, 23

[68] M. Imai, T. Kanda, T. Ono, H. Ishiguro, and K. Mase. Robot mediated round table: Analysis of the effect of robot's gaze. In *Robot and Human Interactive Communication, 2002. Proceedings. 11th IEEE International Workshop on*, pages 411–416, 2002. doi:10.1109/ROMAN.2002.1045657. → pages 22

[69] International Standards Organization. ISO 10218-1:2011 – robots and robotic devices – safety requirements for industrial robots. Technical report, ISO, 2011. URL http://www.iso.org/iso/catalogue{_}detail?csnumber=51330. → pages 85

[70] W. Ju and D. Sirkin. Animate objects : How physical motion encourages public interaction. In P. Thomas, H. Per, and O.-K. Harri, editors, *Persuasive Technology: Lecture Notes in Computer Science*, volume 6137, pages 40–51. Springer, Berlin, Heidelberg, 2010. doi:10.1007/978-3-642-13226-1_6. → pages 10

[71] W. Ju and L. Takayama. Approachability: How people interpret automatic door movement as gesture. *International Journal*, 3(2):1–10, Aug 2009. URL http://www.ijdesign.org/ojs/index.php/IJDesign/article/view/574/244citeulike-article-id:6390837. → pages 10, 11

[72] W. Ju and L. Takayama. Should robots or people do these jobs ? a survey of robotics experts and non-experts about which jobs robots should do. In *International Conference on Intelligent Robots and Systems*, pages 2452–2459, San Francisco, CA, 2011. IEEE/RSJ. ISBN 9781612844558. → pages 8

[73] P. H. Kahn, A. L. Reichert, H. E. Gary, T. Kanda, H. Ishiguro, S. Shen, J. H. Ruckert, and B. Gill. The new ontological category hypothesis in human- robot interaction. In *Proceedings of the 6th international conference on Human-robot interaction - HRI '11*, pages 159–160, Lausanne, Switzerland, 2011. ACM. → pages 105

[74] S. Kajikawa, T. Okino, K. Ohba, and H. Inooka. Motion planning for hand-over between human and robot. In *Proceedings of the International Conference on Intelligent Robots and Systems-Volume 1 - Volume 1*, IROS '95, pages 193–, Washington, DC, USA, 1995. IEEE Computer Society. ISBN 0-8186-7108-4. → pages 17

[75] W. Karwowski, T. Plank, M. Parsaei, and M. Rahimi. Human perception of the maximum safe speed of robot motions. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 31(2):186–190, Sep 1987. ISSN 1071-1813. doi:10.1177/154193128703100211. → pages 85

[76] T. Kazuaki, O. Motoyuki, and O. Natsuki. The hesitation of a robot: A delay in its motion increases learning efficiency and impresses humans as teachable. *Human-Robot Interaction (HRI), 2010 5th ACM/IEEE International Conference on*, 8821007:189–190, 2010. doi:10.1109/HRI.2010.5453200. → pages 37

[77] T. Kellner. Rethink Robotics Is Freeing Robots From Their Cages, 2016. URL http://www.gereports.com/post/107898198705/rethink-robotics-is-freeing-next-gen-robots-from/. Accessed on 2016-11-14. → pages 1

[78] S. M. Khansari-Zadeh and A. Billard. Learning stable non-linear dynamical systems with gaussian mixture models. *Transactions on Robotics*, 27(5):943—-957, 2011. ISSN 1552-3098. doi:10.1109/TRO.2011.2159412. → pages 57

[79] S. Kim, A. Shukla, and A. Billard. Catching objects in flight. *IEEE Transactions on Robotics*, PP:1–17, 2014. ISSN 1552-3098. doi:10.1109/TRO.2014.2316022. → pages 57

[80] N. Kirchner, A. Alempijevic, and G. Dissanayake. Nonverbal robot-group interaction using an imitated gaze cue. In *Proceedings of the 6th International Conference on Human-robot Interaction*, HRI '11, pages 497–504, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0561-7. doi:10.1145/1957656.1957824. → pages 18, 21, 22

[81] C. L. Kleinke. Gaze and eye contact: a research review. *Psychological bulletin*, 100(1):78–100, Jul 1986. ISSN 0033-2909. → pages 16, 18

[82] K. L. Koay and E. A. Sisbot. Exploratory study of a robot approaching a person in the context of handing over an object. In *Symposium on Multidisciplinary Collaboration for Socially Assistive Robotics*, pages 18–24, 2007. → pages 16, 17, 22

[83] D. Kulic and E. Croft. Physiological and subjective responses to articulated robot motion. *Robotica*, 25(01):13, Aug 2006. ISSN 0263-5747. doi:10.1017/S0263574706002955. → pages 10

[84] Y. Kuno, K. Sadazuka, M. Kawashima, K. Yamazaki, A. Yamazaki, and H. Kuzuoka. Museum guide robot based on sociological interaction analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 1191–1194, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-593-9. doi:10.1145/1240624.1240804. → pages 18

[85] D. Leavens, J. Russell, and W. Hopkins. Intentionality as measured in the persistence and elaboration of communication by chimpanzees (pan troglodytes). *Child development*, 76(1): 291–306, 2005. → pages 40

[86] M. K. Lee, J. Forlizzi, S. Kiesler, M. Cakmak, and S. Srinivasa. Predictability or adaptivity?: Designing robot handoffs modeled from trained dogs and people. In *Proceedings of the 6th International Conference on Human-robot Interaction*, HRI '11, pages 179–180, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0561-7. doi:10.1145/1957656.1957720. → pages 15, 16, 18, 22, 100

[87] L. E. Levin. Kinetic dialogs in predator-prey recognition. *Behavioural processes*, 40(2):113–20, Jul 1997. ISSN 0376-6357. → pages 76, 77

[88] A. Lim, T. Ogata, and H. G. Okuno. Towards expressive musical robots: a cross-modal framework for emotional gesture, voice and music. *EURASIP Journal on Audio, Speech, and Music Processing*, 3, 2012. ISSN 1687-4722. doi:10.1186/1687-4722-2012-3. → pages 41

[89] C. Liu, C. T. Ishi, H. Ishiguro, and N. Hagita. Generation of nodding, head tilting and eye gazing for human-robot dialogue interaction. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '12, pages 285–292, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1063-5. doi:10.1145/2157689.2157797. → pages 18, 23

[90] K. E. Lochbaum. *Using Collaborative Plans to Model the Intentional Structure of Discourse*. PhD thesis, Harvard University, Cambridge, MA, USA, 1995. UMI Order No. GAX95-14804. → pages 1

[91] J. Mainprice, M. Gharbi, T. Simon, and R. Alami. Sharing effort in planning human-robot handover tasks. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, pages 764–770, Sept 2012. doi:10.1109/ROMAN.2012.6343844. → pages 17

[92] I. Marlow. The boomer shift: Japan's bold steps, Nov 2015. URL http://www.theglobeandmail.com/globe-investor/retirement/retire-planning/how-japan-is-coping-with-a-rapidly-aging-population/article27259703/. Accessed on 2016-11-03. → pages 1

[93] J. G. Martin and W. Strange. The perception of hesitation in spontaneous speech. *Perception & Psychophysics*, 3(6):427–438, 1968. ISSN 0031-5117. doi:10.3758/BF03205750. → pages 37

[94] Y. Matsusaka, S. Fujie, and T. Kobayashi. Modeling of conversational strategy for the robot participating in the group conversation. In *INTERSPEECH*, 2001. → pages 18

[95] Max Planck Institute for Psycholinguistics. ELAN — The Language Archive. URL http://tla.mpi.nl/tools/tla-tools/elan/. Accessed on 2016-11-03. → pages 29

