

UNIVERSITY OF CALGARY

Exploring Socially Appropriate Nonverbal Robotic Interruption

by

Paul Saulnier

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE
DEGREE OF MASTER OF SCIENCE

DEPARTMENT OF COMPUTER SCIENCE

CALGARY, ALBERTA

JANUARY, 2011

© Paul Saulnier, 2011

UNIVERSITY OF CALGARY
FACULTY OF GRADUATE STUDIES

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies for acceptance, a thesis entitled "Exploring Socially Appropriate Nonverbal Robotic Interruption" submitted by Paul Saulnier in partial fulfilment of the requirements of the degree of Master of Science.

*Supervisor, Dr. Ehud Sharlin
Department of Computer Science*

*Co-supervisor, Dr. Saul Greenberg
Department of Computer Science*

*Dr. Reda Alhajj
Department of Computer Science*

*Dr. Michele Jacobsen
Faculty of Education*

Date

Abstract

Robots are becoming more common and pervasive in our everyday environments, increasing the level of interaction between robots and people. For robots to interact with people in ways that are socially appropriate, both robot and human must understand each other's behaviours and respond accordingly.

In this thesis, we address a narrowed down subset of the above goal: exploring the process of robotic interruption. Our primary contribution is the first (to the best of our knowledge) academic exploration of nonverbal interruption in human-robot interaction. We describe a methodological process for designing minimal robot behaviours for social interruption based on human-human interruption observations, then realized these behaviours on a robot via Wizard of Oz methodologies and robotic interaction implementations, and then designed an evaluation of those behaviours in a set of pilot studies and a final user study. We found that people were able to interpret robot behaviour as interruptions, and we identified the dominant cues people used to relate robotic behaviour to interruption urgency.

This thesis also makes two other secondary contributions. The first is a contribution of a simple yet powerful methodology for probing interruption in HRI. We observe human behaviour in an interruption context, prototype a robot's behaviour using the nonverbal physical behavioural cues observed in the human behaviour, and evaluate how people interpret these cues when used by a robot. The second is the first (to the best of our knowledge) research exploration of bioelectric signal interfaces in implicit human-robot interaction, where the robot is programmed to react to the person's implicit emotional state rather than to direct control.

Publications

Some of the materials, ideas, and figures in this thesis have previously appeared in the following publications.

P. Saulnier, E. Sharlin, and S. Greenberg. Using Bio-electrical Signals to Influence the Social Behaviours of Domesticated Robots. In *Adjunct Proc. Human Robot Interaction (Late Breaking Abstracts) - HRI'09*, page 2 pages plus poster, San Diego, California, March 11-13 2009. Late-breaking abstract. Includes poster shown at conference.

P. Saulnier, E. Sharlin, and S. Greenberg. Exploring interruption in HRI using Wizard of Oz. In *DVD Proc. 5th ACM/IEEE Int'l Conf on Human-Robot Interaction - HRI'2010*, page 2 pages, Osaka, Japan, March 2-5 2010. IEEE/ACM. Late Breaking Report.

Acknowledgements

I would like to acknowledge the contributions of those people who supported me in the creation of this thesis:

Thank you to my supervisor, Dr. Ehud Sharlin, who encouraged me to pursue my master's degree. You have taught me not only how to do research, but also how to face many challenges in life with the correct attitude.

Thank you also to my co-supervisor, Dr. Saul Greenberg. Your advice, your insight and your guidance have been invaluable.

Thank you to the members of the Interactions Lab for your support and critique through the preliminary design stages of my user study. Thank you also to June au Yeung for her assistance in administering my user studies.

Thank you to my parents, Ruth and Terry Saulnier, who have encouraged and supported me in all that I have ever done.

Finally, I would like to acknowledge the Natural Sciences and Engineering Research Council of Canada (NSERC), the Alberta Informatics Circle of Research Excellence (iCORE) and the University of Calgary for their financial support.

Table of Contents

Approval Page	ii
Abstract	iii
Publications	iv
Acknowledgements	v
Table of Contents	vi
List of Tables	ix
List of Figures and Illustrations	x
CHAPTER ONE: INTRODUCTION	1
1.1 Background	1
1.2 Motivation	2
1.3 Research Approach	4
1.4 Thesis Contributions	6
1.5 Thesis Overview	7
CHAPTER TWO: RELATED WORK	9
2.1 Social Human-Robot Interaction	9
2.2 The Concept of Interruption	13
2.3 Identifying Disruptive Interruption Behaviour	17
2.4 Interruption in Human-Computer Interaction	18
2.5 Interruption in Human-Robot Interaction	21
2.6 Summary	26
CHAPTER THREE: EXPLORATION OF IMPLICIT EMOTIVE-BASED HUMAN-ROBOT INTERACTION	28
3.1 Background	29
3.2 The System	30
3.2.1 Direct Explicit Control	31
3.2.2 Behavioural Implicit Control	32
3.3 Discussion	34
3.3.1 Limitations	35
3.3.2 Adapting Rules of Behaviour	35
3.3.3 Public Perception	37
3.3.4 Possible Applications	37
3.4 Summary	38
CHAPTER FOUR: DESIGN APPROACH AND USER STUDY	39
4.1 Observational Study: Identifying Robot Interruption Behaviour	40
4.1.1 Methods	40
4.1.2 Results	43
4.1.3 Discussion	46
4.2 1 st Pilot (Robot Interruption User Study)	49
4.2.1 Selection of Physical Form	49
4.2.1.1 Basic Design	50
4.2.1.2 Feasibility of Human and Animal Forms	50
4.2.2 Implementation	52

4.2.2.1 Robot Platform.....	52
4.2.2.2 Robot Control	53
4.2.3 Methods	53
4.2.3.1 Designing Robotic Interruption Behaviour.....	54
4.2.3.2 Experimental Procedures	55
4.2.4 Results and Discussion	56
4.3 2 nd Pilot (Robot Interruption User Study).....	57
4.3.1 Redesigning Robot Behaviours	57
4.3.2 Description of Specific Robot Behaviour.....	59
4.3.2.1 Null Base Case.....	59
4.3.2.2 Base Case 1 (Episodes 1A to 1F).....	60
4.3.2.3 Base Case 2 (Episodes 2A to 2D).....	60
4.3.3 Implementation.....	61
4.3.3.1 Robot Platform.....	61
4.3.3.2 Controller Station.....	61
4.3.4 Methods	64
4.3.4.1 Data Collection	64
4.3.4.2 Experimental Procedures	65
4.3.5 Results and Discussion	69
4.4 Main Study (Robot Interruption User Study)	69
4.4.1 Participants	69
4.4.2 Implementation.....	70
4.4.3 Methods	70
4.4.4 Results and Discussion	70
4.5 Summary	70
CHAPTER FIVE: RESULTS OF THE USER STUDY.....	72
5.1 Quantitative.....	72
5.1.1 Selection of Statistical Analysis Model.....	72
5.1.2 Identifying Significant Robot Behavioural Cues	73
5.1.3 Means	75
5.1.4 Summary of Quantitative Statistical Analysis.....	76
5.2 Qualitative Analysis Methodology	77
5.2.1 Transcription Log Collection	77
5.2.2 Analysis	79
5.3 Qualitative Results.....	82
5.3.1 The Null Base Case: Impressions of the Robot.....	83
5.3.2 Episode 2D: fast, erratic gaze, close proximity, rotating	84
5.3.3 Episode 1A: slow, direct gaze, far from doorway	86
5.3.4 Episode 2A: slow, direct gaze, close proximity, rotating body language	87
5.3.5 The Robot as a Social Being	88
5.3.6 The Robot as a Machine	88
5.3.7 Politeness When Interrupting	89
5.3.8 Familiarity with Robot	90
5.4 Summary	90
CHAPTER SIX: DISCUSSION	91

6.1 Significance of Robotic Behavioural Cues	91
6.1.1 Impact of Ambient Noise	93
6.2 Social View of the Robot.....	93
6.2.1 Politeness and Emotional Response	94
6.2.2 Suspension of Disbelief.....	95
6.3 Inappropriate behaviour	96
6.4 Appropriate Timing of Interruptions	98
6.5 Emotive-based Coordination of Interruption.....	100
6.5.1 Modifying Bounded Deferral	100
6.5.2 Feedback Mechanism	100
6.6 Summary	102
CHAPTER SEVEN: CONCLUSIONS AND FUTURE WORK.....	103
7.1 Research Question, Revisited	103
7.2 Thesis Contributions, Revisited.....	104
7.3 Future Work	105
7.3.1 Investigating Additional Interruption Behaviour Cues	105
7.3.2 Coordinated Interruption	106
7.3.3 Interruption in Different Environments.....	108
7.3.4 Alternate Physical Forms	109
7.4 Autonomous Implementation	110
7.5 Generalizing to Other Cultures and Environments.....	111
7.6 Final Words.....	112
APPENDIX A: STUDY MATERIALS.....	113
A.1. Informed Consent Form.....	114
A.2. Setting	116
A.3. Experimental Protocol	118
A.4. Conversational Topics	120
A.5. Interview Questions	121
A.6. Robot Interaction Descriptions	122
A.7. Instructions for Study Administrator Tasks.....	124
APPENDIX B: TEDDY BEAR ROBOT FORM.....	126
B.1. Implementation	126
B.2. Discussion	128
REFERENCES	130

List of Tables

Table 4.1: Raw observations of robot actors' behaviour during the observational study.....	45
Table 4.2: Common observations (to 2 or 3 out of 3 robot actors) behaviour during the observational study.	46
Table 4.3: Definition of Behavioural Episodes by Cue used in the 2 nd Pilot. (proximity to person, gaze/head movement, body language, speed of motion)	57
Table 5.1: Significance of Speed of Motion	74
Table 5.2: Statistical Significance of Gaze and Head Movement	74
Table 5.3: Statistical Significance of Proximity	75
Table 5.4: Statistical Significance of Body Language (using rotation)	75
Table 5.5: Interruptedness Means by Factor.....	76

List of Figures and Illustrations

Figure 2.1: A person plays with a Sony AIBO	9
Figure 2.2: Museum visitors interact with Sage (Nourbakhsh et al., 1999)	10
Figure 2.3: PRoP Personal Roving Presence (Paulos and Canny, 1998).....	11
Figure 2.4: Elderly people and a caregiver interact with the Paro seal robot. (Wada 2005)	12
Figure 2.5: A robot approaches a person in a shopping mall (Satake et al. 2009)	21
Figure 2.6: Onlookers observe robots in a train station (left), and a person walks past the robot without observation (right). (Hayashi et al. 2007).....	22
Figure 2.7: A robot approaches a person in a seating context. (Dautenhahn et al. 2006)	24
Figure 3.1: Our NIA-Roomba brain-robot interface prototype	30
Figure 3.2: The OCZ NIA output channels	31
Figure 3.3: Our custom GUI which graphically displays the person's stress level.	33
Figure 4.1: A robot actor looks inside the office from a distance to interrupt unobtrusively (scenario 1).....	40
Figure 4.2: An actor kicks the person's chair while attempting to interrupt him.	47
Figure 4.3: The robot	52
Figure 4.4: The robot interrupts a participant in a design critique session.	55
Figure 4.5: The robot's four distinct motion paths (labelled I, II, III, IV).....	59
Figure 4.6: Remote control station used for robot control and transcription.....	61
Figure 4.7: A view of the user interface used to invoke the pre-programmed behaviour macros.	62
Figure 4.8: The transcription tool.	63
Figure 4.9: A partial view of the Interruptedness Metre used by participants to rank	64
Figure 4.10: The robot enters the office to interrupt a participant.....	66
Figure 5.1: Sample of raw data generated from transcription logs. The bottom two lines have the participant identifiers inserted.....	78

Chapter One: Introduction

1.1 Background

Robots are becoming more common and pervasive in our everyday environments. As this process continues, we can expect to see more instances of robots interacting with people, and more instances of robots interacting with people in a socially acceptable manner. In Japan, for example, an ageing and declining population has led researchers to explore the use of domestic robots as a means to address an expected shortage of human workers (Mori and Scearce, 2010). As one instance of this, Professor Hiroshi Kobayashi (Tokyo University of Science) created a humanoid receptionist robot that, while immobile, performs the function of greeting office visitors and providing basic introduction information within some predefined limits (Hashimoto et al., 2007). More generally, robots have already entered the public consciousness as commercially available highly sophisticated robotic pets (e.g., the Sony AIBO) and even as automatic vacuum cleaners (e.g., the iRobot Roomba).

For robots to work in such social settings, both robot and human must understand each other's behaviours and respond accordingly. This is not yet something that we, as interaction designers, fully understand how to do. The context of communication between people, let alone people and robots, is a complex phenomenon. A large variety of behaviours realize the act of human communication, including language, tone of voice, gesture, posture, body movements, spatial orientation, physical proximity, eye contact, and facial expression (Riley, 1976). It is unrealistic (at least for now) for a robot to fully exhibit this behavioural richness. Our research should be viewed within this larger context of providing social modalities and characteristics to robotic interfaces.

Our own overarching general question is: are there *minimal nonverbal behavioural cues* that robots can exhibit to communicate their internal state, and are those cues understandable by people? By *minimal*, we mean that we are interested in determining behaviours that rely on only a few simple physical capabilities present (or that can be easily added to) most robots.

In this thesis, we primarily address a narrowed down subset of the above goal: exploring the process of robotic interruption. For people, interruptions are a normal part of our daily life; we are all well accustomed to being interrupted by others, and are generally quite proficient in interrupting another person when we need to. We change our behaviours to initiate an interruption, where the particular behaviour we exhibit informs the person being interrupted about the importance and urgency of a situation. Our actions are based on our expectations of how others will understand, interpret and ultimately respond to our interruption behaviours.

Interruption behaviour, while natural for humans, has not yet been a primary focus of human-robot interaction research. We believe that exploring the application of socially acceptable and appropriate robotic interruption has the potential to enhance the social experience that people will have with future robots by adding an extra layer of communication that can be used in human-robot interaction.

1.2 Motivation

Designing comprehensible robotic behaviours that are capable of initiating and tuning a socially acceptable interruption will be extremely important as robots increasingly interact with people. Robots will have to communicate different information to people following different levels of urgency. As such they may have to interrupt the person in a

contextually meaningful and appropriate way, where the person can respond to that interruption accordingly. This leads to our primary research question:

Are there *minimal nonverbal behavioural cues* that robots can exhibit to communicate interruption urgency, and are those cues understandable by people?

Arguably, some classes of robots will be capable of using verbal communication when interrupting users. Indeed, verbal communication is usually an important component, often the most important component, of the way people interrupt each other. However, we argue that interruption between humans, verbal or not, often relies on a fundamental layer of physicality, which is nonverbal and involves spatial movement, interpersonal distance, gaze, etc.,. Thus, we believe that physicality is an important standalone layer of interaction. We hypothesise that for robots to be able to interrupt humans in a socially acceptable manner, designers of social robotic interfaces will need to master and apply this fundamental physical layer of interruption regardless of the type of modality they use. Lessons learned from the study of the physicality aspect of interruption will be applicable to almost any robotic interface, verbal or not.

Mastery of the interruption layer in human-robot interaction (HRI) does not need to begin from scratch. There is already considerable past research on interruption and interruptibility in human-computer interaction (HCI). McFarlane and Latorella (2002) identified different types and methods of interruptions as well as the ways humans respond to them, and the different phases that occur as the interruption is carried out. Their models enable adapting a specific interruption to specific circumstances as determined by factors such as message urgency, the person's interruptibility state, etc.

Fogarty (2006) applied this theory to the design of contextually-aware desktop computer software that decides when and how to deliver information to its users. On a desktop system, interruptions are usually in the form of visual and audio alerts. While simple, there are significant tradeoffs in displaying an alert or other indicator in a way that balances interruption from distractions from disruption.

While robots can possess the computing power of desktop computers, they also have other abilities of physicality at their disposal that they can exploit as interruption cues. They can adapt their spatial position to create an ambient status indicator that exists at the just detectable edge of a person's view. They can move towards a person. If equipped with the right attachments, a robot can make facial or hand gestures the way a person would.

Clearly, it is currently feasible to consider robots that interrupt people. However, as far as we know, there is still no robot behaviour model that directly addresses interruption. We hope that our research will provide a baseline exploration that will inform future exploration of interruption in social human-robot interaction.

1.3 Research Approach

Our research questions recognize minimal nonverbal behavioural cues as an important yet largely unexplored layer of interaction between humans and robots. As mentioned, our primary question is:

Are there *minimal nonverbal behavioural cues* that robots can exhibit to communicate interruption urgency, and are those cues understandable by people?

Our research did not begin with this well-formed question. Initially, our focus was *emotive-based human-robot interaction*, where we tackled the broader question of how a

robot could understand human emotions, and – based on that understanding – respond in an appropriate manner. Specifically, we interpreted one class of signals provided by a brain-computer interface as suggestive of emotional stress, and used that as an emotional parameter to implicitly influence (but not directly control) robot behaviour. In our example, the robot was programmed to avoid contact with a person when that person was ‘stressed’ in an effort to minimize disruption to that person. The results of this initial study are reported in Chapter 3.

While the robot’s emotion-sensing abilities had very severe limitations, this exploration informed our thinking and shifted our research question towards ways of adapting a robot’s behaviour in socially appropriate ways, specifically toward coordinating interruption. Our review of the state of the art revealed both motivation and opportunity for new research that explicitly explores how robots can interrupt people in ways that balance the person’s comfort and the need to interrupt. As plenty of past research describes how computers can interrupt people, we shifted our focus to a previously unexplored area: identifying *minimal nonverbal behaviour cues* that could be exhibited by a robot in an interruption context, and understanding how people interpret these cues. We had found this question to be valuable and important, though be much narrower and more tractable than our original explorations on how a robot can infer a person’s mental state.

With many different behaviour cues usable by robots, it is important to be able to identify which ones are potentially effective at coordinating an interruption and conveying context and reasoning for the interruption. Our approach to this task consists of a set of interface designs and studies, presented in Chapters 4 to 6. For requirements

gathering, we started with observations of interruption between humans, with a human actor attempting to interrupt other humans while being constrained to use only a set of behavioural cues that could be mimicked by a simple nonverbal robot. Next, we programmed a robot to exhibit these social nonverbal cues, and tested their feasibility in two separate pilot user evaluations. Finally, we performed an extensive user study of robotic nonverbal interruption across interruption scenarios. The results from this user study provide the basis for an in-depth theoretical discussion of robot behaviour in an interruption context.

1.4 Thesis Contributions

This thesis makes the following three contributions, one primary and two secondary.

Primary:

To the best of our knowledge, the first academic exploration of nonverbal interruption in human-robot interaction.