[96] R. Maxion and A. DeChambeau. Dependability at the user interface. In *Twenty-Fifth International Symposium on Fault-Tolerant Computing. Digest of Papers*, pages 528–535. IEEE Comput. Soc. Press, 1995. ISBN 0-8186-7079-7. doi:10.1109/FTCS.1995.466944. → pages 37

[97] S. Merlo and P. A. Barbosa. Hesitation phenomena: a dynamical perspective. *Cognitive processing*, 11(3):251–61, Aug 2010. ISSN 1612-4790. doi:10.1007/s10339-009-0348-x. → pages 37

[98] J. Millar and A. Moon. How to engage the public on the ethics and governance of lethal autonomous weapons. In *We Robot*, Miami, FL, 2016. → pages 8

[99] G. Milliez, R. Lallement, M. Fiore, and R. Alami. Using human knowledge awareness to adapt collaborative plan generation, explanation and monitoring. *ACM/IEEE International Conference on Human-Robot Interaction*, 2016-April:43–50, 2016. ISSN 21672148. doi:10.1109/HRI.2016.7451732. → pages 10

[100] R. Mitchell. Human drivers will bully robot cars, says CEO of Mercedes-Benz USA, 2016. URL http://www.latimes.com/business/la-fi-hy-live-updates-2016-la-auto-show-human-drivers-will-bully-robot-cars-1479247249-htmlstory.html. Accessed on 2012-12-10. → pages 75, 96

[101] S. Mohammad Khansari-Zadeh and A. Billard. A dynamical system approach to realtime obstacle avoidance. *Autonomous Robots*, 32(4):433–454, 2012. ISSN 09295593. doi:10.1007/s10514-012-9287-y. → pages 57

[102] A. Moon, P. Danielson, and H. F. M. Van der Loos. Survey-based discussions on morally contentious applications of interactive robotics. *International Journal of Social Robotics*, pages 1–20, Nov 2011. ISSN 1875-4791. doi:10.1007/s12369-011-0120-0. → pages 8

[103] A. Moon, C. A. C. Parker, E. A. Croft, and H. F. M. Van der Loos. Did you see it hesitate? - Empirically grounded design of hesitation trajectories for collaborative robots. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1994–1999, San Francisco, CA, Sep 2011. IEEE/RSJ. ISBN 978-1-61284-454-1. doi:10.1109/IROS.2011.6048228. → pages 38, 44, 45

[104] A. Moon, C. A. C. Parker, E. A. Croft, and H. F. M. Van der Loos. Design and impact of hesitation gestures during human-robot resource conflicts. *Journal of Human-Robot Interaction*, 2(3):18–40, Sept. 2013. ISSN 2163-0364. doi:10.5898/JHRI.2.3.Moon. → pages xii, 38, 42, 44, 45, 57, 80, 82, 85

[105] A. Moon, D. M. Troniak, B. Gleeson, M. K. Pan, M. Zheng, B. A. Blumer, K. MacLean, and E. A. Croft. Meet me where I'm gazing: How shared attention gaze affects human-robot handover timing. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 334–341, Bielefeld, Germany, 2014. ACM/IEEE. ISBN 9781450326582. doi:10.1145/2559636.2559656. → pages iv, 15, 75, 95

[106] A. Moon, E. Calisgan, C. Bassani, F. Ferreira, F. Operto, G. Veruggio, E. A. Croft, and H. F. M. Van der Loos. The Open Roboethics Initiative and the elevator-riding robot. In R. Calo, A. M. Froomkin, and I. Kerr, editors, *Robot Law*, chapter 6, pages 131–162. Edward Elgar Publishin, Cheltenham, UK, Northampton, MA, USA, 2016. → pages 12

[107] Y. Moon and C. Nass. How "real" are computer personalities?: Psychological responses to personality types in human-computer interaction. *Communication Research*, 23(6):651–674, Dec 1996. ISSN 0093-6502. doi:10.1177/009365096023006002. → pages 62, 80, 86

[108] A. Mörtl, T. Lorenz, and S. Hirche. Rhythm patterns interaction - synchronization behavior for human-robot joint action. *PLoS ONE*, 9(4):e95195, Apr 2014. ISSN 1932-6203. doi:10.1371/journal.pone.0095195. → pages 15

[109] K. P. K. Murphy. *Dynamic Bayesian networks: representation, inference and learning*. Ph.d., University of California Berkeley, 2002. → pages 41

[110] B. Mutlu. *Designing Gaze Behavior for Humanlike Robots*. PhD thesis, Pittsburgh, PA, USA, 2009. AAI3367045. → pages 16, 18

[111] B. Mutlu, J. Forlizzi, and J. Hodgins. A storytelling robot: Modeling and evaluation of human-like gaze behavior. In *2006 6th IEEE-RAS International Conference on Humanoid Robots*, pages 518–523, Dec 2006. doi:10.1109/ICHR.2006.321322. → pages 18

[112] B. Mutlu, T. Shiwa, T. Kanda, H. Ishiguro, and N. Hagita. Footing in human-robot conversations: How robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction*, HRI '09, pages 61–68, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-404-1. doi:10.1145/1514095.1514109. → pages 18

[113] B. Mutlu, F. Yamaoka, T. Kanda, H. Ishiguro, and N. Hagita. Nonverbal leakage in robots: communication of intentions through seemingly unintentional behavior. *ACM/IEEE International Conference on Human-Robot Interaction*, pages 69–76, 2009. → pages 11

[114] A. Netick and S. T. Klapp. Hesitations in manual tracking: A single-channel limit in response programming. *Journal of Experimental Psychology: Human Perception and Performance*, 20 (4):766–782, 1994. ISSN 0096-1523. doi:10.1037/0096-1523.20.4.766. → pages 38, 77

[115] S. Nikolaidis, A. Kuznetsov, D. Hsu, and S. Srinivasa. Formalizing human-robot mutual adaptation: A bounded memory model. *ACM/IEEE International Conference on Human-Robot Interaction*, 2016-April:75–82, 2016. ISSN 21672148. doi:10.1109/HRI.2016.7451736. → pages 13

[116] Y. Ogai and T. Ikegami. Microslip as a simulated artificial mind. *Adaptive Behavior*, 16(2-3): 129–147, Apr 2008. ISSN 1059-7123. doi:10.1177/1059712308089182. → pages 38

[117] N. Oliver, A. Garg, and E. Horvitz. Layered representations for learning and inferring office activity from multiple sensory channels. *Computer Vision and Image Understanding*, 96(2): 163–180, Nov 2004. ISSN 10773142. doi:10.1016/j.cviu.2004.02.004. → pages 41

[118] Open Roboethics initiative. Results: Will autonomous cars create more jobs?, 2014. URL http://www.openroboethics.org/results-will-autonomous-cars-create-more-jobs/. Accessed on 2016-11-09. → pages 8

[119] Open Roboethics initiative. The two factors in bath by robots: Privacy and control, 2014. URL http://www.openroboethics.org/the-two-factors-in-bath-by-robots-privacy-and-control/. Accessed on 2016-11-09. → pages

[120] Open Roboethics initiative. The ethics and governance of lethal autonomous weapons systems : An international public opinion poll. Technical report, Open Roboethics initiative, Vancouver, BC, Canada, 2015. → pages 8

[121] Oxford University Press. Definition of negotiation in english, 2013. URL https://en.oxforddictionaries.com/definition/negotiation. Accessed on 2013-10-13. → pages 35

[122] Oxford University Press. Definition of persistence in english, 2013. URL http://oxforddictionaries.com/definition/english/persistence?q=persistency#persistence__5. Accessed on 2013-10-13. → pages 40

[123] A.-L. Pais Ureche and A. Billard. Learning bimanual coordinated tasks from human demonstrations. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*, pages 141–142, 2015. ISBN 9781450333184. doi:10.1145/2701973.2702007. → pages 57

[124] M. L. Patterson. A sequential functional model of nonverbal exchange. *Psychological Review*, 89(3):231–249, 1982. → pages 18

[125] A. Peters. *Spatial coordination - human and robotic communicative whole-body motions in narrow passages*. PhD thesis, 2012. → pages 37

[126] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng. ROS: an open-source Robot Operating System. In *ICRA Workshop on Open Source Software*. IEEE, 2009. → pages 64

[127] M. Ragni, A. Rudenko, B. Kuhnert, and K. O. Arras. Errare humanum est: Erroneous robots in human-robot interaction. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 501–506. IEEE, Aug 2016. ISBN 978-1-5090-3929-6. doi:10.1109/ROMAN.2016.7745164. → pages 104

[128] K. Reed, M. Peshkin, M. J. Hartmann, M. Grabowecky, J. Patton, and P. M. Vishton. Haptically linked dyads: are two motor-control systems better than one? *Psychological science*, 17(5): 365–6, May 2006. ISSN 0956-7976. doi:10.1111/j.1467-9280.2006.01712.x. → pages 8

[129] R. W. Reeder and R. A. Maxion. User interface defect detection by hesitation analysis. *Proceedings of the International Conference on Dependable Systems and Networks*, 2006: 61–70, 2006. doi:10.1109/DSN.2006.71. → pages 37

[130] B. Reeves and C. Nass. *The Media Equation: How People Treat Computers, Television, and New Media like Real People and Places*. Cambridge University Press, 1996. ISBN 157586052X. URL http://www.amazon.com/The-Media-Equation-Television-Information/dp/1575860538. → pages 10, 79

[131] C. Rich, B. Ponsleur, A. Holroyd, and C. L. Sidner. Recognizing engagement in human-robot interaction. In *Proceedings of the 5th ACM/IEEE International Conference on Human-robot Interaction*, HRI '10, pages 375–382, Piscataway, NJ, USA, 2010. IEEE Press. ISBN 978-1-4244-4893-7. → pages 18

[132] P. Rober. Some hypotheses about hesitations and their nonverbal expression in family therapy practice. *Journal of Family Therapy*, 24(2):187–204, 2002. ISSN 1467-6427. doi:10.1111/1467-6427.00211. → pages 37

[133] S. Robotics. Softbank robotics: Who is pepper? URL https://www.ald.softbankrobotics.com/en/cool-robots/pepper. Accessed on 2016-11-14. → pages 1

[134] H. Romat, M. Williams, X. Wang, B. Johnston, and H. Bard. Natural human-robot interaction using social cues. In *Human-Robot Interaction*, pages 503–504, 2016. ISBN 9781467383707. → pages 10

[135] B. Roth. Foreword. In B. Siciliano and O. Khatib, editors, *Springer Handbook of Robotics2*, pages v–ix. Springer-Verlag, New York, 2008. ISBN 9783319325507. → pages 7

[136] A. Sandy. Socially aware computation and communication. *Computer*, 38(3):33–40, 2005. → pages 41

[137] S. Schaal. The New Robotics – towards human-centered machines. *HFSP Journal*, 1(2): 115–126, Jul 2007. ISSN 1955-2068. doi:10.2976/1.2748612. → pages 8

[138] M. Scheutz. The inherent dangers of unidirectional emotional bonds between humans and social robots. In P. Lin, K. Abney, and G. Bekey, editors, *Robot ethics: the ethical and social implications of robotics*, chapter 13, pages 205–221. MIT Press, Cambridge, Massachusetts, 2012. → pages 105

[139] E. E. Shriberg. Phonetic consequences of speech disfluency. In *In Proceedings of the International Congress of Phonetic Sciences (ICPhS-99)*, pages 619–622, 1999. doi:10.1.1.52.2699. → pages 37

[140] C. L. Sidner, C. Lee, C. D. Kidd, N. Lesh, and C. Rich. Explorations in engagement for humans and robots. *Artif. Intell.*, 166(1-2):140–164, Aug 2005. ISSN 0004-3702. doi:10.1016/j.artint.2005.03.005. → pages 18, 23

[141] G. Skantze, A. Hjalmarsson, and C. Oertel. Exploring the effects of gaze and pauses in situated human-robot interaction. In *SIGdial Meeting on Discourse and Dialogue*, pages 163–172. Association for Computational Linguistics, 2013. → pages 18, 23

[142] J. S. Smith, C. Chao, and A. L. Thomaz. Real-time changes to social dynamics in human-robot turn-taking. In *IEEE International Conference on Intelligent Robots and Systems*, volume 2015-Decem, pages 3024–3029, 2015. ISBN 9781479999941. doi:10.1109/IROS.2015.7353794. → pages 37, 38

[143] Statistics Canada. Canada's population estimates: Age and sex, july 1, 2015. Technical report, Statistics Canada, 2015. URL http://www.statcan.gc.ca/daily-quotidien/150929/dq150929b-eng.htm. → pages 1

[144] M. Staudte and M. Crocker. The utility of gaze in spoken human-robot interaction. In *Workshop on Metrics for Human-Robot Interaction*, pages 53–39, 2008. → pages 18

[145] K. Strabala, M. K. Lee, A. Dragan, J. Forlizzi, S. S. Srinivasa, M. Cakmak, W. Garage, and V. Micelli. Towards seamless human-robot handovers. 2012. → pages 15, 16, 17

[146] K. W. Strabala, M. K. Lee, A. D. Dragan, J. L. Forlizzi, and S. S. Srinivasa. Learning the communication of intent prior to physical collaboration. In *International Symposium on Robot and Human Interactive Communication*, pages 968–973, Paris, France, 2012. → pages 15, 16, 18, 19, 100

[147] K. W. Strabala, M. K. Lee, A. D. Dragan, J. L. Forlizzi, S. Srinivasa, M. Cakmak, and V. Micelli. Towards seamless human-robot handovers. *Journal of Human-Robot Interaction*, 2 (1):112–132, Mar 2013. ISSN 21630364. doi:10.5898/JHRI.2.1.Strabala. → pages 5

[148] C. Suda and J. Call. What does an intermediate success rate mean? an analysis of a piagetian liquid conservation task in the great apes. *Cognition*, 99(1):53–71, Feb 2006. ISSN 00100277. doi:10.1016/j.cognition.2005.01.005. → pages 38

[149] D. S. Syrdal, K. Dautenhahn, K. Koay, and M. Walters. The negative attitudes towards robots scale and reactions to robot behaviour in a live human-robot interaction study. In *23rd Convention of the Society for the Study of Artificiel Intelligence and Simulation of Behaviour, AISB*, pages 109–115, Edinburgh; United Kingdom, 2009. ISBN 1902956850. → pages 104

[150] L. Takayama, W. Ju, and C. Nass. Beyond dirty, dangerous, and dull: What everyday people think robots should do. In *Proc. of Human-Robot Interaction (HRI)*, Amsterdam, NL, 2008. ACM/IEEE. doi:10.1145/1349822.1349827. → pages 8

[151] L. Takayama, D. Dooley, and W. Ju. Expressing thought: improving robot readability with animation principles. In *Proceedings of the 6th international conference on Human-robot interaction - HRI '11*, page 69, New York, New York, USA, 2011. ACM Press. ISBN 9781450305617. doi:10.1145/1957656.1957674. → pages 11

[152] C. Tennant, S. Howard, B. Franks, M. W. Bauer, S. Stares, P. Pansegrau, M. Styśko-Kunkowska, and A. Cuevas-Badallo. Autonomous vehicles - negotiating a place on the road: Executive summary. Technical report, London School of Economics and Political Science Department of Psychological and Behavioural Science; City University, London Department of Sociology, 2016. URL http://www.lse.ac.uk/newsAndMedia/PDF/AVs-negociating-a-place-on-the-road-1110.pdf. → pages 75