To the best of our knowledge, this thesis represents the first attempt to explicitly justify and explore robots that interrupt people using nonverbal behavioural cues in a socially meaningful and acceptable manner.

Secondary:

A methodology for probing interruption in HRI.

We contribute a simple yet powerful methodology to observe human behaviour in an interruption context, prototype a robot's behaviour using the nonverbal physical behavioural cues observed in the human behaviour, and evaluate how people interpret these cues when used by a robot.

A research exploration of bioelectric signal interfaces in implicit human-robot interaction, where the robot is programmed to react to the person's implicit emotional state rather than to direct control.

To the best of our knowledge, the (vast) prior work in this domain concerns only direct brain-robot control. As far as we know, our emotional state exploration, while rudimentary, is the first attempt to justify and prototype the use of a brain-computer interface to infer a person's implicit emotional state and to mediate a robot's behaviour as a consequence.

1.5 Thesis Overview

The subsequent chapters of this thesis describe in detail the research effort and contributions outlined above. In Chapter 2, we explore the research domain that motivates this thesis and provide background on work related to our efforts, and the current state-of-the-art. In Chapter 3, we prototype a robot – a Roomba - that uses an estimation of a person's emotional state, as sensed by a brain interface, to mediate its behaviour to be less disruptive to the person, and thus more socially appropriate. This earlier work inspired our main research question, addressed in the following chapters. In Chapter 4, we describe our methodological process for designing and evaluating minimal nonverbal physical robot behaviours for social interruption. This process comprises four elements: (a) an observational study to see how people improvise their behaviour to interrupt others using a subset of nonverbal cues over scenarios that vary in urgency, (b) a design critique of these behaviours when mimicked by a robot, (c) a robotic implementation of the behaviours which are triggered and somewhat controlled by a human operator, and (d) a robot interruption user study, where we expose people to these

robotic behaviours, and gather their reaction and interpretation of those behaviours. In Chapter 5, we present the quantitative and qualitative results of the robot interruption user study, and describe the methodology that we used to analyze these results. In Chapter 6, we discuss the user study results, and explore the insights gained from them. Finally, in Chapter 7, we summarize our final conclusions from this research and discuss possible future work. Appendix A includes additional user study materials for the robot interruption user study described in Chapter 4. In Appendix B, we explore our limited implementation of a teddy bear robot form.

Chapter Two: Related Work

This thesis is concerned with designing robot behaviour that can be used to interrupt in socially appropriate ways. Specifically, we focus on our research question: are there *minimal non-verbal behavioural cues* that robots can exhibit to communicate interruption urgency, and are those cues understandable by people?

Being able to do so will be increasingly important in future robot implementations. In this chapter, we explore several themes related to this concern: the emerging field of social human-robot interaction, the concept of interruption itself, a discussion of disruptive interruption behaviour, previous research involving interruption in human-computer interaction, and previous robot applications that use robots to initiate interactions with people. We close by situating the goals of this thesis within the context of HRI for interruptibility.

2.1 Social Human-Robot Interaction

Defining social human-robot interaction is a difficult task. There are many opinions about what qualifies as a social interaction between humans and robots, and many applications that involve some social connection between humans and robots.



Figure 2.1: A person plays with a Sony AIBO

(from <http://www.cse.unr.edu/~monica/Students/students.html>)

Traditionally, robots have been designed for various applications that require very little, if any, interaction with people (Breazeal 2004). Examples include exploring planets, sweeping mine fields, or assembling components in factory assembly lines. In these types of applications, the robot operates autonomously as a sophisticated tool that may be remotely controlled by a human operator. Service robots that vacuum floors or mow lawns may have the additional element of sharing a physical environment with people, though in these cases (from the robot's point of view) people are treated more as obstacles to navigate around.

Recent applications of robots have been developed that possess the ability to interact directly with people in more direct and engaging ways. Commercial applications of robotic toys, such as Sony's AIBO robotic dog (shown in Figure 2.1), adapt and change their behaviour as children (or curious robotics researchers) play with them. It is able to hear and recognize its name as well as dozens of verbal commands, sense touches, and see pink objects. It is able to act autonomously, though its behaviour can also depend on its interaction with a person. It is able to learn and can be trained by a person albeit



Figure 2.2: Museum visitors interact with Sage (Nourbakhsh et al., 1999)

within limitations.

Robot tour guides are also being used in public places to interact with the general public. For example, Sage is a full-time autonomous robotic ‘staff member’ at the Carnegie Museum of Natural History. As visible in Figure 2.2, it interacts with visitors, where it provides them with both entertainment and information (Nourbakhsh et al., 1999). Because it communicates directly with people and moves with them in the same physical space, it also features functional obstacle avoidance and a navigation system.

Another HRI application area is teleconferencing. A robot becomes the surrogate of a distant person, where it embodies the physical and social presence of the distant person in the local location. Figure 2.3 illustrates the PRoP robot which the remote person uses (via teleoperation) to socially interact with others. By moving its head and hands, the person (via the robot) can make eye contact and even shake hands (Paulos and Canny, 1998).




Figure 2.3: PRoP Personal Roving Presence (Paulos and Canny, 1998)

HRI research also covers robots that assist people more as collaborators, personal assistants, or pets rather than merely sophisticated tools. For example, Paro is a commercially-available robot that is designed to provide some of the health related benefits of animal therapy (e.g., Figure 2.4). As a robot, Paro is useful in environments such as hospitals and extended care facilities where live animals pose treatment or logistical difficulties (Wada et al., 2005).

These examples show increasing levels of social interaction between humans and robots, which at the same time creates opportunities to design robots that behave in socially appropriate ways. As this integration deepens, it becomes more important to consider social aspects of interactions between humans and robots.

Some may be sceptical that social interaction is even possible with a machine. Nass et al. (1995) countered this scepticism, where they found that when machines are endowed with even minimal personality-like characteristics, people will respond to them as if they have personalities. In fact, not only is this possible, but it is advisable to take



Figure 2.4: Elderly people and a caregiver interact with the Paro seal robot. (Wada 2005)

this social interaction into consideration under many circumstances. The satisfaction derived from any session involving a person and a machine can be strongly affected by the interaction of the individual's personality and the machine's personality.

Thus, there is value in exploring not only how robots can interact socially with people, but also how they can do so in a way that is appropriate and comfortable for people in social interactions. This is still a broad area, so we now narrow our focus to situations where a robot must initiate an interaction with a person when that person is engaged in tasks unrelated to the robot. In essence, the robot must appropriately perform an interruption. First, we briefly review what is meant by interruption.

2.2 The Concept of Interruption

A high level definition of human interruption is offered by McFarlane and Latorella (2002) as “the process of coordinating abrupt changes in people's activities”. Interruption has been studied scientifically for decades, and includes many factors. Most will not be pursued further in this thesis, although we give a sampling below that motivates our pursuit of interruption behaviour in robots.

Scientific study of interruption began with the classic experiments of Zeigarnik (1927) that positioned interruption as observable phenomena that can change a person's abilities. She showed that people being interrupted during a task would recall the interrupted task more readily than a task that had not been interrupted, thus giving an early glimpse into how interruptions affect human behaviour. Other effects of interruption have also been documented since the Zeigarnik studies.

The value of an interruption depends upon the way the interruption is done, the context of that interruption (e.g., its urgency), and its effect on the person. As a result,

interruptions can range from highly valuable to quite annoying. That is, interruptions themselves are not inherently bad: they are a normal part of communication. They are only 'bad' if there is a poor match between the needs of the interrupter and interrupted, and if the two have not coordinated their communication appropriately. Cohen (1980), for example, found that personal stress can be induced from unpredictable and uncontrollable interruptions. Interruptions can also cause people to become less efficient, more error prone or even both (Kreifeldt and McCarthy, 1981).

People do, however, have natural abilities to adapt behaviour to accommodate interruptions. Flexibility in task performance can reduce the effects of interruption. In some cases, such as those where the primary tasks being interrupted are simple and unchallenging, interruptions can actually enhance performance. Horvitz et. al suggest that this is possible as a result of the person using their unused cognitive capacity with things not related to the central task. Decreased performance does result from interruptions to more complex tasks due to higher demands on cognitive capacity (Horvitz et al., 2005). Thus, it is clear that there are times when an interruption can be a positive experience, and when it can be a hindrance.

These lessons above are particularly relevant to the research presented in this thesis, as they must be considered as factors when designing robot behaviour to appropriate interrupt people. That is, for 'positive' interruptions, a robot must interrupt in appropriate ways and at appropriate times.

The very mechanics and methods of interruption have been scientifically identified and scrutinized. McFarlane and Latorella (2002) provide both the interruption management stage model (IMSM), and a taxonomy covering several major dimensions of

the human interruption problem, some of which we will focus on, as we feel they are particularly relevant to this thesis.

First, their taxonomy identifies the *source of interruption* as a dimension. Sources can be distinguished between *internal interruptions*, which are activities that people perform outside their focus of conscious attention (such as a hiccup), and *external interruptions*, which are side effects of activities that people delegate to other entities such as a computer or other object. The focus of this thesis is the robot as the source of external interruption.

Second, *methods of coordination* address several problems that can cause a failed or inconvenient interruption (McFarlane and Latorella, 2002). One problem in the IMSM is oblivious or unintentional dismissal of an interruption, whereby the person does not notice an interruption. This highlights the importance of designing interruptions that effectively attract the attention of a person and communicate an appropriate level of importance that matches the situation. If done well, a person can intentionally (but appropriately) dismiss an interruption, or place it “on hold” so that focus can remain on a more important task.

McFarlane et.al’s taxonomy also provides four ‘design solutions’ to coordinate the interruption of a person: (a) immediate interruption, (b) negotiated interruption, (c) mediated interruption, and (d) scheduled interruption. These solutions vary in what options and control are given to the person to deal with the interruption. To explain these differences, we use an example of a person seated in an office engaged in a conversation with another person that is periodically interrupted by a robotic office secretary. An immediate solution would have the robot stop the conversation and require the person to

interact with it immediately. A negotiated solution would have the robot announce its desire to interrupt, and then allow the person to deal with the interruption or defer it until later. A mediated solution allows the robot to indirectly interrupt through a mediator object, such as a PDA, that decides when and how to perform an interruption. A scheduled solution would restrict the robot's interruptions to a prearranged schedule, such as once every 30 minutes or every day at 9 o'clock and 1 o'clock. Thus, a robot that is designed to interrupt a person under varied circumstances could be designed to coordinate its interruption attempt in different ways.

Researchers have found that the method of expressing an interruption can mitigate possible negative effects on performance. Obermayer and Nugent (2000), for example, suggest that the importance and urgency of an interruption's message should define the degree of attention-getting cues. For example, an urgent or important interruption may use behaviour that prevents a person from continuing with their current task until they address the interruption (e.g., a modal dialog box). A less urgent or less important message may use behaviour that allows the person to ignore or defer the interruption until a more convenient time (e.g., an alert at the side of the screen). Thus, any system, robots included, designed to interrupt people for a range of different reasons should be able to tailor its interruption behaviour for the context.

Another possible way to mitigate the negative effects on performance is careful choice of *channel of conveyance*, which is the factor of McFarlane's taxonomy that refers to the channel of communication used to communicate with a person (e.g. visual, auditory, etc.). This selection can affect how long it takes for a person to notice an interruption, and also what negative effects result from the interruption. For example,

Taylor (1989) found that visual channels are useful for providing spatial information to aircraft pilots, and voice channels are better used when a pilot is using their eyes for another task. However, Taylor also found people are very sensitive to deficiencies in the designs of voice systems, which increases the amount of serious attention required when engineering them. Thus, the choice of communication channel, along with potential effects on performance, must be considered carefully when designing interruptions to avoid an interruption that is needlessly disruptive.

2.3 Identifying Disruptive Interruption Behaviour

With many methods of interruption possible, it is useful to think about how they will affect a person. Obviously, we should avoid negative effects on the person's performance, while stressing the positive effects if possible.

McFarlane et. al's taxonomy describes various effects of interruption. They explain that interruptions can temporarily inhibit a person's ability to perform post-interruption tasks. They can also cause people to make mistakes, reduce their efficiency, and increase their stress.

However, interruptions do not always have negative effects. The classic interruption experiments by Zeigarnik (1927) mentioned earlier in Section 2.2 are examples of how an interruption can have a positive effect; in that case, participants were able to recall interrupted tasks more readily than non-interrupted tasks. In other cases, the effects of the interruption may be negligible. Chapanis and Overbey (1974), for example, found that while interruptions changed the way that participants chose to accomplish a task, the actual performance time was not affected. Clearly, an interruption can have a

wide range of effects, so it is important to know what effect an interruption will have, particularly if it will have a negative effect.

A study by Gillie and Broadbent (1989) explored how some interruptions are disruptive and how some are not. Specifically, they were interested in the factors of interruption length, similarity of the interruption to the task being interrupted, and complexity of the interruption. They concluded that their participants were surprisingly unaffected by the timing or length of the interruption. They were, however, affected by interruption tasks which were similar to the task being interrupted, and also by complex tasks that made large demands on working memory. Thus, a system that is designed to interrupt people appropriately must consider the task being interrupted very carefully, and have some method of comparing it to the interruption task. This information can be combined with the importance and urgency of the interruption to determine whether or not it is appropriate to perform an interruption.

In the next section, we discuss possible ways to manage interruption.

2.4 Interruption in Human-Computer Interaction

As a part of normal operation, computers often have to present alerts to the user, possibly interrupting them from another task. There are various approaches at mediating interruptions so that they are less disruptive in the context of a computing environment, for example, in the workplace. While people need full concentration to make good decisions and perform well, people are also expected to collaborate with other people, monitor dynamically changing variables in an information environment and supervise other tasks that may be occurring in the background. While it is possible to create

systems to automate these tasks, such systems may need to raise alerts of their own during normal operation. Thus, a means to manage and triage interruption is necessary.

Horvitz et al. (2005) investigate the concept of *bounded deferral* as a method of potentially providing a calmer computing environment with less disruption. The idea is simple: if the user is busy when the system needs to interrupt them with an alert, the system will delay the alert up to a predetermined maximum of time. If the user transitions to a non-busy state before this maximum time elapses, the alert is delivered. If the user remains busy, the alert is delivered once the maximum waiting time has passed. The end result is a small cost of delayed awareness in exchange for a calmer working atmosphere.

However, this work does not provide any results to demonstrate that bounded deferral is actually providing a calmer work environment. Indeed, Gillie and Broadbent's (1989) work shows us that the timing of an interruption does not actually increase or decrease disruptiveness, though it is conceivable that any interruption is going to be less disruptive if it can be performed when the participant is not busy with any task at all.

Storch (1992) explores whether the style of computer user interface used by a person affects their performance following an interruption. In her study, each participant performed data entry using either a graphical user interface with a mouse and screen buttons or a text-based interface with tab and function keys. Each participant was interrupted three times with three different forms: a telephone call, an in-person visitor, and a modal on-screen message. Following each interruption, the participant's performance, as measured by the number of correct data fields entered with errors subtracted, and eye movement were recorded for two minutes.

While the study actually showed that the type of user interface used for data entry had no significant effect on the post-interruption performance of the participant, several interesting differences did emerge regarding the effects of different interruption methods. Participants found the on-screen interruption messages to be most disruptive, as measured by level of performance, while the telephone call was surprisingly not disruptive at all. The in-person visitor was somewhat disruptive, though not as disruptive as the on-screen message.

Storch argued that the nature of the screen interruption sheds light on why it was most disruptive. While all three types of interruptions occurred abruptly as a surprise to the participant, only the screen interruption locked the participant out from the main task. The participant likely had previous life experience with mediating interruptions via telephone and in-person, but probably had less or no experience with the specific design of the on-screen interruption used in the experiment, thus making it more disruptive. Alternatively, since both the main task and screen interruption were mediated through the computer, this similarity may have impacted the level of disruption.

The Storch study presented many interesting lessons to consider when designing systems that interrupt people, as the different abilities of interruption systems by nature may affect how disruptive they will be to a person. For example, a computer's primary method of interacting with a person is through its screen, which was found to be most disruptive in this study. Of course, other form factors of a computer – such as a robot – can significantly change the nature of the interruption. This will be explored next.

Robots, the focus of this thesis, are social players, very much like humans. Because of their physicality, they primarily have the ability to interrupt people “in-

person”, as a person would. Thus, a robot will potentially be interpreted more like a person would be in an interruption context, rather than a machine.

2.5 Interruption in Human-Robot Interaction

The concept of interaction between humans and robots is not new, and plenty of research explores various types of human-robot interaction. We believe that the definition of human interruption proposed by McFarlane and Latorella (2002) as “the process of coordinating abrupt changes in people’s activities” is also applicable to a robot capable of physically interacting with a person in some way. Still, there is room for multiple interpretations of McFarlane and Latorella’s definition. For example, imagine a scenario where a robot presents itself at a person’s office door. The person notices the robot, but chooses to ignore it. McFarlane’s definition seemingly excludes this scenario as no ‘abrupt change’ of activity occurs, yet most people would still consider this an interruption. In this case, the change in the person’s mental state as they notice the robot could be considered “abrupt” enough to classify as an interruption. For the purposes of this thesis, we will refer to any interaction between a human and a robot that results in an abrupt change of a person’s activity, or mental state, as an interruption.

There are many examples of research in HRI that involve interactions between



Figure 2.5: A robot approaches a person in a shopping mall (Satake et al. 2009)

humans and robots that are effectively interruptions as well. However, no prior research explicitly focused on the interruption aspect of the interaction. Thus, the unique circumstances of interruption in HRI are unexplored. Nonetheless, previous research does provide useful lessons generalizable to interruption.

Satake et al. (2009) created a robot that approaches visitors to a shopping mall. The robot, shown in Figure 2.5, roams around a predetermined area looking for people within a close distance to approach. Once a person is identified, the robot takes the shortest path at a constant predetermined speed to reach a position close enough to initiate a conversation while gazing toward the person. If the person responds to the robot, then it recommends local shops to visit. Otherwise, the robot continues on its path and looks for another person to approach.

Let us reinterpret this work in interruption terms. The robot uses spatial locality and gaze to attempt to attract the attention of and thus interrupt a person. If the interruption is accepted by the person, then the robot communicates verbally. In this shopping scenario, the interruption is neither important nor critical. If the robot fails to get the person to respond, it disregards the attempt, and looks for another person.



Figure 2.6: Onlookers observe robots in a train station (left), and a person walks past the robot without observation (right). (Hayashi et al. 2007)

Although the path and specific motion path of the robot will differ for each interruption attempt, the methods of interruption used and interruption message (e.g. recommendation of shops) remain the same. This work provides excellent motivation for exploring whether varied speed or robot movement has an effect on how people perceive the interruption, and how to properly interrupt people in a wide range of situations.