[153] The IEEE Global Initiative for Ethical Considerations and in Artificial Intelligence and Autonomous Systems. Ethically aligned design: A vision for prioritizing human wellbeing with artificial intelligence and autonomous systems. Technical report, IEEE Standards Association, 2016. → pages 105

[154] M. Tomasello, M. Carpenter, J. Call, T. Behne, and H. Moll. Understanding and sharing intentions: the origins of cultural cognition. *The Behavioral and brain sciences*, 28(5):675–735, Oct 2005. ISSN 0140-525X. doi:10.1017/S0140525X05000129. → pages 8, 11

[155] P. D. Tremoulet and J. Feldman. The influence of spatial context and the role of intentionality in the interpretation of animacy from motion. *Perception & psychophysics*, 68(6):1047–58, Aug 2006. ISSN 0031-5117. → pages 10

[156] V. Tsiaras, C. Panagiotakis, and Y. Stylianou. Video and audio based detection of filled hesitation pauses in classroom lectures. In *17th European Signal Processing Conference*, pages 834–838, Glsgow, Scotland, 2009. → pages 37

[157] A. Veiga, S. Candeias, D. Celorico, J. Proenca, and F. Perdigao. Towards automatic classification of speech styles. In H. Caseli, A. Villavicencio, A. Teixeira, and F. Perdigao, editors, *Lecture Notes in Artificial Intelligence: Computational Processing of the Portuguese Language, PROPOR 2012*, pages 421–426. Springer-Verlag, 2012. ISBN 978-3-642-28885-2. → pages 37

[158] R. Vertegaal and Y. Ding. Explaining effects of eye gaze on mediated group conversations:: Amount or synchronization? In *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work*, CSCW '02, pages 41–48, New York, NY, USA, 2002. ACM. ISBN 1-58113-560-2. doi:10.1145/587078.587085. → pages 18

[159] Vicon Motion Systems Ltd. VICON, 2013. URL http://www.vicon.com/. Accessed on 2016-12-10. → pages 41

[160] A. Vinciarelli, M. Pantic, and H. Bourlard. Social signal processing: Survey of an emerging domain. *Image and Vision Computing*, 27(12):1743–1759, Nov 2009. ISSN 02628856. doi:10.1016/j.imavis.2008.11.007. → pages 37, 41

[161] M. Von Cranach and J. H. Ellgring. The perception of looking behaviour. In *Social Commununication and Movement*. Academic Press, London, 1973. → pages 18

[162] E. O. Wilson. *The Insect Societies*. Belknap Press of Harvard University Press, 1971. ISBN 9780674454903. → pages 8

[163] C. J. Wong, Y. L. Tay, R. Wang, and Y. Wu. Human-robot partnership: A study on collaborative storytelling. *ACM/IEEE International Conference on Human-Robot Interaction*, 2016-April: 535–536, 2016. ISSN 21672148. doi:10.1109/HRI.2016.7451843. → pages 8

[164] E. S. A. S. M. H. Yang Liu. Comparing HMM, maximum entropy, and conditional random fields for disfluency detection. In *Proceedings of the International Congress of Phonetic Sciences*, pages 619–622, 1999. → pages 37

[165] T. Yokoi and K. Fujisaki. Hesitation behaviour of hoverflies sphaerophoria spp. to avoid ambush by crab spiders. *Die Naturwissenschaften*, 96(2):195–200, Feb 2009. ISSN 0028-1042. doi:10.1007/s00114-008-0459-8. → pages 38

[166] Yotel. Yotel, 2016. URL http://www.yotel.com/. Accessed on 2016-12-21. → pages 7

[167] M. Zheng, A. Moon, B. Gleeson, D. Troniak, M. Pan, B. Blumer, M. Meng, and E. Croft. Human behavioural responses to robot head gaze during robot-to-human handovers. In *2014 IEEE International Conference on Robotics and Biomimetics, IEEE ROBIO 2014*, 2014. ISBN 9781479973965. doi:10.1109/ROBIO.2014.7090357. → pages iv, 15

[168] M. Zheng, A. Moon, E. Croft, and M.-H. Meng. Impacts of robot head gaze on robot-to-human handovers. *International Journal of Social Robotics*, 7(5), 2015. ISSN 18754805 18754791. doi:10.1007/s12369-015-0305-z. → pages iv, 15

[169] D. Zöbel. The Deadlock problem. *ACM SIGOPS Operating Systems Review*, 17(4):6–15, Oct 1983. ISSN 01635980. doi:10.1145/850752.850753. → pages 82

Appendix A

Supplementary Investigations

As discussed in the main body of this thesis, there are studies that supplement the contents of the thesis. Study 6 outlined in this Appendix is one such study the provides additional support for the impact robot use of gaze cues can have on HR handovers (see Section A.1). In addition, presented in Section A.2 are investigations conducted on samples of human hesitation motions. These investigations helped better understand the dataset, although they did not significantly contribute to the design of artificial negotiative hesitation gestures of robots.

A.1 Study 6: Impact of Gaze on Non-naïve Handover Behaviour

This section presents a summary of a handover study that followed Studies 1 and 2 presented in Chapter 3. While the results from Study 2 (Section 3.4) suggest that a robot's use of gaze can influence a human's behavioural decision on when and where to reach for the offered subject, the study was focused on naïve human responses to HR handovers. Study 6, presented in this section, investigates whether such impact of the robot's use of gaze persists in non-naïve handover responses. One of the aims of this study was to reaffirm that the robot's use of gaze during handovers positively helps the two agents interweave subplans about *when* the handover should occur, thereby increasing fluency of the interaction. While the results of Study 2 provide evidence of such impact of gaze, the scope of experimental result is limited to naïve participants' responses. Study 2 by itself is also insufficient to assert that the timing of the robot's use of shared attention gaze or face gaze affects the timing of reaching behaviours elicited in the human receiver. Results from Study 6 sheds light on this relationship. In particular, this study investigates whether earlier timing of robot's use of face gaze during a robot-to-human handover elicits earlier reaching behaviours on the human receiver.

The presented study was led by Minhua Zheng in collaboration with the author. It is only briefly outlined here in order to highlight the implications of the findings relevant to this thesis[1].

[1] More detailed contents of this section has been published in the International Journal of Social Robotics, in which AJung Moon is a co-author: Zheng, M., Moon, A., Croft, E. A., Meng, M. Q. H., Impact of robot head gaze on robot-to-human handover, International Journal of Social Robotics. Minhua Zheng and AJung Moon collaboratively designed the *in situ* experiment, while the experiment and majority of the analysis was conducted by Minhua Zheng with the supervision of AJung Moon and Elizabeth A. Croft.

A.1.1 Experimental Procedure and Hypotheses

Three types of gaze patterns are implemented:

Attn This condition is the same as the Attn condition employed in Study 2. The robot moved its head smoothly and quickly from the *ready position* to the *handover position*. It stayed gazing at the *handover position* afterwards, providing a prolonged shared attention gaze. Given that this was the best performing gaze pattern in Study 2, it is used as a baseline to compare the two new conditions (Face and LongFace-Attn) not tested previously.

Face The robot moved its head from the *ready position* to the participant's face, and remained gazing at the participant's face to provide a prolonged face gaze.

LongFace-Attn Similar to the Face condition, the robot shifted its gaze from the *ready position* to the participant's face. However, afterwards, the robot moved its gaze from the participant's face to the *handover position* to arrive at the handover position approximately at the same time the gripper arrived at the *handover position*.

In all of the gaze condition, the robot used an average velocity of 90 deg/s to shift its gaze from one point to another.