Another study by Hayashi et al. (2007) explores how a robot can attract attention to itself in a public space, such as a train station (see Figure 2.6). Two robots were situated in the train station and engaged in a staged conversation with each other about station information and travel as people watched. This format is similar to many television news programs that have several anchors conversing with one another. They found that this passive-social medium was an effective means for robots to attract the attention of people, and notably more effective than a single robot simply talking to itself. Many people, especially those who did not appear to be busy, were attracted to the robots and stood watching them. People who were busy with their travels simply ignored or did not notice the robots, and thus were not disturbed by their behaviour. In fact, the majority of people passing through the train station ignored the robots. However, others were very interested in the robots, stopping to talk about them, touch them, and looking at them for a long time. The methodology used in this study could be useful when designing robots to interrupt people, particularly if it is desirable to interrupt only when the person is not busy.

While these two studies looked at robots interacting with any people that actually noticed them, other studies focused on robots designed to interact with specific people.

Dautenhahn et al. (2006) examined different ways in which a robot companion

approaches a seated person in a helping context (e.g. Figure 2.7). The person was seated in a simulated living room with a chair, two tables and a television. The remote was hidden from the person, so they were required to ask the robot to fetch it. Once requested, the robot brought the remote, approaching the person from the front, the left or the right position. Participants preferred and were most comfortable with the robot approaching from the left or the right, and least preferred approaches from the front. This showed that a detail as simple as motion path planning relative to the person should be taken into account to maximize comfort. This concept could be applied to interruption scenarios that require a robot to approach a person.

In a more complex scenario of proximity to people, Yamaoka et al. (2008) described a model for a robot to appropriately control its position i.e., proximity to the person and proximity to an object being presented to the person. This was based on the notion that people establish a joint view toward a target object, and that robots should do the same. The robot stands at a position that considers the positions of both the listener and object to optimize the listener's field view, and establish a joint view. Again, we believe these lessons are applicable to interruption scenarios due to their consideration of proximity not only to a person, but also of an object of interest to the person. Indeed,



Figure 2.7: A robot approaches a person in a seating context. (Dautenhahn et al. 2006)

Yamaoka et. al. expected this capability to be important in future applications of robots acting as shopkeepers presenting products to customers, or museum guides presenting information to visitors. Both of these applications will involve some element of interruption if the robot is designed to approach customers before waiting to be asked.

As workers, robots are already used in many different types of working environments. Mutlu and Forlizzi (2008) examined the effects of integrating robot workers into different units of a hospital including support units, such as laundry collection, post-partum units, and medical units such as surgical and oncology units. The robots were designed to perform tasks such as collecting linen from outside patient rooms with assistance from hospital staff and return empty food trays to the kitchen. The differences in how these work environments operated significantly affected the way the staff perceived the robot. The acceptability of the robot's behaviour depended very much on the ward and on the situation.

The support units benefited the most from the robots bringing linen directly to them, but this benefit came at the expense of more work for the medical and post-partum units because the robot lacked the ability to operate completely autonomously. Nurses and housekeepers were required to load linen onto the robot's cart, a new task that was not necessary when laundry collection staff collected the linen themselves. Thus, nurses worried that they would be interrupted from their primary health care duties in order to assist the robot.

Staff that normally had no reason to interact with the robot worried about the robot's disruption in the normal flow through hallways. Many informants had stories of collisions with the robot causing physical pain, due to it being unable to always detect

obstacles. Elders felt disrespected by the robot as it would always take precedence when moving through a hallway. Nurses worried about what would happen in an emergency if someone needed to be rushed to the delivery room and the robot was blocking the path while pausing to determine the best way to go around an obstacle.

In the oncology ward, nurses are required to spend more time with their patients and thus have lower tolerances for interruptions. The robot was designed to approach the nurses and other staff, and continually announce its presence until acknowledged, regardless of circumstances, such as the interruptibility level of the staff member. This behaviour was so disruptive to some staff members that physical abuse toward the robot sometimes occurred.

Thus, Mutlu and Forlizzi found (albeit indirectly) that interruptibility plays a key role on how the robot may disrupt work flow. If these interruptions result in negative effects such as stress or even errors, they run the risk of having the robot's advantages being outweighed by its disadvantages. This work, as with others mentioned earlier, again suggests that we need to explore methods of interruption, e.g. in a hospital where staff can prioritize the robot's interruption and deal with it when it is appropriate to do so without the expense of higher priority tasks.

Interruption is a concern in HRI. The lessons learnt from even this small amount of related work show many opportunities for future exploration, and motivate a need for designing appropriate interruption behaviour.

2.6 Summary

Although robots began as mere tools for helping people to do certain tasks, they are now capable of assisting people more as collaborators, as personal assistants and pets. As

robots continue to evolve as social entities, their ability to interact with people in socially appropriate ways will become more important. One important type of social interaction between people and robots is interruption, which is a phenomenon that can be managed in many different ways that are more or less suitable for different contexts.

However, interrupting people appropriately has not been an explicit design focus of current robots. In this thesis, we will focus on a subset of this, where we consider the minimal non-verbal behavioural cues that robots can exhibit to communicate interruption urgency, and ask if those cues are understandable by people. The related work explored in this chapter highlights many situations where robots interact with people in ways that reveal useful lessons that can be used when designing this behaviour.

In Chapters 4 to 6, we build on the related work of this chapter through the design, evaluation, and analysis of a set of robot interruption studies. However, we begin with somewhat of a side story in Chapter 3, where we present our earlier work. This work is relevant as it inspired our main research question and directed our focus to interruption.

Chapter Three: Exploration of Implicit Emotive-based Human-Robot Interaction

When beginning the research presented in this thesis, our focus was quite general. We were exploring emotive-based human-robot interaction ideas, and did not focus explicitly on interruption. Rather, our initial work investigated the differences between explicit vs. implicit robot control by having the robot react to human ‘stress’ as detected by a crude commercial ‘brain-computer’ input device. As this chapter will reveal, we implemented this interface, and had people use it when trying to control the robot’s speed of motion directly via the brain-computer input device. This direct explicit mapping between the brain interface and the robot locomotion proved difficult to do reliably. Next, we designed the robot to interpret one class of input (suggestive of emotional stress) as an emotional parameter to influence (but not directly control) robot behaviour. In this case, the robot would implicitly react to human stress by staying out of the person’s way. The idea is that the robot would use this interaction modality in order to be less disruptive to a person during a stressful situation.

After completing the system described in this chapter, we realized that we had created a system where the robot tries to avoid ‘interrupting’ a person, and began to think more deeply about interruptibility. We also realized that our original research goal was too broad for an initial exploration, i.e., we would have had to research methods for detecting human emotions, as well as designing robots that could behave appropriately after sensing that emotion. Consequently, we narrowed the primary research goal to the form stated in Chapter 1, where we considered robot behaviour in terms of interruption.

This chapter summarizes this early work on emotive-based human-robot interaction.¹ As will be evident, this meant that a portion of this work did not focus directly on interruption. However, components of this early work provided insight and directly affected our exploration of interruption in human-robot interaction.

3.1 Background

The concept of ‘brain-computer interaction’ involves one or two-way communication between a human brain, and an external computer device. We are particularly interested in using such a system to direct robot behaviour (e.g., McFarland and Wolpaw, 2010). This method of control is not the realm of pure science fiction: Millan et. al. (2004), as one example, demonstrated how two people could successfully move a robot between rooms using an EEG-based interface that recognized three mental states.

In our effort, we focused on using a very limited type of brain-computer interaction, controlling an iRobot® Roomba® vacuum cleaner robot. In particular, we used the OCZ NIA™ neural impulse actuator (OCZ, 2008), an off-the-shelf, low-cost commercial interface designed for video game use that is worn on the user’s head and reads bioelectrical signals. We mapped this input device to the Roomba and constructed two methods of allowing people to control the Roomba. The first was *direct control*, where a person controlled their bioelectrical signals to directly affect robot speed. The second was *behavioural control*, where a person’s emotional state was inferred based on their bio-electrical signal state, and the robot adjusted its behaviour to fit that person’s

¹ Portions of this chapter were published previously in a modified form in Saulnier et al., 2009

emotional state. Thus, the robot was able to modify its own behaviour to avoid interrupting the person.

3.2 The System

The OCZ NIA™ consists of a headband worn by the person on their forehead (Figure 3.1). The manufacturer claims that the device reads “bioelectric signals that are amplified, digitized and further de-convoluted into computer commands”, where these bioelectric signals are collectively generated by facial muscles, eyes, and the brain (OCZ, 2008). It is somewhat unclear what signals are actually being read, but one reviewer (TechRadar, 2008) suggests that its sensors read skin biopotentials, i.e., small electrical changes on the surface of the skin. Software calibration must be done before every use, and the software displays live measurements for muscle tension, eye glancing as well as alpha and beta waves (Figure 3.2). The device’s software convolutes its sensor reading into various applications, such as keystrokes that can be mapped to computer games brain interaction.



Figure 3.1: Our NIA-Roomba brain-robot interface prototype

We were able to customize this convolution to our own uses, i.e., to control the iRobot® Roomba® through its API. However (and in spite of manufacturer claims inherent in the software display, Figure 3.2), we found the OCZ NIA™ capabilities as a ‘brain interface’ to be quite limiting, where we could only capture muscle tension data reliably (Figure 3.2, yellow bar) and generally failed to interpret the “brain signals” in a meaningful way. In spite of these shortcomings, we were able to use muscle tension input (a) as commands to directly control robot speed, and (b) for roughly inferring human emotions, such as stress, to mediate a robot’s behaviour to avoid interruptions at inappropriate and undesirable times.

3.2.1 Direct Explicit Control

We initially mapped the OCZ NIA™’s muscle tension readings to robot speed as an exploration exercise into the viability of the brain-computer interface for controlling a

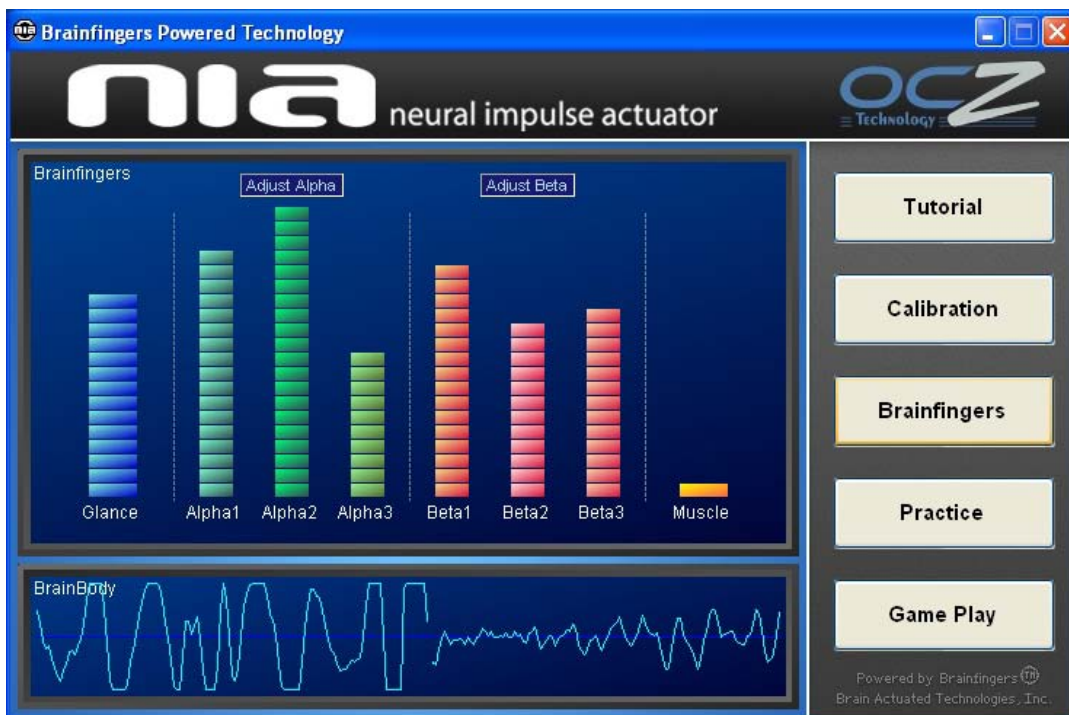


Figure 3.2: The OCZ NIA output channels

robot. With minimal training and practice, we were able to increase and relax muscle tension as desired to control the robot's speed in real time using the control software we developed. We did this by converting the continuous bioelectric signal input into four discrete values, where each value corresponded to a specific robot speed in the forward direction. Thus, a person would consciously change their muscle tension to change robot speed, using methods such as jaw clenching, or eyebrow tensing.

While this mapping worked, its utility was found to be limited. We found that using muscle tension alone is not sufficient. Not enough reliable input parameters were available to effectively control the Roomba vacuum cleaning robot. Speed is just one aspect; other critical control factors include direction, rotation, operation (e.g., clean / don't clean), and so on.

Even though we have found the capabilities of the OCZ NIA to be limited (and our system was admittedly a crude but working prototype), we did find that muscle tension as measured by this device was usable as input. We prototype our application of muscle tension in the next section.

3.2.2 Behavioural Implicit Control

In our next exploration, we wanted to see if we could infer a rough estimation of the person's emotional state from muscle tension reading, and then use this rough estimation to influence (rather than directly control) robot behaviour. We attempted to demonstrate that this subtle difference may enable a robot to autonomously adapt its behaviour in ways that (ideally) are more socially appropriate and less disruptive to people, based on their emotional state.

Our control software interprets the muscle tension reading as crudely suggestive of one's stress level: the more muscle tension, the more stress is inferred. We polled the muscle tension readings ten times per second, and then averaged all values over five seconds (i.e., calculate average of the last 50 readings). This average reading was displayed graphically on a GUI (Figure 3.3) and used as input to adjust robot behaviour.

Two distinct robotic behaviours corresponding to two extreme emotional states, either relaxed or stressed, are triggered when the stress reading crossed a threshold. Robot actions are then influenced by these stress readings. When a person shows high stress (~levels 3 & 4), the robot assumes the user is stressed, enters its cleaning mode but moves away from the user so as not annoy them. When a person is relaxed (~level 1), the

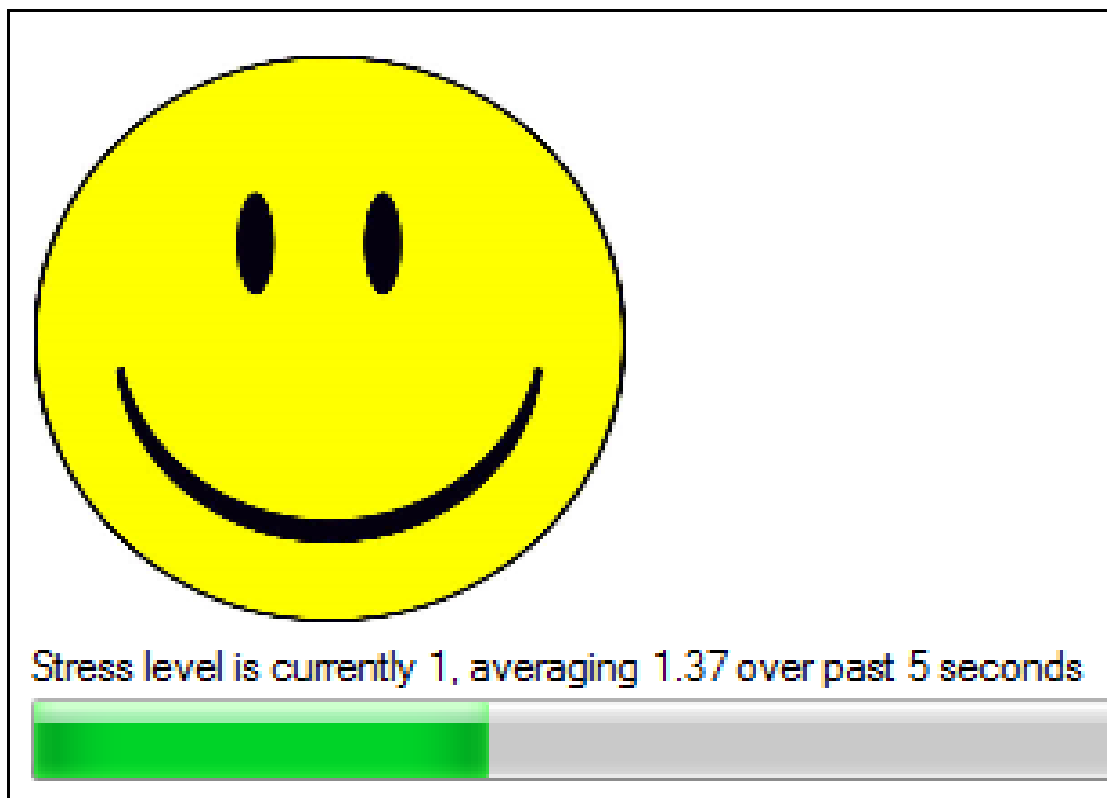


Figure 3.3: Our custom GUI which graphically displays the person's stress level.

robot (if cleaning) approaches the person and then stops, simulating a pet sitting next to its owner.

Our brain-robot implicit interfaces operated, albeit not perfectly. It demonstrated the use of bio-electrical signals as crude, implicit input with the assumption that it can be suggestive of emotional state. Indeed, this type of mapping may not be valid. While far from an ideal vision for a brain-robot interaction, it did roughly simulate the kinds of input we may be able to capture in the future if more powerful devices emerge with the ability to read brain activity reliably and non-invasively. This deeper sensing of a person's emotional state would enable a robot to coordinate interruptions and other behaviours with more accuracy, perhaps timed with consideration to the person's level of engagement in their current task.

3.3 Discussion

In this discussion, we explore several elements of this project. First, we discuss the limitations of the technology we used, and how these impacted its usefulness. Second, we discuss methods of using emotion-sensing to adapt the robot's behaviour effectively while attempting to avoid undesirable side-effects from emotion sensing errors. Third, we discuss how the general public perceived our concept of emotion-sensing robots through the intense media coverage we received. Finally, we discuss potential applications of emotion-sensing robots to be explored once available technology makes their development feasible.

3.3.1 Limitations

To start our discussion, we will describe the limitations that we encountered which we felt posed issues for any designer trying to incorporate human emotion as an input to influence robot's behaviour.