In Study 6, the robotic platform as well as the robot's arm motion during handovers remained the same as in Study 2. Only the experimental gaze cues were changed in order to investigate the following hypotheses:

1. Since the use of face gaze seems to trigger gaze behaviours in humans, the Face and LongFace-Attn conditions, both exhibiting face gaze early in the handover process, will elicit an earlier reach time from the participants than the Attn.
2. Given the participants' tendency to prefer the Turn (that includes a face gaze) over the Attn condition in Study 2, Participants will perceive the Face and LongFace-Attn conditions to be more favorable than the Attn condition.

The timeline of the robot's gripper and head gaze behaviours for the conditions are illustrated in Figure A.1. The robot's arm and gripper behaviours remained consistent across the conditions, but only the robot's head motion varied according to the three conditions above. Just like the procedure used in Study 2, in all experimental conditions, when the robot gripper moved from the *grasp position* to the *ready position*, the robot head tracked the gripper; when the gripper moved from the *ready position* to the *handover location*, the robot head moved according to the conditions.

While gathering a large quantity of samples was important in the between-subjects comparison of Study 2, Study 6 focuses on within-subject comparisons with smaller number of subjects in order to investigate non-naïve handover responses occurring after a set of repeated handovers. A balanced incomplete block design $(v = 3, b = 3, k = 2, r = 2, \lambda = 1)$ [8] was used with six ordered pair of conditions. The condition orders were randomized across participants. Recruitment of participants involved emails, web advertisements and posters with a \$5 monetary incentive. The experiment took approximately 15 minutes per participant.

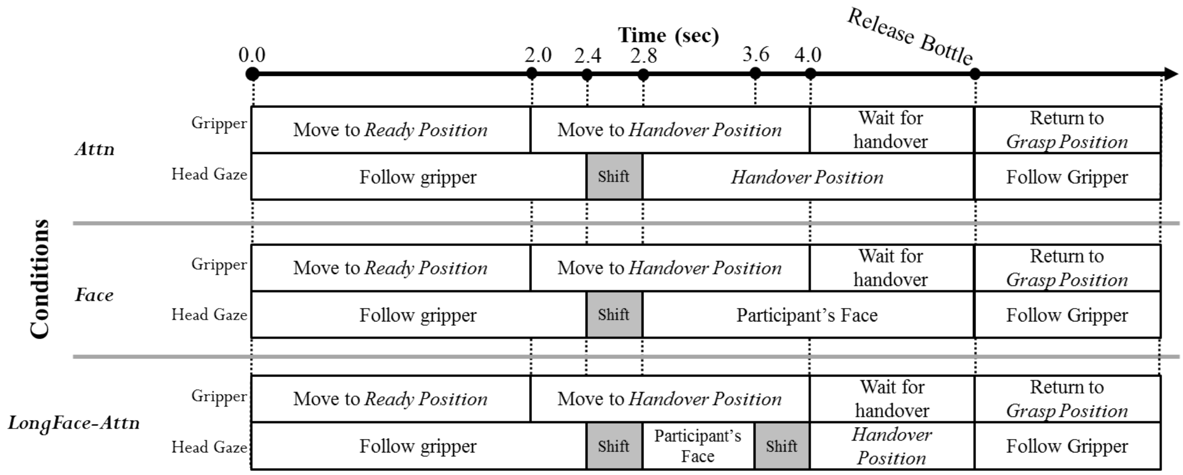

Figure A.1: Timeline of the robot's gripper motion and head gaze for the conditions. The release bottle time depends on the participant's behaviour because the robot was programmed to release the bottle when the participant took it. (©2014 IEEE)

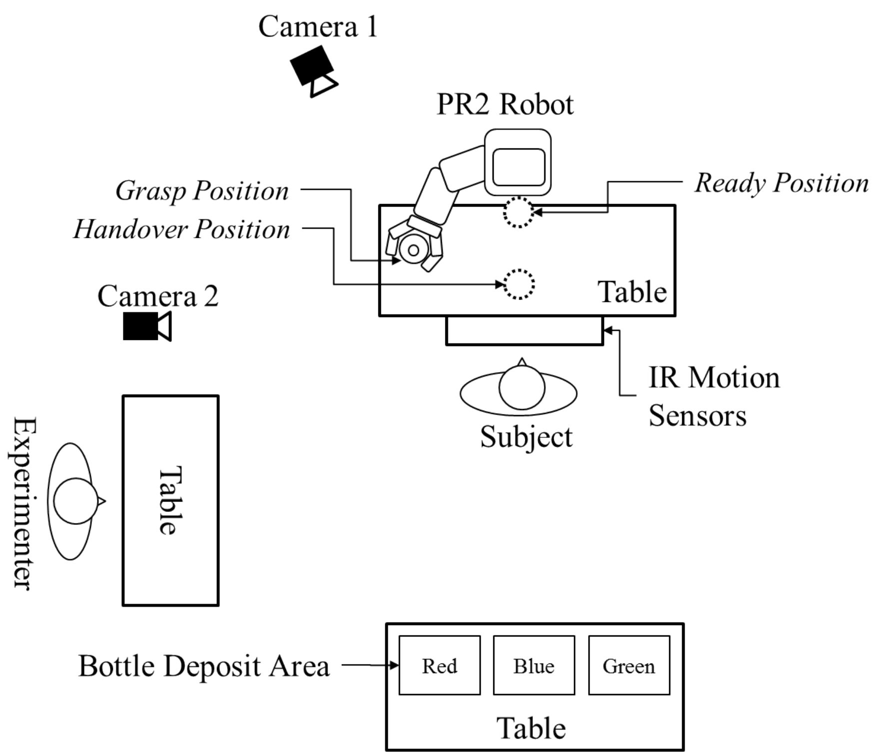

Figure A.2: Study 6 experiment set-up. The participant was instructed to take the bottle from the robot whenever s/he felt it was the right time to do so. After receiving the bottle, the participant took the bottle to a table approximately two meters behind him/her where three tubs were located. (©2015 Springer)

Participants read and signed a consent form before the experiment, and filled in a pre-questionnaire dedicated to collect demographic information. Then the experimenter briefly introduced the experimental procedure to the participants. The experiment included six sessions of handovers. Each session consisted of two handovers and a questionnaire directed to compare the two handovers. The first three sessions included all three pairs of the conditions, and the last three sessions were a repeat of the first three sessions with the order of two handovers reversed in each pair. In total, each participant completed 12 handovers (four trials for each condition) and six questionnaires.

In each handover, the robot lifted a filled water bottle from a table located between the robot and the participant, and handed it to the participant. The robot used gaze cues according to the condition assigned to the trial. The participant was instructed to take the bottle from the robot whenever s/he felt it was the right time to do so. After receiving the bottle, the participant took the bottle to a table approximately two meters behind him/her where three tubs were located. The tubs were labeled red, green and blue, respectively (Figure A.2). The water bottle was also labeled with one of the three colors. The participant was asked to empty the bottle into a tub that match the colour of the water bottle, put the empty bottle into a bin, and return to the robot to start the next handover. The water pouring activity served as a washout between handovers, and distracted the participant while the experimenter readied the robot for subsequent handovers.

After completing two handovers, the participant compared the them on three subjective metrics (*likeability*, *anthropomorphism* and *timing communication*) by selecting either the first or the second handover to the following questions:

1. Which handover did you like better?
2. Which handover seemed more friendly?
3. Which handover seemed more natural?
4. Which handover seemed more humanlike?
5. Which handover made it easier to tell when, exactly, the robot wanted you to take the object?

Similar to Study 2, the participants could optionally provide additional comments.

Questions 1 and 2, inspired by the Godspeed questionnaire [14], provided a measure of *likeability* (Cronbach's $\alpha = 0.83$). Questions 3 and 4, also inspired by [14], measured *anthropomorphism* (Cronbach's $\alpha = 0.91$). Question 5, echoing one of the questions from Study 2, measured *perceived timing communication*.

At the end of the experiment, an experimenter conducted a short structured interview with the participant. The first question was "Did you notice any difference between two handovers in each session?" If the participant answered "Yes", the experimenter asked the participant to describe the difference(s). If the participant mentioned the difference in the robot's gaze pattern or head motion by using words such as "looking at", "head", or "gaze", the experimenter asked about his/her opinion on or preference for the gaze patterns or head motions.

A.1.2 Results

Of the 30 participants recruited, only one participant's data was rejected as an outlier (*reach time* outside 3.6 SD). Hence, data from only 29 participants' (17 male, 12 female; age $M = 24.8, SD = 3.31$) was used in the following analysis. Unsurprisingly, and echoing the findings from Study 2, there is a training effect in *reach time* in the first and second handovers ($t(28) = 1.60$, one-tailed $p = 0.06$). However, this disappears in later trials. Since the focus of this study is in non-naïve handover responses, the analysis only includes measures from the last six HR handovers out of a total of twelve handovers performed per participant. This effectively includes only the third and fourth times each gaze condition was presented to the participant. Results from a repeated-measures ANOVA confirm the findings from Study 2 that gaze condition significantly affects participant's *reach time* during handovers. Post-hoc analyses with Bonferroni method suggest that participants reach for the object significantly earlier in the Face condition ($M = 3.49, SD = 0.34$) compared to Attn ($M = 3.62, SD = 0.42$) as well as LongFace-Attn ($M = 3.61, SD = 0.44$).

Bradley method [9] was used to analyze the paired ranking data for *likeability*, *anthropomorphism*, and *timing communication* measures. The results indicate that the subjective reports vary significantly across gaze conditions ($L > 5.99$ for all three measures)[2]. The ranking of the conditions for the three measures are as follows:

Likeability $LongFace - Attn(0.45) > Face(0.38) > Attn(0.17)$

Anthropomorphism $Face(0.41) > LongFace - Attn(0.39) > Attn(0.20)$

Timing communication $LongFace - Attn(0.48) > Face(0.30) > Attn(0.22)$

A.2 Supplementary Investigations on Human Hesitations

This section provides investigations and results that supplement the work described in Chapter 4.

As reported in Chapter 4, the author used Shooting Algorithm to find saliency of 75 trajectory features studied. The features with a non-zero weight from the Shooting Algorithm are presented in Table A.2. Figure A.3 demonstrates the regularization path of the Shooting Algorithm applied to the four sets of motion samples tested ($N_{b2} = 384$, $N_{b1} = 596$, $N_{ub2} = 1898$, $N_{ub1} = 2004$). The optimum λ value found through this approach 16 for all sample sets and is summarized in Table A.1.

Those that exhibited significant mean differences in t-test or Welch test in addition to having a non-zero weight from the Shooting Algorithm were selected for analysis with SVM. Table A.3 presents the full results of the t-tests. The analysis involving the SVM is outlined in Section A.2.

In order to identify trajectory features that are different in reach and negotiative hesitation gestures, the author developed a number of SVMs using combinations of trajectory features. In addition, the values of ν was varied to tune the models. With the value of $\nu = 0.25$ for the sample set, $N_{b1} = 596$, a total of 82 (9.15%) of support vectors were used when three features ($\max(\dot{d}_1(t))$, $\mu_{\dot{\delta}(t)}$, and $A_{\alpha_1(t)}$) were employed for an SVM. Figure A.4 illustrates the ROC curve of this model.

[2] Critical $X^2(2, 0.05) = 5.99$, with *likeability* ($L = 11.03$), *anthropomorphism* ($L = 6.82$), and *timing communication* ($L = 6.61$).

Table A.1: Optimum λ values obtained from Shooting Algorithm

Sample Set	λ
$N_{b1} = 596$	16
$N_{ub1} = 2004$	16
$N_{b2} = 384$	16
$N_{ub2} = 1898$	16

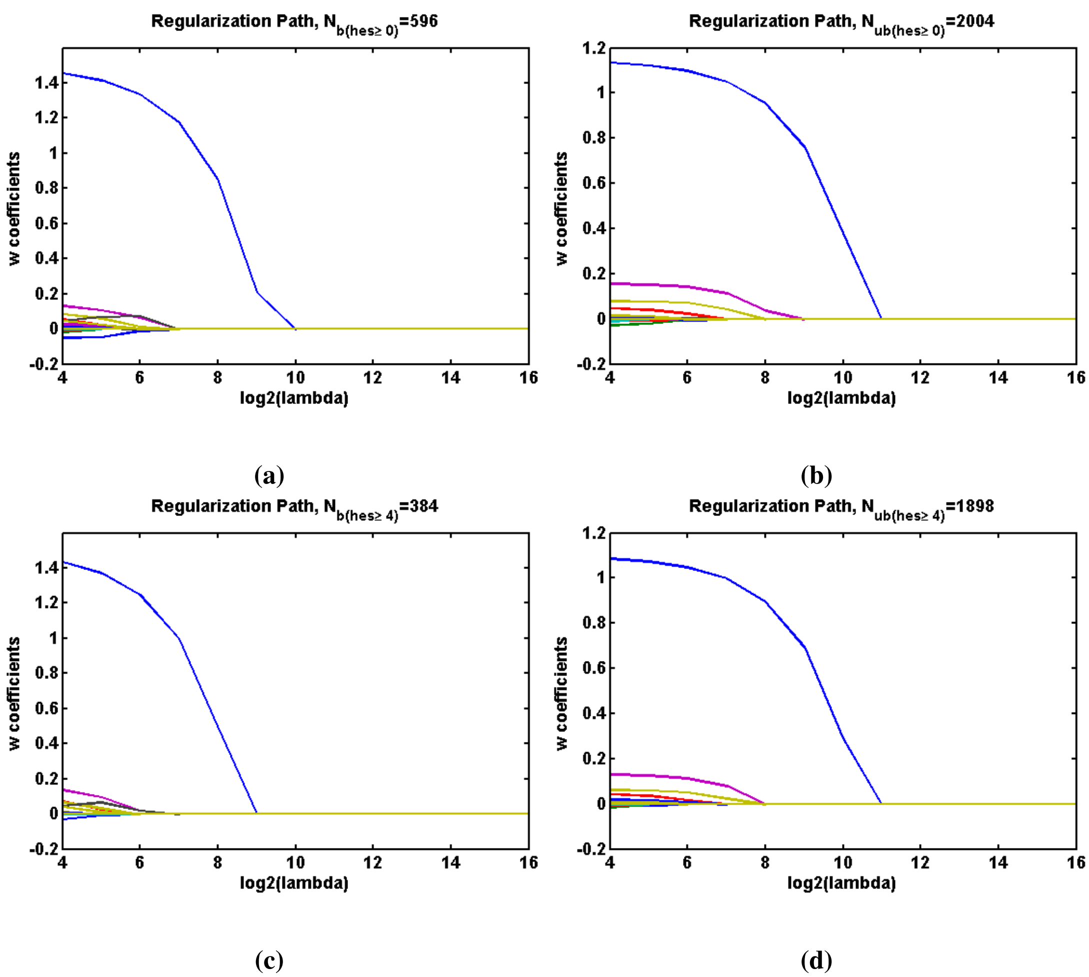

Figure A.3: Regularization path for Shooting Algorithm applied to the four sets of motion samples tested. The top blue line indicates the bias variable in all cases.