First and foremost, the ability of the current technology we used to accurately read or infer a person's brain activity emotional state is limited. The device we used did not provide any input reliably, except for muscle tension, so inferring emotion from this data was very questionable. While some specialized devices in use in medical fields can read brain activity more reliably, many are invasive, requiring physical implants into a person's brain. This approach, regardless of reliability and robustness, is not practical for many human-robot applications.

Second, we feel that the current requirement to wear a physical input device, invasive or non-invasive, can be cumbersome. The device we used requires the person to wear a headband securely on their head, which is very awkward socially and uncomfortable physically.

3.3.2 Adapting Rules of Behaviour

Our behavioural prototype demonstrated implicit control of a robot, where the robot reacts to muscle tension input rather than direct control. The person does not have to do anything (except wear the input device). The robot's behaviour – not low level actions – is altered based on its perception of the muscle tension, which we believe can be viewed to some extent as an indicator for emotional stress. This technique, if fully implemented, opens up a wide variety of applications, e.g., household worker robots that do not disturb their owners, or companion robots that provide comfort when a person is sad. However,

these applications must also be able to gracefully handle any errors made in the judgement of a person's emotional state. In other words, the worst case scenario of incorrect emotion sensing must result in a minimum of undesirable robot behaviour, and other forms of input should be used to supplement the emotional input when possible to mitigate the risks of incorrect emotional data readings. To reduce the negative impact of incorrect emotion sensing, appropriate robot behaviour can be provided as simple and perhaps anthropomorphic characteristics. These traits can be designed to be robust and less sensitive to input errors. That is, in this example none of the actions of the robot are 'bad ones'. At its worst, if the robot incorrectly assumes a relaxed state when the person is in fact stressed, the consequences are small. Our example illustrated how the Roomba would work away from stressed people so as not to annoy them. Even if the Roomba were wrong and adapted the wrong behavioural traits, this could be designed so it will not be considered bad behaviour.

Another way to reduce the impact of errors in emotion sensing, at least in implementations like ours, is to augment this data with information about the person's emotional state from other sources. For example, facial gesture recognition, heart beat monitoring, and blood pressure monitoring could all be used in addition to a bioelectrical or brain-computer interface. Of course, the addition of more monitors requires more sensory equipment, along with better extraction and recognition algorithms. These additions make the whole solution more physically cumbersome for the person and more demanding computationally for the robot.

3.3.3 Public Perception

Following a public presentation and publication of our work, many media outlets around the world followed up with coverage of the project with TV and radio interviews, newspaper articles, and over a dozen blog posts. The aspects of our research highlighted by them, and their subsequent re-interpretations uncovered several interesting details regarding how people feel about robots that read and react to their emotions.

The general reaction was overwhelming positive. This response was most visible on high-profile gadget blogs Engadget (Stevens, 2009) and Gizmodo (Fallon, 2009), among others. Their headlines made light of the robot reading the human's emotion scenario by saying "University of Calgary researchers teach little robots to be scared of angry humans" (Engadget) and "Modified Roomba Detects Stress, Runs Away When It Thinks You Might Abuse It" (Gizmodo). Dozens of comments by readers for these showed positive interest in the idea. None seemed fearful. Many people specifically imagined servant robots that would avoid people when they were angry, thus to prevent disturbing them.

Based on this media coverage, there's no way to know if people are truly interested in interacting with emotion-sensing robots. However, the positive response is encouraging, as technology may improve enough in the future to allow applications of emotion-sensing robots to be realized.

3.3.4 Possible Applications

Our initial prototype implementation is potentially applicable in the improvement of house servants (e.g., a robot vacuum cleaner). Above, we discussed an implementation that senses stress, but other emotions could be used as well. Happiness could be used as a

trigger to notify the person of anything requiring physical action from the person, such as the emptying of a full vacuum cleaner dust bin, or cleaning of the brushes. Anger could be used as a trigger to continue cleaning, but to do so quietly far away from the person. In these cases, the robot uses the person's emotional state to determine how it can behave in a socially appropriate way, i.e., not to annoy the person when they are busy, and ensure that maintenance tasks are performed when the person is likely to find them the least disruptive.

3.4 Summary

In this chapter, we have summarized our early work on emotive-based human-robot interaction. As far as we know, we were amongst the very first to pursue bioelectric signal interfaces in implicit human-robot interaction, where we programmed a robot to react to human 'stress' as detected by a very crude input device

After completing the system described in this chapter, we realized that we had created a system where the robot tries to avoid 'interrupting' a person. We also thought that our scope was too broad and may not allow us to reach meaningful results. Thus this early work led us to further focus our research toward interruption, which led to the primary research goal of Chapter 1. We have recast this work – albeit after the fact – as the first part of a two problem, consisting of how does a robot infer what a person is doing, and based on that modify its behaviour to avoid interrupting a person. In the subsequent chapters of this thesis, we continue with the second part of this interruption problem: how should a robot exhibit interruption behaviour when it needs to act in a socially appropriate way?

Chapter Four: Design Approach and User Study

Our primary research question is: are there *minimal nonverbal behavioural cues* that robots can exhibit to communicate interruption urgency, and are those cues understandable by people? To answer this question, the remainder of this thesis will focus on our efforts of identifying the behavioural cues that a robot can use to indicate interruption to a person, and on understanding how a person will interpret these cues in terms of interruption urgency and importance. These insights will likely prove valuable when designing interruption behaviour that is minimally disruptive while still adequate to convey urgency and importance appropriately within a given context.

To address our research question, we followed a four-step methodological design and evaluation process.

1. We observed how people use a subset of nonverbal cues to interrupt other people over scenarios that vary in urgency.
2. We used these observed behavioural cues to design and critique robotic behaviours.
3. We implemented these behavioural cues on a mobile robot, where the behaviours are triggered and somewhat controlled by a human operator using a Wizard of Oz methodology (Dahlbäck et al., 1993).
4. We conducted a lab study (following two pilot studies), where we exposed people to these robotic behaviours, and gathered their reaction and interpretation of those behaviours.

This chapter discusses the set of four studies that constituted our design process. We first discuss an observational study, where we asked a small number of people to try to

interrupt others using only a minimal set of physical behaviour cues. The purpose of this study was to see how human actors would use these cues and thus get a general sense of how they can be applied using robots in an interruption context. Next, we discuss our three robot interruption user studies, comprising two pilot experiments and a main study, where we exposed people to robotic behaviours designed in part from the results of the observational study, and evaluated the participants' interpretation of those behaviours.

4.1 Observational Study: Identifying Robot Interruption Behaviour

Our overall goal considers *minimal nonverbal behavioural cues* that a robot could use to interrupt people. By this we mean that we are interested in determining interruption behaviours that rely on only a few simple physical capabilities present (or that can be easily added to) most robots. We deliberately exclude verbal cues for now.

4.1.1 Methods

As a starting point, we hypothesized three general physical robot abilities as being likely



Figure 4.1: A robot actor looks inside the office from a distance to interrupt unobtrusively (scenario 1)

candidates for nonverbal interruption cues: physical position (e.g., location on floor in relation to people or objects), speed and type of movement (e.g., standing, walking, running, rotating body, etc.), and gaze. We chose these abilities as we felt they could be easily mimicked by most mobile robots with minimal effort.

We recruited three people as “robot actors” and asked them to act through five situational interruption scenarios. For each scenario, the actor had to interrupt two people who were engaged in a meeting inside an office with an open door (e.g., Figure 4.1). One of these two people inside the office was identified as the main “interruption subject”. The actor was asked to improvise interruption behaviour appropriate to varying urgency levels, with a focus on the one interruption subject, using only nonverbal interruption cues. Furthermore, we instructed the robot actors to leave if no acknowledgement of their actions was provided after 10 to 15 seconds. An element kept secret from them was that the people they were interrupting were informed of the coming interruption, and instructed to ignore the robot actors for at least 10 seconds, to allow the experimenters to have enough time to observe the robot actors’ behaviour, and to investigate the way they employed nonverbal interruption cues

The process followed for each of the five situational scenarios consisted of two steps: (a) verbally brief the robot actor on the details of the scenario, and (b) observe the robot actor as she or he interrupts two people inside the office using behaviour improvised on the spot. We did not provide any feedback to the robot actors’ regarding their behaviour at any point during the study.

The robot actors were instructed to use only the three physical robot abilities that we hypothesized earlier. Specifically, we told the robot actors at the beginning of the study to use only the following factors into their behaviour:

- Physical position. (We suggested that actors limit standing to designated spots, e.g., outside the doorway, at the doorway, inside the office, or wherever is appropriate.)
- Speed and type of movement (i.e. standing, not moving, walking, running, pacing, rotating, etc.)
- Eye glancing (emphasis on gaze)

We also explicitly defined the types of behaviour that we did *not* want the robot actors to use, as we felt they could be beyond the capabilities of a robot, and beyond our design goals. Our guidelines were:

- Voice, sound, physical props, hands, arms, and legs cannot be used (i.e. no knocking, waving, kicking, etc.)
- Don't do anything that isn't naturally human
- Communicate only with person to be "interrupted" (i.e., "interruption subject"). Any other guests are only secondary.

The five situational scenarios provided to the robot actors in the study ranged a spectrum from typical time-insensitive non-urgent matters, to important time-sensitive matters, and to extreme emergency situations that required immediate attention and

action. We thought that the use of both typical and extreme scenarios would allow a wide range of matching typical and extreme behaviour to emerge. The scenarios were:

1. Earlier in the day, the interrupted person was wondering (for personal interest) when Calgary was established as a town. You have discovered that this date is 1884, so you want to tell them about it right now if possible.
2. You are supposed to be meeting with the interruption subject right now about an important issue, and you cannot proceed until after you meet. However, you have enough time and patience to wait until they are finished with their current meeting.
3. Same as scenario 2, except you have no spare time to wait and the meeting must begin as soon as possible.
4. Something important has come up, and the University President needs to see the interruption subject immediately. If the President waits too long, he or she will get annoyed and terminate the interrupted person.
5. The building is on fire.

4.1.2 Results

We videotaped the behaviour of the three robot actors acting through each of the five situational scenarios using a camera located inside the office. This produced 15 scenes that we then reviewed to identify characteristic behavioural trends. We saw that our robot actors improvised with a range of nonverbal physical behaviours adapted to the given scenarios.

Table 4.1 describes our raw observations of the behaviours used by the three robot actors' for interruption in the different scenarios. All of the robot actors used some

behaviour that we told them to avoid at least once, though the majority of the behaviours remained within the guidelines that we described earlier.

There was clearly some variance across our three robot actors with regards to the behaviour they used for the situational scenarios. However, we also observed many cues of behaviour that were common to two or three of the three robot actors. Table 4.2 summarizes the common cues, and exclude cues that were different for all three robot actors. Using this table, we can see a consensus of interruption behaviour used for each of the five scenarios.

For less urgent, less important scenarios (e.g., scenarios 1 and 2), the robot actors used slow, non-disruptive behaviour. In these cases, the robot actor ‘peaked into’ the office from a distance outside the doorway (e.g., Figure 4.1), perhaps to see if the people inside the office were busy, and if it would be possible to interrupt without disrupting a more important task. After receiving no immediate acknowledgement, the robot actor left without waiting. While it is possible for the people seated in the office to notice the actor, it is also possible and appropriate to ignore him or her if desired.

As urgency increased (e.g., scenarios 3 & 4), the robot actors generally used more disruptive behaviour. They maintained close interpersonal proximity with the people inside the office and only left when their interruption was acknowledged and addressed.

For the most urgent scenario (i.e., scenario 5, the fire), the robot actors used their most disruptive behaviour. One robot actor entered the office running, circled around the people inside a few times before kicking the person’s chair and legs until acknowledged, as shown in Figure 4.2 (despite being told that kicking was not an allowed behaviour).

Table 4.1: Raw observations of robot actors' behaviour during the observational study.

Scenario	Actor	Observations of Robot Actor's Behaviour
1	1	Slow approach, away from doorway (~50cm), looking toward interruption subject, leaves quickly without acknowledgement, office people later mention interruption
	2	Slow approach, away from doorway (~80cm), looking toward interruption subject, slight leaning during glancing, leaves quickly without acknowledgement
	3	Slow approach, inside office, next to interruption subject (~40cm), looking at interruption subject until acknowledged
2	1	Slow approach, at the doorway, looks toward interruption subject, begins to leave, pauses for a second, then leaves out of view, office people describe interruption as a "breeze" passing through
	2	Slow approach, at the doorway, looks at both people in office, begins to leave, pauses, small shrug, then leaves without acknowledgement
	3	Slow approach, at the doorway, looking at laptop which office people are using, waits a few seconds before leaving without acknowledgment
3	1	Slow approach, slightly inside doorway (~10cm), tapping hand on body and looking directly at both people inside the office, does not leave until acknowledged
	2	Noticeably faster approach, inside doorway (~50cm), very close to interruption subject, few taps of foot on floor, slight pacing, no direct eye contact with either person, does not leave until acknowledged
	3	Noticeably faster approach, inside office, next to interruption subject, kicking chair and then looking at interruption subject until acknowledged after a few seconds
4	1	Noticeably faster approach, inside doorway (~40cm), very close to interruption subject (~40cm), tapping both hands on body, looking at interruption subject, does not leave until acknowledged
	2	Noticeably faster approach, inside doorway (~40cm), very close to interruption subject (~40cm), tapping feet on ground, kicking ground, looking at interruption subject, does not leave until acknowledged
	3	Slow approach (similar to scenario 1&2 for this actor), inside office between two office people, looking at interruption subject, does not move until acknowledged
5	1	Much faster approach, inside doorway (~30cm), next to interruption subject, one small jump, gesturing head toward doorway, no direct eye contact with interruption subject, does not leave until acknowledged
	2	Much faster approach, running into office and around interruption subject, jumping and rotating of body, leaves office once acknowledged
	3	Noticeably faster approach (same as scenario #3 for this actor), inside office, between office people, moving head back and forth, gesturing head toward doorway, no direct eye contact with interruption subject, leaves when acknowledged

In these cases, the behaviour was certainly more noticeable to the people inside the office, who were unable to continue their conversation because of the interruption.

Despite some variance across a small number of robot actors and the use of a few behaviours outside the guidelines we gave, we considered these results to be helpful in developing a robotic vocabulary of interruption cues. .

4.1.3 Discussion

If we examine the common elements of behaviour used by the robot actors in Table 4.2, we can define a list of four nonverbal physical behavioural cues that are both easily applied to mobile robots and representative of the behaviour used by the robot actors. We describe the four nonverbal physical behavioural cues below.

Proximity to Person. The robot actors generally used positions of proximity that corresponded to three thresholds: outside the doorway, at the doorway, and next to the person to be interrupted inside the office. A robot can move around in a physical space

Table 4.2: Common observations (to 2 or 3 out of 3 robot actors) behaviour during the observational study.

Scenario	Observations of Robot Actor's Behaviour
1	Slow approach, away from doorway, looking toward interruption subject, leaves quickly without acknowledgement
2	Slow approach, at the doorway, looks toward interruption subject, begins to leave, pauses for a second, then leaves out of view
3	Fast approach, inside the office, foot tapping or kicking chair, different eye contact for all 3 actors, does not leave until acknowledged
4	Fast approach, inside office, very close to interruption subject, tapping hands or feet, looking at interruption subject, does not leave until acknowledged
5	Fast approach, inside office, next to interruption subject, gesturing head toward doorway, no direct eye contact with interruption subject, variable body movement (e.g., jumping, rotating), does not leave until acknowledged

and, as a consequence, they can move close or away from a person. We hypothesise that proximity contributes to the interruption cue, i.e., the closer the robot is to the person, the more interruptive the robot will appear to be.

Gaze and Head Movement. The robot actors generally directed their gaze toward the person being interrupted during the least urgent scenarios, although this gaze became less directed as the urgency of the scenarios increased. For the more urgent scenarios, the robot actors used their head to gesture towards the doorway and did not hold their gaze to any person or object in the office. Thus, we hypothesise that both gaze and movement can be collapsed into a single behavioural cue of head movement for robotic use within the context of an interruption. Direct gaze can be used in low urgency situations to address a single person, while movement of the head without direct gaze could be used in high urgency situations. To do this, a robot can be equipped with a simple movable ‘head’ that has a clear directional ‘gaze’, as indicated by its position on the body of the



Figure 4.2: An actor kicks the person’s chair while attempting to interrupt him.

robot. Alternatively, a robot can be equipped with ‘eyes’ that indicate gaze direction. This can be as simple as two eyes painted on a robot’s head, or eyes that can move independent of the head.

Body Language. Despite being told not to use physical abilities such as kicking and rotation, the robot actors sometimes used them anyway for the most urgent scenarios. While it is difficult for robots to kick things, simple body rotation is easy. A robot can exhibit simple body language by the way it rotates or can use persistent body movement to show impatience it is waiting too long for a person to acknowledge its actions.

Speed of Motion. As urgency increased, the robot actors increased the speed of their approach to the office and overall movement. A robot can also vary its speed of motion. We hypothesize, based on the behaviour of the robot actors that increased motion contributes to an increased sense of urgency. For example, a fast speed can imply a more urgent or important situation, while a slow speed implies more calm, less urgent circumstances.

Clearly, a robot can be designed with many different nonverbal behavioural cues that can be used to modify behaviour in ways that facilitate communication and interruption. However, it may not always be useful for a robot to use all of these cues, as some may be redundant or less effective than others for the same purpose. Thus, there is value in identifying which cues are usable in the context of interruption, and which ones do not facilitate any communication. Our observational study approach is far from exhaustive; rather, it serves to preliminarily identify some nonverbal behavioural cues used by humans in an interruption context that can also be used by a robot.

In the next section, and for the rest of this chapter, we explore our robot interruption user studies which build upon our human-robot observational study, and attempt to preliminarily verify the comprehensibility of robot behaviour in an interruption context.

4.2 1st Pilot (Robot Interruption User Study)

The *Robot Interruption User Study* was designed to test the degree to which particular *minimal nonverbal behavioural cues* used by robots to communicate interruption urgency are understandable by people. The study comprised two pilot experiments and a main study. All revolved around having the participant seated and engaged in some task which is unexpectedly interrupted by the robot. The robot's interruption behaviours were limited to the particular combinations of minimal nonverbal behavioural cues (proximity to person, gaze and head movement, body language, and speed of motion) that emerged from the observational study. Following the attempted interruption, the participant was asked to describe their interpretation and understanding of the robot behaviour.

4.2.1 Selection of Physical Form

We begin with our choice of the robot's physical form for the 1st pilot of the robot interruption user study. When deciding which form was appropriate to study in the context of interruption, our main considerations initially were to try to limit ourselves to fairly basic capabilities common to many robots, arguably to almost any robot, and to identify how those capabilities could implement at least a rudimentary aspect of the behavioural cues identified in the observational study.

4.2.1.1 Basic Design

To implement the previously listed behavioural cues, our robot had to meet the following requirements.

First, our interruption robot should be capable of using *different speeds of motion*. We considered two different speeds: a “slow” speed that is quite a bit slower than human walking speed, and a “fast” speed, that is equivalent to a brisk pace.

Second, the robot should have some form of *movable head or eyes*. Both could be used to address a person directly. In cases where a robot approaches a group of people, gaze would be used by the robot to indicate that it was addressing a particular person within the group.

Given the above, we also believe that the size and height of a robot should be large enough to make the interruption cues salient. We considered a robot that is the height of a small child as reasonable, as it is large enough to be easily visible but not so large as to be threatening.

4.2.1.2 Feasibility of Human and Animal Forms

Our assumption is that people will – at least in part – consider a robot as a social entity. Thus if the robot displays a particular interruption behaviour, they would interpret it somewhat similarly to human interruption. As people are generally most familiar with receiving interruptions from other people or their pets, it may appear that the robot form factor should be either a human form (through use of an android or humanoid robot, via anthropomorphism), or animal form (via zoomorphism). However, in practice, the use of these forms may not be ideal, may be extremely difficult to implement and – most importantly – may not even be necessary.

It is very difficult to model a robot following a convincing human or animal form, and to maintain an acceptable level of realism. People have very complex expectations of how people and animals behave, and current robots fail to achieve a level of realism that makes them indistinguishable from actual people or animals, although some do come close (Nishio, Ishiguro, and Hagita, 2007). Many attempts fall into the ‘*uncanny valley*’, which theorizes that humans will become emotionally empathic of a robot up until the point at which the realism of the robot’s appearance and motion pass a certain threshold that causes the human response to shift from empathy to revulsion (Mori, 1970). We have made an informed decision not to explore robotic form, and are not looking to make a contribution to the field of animatronics. In Appendix B, however, we explore our limited implementation of a teddy bear robot form as a side project.

Our main focus in this pilot was a first look at robotic interruption through nonverbal cues. Thus, we are interested in exploring robot forms that do not attempt to represent any human or animal form with any level of realism. We decided to use a minimalist robot that does not resemble animals or humans at all, but is still capable of interrupting people. This approach greatly minimizes the risk of uncanny valley effects. Indeed, there are many forms that go beyond a minimalist form without falling into the uncanny valley. As we are interested in finding minimal nonverbal behavioural cues that are usable by robots of all types, including all physical forms, we thus feel that a minimalist form is most suitable. We describe our particular robot form in more detail in the next section.

4.2.2 Implementation

4.2.2.1 Robot Platform

The robot platform (Figure 4.3) we used for our pilot study consists of a Mobile Robots Inc. Pioneer 3-DX base (the red wheeled portion of Figure 4.3) with a custom body and head we have added on top. The base is motorized, and is capable of moving faster than human walking speed. It can also carry heavy loads, although we just used it to carry a wireless laptop (used to control the robot via a serial connection) hidden in the body. The custom ‘body’ consists of a plastic container used to increase the height of the robot, covered by a t-shirt to reduce the robot’s mechanical appearance without going as far as directly anthropomorphizing it. The robot’s ‘head’, used to portray head movement and gaze, is just a small cardboard box affixed to a computer-controllable servo motor, which lets us rotate the head left/right and up/down. The head does not include any facial markings such as eyes. However, it does have a clear directional forward position based on the way it is positioned on the body. We used this minimalist design because we wanted to rely on only a generic shape and a few simple physical behavioural capabilities that are present in (or that can be easily added to) most robots.



Figure 4.3: The robot

4.2.2.2 Robot Control

The robot is remotely controlled using a Wizard of Oz methodology (Dahlbäck et al., 1993) by a study administrator sitting outside the room and out of view of the evaluator. The study administrator was equipped with a standard gamepad controller to issue commands using joysticks to control the robot's direction and speed, and buttons corresponding to short pre-programmed sequences of robot behaviour (e.g., head shaking). The controller was connected to a computer which is connected wirelessly to the on-board laptop running custom C#/C++ software, which itself is connected to the robot base via its direct serial connection. All sensory monitoring (including position awareness and distance travelled) also occurs on the laptop via this serial connection.

4.2.3 Methods

We developed a first suite of robot interruption behaviours based on the behavioural features mentioned in Section 4.2.1.1. To test whether these were reasonable, we evaluated these behaviours in a design critique session² as we were aware that we likely wouldn't get this 'right' the first time. We conducted the design critique using the scenario walkthroughs from the observational study that included two participants recruited from our laboratory. Our participants were asked to help us test, discuss and critique the general suitability of the chosen interruption behaviours when applied to robots, as well as the robustness and nuances of the technical implementation of the robot and its controller software. While somewhat informal and very limited in scope, these sessions helped us discover 'big effects', i.e., where our behaviours and/or

² Portions of this chapter were published previously in a modified form in Saulnier et al., 2010

implementation had serious problems. In the following subsections, we will describe the design process that we used for this pilot study and then explore its results.

4.2.3.1 Designing Robotic Interruption Behaviour

We used our observations of interruption behaviour expressed by the robot actors in the observational study to design and program robot behaviours. Proximity and speed of motion were easily reproduced as rudimentary robotic behaviour within the capabilities of our robot base. Gaze/head movement were collapsed into head movement, as implemented by the servo motor and controlled via remote controls. Subtle body language, such as slight leaning in one direction, could not be implemented due to robot limitations. However, we introduced more blatant body language behaviour (e.g., fidgeting) as body rotation.

These behaviours were combined as part of the interruption flow. For example, for situations with low urgency and importance, the robot was programmed to move slowly, with fluid head movement and no body language. Similar to a person walking by an office, the robot would move to a position where it could observe the person it wished to interrupt (and thus presumably be seen by that person) but would not approach the person. For situations with high urgency and importance, the robot was programmed to move very quickly, with erratic head and body movement (i.e., rotation), all designed to be as disruptive as possible to the user.

4.2.3.2 Experimental Procedures

In this pilot study, each participant was seated in a room alone with their laptop computer and instructed to work on a task of their choice. All attempts were made to minimize distractions in the room so that the participant could focus on their task. The participant was seated so that the doorway to the room was visible to their right. This doorway was then used by the robot to approach and optionally enter the room to facilitate an interruption that could attract the attention of the participant. Participants were aware that they would be interrupted by the robot at undisclosed intervals, but they were not aware that the robot was being remotely controlled.

Figure 4.4 shows one low-urgency interruption scenario, where the robot simply passes by outside the door without entering the room. Following each interaction with the robot, the study administrator entered the room to discuss and critique the interaction with the participant before returning to control the robot again.



Figure 4.4: The robot interrupts a participant in a design critique session.

4.2.4 Results and Discussion

Generally, participants were able to discern the meaning and level of urgency and importance from at least some of the robot's behaviour, which was encouraging to us.

“[Urgency was] not so huge. Because it kind of stopped as I was looking at it... If it wanted to do more, it would've kept rolling toward me. But it stopped.”

(1st Pilot participant 2, low urgency behaviour)

“I would think that there is some kind of emergency... like someone is having a heart attack or something is on fire.”

(1st Pilot participant 1, high urgency behaviour)

These initial design critique sessions proved valuable at identifying not only successes, but fundamental problems in the design of our robot behaviours. They also helped us identify technical issues that caused the robot to malfunction or produce undesirable behaviours. For example, the study administrator had to manually control the robot at all times, which proved difficult to do and thus compromised the reliability and repeatability of the robot's behaviour.

This pilot also revealed problems that we would face in a full study, in particular, collecting meaningful quantitative data. Specifically, there was no way in our pilot to determine, aside from the interview, how participants interpreted specific elements of the robot's behaviour without relating them to the overall experience. For example, how did the speed of motion, type of head movement, or physical proximity position influence the sense of interruption? Nor was there any way to quantitatively identify the degree to which these behavioural cues were effective at conveying information such as urgency and importance in a particular interruption scenario. Another problem was the ability for the participant to choose their own task during the study, which could affect the robot's

ability to interrupt. For example, it may be more difficult to interrupt a person who is deeply engaged in a writing task than a person browsing the web.

As we will present now, this led to the design of a more sophisticated experiment.

4.3 2nd Pilot (Robot Interruption User Study)

To address the deficiencies uncovered in the first pilot user study, we conducted a second pilot that incorporated a series of major changes to the experimental design. First, we will describe the most major change which involved the design of the robot's behaviour. Second, we will discuss changes to our implementation which were mostly software based. Finally, we will discuss our revised data collection methods and experimental procedures.

4.3.1 Redesigning Robot Behaviours

We replaced the five initial situational scenarios from the 1st Pilot with ten robotic interruption behavioural episodes using different *combinations of four behavioural cues* (proximity to person, gaze/head movement, body language, speed of motion) as summarized in Table 4.3. These episodes were designed using combinations of cues such

Table 4.3: Definition of Behavioural Episodes by Cue used in the 2nd Pilot. (proximity to person, gaze/head movement, body language, speed of motion)

		<i>slow speed of motion</i>		<i>fast speed of motion</i>	
<i>proximity to person</i>	<i>body language</i>	<i>direct gaze</i>	<i>erratic gaze</i>	<i>direct gaze</i>	<i>erratic gaze</i>
<i>far from doorway</i>	<i>none</i>	1A		1B	
	<i>rotation</i>				
<i>at the doorway</i>	<i>none</i>	1C		1D	
	<i>rotation</i>				
<i>next to participant</i>	<i>none</i>	1E		1F	
	<i>rotation</i>	2A	2C	2B	2D

that any one episode had another matching episode that differed by exactly one behavioural cue. The purpose of this change was to tease out effects of particular behavioural cues, i.e., to see if there was a statistical difference between a behaviour that included or did not include a particular cue. Thus, the behavioural episodes, labelled 1A-F, and 2A-D, no longer corresponded directly with any situational scenario (although some are suggestive of them). For example, episodes 1A and 1B in Table 4.3 both had the robot moving to just outside the doorway and gazing at the participant, but they differ in the speed of motion used. This approach thus enables any differences in the participant's interpretation between two episodes to be feasibly attributed to the single behavioural cue that differs between the two episodes. We also included an additional episode, which we called the "null base case", where a robot would do an action that (we believed) had nothing to do with interruptions: slow movement outside the office without any direct interaction with the participant. The purpose of this case is to serve as a baseline to compare all other episodes to, regardless of their design. We considered this to be a valid baseline as it provides the participant's interpretation of the *minimum* level of interruption behaviour; we would generally expect all other episodes that use more interaction with the participant to be interpreted as more interruptive. The use of this baseline enables more accurate statistical analysis to occur, which we will discuss in further detail in Chapter 5.

We do not rate these episodes by urgency – this is what participants would do through their interpretation of robot behaviour. However, we hypothesise that episodes with fast speed, close proximity, direct gaze, and body language using rotation (Table

4.3, bottom right) would be higher in interruption magnitude than those with opposite values (Table 4.3, top left).

While the five situational scenarios used in our first pilot were presented in an increasing order of magnitude, the ten cue-based episodes were presented in a scrambled order to the participant, following the null base case episode at the beginning.

4.3.2 Description of Specific Robot Behaviour

Here, we describe the specific behaviour used by the robot for each of the behavioural episodes introduced in the previous section and summarized in Table 4.3. Each of the eleven episodes fits into one of three *base cases*. The base cases are differentiated by which behaviour cues are varied over different episodes, and which ones are constant.

Below, we describe each of the base cases in detail.

4.3.2.1 Null Base Case

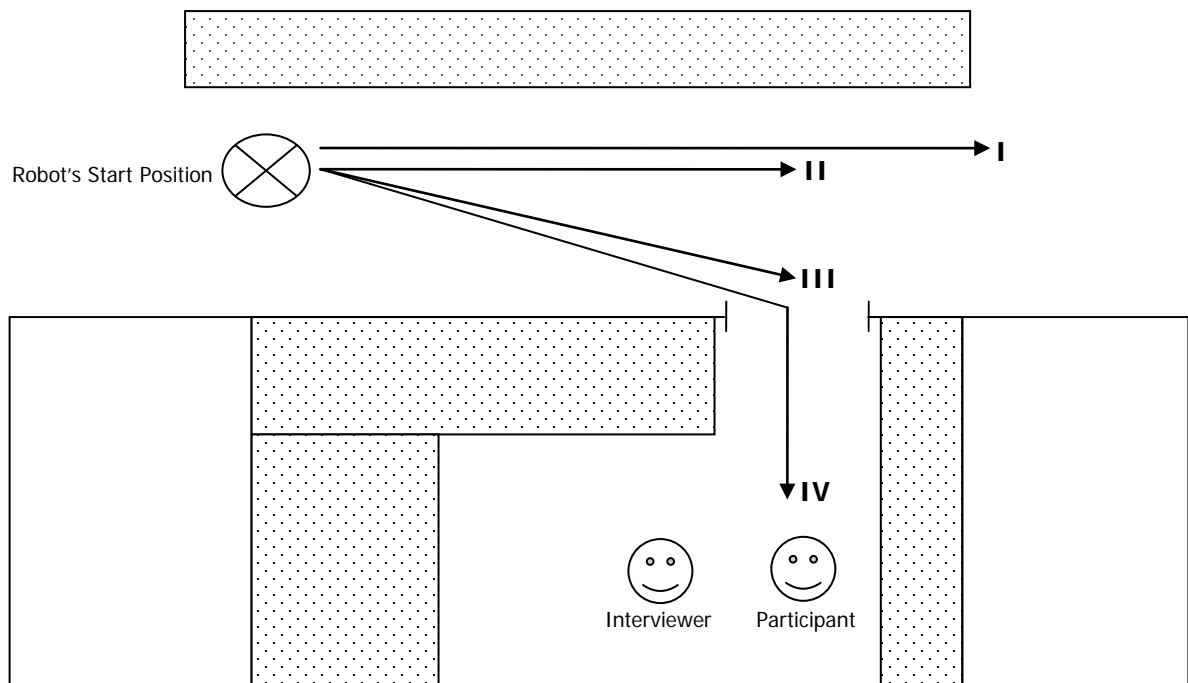


Figure 4.5: The robot's four distinct motion paths (labelled I, II, III, IV).

The null base case, comprising a single behavioural episode, is presented to the participant before all of the other episodes in both phases of the user study. This episode consists of the robot rolling past the doorway (Figure 4.5, motion path I) at slow speed until it is out of view of the participant, where it waits 15 seconds, and then returns using the same path. The robot does not stop, move its head, or otherwise interact with the participant.

4.3.2.2 Base Case 1 (Episodes 1A to 1F)

Base case 1 comprises six variations with varied *speed of motion* and *proximity to participant* positions. Motion paths II (far from doorway), III (at the doorway), and IV (next to the participant) in Figure 4.5 correspond to the different proximity to participant positions used by the robot. For each base case 1 episode, the robot approaches the defined proximity position using one of the three motion paths (II, III, IV) at either slow or fast speed, looks at the participant directly with no head movement, waits 15 seconds, and then returns to its starting position. No rotational body language is used.

4.3.2.3 Base Case 2 (Episodes 2A to 2D)

Base case 2 comprises four episodes with varied *gaze/head movement*, and *speed of motion*. All four episodes use close proximity to the participant corresponding to motion path IV in Figure 4.5 using either slow or fast speed of motion. Once at this position, the robot either moves its head erratically or gazes directly at the participant. It stays in this position for 15 seconds, and then leaves the room to return to its starting position. Rotational body language is always used.

4.3.3 Implementation

4.3.3.1 Robot Platform

The robot platform used for the 2nd Pilot was same as the one used for the 1st Pilot (see Section 4.2.2.1).

4.3.3.2 Controller Station

The controller station used for the 2nd Pilot allowed the study administrator to serve two purposes (a) to remotely control the robot's behaviour using fully automated *pre-programmed behaviour macros* that ran autonomously, significantly reducing the amount of reliance on manual remote controlled behaviour, and (b) to record relevant participant comments using a *transcription tool* that we designed to enable easier analysis of the qualitative remarks made by participants during the study. We describe both of these




Figure 4.6: Remote control station used for robot control and transcription

elements of the controller station below.

The station (Figure 4.6) comprised a standard laptop with a second monitor, and a wireless router that linked the laptop with the robot. The controller station was positioned so that the robot was always within the study administrator's view, except when it entered the office. Participants could not see the controller station (or the study administrator) from within the office.