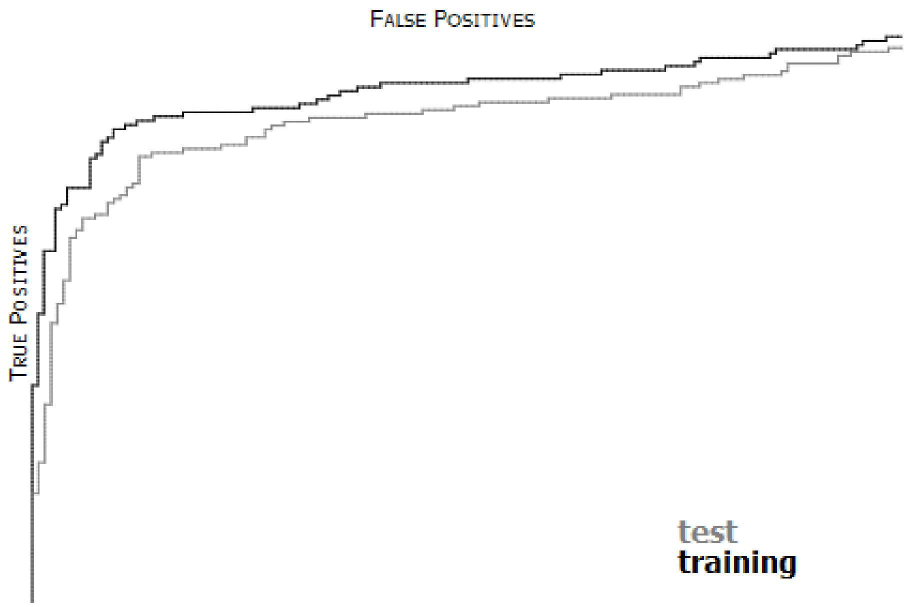

Figure A.4: ROC curve of one of SVM models using $\max(\dot{d}_1(t))$, $\mu_{\dot{\delta}(t)}$, and $A_{\alpha_1(t)}$. 9.15% of the samples were used as Support Vectors.

As mentioned in Section 4.4.3 two features ($\max(\dot{d}_1(t))$ and $\mu_{\dot{\delta}(t)}$) stand out from the rest of the features as strong contributors for accurately classifying hesitations from reach motions. An SVM classifier using only these two features for the $N_{b2} = 384$ sample set returns a test accuracy of 85.8% ($CI: 69.4, 87.0$) with 50% test/train ratio, and 82.4% ($CI: 68.8, 92.3$) with 75% test/train ratio ($\nu = 0.3, SV = 67(17\%)$). This SVM model in feature space is shown in Figure 4.6.

A.2.1 Correlation between Features

In addition to identifying trajectory features that can be used to design artificial negotiative hesitation gestures for a robot, the author also sought to identify what trajectory features may be strongly correlated with human perception of *hesitancy* and *persistency*. Taking the *Hesitency* and *Persistency* scores collected from Study 3, the author conducted a linear regression across the 75 features of all hesitation samples.

Many of the features yielded a non-zero correlation coefficient with significance ($p < 0.05$). However, r and R^2 obtained were too small in all of the features to consider them as reliable predictors of perceived *Hesitancy* and *Persistency*. Interestingly, the author found that the highest linear correlation occurs between *Persistency* and *Hesitancy* scores, ($r = -.92$ [CI:-0.93, -0.90], R^2=0.84). Figure A.5 illustrates this relationship. However, this finding, although intriguing, is not useful for the purposes of generating artificial hesitation trajectories. The next highest correlation with *Hesitancy* score was found with $\rho_{\ddot{d}_2(t)}$ ($r = 0.35$ [$CI: 0.24, 0.44$], R^2=0.12), $\rho_{\ddot{\delta}(t)}$ ($r = 0.34$ [$CI: 0.23, 0.44$], R^2=0.11), and $\rho_{\alpha_2(t)}$ ($r = 0.32$ [$CI: 0.21, 0.42$], R^2=0.10). No correlation with *Persistency*, other than *Hesitancy* Scores, yielded $R^2 > 0.1$.

Table A.2: List of features with non-zero (> 0.0001) weights from the shooting algorithm. Here, S represents the number of sample sets out of four tested that have non-zero weights.

	Feature	Sample Set				S
		$N_{b2} = 384$	$N_{b1} = 596$	$N_{ub2} = 1898$	$N_{ub1} = 2004$	
1	$\mu_{\ddot{d}_1(t)}$	0	0	0.001408951	0.003123393	2
2	$\min(\ddot{d}_1(t))$	0	0.000555272	0.00356456	0	2
3	$A_{\ddot{d}_1(t)}$	0.007659523	0.004717826	0.013822345	0.03424624	4
4	$\min(d_1(t))$	0.092830002	0.072352424	0.047765145	0.0524311	4
5	$\rho_{d_1(t)}$	0.081822797	0.092379396	0.065071197	0.080628918	4
6	$\max(\dot{d}_1(t))$	0.15614317	0.146142143	0.129540803	0.145034212	4
7	$\mu_{\dot{d}_1(t)}$	0.034778617	0.036699048	0	0	2
8	$A_{\dot{d}_1(t)}$	0	0	0.015197184	0.042392476	2
9	$\min(\ddot{\alpha}_1(t))$	0	0.002656939	0.002036993	0.005449252	3
10	$\max(\alpha_1(t))$	0	0	0.010195064	0.009812983	2
11	$\rho_{\alpha_1(t)}$	0.008706001	0.003330487	0.022997939	0.019261962	4
12	$A_{\alpha_1(t)}$	0	0	0.00079188	0.012803684	2
13	$\mu_{\delta_1(t)}$	0.051925879	0.0585937	0.008110461	0.020604495	4
14	$\mu_{d\ddot{e}lta_2(t)}$	0	0	0.004541851	0.010461328	2
15	$\mu_{d_2(t)}$	0.043327179	0.046716392	0.014594146	0.026333731	4
16	$\mu_{\delta(t)}$	0.027825126	0	0.015221386	0.015401922	3
17	$\mu_{\alpha_2(t)}$	0	0	0.003210239	0.007110355	2
18	$\min(\alpha_2(t))$	0	0	0.000635747	0.006038327	2
19	$A_{\alpha_2(t)}$	0	0.047078946	0.003604433	0	2
20	$\max(\dot{\alpha}_2(t))$	0	0	0.003323284	0.000152331	2
21	$A_{\dot{\alpha}_2(t)}$	0	0.021096329	0.001312776	0.009001647	3

Table A.3: Number of significant results ($p < 0.05$) found for each feature across the four runs of t- and Welch tests. The multiple runs of Welch tests on the unbalanced data are redundant.

Feature	N_{b1}	N_{b2}	N_{ub1}	N_{ub2}	Feature	N_{b1}	N_{b2}	N_{ub1}	N_{ub2}
$\max(d_1(t))$	0	1	0	0	$\max(d_2(t))$	4	3	4	4
$\min(d_1(t))$	4	4	4	4	$\min(d_2(t))$	4	4	4	4
$\mu_{d_1(t)}$	0	3	0	0	$\mu_{d_2(t)}$	4	4	4	4
$A_{d_1(t)}$	4	4	4	4	$A_{d_2(t)}$	4	4	4	4
$\rho_{d_1(t)}$	4	4	4	4	$\rho_{d_2(t)}$	4	4	4	4
$\max(\dot{d}_1(t))$	4	4	4	4	$\max(\dot{d}_2(t))$	3	4	4	4
$\min(\dot{d}_1(t))$	4	4	4	4	$\min(\dot{d}_2(t))$	4	4	4	4
$\mu_{\dot{d}_1(t)}$	4	4	4	4	$\mu_{\dot{d}_2(t)}$	4	4	4	4
$A_{\dot{d}_1(t)}$	4	4	4	4	$A_{\dot{d}_2(t)}$	3	4	4	4
$\rho_{\dot{d}_1(t)}$	4	4	4	4	$\rho_{\dot{d}_2(t)}$	4	4	4	4
$\max(\ddot{d}_1(t))$	1	2	4	4	$\max(\ddot{d}_2(t))$	3	4	0	0
$\min(\ddot{d}_1(t))$	4	4	4	4	$\min(\ddot{d}_2(t))$	3	4	0	0
$\mu_{\ddot{d}_1(t)}$	4	4	4	4	$\mu_{\ddot{d}_2(t)}$	4	3	4	4
$A_{\ddot{d}_1(t)}$	3	4	4	4	$A_{\ddot{d}_2(t)}$	3	4	0	0
$\rho_{\ddot{d}_1(t)}$	4	4	4	4	$\rho_{\ddot{d}_2(t)}$	4	4	4	4
$\max(\alpha_1(t))$	4	4	4	4	$\max(\alpha_2(t))$	4	4	4	4
$\min(\alpha_1(t))$	4	4	4	4	$\min(\alpha_2(t))$	4	4	4	4
$\mu_{\alpha_1(t)}$	0	0	0	0	$\mu_{\alpha_2(t)}$	0	0	0	0
$A_{\alpha_1(t)}$	4	4	4	4	$A_{\alpha_2(t)}$	4	4	4	4
$\rho_{\alpha_1(t)}$	4	4	4	4	$\rho_{\alpha_2(t)}$	4	4	4	4
$\max(\dot{\alpha}_1(t))$	0	0	0	0	$\max(\dot{\alpha}_2(t))$	4	4	4	4
$\min(\dot{\alpha}_1(t))$	4	4	4	4	$\min(\dot{\alpha}_2(t))$	4	4	4	4
$\mu_{\dot{\alpha}_1(t)}$	1	0	0	0	$\mu_{\dot{\alpha}_2(t)}$	0	0	0	0
$A_{\dot{\alpha}_1(t)}$	4	4	4	4	$A_{\dot{\alpha}_2(t)}$	4	4	4	4
$\rho_{\dot{\alpha}_1(t)}$	4	4	4	4	$\rho_{\dot{\alpha}_2(t)}$	4	4	4	4
$\max(\ddot{\alpha}_1(t))$	2	3	4	4	$\max(\ddot{\alpha}_2(t))$	1	4	4	4
$\min(\ddot{\alpha}_1(t))$	1	2	4	4	$\min(\ddot{\alpha}_2(t))$	3	4	4	4
$\mu_{\ddot{\alpha}_1(t)}$	1	0	0	0	$\mu_{\ddot{\alpha}_2(t)}$	0	0	0	0
$A_{\ddot{\alpha}_1(t)}$	3	4	4	4	$A_{\ddot{\alpha}_2(t)}$	3	4	4	4
$\rho_{\ddot{\alpha}_1(t)}$	4	4	4	4	$\rho_{\ddot{\alpha}_2(t)}$	4	4	4	4
$A_{\dot{\delta}(t)}$	3	4	4	4	$\max(\delta(t))$	4	4	4	4
$\rho_{\dot{\delta}(t)}$	4	4	4	4	$\min(\delta(t))$	4	4	4	4
$\max(\ddot{\delta}(t))$	0	0	0	0	$\mu_{\delta(t)}$	4	4	4	4
$\min(\ddot{\delta}(t))$	1	1	0	0	$A_{\delta(t)}$	0	0	0	0
$\mu_{\ddot{\delta}(t)}$	0	0	0	0	$\rho_{\delta(t)}$	4	4	4	4
$A_{\ddot{\delta}(t)}$	0	0	0	0	$\max(\dot{\delta}(t))$	4	4	4	4
$\rho_{\ddot{\delta}(t)}$	4	4	4	4	$\min(\dot{\delta}(t))$	4	4	4	4
$\mu_{\dot{\delta}(t)}$	4	4	4	4					

Table A.4: Ratio of the main participant's Euclidean distance to target at zero velocity crossings

Ratio	N	Mean	Std	Var	Range	Min	Max
ZVC2 / ZVC1	179	1.0304	0.3841	0.1475	3.5800	0.2586	3.8385
ZVC3 / ZVC2	112	0.9719	0.2611	0.0682	2.0391	0.3219	2.3609
ZVC4 / ZVC3	56	1.1273	0.3148	0.0991	1.6052	0.7338	2.3390
ZVC5 / ZVC4	28	1.0350	0.1352	0.0183	0.6157	0.8194	1.4351
ZVC6 / ZVC5	20	1.0755	0.4036	0.1629	2.0191	0.5527	2.5719
ZVC7 / ZVC6	10	1.0459	0.1052	0.0111	0.2880	0.9236	1.2116
ZVC8 / ZVC7	7	0.9707	0.0524	0.0027	0.1316	0.8751	1.0067
ZVC9 / ZVC8	4	0.9240	0.1528	0.0233	0.3097	0.6949	1.0046
ZVC10 / ZVC9	4	1.0066	0.0248	0.0006	0.0538	0.9889	1.0427
Mean	-	1.02080	0.20380	0.05930	1.18250	0.68540	1.86790
STD	-	0.06100	0.14140	0.06340	1.20530	0.25910	0.97660
VAR	-	0.00370	0.02000	0.00400	1.45270	0.06710	0.95380
Min	-	0.92400	0.02480	0.00060	0.05380	0.25860	1.00460
Max	-	1.12730	0.40360	0.16290	3.58000	0.98890	3.83850

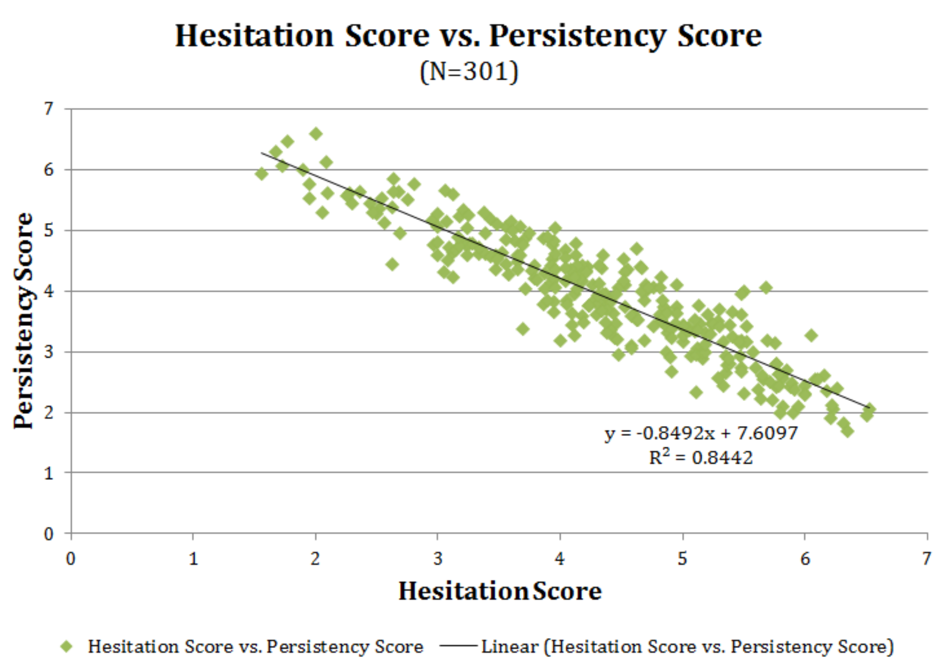

Figure A.5: Correlation between *Hesitation* and *Persistency* scores obtained from the Mechanical Turk survey (Study 3).