Robot Behaviour Remote Control. Custom software on the controller station's laptop was primarily used to issue high level commands to the robot that triggered predefined macros. These macros in turn executed particular robot behaviour. Each behavioural episode had exactly one corresponding macro. The use of pre-programmed macros for the robot's behaviour relieved the study administrator of the need to manually control the robot's behaviour at all times, and ensured that all participants would observe nearly identical robotic behaviour. Figure 4.7 shows the user interface used to invoke the

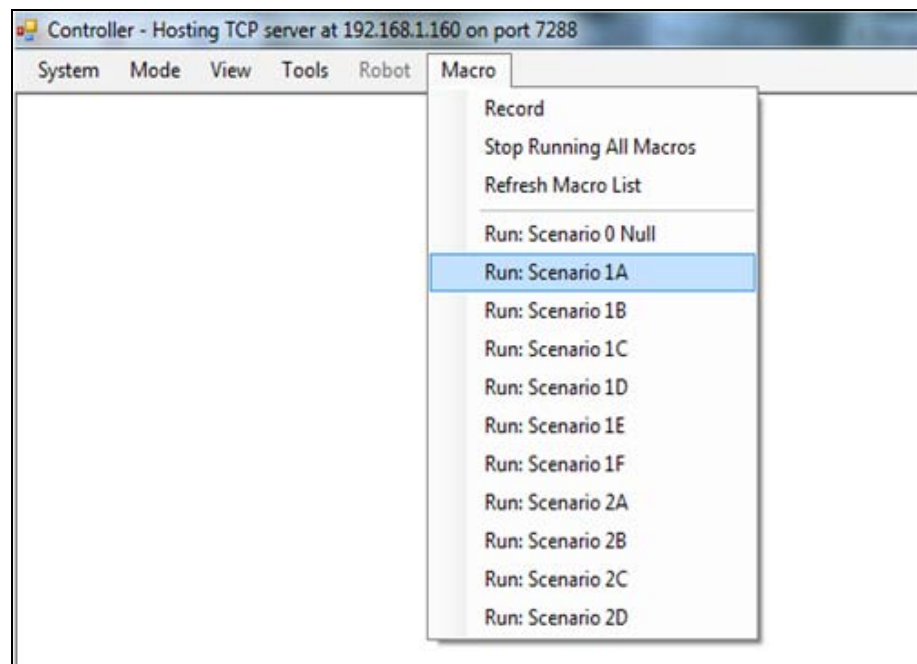
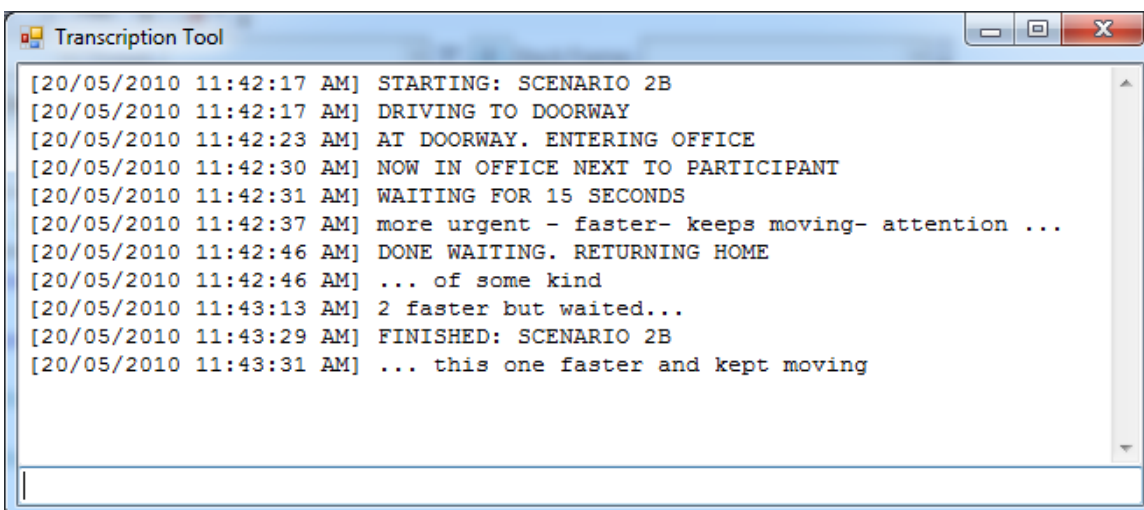


Figure 4.7: A view of the user interface used to invoke the pre-programmed behaviour macros.

behaviour macros (behavioural episodes were labelled as “scenarios” in the interface).

Manual positioning controls were also provided to allow the administrator to reposition the robot to its start location, as marked on the floor, as the pre-programmed macro behaviour did not always do so precisely after completing a behavioural episode. The amount of distance that the robot is asked to travel was not guaranteed by the robot’s programming interface to be equal to the amount that the robot would *actually* travel. Thus, the two would often differ. These differences, while minor (~ 5-10cm), could become significant if left unchecked. Thus, the study administrator was instructed to ensure that the robot always began each behavioural episode from the same starting position marked on the floor.

Transcription Tool. The software also included a transcription tool we designed, which enabled logging of participant comments with timestamps and other high level events sent back to the station by the robot. This tool, which ran on the study administrator’s secondary computer monitor, was integrated into the robot’s controller software. This enabled the resulting transcript to be augmented automatically with



```

Transcription Tool
[20/05/2010 11:42:17 AM] STARTING: SCENARIO 2B
[20/05/2010 11:42:17 AM] DRIVING TO DOORWAY
[20/05/2010 11:42:23 AM] AT DOORWAY. ENTERING OFFICE
[20/05/2010 11:42:30 AM] NOW IN OFFICE NEXT TO PARTICIPANT
[20/05/2010 11:42:31 AM] WAITING FOR 15 SECONDS
[20/05/2010 11:42:37 AM] more urgent - faster- keeps moving- attention ...
[20/05/2010 11:42:46 AM] DONE WAITING. RETURNING HOME
[20/05/2010 11:42:46 AM] ... of some kind
[20/05/2010 11:43:13 AM] 2 faster but waited...
[20/05/2010 11:43:29 AM] FINISHED: SCENARIO 2B
[20/05/2010 11:43:31 AM] ... this one faster and kept moving

```

Figure 4.8: The transcription tool.

timestamps and high level events regarding the robot's behaviour (e.g. "the robot is now approaching the doorway", etc.). Figure 4.8 shows the user interface used by the study administrator. Since the robot controls are manipulated using only the computer's mouse, all keyboard input is redirected to the transcription tool, even if the transcription tool's window does not have focus or is minimized. All text entry by the study administrator is saved to the hard disk automatically, which prevents data loss due to unexpected termination of the controller software caused by unanticipated bugs.

4.3.4 Methods

4.3.4.1 Data Collection

Qualitative comments made by participants during the study were captured in both video recordings, and in textual notes taken by the study administrator in real time using the transcription tool. These notes were used to assist processing of the video recordings by allowing comments of interest in the text notes to be quickly found in the videos using the recorded timestamps.

We wanted to collect quantitative as well as qualitative data, to enable statistical comparison of differences between people's interpretations. Consequently, we added a second phase to the study to focus on quantitative data collection. In this new second



Figure 4.9: A partial view of the Interruptedness Metre used by participants to rank interruptedness in the study's second phase.

phase, the robot still interrupts the conversation between the interviewer and the participant for each episode, as with the first phase. In this second phase, however, the interviewer stops the conversation after each interruption and asks the participant to quantitatively rank the interruption behaviour through the use of a custom ranking device we called 'Interruptedness Metre' (Figure 4.9). The participant would order the sequences they saw from least interruptive (left) to most interruptive (right). Rankings, which were translated from their relative position on the metre to a continuous scale from 0 to 100, form the participant's subjective measure of how they interpreted the robot's behaviour. A higher ranking corresponded to a higher level of interruptedness, while a lower ranking corresponded to a lower level of interruptedness. Qualitative remarks are also collected in Phase 2, where we ask participants questions relating to their choice of quantitative ranking.

After adding the second phase, we then shortened the first phase to use only four representative behavioural episodes (including the null base case at the beginning), while the second phase has all episodes. We did this to save time, as running all episodes in both phases would have led to excessively long study sessions.

4.3.4.2 Experimental Procedures

We conducted the 2nd Pilot using four participants. None had participated in our previous pilot. The study consists of an introduction outside the office, two phases of robot behaviour inside the office, followed by a brief interview period for general reflection.

The introduction consisted explanations about the general format of the study. Participants were told that the researchers are interested in exploring how people interpret the robot's behaviour in an interruption context. No further details were given regarding

any approach used by the robot to interrupt, nor that it was pre-programmed and being controlled by the study administrator.

Both phases of robot interaction with the participant were qualitative, while the second was also quantitative. Instead of having participants choose their own task (which could affect how attentive they might be toward the robot), the participant and the interviewer (an actor) were seated in an office (e.g., Figure 4.10) for both phases, having a conversation about topics unrelated to robots or the user study. Both the participant and interviewer had a clear view of the open doorway. This seating arrangement allowed the participants to see and comment on the robot's behaviour, but also enabled them to ignore the robot's behaviour if they were too "busy" with the conversation with the interviewer. The conversation typically began with small talk about school and continued with the participant's interests that emerged; specific questions used to seed the conversation are listed in Appendix A.



Figure 4.10: The robot enters the office to interrupt a participant.

While the conversation occurred, the robot underwent attempts to interrupt the participant using a series of minimal nonverbal behavioural episodes as mentioned previously and summarized in Table 4.3. All episodes began with the robot out of view outside the office. The two phases differed in the particular set of episodes used and whether a verbal interview or ranking by the participants occurred once the robot had completed an episode.

The primary purpose of Phase 1 was to gather qualitative and unbiased reactions to interruption. This phase comprised four pre-programmed behavioural episodes (Null, 2D, 1A, and 2A, as defined in Table 4.3) initiated by the study administrator. These specific episodes were selected to representatively exhibit a wide range of robot behaviour. The order of the episodes used following the null case was randomly generated, and was the same across all participants.

To reduce predictability of when an interruption may occur, each interruption attempt was separated by a short delay of a few minutes. During each attempted interruption, the interviewer encouraged the participant to talk about his or her reaction, i.e., the methodology followed that of constructive interaction think-aloud. After the robot completed its episode and was out of view, the interviewer asked the participant questions about the interruption to get a sense of how they understood the behaviour, e.g., how they would describe their experience with the robot, and what circumstances they thought would have led the robot to behave this way. Although the participants were asked to think about how they would interpret the robot's behaviour in an interruption context, they were also invited to discuss their feelings about the robot and its behaviour in general, e.g., what emotions, if any, the robot seemed to be conveying, how annoying

the robot was, etc. We felt that these initial interactions with the robot provided an unbiased participant reaction, as they were the first experiences that the participants had with the robot, and no expectations about “right or wrong” answers were expressed. Everything the participant said was recorded in real time by the study administrator.

The primary purpose of Phase 2 was to have participants quantitatively rank the how “interrupted” they felt by the robot during each behavioural episode. In this phase, the robot progressed through all ten pre-programmed behavioural episodes in a random order that was kept the same for all participants, following the null case episode (Null, 1F, 2C, 2B, 2D, 1C, 2A, 1D, 1E, 1B, 1A, as defined in Table 4.3), with little delay between them. Following the robot’s completion of each episode, the participant was asked to rank how interrupted he or she felt by the robot by placing a marker on our custom ‘Interruptedness Metre’ (e.g., Figure 4.9). There were eleven markers total, each corresponding to one specific robotic behavioural episode seen by the participant. The participant was informed that markers could be placed anywhere on the ranking device, but they could not overlap. To gain additional reaction, the interviewer asked participants to explain their choice of spot for the marker as they placed it on the Interruptedness Metre.

After both phases, participants were interviewed for their final impressions and thoughts. Overall, the primary goal in data collection is to gain insight into the participants’ understanding of the robot’s behaviour, thus addressing our primary research question in identifying a minimal set of nonverbal behaviour cues that are understandable by people in an interruption context.

4.3.5 Results and Discussion

We found that the changes to the format and implementation resolved the technical deficiencies of the 1st Pilot. Our preliminary analysis of this 2nd Pilot indicated that changing the robot's behaviour to eleven cue-based behavioural episodes was sufficient to permit quantitative analysis that answers our primary research question, i.e., to find *minimal nonverbal behavioural cues* usable by robots to communicate interruption urgency in ways are understandable by people. The qualitative data enriched this information. For example, our pilot participants expressed different opinions for different robotic behavioural episodes.

“[The robot's behaviour] triggered my attention that something is wrong. [The robot] is signalling a warning, or an emergency.”
(2nd Pilot participant 1, episode 2D, fast movement)

“The robot is calm and happy trying to tell me something. It's trying to get my attention. It's looking at me.”
(2nd Pilot participant 2, episode 2A, slow movement)

4.4 Main Study (Robot Interruption User Study)

Following a successful 2nd Pilot, we were ready to proceed with our main study which closely paralleled the design of the 2nd Pilot. Appendix A includes additional study materials for the main study that are not included in this chapter, such as consent forms.

4.4.1 Participants

Twenty participants were recruited for the main study through mailing lists at the University of Calgary. Although no particular groups were targeted, participants were a nearly equal mix of male and female graduate students with varied ethnic backgrounds, many of whom were members of the Faculty of Engineering, with ages ranging from 20

to 30. Participants received \$15 in compensation. Each study session was approximately 45 – 60 minutes long.

4.4.2 Implementation

The materials, including the robot platform and controller station, used for the main study were the same as those used for the 2nd Pilot (see Section 4.3.3).

4.4.3 Methods

The methods used for the main study do not include any substantial changes from those used for the 2nd Pilot (see Section 4.3.4), except those required to properly conduct a formal in accordance with university guidelines. The introduction of the main study was extended to include formalities such as signing of the consent form (see Appendix A), and dispensing of compensation in addition to the explanations about the general format of the study.

4.4.4 Results and Discussion

We discuss the results from this main study in Chapter 5, and discuss their meaning in further detail in Chapter 6.

4.5 Summary

In this chapter, we explored the approach we used in addressing our primary research question: are there *minimal nonverbal behavioural cues* that robots can exhibit to communicate interruption urgency, and are those cues understandable by people?

First, we identified a set of four minimal behavioural cues, and specific behaviour based on these cues usable in a series of situational interruption scenarios in a human observation study. Second, we designed and critiqued the implementation of potential robotic interruption behaviour that mimicked the observed behaviour using an actual

robotic interface in our 1st Pilot. Third, we redesigned the robot's behaviour and experimental design in the 2nd Pilot to resolve various issues that emerged during the 1st Pilot. Finally, we conducted our main study once we verified the efficacy of the robot behaviours and the human understanding of them.

In the next chapter, we describe our analysis procedures and summarize the quantitative and qualitative results of the main robot interruption user study.

Chapter Five: Results of the User Study

This chapter describes the quantitative and qualitative results from the main user study described in Chapter 4. First, we describe the process we followed to select a viable statistical model to analyze our quantitative results, followed by the results of this analysis. Second, we present the process we used to collect and analyze our qualitative results. Finally, we explore and summarize our qualitative results and prepare a foundation to discuss their meaning in Chapter 6.

5.1 Quantitative

5.1.1 Selection of Statistical Analysis Model

Prior to this study, I had no formal statistical experience except through a course that introduced basic concepts of t-tests and ANOVA. I thus consulted with a statistician³ about the best test to use to match my study conditions.

To review, I had 19 participant data sets for 11 different trials (i.e., behavioural episodes) that were kept the same for all participants. The null base case episode (as defined in Chapter 4) was designed to be a covariate, or independent variable. While debatably not entirely “null”, we expected its results to act as a baseline for comparison with all other results.

Initially, I considered ANOVA, but it did not emerge as the best test for three reasons.¹ First, only nineteen participant data sets are available for eleven different behavioural episodes. For more ideal ANOVA analysis, the data would consist of either more participant data sets or fewer behavioural episodes. Second, the participants were

³ Personal communication with Gisela Engels, Senior Stat Consultant, Information Technologies, University of Calgary

expressing their opinion of the robot's behaviour as a percentage on the Interruptedness Metre, which amounts to a ranking. This makes ANOVA less useful for analysis. Finally, ANOVA would not account for the baseline as a covariate

Instead, we decided upon a *linear mixed model* for our analysis. With this model, one can choose whether or not to use a covariate. We attempted to run this model with and without using the null base case as a covariate. A comparison of the -2 Restricted Log Likelihood and Akaike's Information Criterion (AIC) revealed that it is better to use a covariate in this data set ¹, according to the statistician who I consulted.

5.1.2 Identifying Significant Robot Behavioural Cues

The robot's behaviours (summarized in Chapter 4 in Table 4.3) were designed to enable statistical analysis that identified which of the robot's behavioural cues (proximity to person, gaze and head movement, body language, and speed of motion) actually had a statistically significant impact on the interruptedness felt by a person due to the robot's behaviour.

Tables 5.1 through 5.4 summarize the statistical significance of each individual cue of the robot's behaviour used in the study as well as the interaction between cues. P-values are considered statistically significant based on a threshold of $p < 0.05$. These significant values are distinguished using bold text in the tables.

Table 5.1 summarizes the effect of speed of motion, interacting with gaze/head movement and proximity. The speed used by the robot for its spatial motion as well as its head movement was either slow or fast. As Table 5.1 indicates, speed of motion was significant only when the robot was situated next to the participant. When the robot was

located at the doorway of the office or outside the doorway, no significant impact was observed.

Table 5.1: Significance of Speed of Motion

<i>Gaze</i>	At participant				Erratic
<i>Proximity</i>	Far from Doorway	At Doorway	Next to Participant		
<i>Episodes</i>	1A & 1B	1C & 1D	1E & 1F	2A & 2B	2C & 2D
<i>P-Value</i>	0.139	0.360	0.025	0.010	0.006

Table 5.2 summarizes the effect of gaze and head movement, which interacts with speed. The gaze suggested by the robot's head movement was either directly focused on the participant, or erratic movement where the head was constantly moving in all directions, as described in Chapter 4. The data in the table indicates that head movement/gaze had no statistically significant impact.

Table 5.2: Statistical Significance of Gaze and Head Movement

<i>Speed</i>	At Slow Speed	At Fast Speed
<i>Episodes</i>	2A & 2C	2B & 2D
<i>P-Value</i>	1.000	0.996

Table 5.3 summarizes the effect of proximity to person, interacting with speed. The robot used three proximity positions: next to the participant, at the doorway, and outside the doorway. The data shows that there was no statistically significant difference between being at the doorway or far from the doorway. However, there was a significant difference between being far from the doorway and being next to the participant. When comparing positions at the doorway and next to the participant, there was only a significant difference at fast speed, and not slow speed.

Table 5.3: Statistical Significance of Proximity

<i>Speed/Proximity</i>	At Slow Speed	At Fast Speed
<i>Far from Doorway vs. At doorway</i>	1.000	1.000
<i>At Doorway vs. Next to Participant</i>	0.050	p<0.001
<i>Far from Doorway vs. Next to Participant</i>	0.012	p<0.001

Table 5.4 summarizes the effect of body language using rotation, which interacts with speed. For some of the behavioural episodes, the robot rotated its body in place while stopped, while it used no body movement when stopped for other episodes. The data shows that this cue was not statistically significant.

Table 5.4: Statistical Significance of Body Language (using rotation)

<i>Speed</i>	At Slow Speed	At Fast Speed
<i>Episodes</i>	<i>2A & 1E</i>	<i>1F & 2B</i>
<i>P-Value</i>	1.00000	1.00000

5.1.3 Means

Table 5.5 presents the mean rankings collected from the Interruptedness Metre. To review, the rankings are translated from their relative position on the metre to a continuous scale from 0 to 100. The means shed light on the magnitude and direction of differences for the robotic behavioural cues that proved significant: speed of motion, and proximity to the participant. For proximity, the significant differences in interruptedness are around 25-30 when comparing positions of far from the doorway to next to the participant (see Table 5.5, 19.3 far from doorway vs. 43.8 next to participant, and 34.6 far from doorway vs. 64.0 next to participant).

For speed of motion, the differences in interruptedness between slow and fast when the robot is next to the participant are not only statistically significant, they are also large: around 20 each (see Table 5.5, bottom row is close proximity, 46.2 slow vs. 63.7 fast, and 51.2 slow vs. 72.2 fast).

As mentioned, our statistical analysis indicates that the differences of means for gaze/head movement and body language are not significant.

Table 5.5: Interruptedness Means by Factor

<i>Proximity Position</i>	<i>Body Language</i>	<i>at slow speed</i>		<i>at fast speed</i>	
		<i>direct gaze</i>	<i>erratic gaze</i>	<i>direct gaze</i>	<i>erratic gaze</i>
<i>Far from doorway</i>	<i>none</i>	19.3		34.6	
	<i>rotating</i>				
<i>At the doorway</i>	<i>none</i>	23.9		33.4	
	<i>rotating</i>				
<i>Next to participant</i>	<i>none</i>	43.8		64.0	
	<i>rotating</i>	46.2	51.2	63.7	72.2

5.1.4 Summary of Quantitative Statistical Analysis

Pairwise comparisons of the quantitative results for robotic behavioural episodes do not show any statistical significance for gaze/head movement or body language. Speed of motion is only significant when the robot approaches a person at close proximity. The proximity position used by the robot is also significant, at least when comparing the robot's position inside the office to outside the office, i.e., whether the robot appears at the doorway or far from the doorway does not matter.

5.2 Qualitative Analysis Methodology

Following our quantitative results from Phase 2 which provide numerical measurements, we focused on qualitative-oriented exploration of the experiences and observations by participants.

Phase 1 was primarily qualitative with no quantitative data collection whatsoever. Generally, all participants made comments about the robot and its behaviour as it appeared. Once the robot completed its episode and left the room, the interviewer asked questions (described in Chapter 4) to gain deeper insight of the participant's understanding of the robot's behaviour. In Phase 2, the interviewer questioned the participants' choice of ranking using the Interruptedness Metre.

5.2.1 Transcription Log Collection

All comments made by the participants during both phases were recorded in real-time by the study administrator using our custom Transcription Tool software (described in Chapter 4). The records produced by this tool included timestamps (e.g., "11/06/2010 1:37:55 PM") to allow cross referencing with the recorded video and high level event notifications inserted by the robot controller software (e.g., "the robot is now at the doorway").

For each participant, the software generated a log file in plain text format (e.g., Figure 5.1) using the participant identifier as the file name. To ensure accurate time stamps, the high level events often appeared in the middle of a participant's comment. Ellipsis symbols (i.e., three periods in a row) were inserted automatically into the logs to indicate that this had occurred, and that the participant's comment was continuous, even if it was broken up into multiple lines. The tool did not allow the study administrator to

manually type an ellipsis, so all instances of this symbol appearing in the log files are automatically inserted.

First, we modified these log files to include the participant identifiers into each line (e.g., Figure 5.1, bottom), as the software did not do this automatically. Next, the comments were transferred to word-processing software (via copy and paste) where they were grouped by behavioural episode (originally using “scenario” as a label instead of “episode”) while still retaining all comments, participant identifiers, high level events, and timestamps to enable easy reference with the original transcription logs as necessary. Through this transformation process, no participant comments were removed or modified.

The separation by behavioural episode was done by including all comments made

```
[11/06/2010 1:36:56 PM] STARTING: SCENARIO 2A
[11/06/2010 1:36:57 PM] DRIVING TO DOORWAY
[11/06/2010 1:37:06 PM] AT DOORWAY. ENTERING OFFICE
[11/06/2010 1:37:15 PM] NOW IN OFFICE NEXT TO PARTICIPANT
[11/06/2010 1:37:15 PM] WAITING FOR 15 SECONDS
[11/06/2010 1:37:30 PM] DONE WAITING. RETURNING HOME
[11/06/2010 1:37:48 PM] again he came here his movement wa...
[11/06/2010 1:37:55 PM] FINISHED: SCENARIO 2A
[11/06/2010 1:37:55 PM] ... sn't very fast - didn't interrupt me very much
[11/06/2010 1:38:36 PM] if he bump me i would get very interrupted
[11/06/2010 1:38:57 PM] in an office i can see a lot of movement so i don't find it very
interrupted
[11/06/2010 1:39:30 PM] STARTING: SCENARIO 1D

[P17 11/06/2010 1:36:56 PM] STARTING: SCENARIO 2A
[P17 11/06/2010 1:36:57 PM] DRIVING TO DOORWAY
```

Figure 5.1: Sample of raw data generated from transcription logs. The bottom two lines have the participant identifiers inserted.

after the beginning of one episode, but before the beginning of the next episode. For example, in Figure 5.1 at the top, all comments are pertaining to episode 2A because they are made before the robot begins episode 1D. The timestamps shown refer to the *start time* of the comment appearing adjacent to it, and serve to enable cross-referencing of the comments with the video recordings.

While these transcription logs proved to very useful and viable for our analysis, we did frequently refer back to the video recordings to verify and improve the completeness of some parts of the logs that we felt required deeper analysis. For example, the videos were used to add additional context to a participant's comments, such as what the interviewer had specifically asked or said previously, as the study administrator only transcribed what the participants said, not the interviewer. In other cases, the video recordings allowed us to correct grammatical mistakes or restore connecting words not recorded by the study administrator. In some cases, the study administrator could not hear some of the participant's comments, thus leaving a portion of the transcription log blank. Thus, the video recordings enabled us to fill in these blanks.

5.2.2 Analysis

Our analysis began with a process akin to open coding (Strauss and Corbin, 1998) and affinity diagramming where the purpose was to synthesize categories based on similarities between the collected data. We focused on exploring the qualitative comments and taped interviews to find themes and present them in summaries. The exact procedure conducted, from a high level, consisted of three parts: read the recorded comments, cluster quotes into related themes using the open coding methodology, and expand each theme into paragraphs of discussion.

First, we read all of the recorded comments to gain an insight of the range of comments received and began to develop impressions of possible concepts and trends present in the data. Early on, large trends emerged along with interesting and unique interpretations from smaller subsets of participants. The trends that emerged came from our own logical reasoning and interpretation based on related work (Chapter 2) and our experience in conducting this user study.

Second, we clustered the quotes into related themes. To do this, we identified groups of comments from different participants that seemed to express similar interpretations of the same robot behaviour. For example, we found that many participants compared the robot to various social entities such as small children and dogs, thus a focused effort was specifically made to collect all of these comments together. The same process was iteratively followed for numerous other themes, until only a minority (< 25%) of unique comments remained as possible outliers.

These remaining comments were themselves collected into a group for reconsideration for inclusion into the other groups, or new ones. Most of these comments were in fact reorganized into new sections that appear later in this chapter. Some were examined more closely, and deemed redundant with other comments already included in a themed section. A small number (< 10%) of other comments didn't seem to add any clear insight to the participant experience, and were thus not discussed further.

Third, from the groups of themed comments, we created high level category headings to describe the comments, and expanded the groups of themes into paragraphs of text. Some emerged from participants (e.g., our "politeness" heading emerged from participants calling the robot "polite"), while others did not emerge directly from

participant quotes. For example, we refer to a group of comments under “Robot as a social entity”, even though the term “social entity” was never used specifically by any participant. Having said this, the comments made by participants make reference to various individual examples of social entities, hence the heading we chose.

Once we finished creating the initial set of themes, we reordered them to present a coherent story of the participant experience of the robot. Only at this point did we reformat the raw comments to remove timestamps, participant identifiers, etc. that were no longer needed, as additional context from the video recordings had already been added as needed. The participant quotes were also shortened to keep only the most significant part (e.g., “eavesdropping” is extracted from “I think the robot was eavesdropping”) so that comments from multiple participants could be summarized in single sentences, as appropriate. To improve the readability of sentences incorporating a large number of participant quotes, redundancies were reduced. For example, a list of quotes like “annoyed”, “distracted”, “disturbed” and “interrupted” could be shortened to a smaller list (e.g. just “annoyed” and “distracted”) without losing any significant meaning.

These quote groupings have the advantage of providing a wide breadth of participant experience with many descriptive words derived directly from the participants, at the expense of wider and more complete context for any individual participant’s experience. To balance this, a representative set of complete participant comments comprising two or three sentences are included in our discussion at certain points to give a better impression of the participant experience. These were chosen based on their relevance to our discussion and themes seen overall, and readability. Our process, while initially akin to open coding (Strauss and Corbin, 1998), was less

structured than the formal method. We assigned categories to virtually all of the recorded comments, but we didn't try to connect them together using a single theory. Instead, our goal was to allow ideas and themes to emerge from the data. This process was similar to affinity diagramming, except that we didn't record or transfer the comments to another medium, but we did follow through the process of finding ideas that seem related and grouping them until no ideas remain.

In the sections that follow, we first group many of the trends that emerged from the four representative robot behaviours used in Phase 1. As these cases were intended to represent a wide range of robot behaviour, we expected to also provide the widest range of participant experience to emerge in these behaviours in a fashion that could be discussed and contrasted effectively. These results are summarized in sections 5.3.1 through to 5.3.4. Following these, we explore the remaining themes and clusters of comments (e.g., the robot as a social being, politeness, etc.) that emerged from the data across all behaviours, not just the first four. These results are explored in sections 5.3.5 through to 5.3.8.

5.3 Qualitative Results

We now summarize and discuss the qualitative comments received for both phases of the study: the first qualitative phase, as well as the second phase which was focused on quantitative feedback but also allowed the participant to provide further qualitative reflections. We begin with participant impressions of each robotic behavioural episode in Phase 1. Following that, we talk more generally about themes regarding particular perceptions people had across both phases.

5.3.1 The Null Base Case: Impressions of the Robot

The null base case episode occurred first, after the participant was in the office for a few minutes, i.e., the robot passed by the office door without any head movement, and did not gaze into the office. During this episode, about half of the participants commented on the robot's behaviour just as it began moving past the doorway; the others just kept talking to the interviewer, and talked about the robot only when asked by the interviewer when the episode ended. Most said they first detected the robot because of its noise, even before it was visible through the doorway. Many described the details of how they observed the robot's behaviour using phrases like "it just passed by", "it's coming", and "it disappeared". The behaviour was "calm", and "not disturbing". One participant said the robot looked as if "it could move faster" than it was. Although all participants clearly noticed the robot, one said it was not "super distracting" and that it got his attention in a "polite way". Another said the robot was "minding its own business" and that it "didn't affect the flow of conversation". While one found the whole behaviour to be "pretty weird", most said nothing at all about the experience being strange or weird.

Participants were asked what they thought the robot was trying to do, or what its intentions were. None felt that the robot was trying to interact directly with them, but opinions of what it was doing varied. Some assigned social presence to the robot similar to a person just passing by on the way to some other location, or pacing about with no specific mission, or in the middle of accomplishing a task such as delivering messages or moving objects around the office. One even compared the robot to a "child, waiting to be noticed".

Participants were quite generous in the social abilities they afforded to the robot, despite it lacking any form of eyes, ears, cameras, microphones, or speakers. Many participants felt the robot was “curious” about their presence, even “spying” or “eavesdropping” on the conversation, because the robot did not “know” who the participant was. One even implied that the robot felt territorial because it was approaching for a “sense of security”.

“I heard its wheels. I had a feeling it was moving. I thought he would come inside [the office], but he didn’t. He was examining the perimeter, becoming familiar with his surroundings, and mapping out objects.” (P08)

Others felt that the robot was responding to louder talking between them and the interviewer. Some were more specific, saying that the robot “heard its name” (despite the robot not having a name during the study) or the word “robot”, and wanted to hear more of the conversation. Many noticed the robot’s lack of active behaviour (other than moving by the doorway), and said the robot was not interacting because they were “not paying attention”.

5.3.2 Episode 2D: fast, erratic gaze, close proximity, rotating

In the next episode, the robot directly approached participants with its most extreme behaviour, where it was active and fast-paced (Table 4.3, episode 2D).

Participants initially described this behaviour using active words such as “weird”, “big”, “racing”, “scared”, “frantic”, “hard to ignore”, “in a rush”, etc. Many participants said they were “annoyed”, “distracted”, “disturbed” and “interrupted” by the behaviour, and unable to continue their conversation with the interviewer. The entrance of the robot into the room was described as “forceful” or comparable to banging on a door. Because

of the robot's faster movement, its motors made more noise, which one participant described as "different" and "huge".

"[The robot] distracted me. He came in quickly and moved... It had to say something to me. It was urgent. It came in forcefully and tried to gain my attention." (P06)

The behavioural cue mentioned most often was speed. Head movement was also mentioned, but to a much lesser extent. Many participants also noted that the robot came into the room (referring to the cue of close proximity) *vs.* the previous behaviour where the robot just passed by the doorway. Very few commented specifically on the robot's body language and movement while inside the room during the whole study, even though it was persistently rotating back and forth for 15 seconds. One specifically said that the closeness of the robot felt more significant than its movement. Another said it was "kind of weird" that the robot was communicating with body language only, and no verbal communication.

"It definitely gets your attention. There is something very important going on. It is still moving, and racing. It can't wait. Like a really important message, fire in the building, emergency, or the boss is calling." (P01)

Almost all participants viewed this behaviour as representative of an "emergency", "something [being] wrong", "someone hurt", or something having "happened". Several participants even identified the emergency as a possible "fire", one saying "probably a fire". One said the robot's behaviour indicated that it was necessary to stop the conversation and move out of the room.

Almost all used words such as “important” or “urgent” to describe the potential reasoning behind the interruption. One said this behaviour would be “rude” if it was used to interrupt an important meeting, but not a casual one.

In summary, it is clear that this behaviour was largely associated with “fire” or “emergency” behaviours. Indeed, one participant said the behaviour would be “inappropriate” for a non-urgent interruption such as a greeting.

5.3.3 Episode 1A: slow, direct gaze, far from doorway

In the next behaviour in Phase I, the robot stood outside the doorway and did not enter the room (Table 4.3, episode 1A). Generally, this behaviour was seen as non-interruptive. In all but one case (where the robot was not even noticed), participants noticed the robot in part due to the noise it was making. Comments described how non-interruptive it seemed, for example it was “not interrupting” because it “did not approach too close, but from a distance”.

“The robot looks like he is still interested. He was peeking inside the room for a short while, perhaps to report back to someone else. He was not interrupting and did not approach too close, but from a distance.” (P02)

Many participants felt that the robot was acknowledging their presence and “noticing” them, e.g., “this time I’m sure it’s noticing us”, because of the “head movement”. A few said the robot was going by, but was stopping to “listen to the conversation”, and that it was “paying attention”. Another said it was “curious” and that it was “eavesdropping a bit” because it “overheard the conversation and was interested”.

Other participants interpreted the robot’s behaviour as something other than interruption-based. One said it doing “periodic checking, in case we need something”.

Another said the robot was “peeking inside the room” and then “reporting back to someone else”.

5.3.4 Episode 2A: slow, direct gaze, close proximity, rotating body language

In the next episode, the robot operated at close proximity (Table 4.3, episode 2A).

Participants had varied impressions.

One participant noted that the robot, like a person, was more interruptive when it entered the room, compared to when it did not enter. Another said the robot seemed to be acting with more “maturity” due to the eye contact, and that it was respectful and more “accustomed to social rules”. One said he was “surprised by the smooth motion”, and that it was “not going crazy”.

“[The robot is] looking at me, addressing me. He needed to tell me but was not impolite. He knew we were having this conversation, didn't move as aggressively. He had softer movement. [The interruption is] important because he approached me.” (P11)

Many participants expressed how they felt emotionally about this interaction, contrasting it to the previous ‘urgent’ behaviour noted in episode 2D. Two participants said that this behaviour “didn’t scare” them. One said that the previous one had lots of “shaking” and required some “getting used to”. Another said the robot was “not very annoying” whereas it was previously “making a lot of noise” and “bothering” him. One participant felt more comfortable, whereas they had been previously worried that the robot might have hurt them in episode 2D. One participant “preferred” this behaviour.

We now turn to more general impressions of the robot across all episodes of both phases.

5.3.5 The Robot as a Social Being

Many participants made comments about the robot as if it were a person. One participant said the robot was like a “real being” because it was showing interest in things, going away and then coming back. Another said it moved and tried to gain attention by “barging” in and moving its head. One felt that the robot was “annoyed” that its space was being intruded on. Another suggested that the robot was actually trying to annoy him or do something funny.

Two participants did compare the robot to non-human entities, such as a dog running up to a visitor when entering a house.

“It’s like when you walk into some people's house and the dog comes to you. It’s an inappropriate greeting. It’s too much. Slower would be better.” (P20)

Another compared the robot to a child entering the room, in a manner that a child might approach his or her parent, to say that someone was annoying them.

5.3.6 The Robot as a Machine

A few participants described the robot as a machine. One said its procedure was “smooth”, because of the “mechanics or software”. Another suggested that the robot was exhibiting certain behaviour because it was “broken” or “damaged”. Yet another felt that the robot was “examining the perimeter, becoming familiar with its surroundings, and mapping out objects”. One said the robot seemed to be “analyzing” them, collecting data, taking pictures, and recording audio. A small number of participants suggested the robot was running through “programmed” behaviour or being controlled by the study administrator.

5.3.7 Politeness When Interrupting

A common theme used in describing the robot's behaviour in many episodes was politeness. Many participants felt the robot had some intention, but that it chose to defer that intention when it noticed that a conversation was in progress. One participant thought the robot wanted to "say something" that was "not important", but that it "changed its mind" because of the conversation, and that it would "come back later". Another felt the robot was coming for a "scheduled meeting", but that it was "waiting outside".

Many participants defined this type of behaviour as either "normal", "better" and "more gentle" compared to the more extreme behaviour used in previous episodes. Another said the robot was "looking for someone" on behalf of someone else, and that it was "trying to say something", but did not say anything because it "didn't want to interrupt" the conversation. Similarly, another participant said the robot was "trying to look for an opening in the conversation" so that it could "add to it".

"It's like someone is here for a meeting. I see it as a messenger or servant. [The robot used] slow motion, more control, further distance. It was more polite." (P13)

A couple of participants interpreted authority as a factor, and compared the behaviour to someone who "is waiting for a superior to finish" and that the matter was not "urgent", as the robot was not "actively catching attention". Another participant saw the robot as a "messenger" or "servant". Another felt the robot was acting as a servant, but for someone else.

While most participants used comments that implied some element of politeness, one participant said the robot was "impolite" because it was just "staying there and staring", though even this person noted that the robot "didn't want to interrupt".

5.3.8 Familiarity with Robot

Two participants commented on their increasing familiarity with the robot across behaviours. One said the robot made a bit of noise and was distracting in a way that was “out of character”, implying a certain familiarity with some behaviour that was “in character” for the robot. Another said he had “seen the robot too many times before” and that he was becoming “more sensitive” to noticing it over time. Another didn’t look at the robot much because it was becoming a “common occurrence”, while another said he was getting “used to” the robot.

5.4 Summary

In this chapter, we explored the results of the main robot interruption user study. We first presented the quantitative results from the Phase 1 which revealed that speed of motion and proximity had a statistically significant impact on the participant’s interruptedness, while gaze/head movement and body language had no impact. Then, we moved on to qualitative comments made by participants during both phases in which participants viewed the robot, in many cases, as a social entity with wide ranging abilities to convey context about an interruption.

In the next chapter, we move onto deeper interpretations of these results, including elements that surprised us, and what lessons have learned from them.

Chapter Six: Discussion

In the last chapter, we summarized the results of our main user study regarding human interpretation of robotic behaviour cues in an interruption scenario. In this chapter, we will explore these study results in more detail. First, we will discuss which behavioural cues used by the robot were significant, and reasons why other cues may not have been significant. Second, we will explore the view of the robot as a social entity that emerged from the comments made by participants. Third, we move on to behaviour used by the robot that was considered to be inappropriate on a social level. This brings us to our final point of discussing how a robot can moderate the timing of its interruption to minimize unwanted disruption.

6.1 Significance of Robotic Behavioural Cues

From the outset, we sought to determine if there are *minimal nonverbal behavioural cues* that robots can exhibit to communicate their internal state, specifically in an interruption context, and if those cues understandable by people. From the results of the robot interruption user study, we have verified that robots can convey urgency about an interruption situation using only basic elements of its physical behaviour.

Our quantitative results statistically show that both speed of motion, and proximity to the person can both provide a range of interruptibility over urgency. This alone is a significant contribution, as it demonstrates simple behavioural cues can be used by any robot capable of physical movement (e.g., a Roomba) to convey interruption context. These cues are even feasible for robot implementations that lack gaze or precise body language abilities, as we surprisingly found these cues to be statistically insignificant in our study.

Despite the lack of quantitative significance for gaze/head movement, anecdotal comments from participants do suggest that some form of eyes, head, or indication of forward direction is useful. Many participants did mention that the robot was looking at them. When the robot was distant or not gazing directly, participants did not feel they were ones that the robot wanted to interrupt. Instead, they seemed confused or that the robot was searching for someone else when far away. These comments suggest that gaze plays a useful role in the coordination of an interruption to identify the person being interrupted, even if gaze does not have a statistically significant impact on how interruptive the robot seems to a person. Fortunately, gaze can be easily added to simple robots, e.g., by ‘painting eyes’ onto its front, thus the effort required to gain the benefits of gaze is low.

Body language, specifically rotation, was the other behavioural cue that we found to be statistically insignificant. Unlike gaze, there were far fewer comments regarding rotation. Many participants ambiguously referred to the robot’s “movement”, a term which can also include the motion of the robot’s head or body, in addition to rotation. Based on these results and comments, we suspect that rotation was simply an unnecessary additive behavioural cue. Indeed, this cue was the least observed of the four cues used in the human-human interruption observational study in Chapter 4. As our robot could not simulate the hand and foot gestures used by some of our actors, rotation was used as a compromise to simulate some level of body language. Thus, we cannot conclude that every possible form of body language is statistically insignificant, though we feel that more complex forms of body language would require a greater engineering effort to gain the possible benefits of body language.

Of course, the use of rotation may have given rise to additional ambient noise from the robot's motors that would not have been present if the robot remained still. This ambient noise, as a product of all of the robot's physical movement, had noticeable impacts of its own, as described below

6.1.1 Impact of Ambient Noise

While statistical analysis of the quantitative data collected from the participants shows that speed and proximity had a significant impact on interruptedness, comments from many participants also mentioned ambient noise. This was an unintended cue in our study that arose as a side effect. The nature of the robot's design causes the noise from its motors to be higher as it moves at a faster speed. Thus, this noise alone may have had some impact on the interruptedness felt by participants, especially when combined with the fast lateral motion of the robot.

Indeed, many participants did note that they heard the robot before they could see it. Thus, ambient noise did play a role in interrupting the participants that we did not isolate entirely. It remains unclear how participants may have ranked interruptedness with fast speed of motion alone without the same level of ambient noise audible during slow speed.

In the next section, we move from discussion of the robot's physical behaviour to the social view of the robot.

6.2 Social View of the Robot

The results of the user study presented in Chapter 5 reveal many different social interpretations of the robot. Some participants viewed the robot for what it is – a robot. They used machine terms to describe its abilities (e.g. sensors). They suggested that

aspects of its behaviour were “programmed”. However, many participants saw the robot as more than just a machine and referred to it as a social being with its own desires, goals, and thought process.

In less than an hour after meeting the robot for the first time, participants were noting that they were already becoming familiar with the robot and its behaviours. Participants compared the robot to various social entities such as pets, small children, and even office workers or messengers. They described the robot as one might describe a *person* using similar behaviour. In these cases, the human comparison was not explicit, but it is easily inferred from the attribution of human characteristics to the robot. For instance, many participants noted that the robot “didn’t want to interrupt”, perhaps as a preference, rather than the result of some computer algorithm. All of these interpretations emerged from the robot’s behaviour, even though the robot’s appearance itself bears little resemblance to any person or animal.

6.2.1 Politeness and Emotional Response

Many participants surprisingly expressed how they felt emotionally about their interaction with the robot. Two participants said that one of the behavioural episodes (1C, slow, direct gaze, at the doorway) “didn’t scare” them. Another said the robot was “not very annoying” whereas it was “making a lot of noise” and “bothering” him during episode 2D (erratic gaze, fast speed, rotating, next to the participant). One participant felt more comfortable during episode 1C, whereas he had been previously worried that the robot might have hurt him in 2D.

The comments by participants showing interpretation of politeness in the robot’s behaviour provide confirmation that a robot can communicate interruption urgency in a

way that minimizes disruption. Thus, the robot is capable of behaving more socially appropriately. Clearly, there are cases, emergencies for example, where being ‘polite’ may not be important, so long as the person understands the message. However, minimizing disruption could be very important in cases where a robot is attempting to interrupt a busy person for a non-urgent or non-important (to the person) matter, e.g., the robot’s batteries will soon need to be recharged.

The participant comments discussed above reveal interpretations of the robot as more than just a machine running code or following remote instructions. They reveal the willingness of people to perceive the robot as a social entity.

6.2.2 Suspension of Disbelief

It is critical to cultivate an atmosphere of open-mindedness and “suspension of disbelief” for the study’s participants. The majority of phenomena being studied in our experiment relied on the implicit assumption that the robotic entity has an internal purpose and intent driving its actions, as opposed to merely being a procedural machine carrying out its programming. That is, unless the human observer perceives the robot as a social entity, there is no reason for them to interpret the robot’s motions as anything more than abstract patterns of movement. Instead, if the observer tries to ascribe intent to the robot’s behaviour, suddenly they are able to make a “leap of faith” and associate the robot’s motions with those social gestures they are already familiar with in daily life; suddenly there is a “why” (which is the social/emotional message) behind the “what” of the robot’s motions.

The comments made by our participants indicate a largely successful attempt in cultivating this atmosphere of open-mindedness and suspension of disbelief. Only two

participants actually suggested the robot was running through “programmed” behaviour or being controlled by the study administrator during the study. We actually expected more participants to make this conclusion, as the robot’s appearance did not imply any possession of adequate sensors or intelligence that would enable fully autonomous operation. The robot followed similar motion paths for each of its behavioural episodes in ways suggestive of pre-programmed code being run (e.g., the robot always entered the office in the precise same way if it was going to approach the participant). The robot did not actively respond to any stimuli initiated by the participant, such as talking or eye contact. Indeed, the robot would occasionally turn around and leave while the participant was staring at it, as the robot only remained stationary in the office for 15 seconds every time.

Instead of commenting on the robot’s limited abilities and inferring that the robot was not actually acting autonomously, participants ascribed social roles to the robot. This is important, as it is critical that people accept robots as social entities if they are expected to coordinate interruptions with them. Of course, for each social entity, there are expectations regarding what is considered to be “appropriate” and what behaviour may be misunderstood or considered “inappropriate” for certain situations.

6.3 Inappropriate behaviour

Many comments also emerged where participants suggested that the robot behaved “inappropriately”, particularly in cases where they could not interpret the behaviour as meaningful.

Of course, this information alone cannot be used to decide if the robot’s behaviour should be considered appropriate or not for all situations. Indeed, the suitability of a

robot's behaviour depends on the context of the interruption including factors such as urgency and importance. Appropriate behaviour to use in an urgent situation may (and probably does) differ from behaviour that should be used in a non-urgent situation. To get a sense of which behaviour is appropriate for each circumstance, we asked the participants which circumstance they thought was shaping the robot's behaviour. The underlying assumption here is that participants will do their best to identify circumstances for which the behaviour is most suitable, if there any at all. No information about the robot's reason for interrupting was communicated to the participant.

Behavioural episode 2D was designed to be the most extreme robot behaviour with the fastest movement at close proximity to the participant. Indeed, participants did find this behaviour to be the most interruptive overall according to the means presented in Table 5.5. However, it is possible that this behaviour was interruptive simply because it did not make sense to the participants. Some participants described this behaviour as "weird" or just plain annoying. Ideally, the robot's behaviour should seem somewhat "normal", especially to people who have been introduced to the robot and its abilities. Others thought that the robot was "broken" or "damaged", which isn't desirable because the participant was not tasked with repairing the robot if it were broken. Thus, it would not be appropriate for the robot to approach them with such a request. Indeed, this behaviour could be appropriate in a working environment where the robot approaches people with problems who have been trained to deal with these situations.

There are situations, however, that may be easily recognizable but "out of place" most of time. Thus, people may have preconceived notions of certain robotic behaviour that is inappropriate in almost any situation. Most participants associated the fast, erratic

head and body movement at close proximity with some type of emergency such as a fire, thus making it inappropriate for any other less critical scenarios to prevent any false alarms. It may be quite problematic if people misinterpret a mail delivery robot to be communicating a fire instead. Of course, ensuring appropriate behaviour includes not only the nature of the robot's actions but their timing as well.

6.4 Appropriate Timing of Interruptions

In our interruption user study, the robot interrupted the participants at unexpected intervals while they were engaged in conversation with the interviewer. Although the interruptions were timed to have a delay of a few minutes between each, they were not timed to coordinate with any specific state of the participant. The interviewer tried to maintain discussion with the participants using topics unrelated to the study, but sometimes the topics so engaging to the participants that they were actually less interested in talking about the robot as it appeared, at least until prompted by the interviewer. Thus, some interruptions sometimes occurred when the participant was more "busy" talking than other times.

Our study did not explicitly examine or isolate this factor, but its presence does raise *timing* as an issue to consider when choosing how to interrupt appropriately. While our robot was designed to interrupt without any regard to the timing (aside from ensuring delays between each interruption), a more socially appropriate approach could moderate the timing of the interruption based on the person's current activity and how interruptible they are.

In Chapter 2, we discussed how Gillie and Broadbent (1989) found that task similarity plays a significant role in the level of disruption caused by an interruption. That

is, an interruption will be more disruptive to tasks of a similar nature than to tasks that are dissimilar. Perhaps the similarity between the interruption and the task being interrupted should determine when an interruption occurs. Thus, we propose that the concept of bounded deferral (Horvitz et al., 2005) can be modified such that an interruption is opportunistically carried out when the participant transitions to a task that is similar to the interruption task as a necessary precondition, instead of a non-busy state. This approach, however, requires the interrupting system to have the ability to determine task type.

Research by Fogarty et al. (2005) has demonstrated that it is already possible, even practical, to determine a person's interruptibility state in a typical office environment. The methods used in this research could be adapted to determine the type of at least some tasks that are engaging a person, in order to make it possible to tailor an interruption so that it is similar to the interrupted task. For example, an audio sensor is used in the Fogarty et al. study to determine whether the person is engaged in conversation or not as one component in the measurement of busy state. If an interruption is designed to use audio or verbal cues, a system using our modified approach to bounded deferral could defer the interruption until the person begins a conversation, as detected by the audio sensor.

However, the bounded deferral approach is only effective when delayed awareness is acceptable. In cases where immediate awareness is required, careful consideration of the method of interruption becomes key at reducing disruption. The interruptibility sensors used in the Fogarty et al. study can still be used. For example, the same audio sensor that detects whether the person is engaged in conversation or not as a busy state measurement could be used to determine that an interruption should use the

auditory channel of conveyance. In this case, the interruption is perhaps less disruptive than an interruption that uses the visual channel of conveyance.

6.5 Emotive-based Coordination of Interruption

In Chapter 3, we discussed how we used an off-the-shelf commercially-available brain-computer interface to infer emotional stress. Although we feel that current technology is too limited to be useful in practical applications outside of research laboratories, the potential of future technology is promising for the application of emotion-sensing robots that coordinate interruption in socially appropriate ways.

6.5.1 Modifying Bounded Deferral

For our prototype, we programmed a Roomba to avoid a person when they are sensed to be emotionally stressed. When coordinating interruption, this same emotion could be used as a cue to delay an interruption, at least up to a certain amount of time, adding another modified approach of bounded deferral (Horvitz et al., 2005) to the one we discussed earlier in this chapter. In this case, the robot waits for the person to transition to a non-stressed emotional state, or until a predetermined amount of time has elapsed, whichever comes first.

6.5.2 Feedback Mechanism

Emotion-sensing could also be used as an implicit feedback mechanism for the robot's behaviour. Specifically, if the robot becomes aware of a change in emotion during its interruption consistently over a number of trials, it could feasibly attribute this change of emotion to the robot's behaviour itself, and not some other task. This is particularly important if the person's becomes scared or angry because of the robot's behaviour.

Indeed, there is a range of behaviour that almost everyone will find to be inappropriate in most situations. However, across different people, there is also a range of behaviour that some consider appropriate, and some do not. Our participants provide strong evidence to support this notion, as there was no single comment or impression that was echoed the same way by everyone. While a majority of participants associated the aggressive 2D behavioural episode with an emergency, some did not. Participant impressions of less aggressive episodes had even more variation. To maximize the chance of an atmosphere that cultivates suspension of disbelief, a social robot could adapt its behaviour toward different people in ways that work best for them, based on past feedback from implicit emotional responses.

As an example, suppose that a robot has an urgent interruption to deliver. To do this, it enters a person's office at fast speed, reaches a position next to the person, and starts moving its head rapidly, frightening the person in the process. This is likely to cause the person to develop a negative opinion of the robot, or worse if anger is also involved (e.g., physical abuse toward robots is not uncommon, as we discussed briefly in Chapter 2). This scenario could limit the robot's ability to effectively interrupt in the future, as the person may become distrusting of the robot and ignore it, or take steps to prevent its ability to interrupt (e.g., keeping the office door closed at all times). To maintain a positive impression with the person, the robot could use slower speed of motion, and instead rely on other behavioural cues to convey information about the interruption.

Clearly, there are many ways to incorporate emotion-sensing into a robot's interruption abilities. We feel that emotion-sensing technology is not yet feasible enough

for this type of application, but future developments could make it less invasive and more subtle. These developments could greatly improve a robot's ability to interrupt people in socially appropriate ways.

6.6 Summary

In this chapter, we discussed various behavioural cues that we tested in interruption scenarios as well as various methods of coordinating the interruptions themselves. First, we identified a few reasons why gaze and body language may not have been statistically significant behavioural cues. Second, we explored the view of the robot as a social entity that emerged by way of emotional responses from participants, and their perception of politeness in the robot's behaviour. This brought us to suspension of disbelief, a concept that we feel is critical to create an atmosphere where people accept the robot as a social entity instead of just a machine. Finally, we discussed the concept of behaviour that is appropriate, in terms of how it makes the participant feel and its timing.

In the next and final chapter, we will revisit our research questions, summarize our conclusions, and discuss possible future work that builds on our contributions.

Chapter Seven: Conclusions and Future Work

In this thesis, we presented an inquiry into the topic of socially appropriate nonverbal robotic interruption. First, we began our exploration with a review of the current state of the art (Chapter 2). Second, we described an early precursor to our work with interruption, a robot prototype that mediates its own behaviour to be less disruptive to people and thus (ideally) more socially appropriate (Chapter 3). Third, we narrowed our focus on interruption. Specifically, we described a methodological process for designing minimal robot behaviours for social interruption based on human-human interruption observations, then realized these behaviours on a robot via Wizard of Oz methodologies and robotic interaction implementations, and then designed an evaluation of those behaviours in a set of pilot studies and a final user study (Chapter 4). We explored the results of the user study (Chapter 5), and then discussed what we believe these results mean (Chapter 6).

7.1 Research Question, Revisited

In recognizing minimal non-verbal behavioural cues as an important yet largely unexplored layer of interaction between humans and robots, this thesis explored the following primary research question raised in Chapter 1.

Are there *minimal nonverbal behavioural cues* that robots can exhibit to communicate interruption urgency, and are those cues understandable by people?

7.2 Thesis Contributions, Revisited

This thesis makes the following three contributions, one primary and two secondary.

Primary:

To the best of our knowledge, the first academic exploration of nonverbal interruption in human-robot interaction.

To the best of our knowledge, this thesis represents the first attempt to explicitly justify and explore robots that interrupt people using nonverbal behavioural cues in a socially meaningful and acceptable manner. We created a design and evaluation process comprising four elements: (a) an observational study to see how people improvise their behaviour to interrupt others using a minimal subset of nonverbal cues over scenarios that vary in urgency, (b) a design critique of these behaviours when mimicked by a robot, (c) a robotic implementation of the behaviours which are triggered and somewhat controlled by a human operator, and (d) a user study, where we exposed people to these robotic behaviours, and gathered their reaction and interpretation of those behaviours. We found that people were able to infer urgency context about an interruption from a robot that is using only minimal non-verbal behavioural cues.

Secondary:

A methodology for probing interruption in HRI.

We contribute a simple yet powerful methodology to observe human behaviour in an interruption context, prototype a robot's behaviour using the nonverbal physical behavioural cues observed in the human behaviour, and evaluate how

people interpret these cues when used by a robot.

A research exploration of bioelectric signal interfaces in implicit human-robot interaction, where the robot is programmed to react to the person's implicit emotional state rather than to direct control.

To the best of our knowledge, the (vast) prior work in this domain concerns only direct brain-robot control. As far as we know, our emotional state exploration, while rudimentary, is the first attempt to justify and prototype the use of a brain-computer interface to infer a person's implicit emotional state and to mediate a robot's behaviour as a consequence.

7.3 Future Work

The research presented in this thesis hints at the potential benefits of applying techniques of nonverbal interruption behaviour to robots. We believe our efforts are novel, and contribute to the domain, however, our work is quite preliminary and thus limited in scope and in depth. There is still plenty of work that needs to be done to fully understand how to develop robots that interrupt people non-verbally in socially appropriate ways. Here, we present a few of the many possible research extensions to our work.

7.3.1 Investigating Additional Interruption Behaviour Cues

In our observational study in Chapter 4, we used a set of four minimal nonverbal behaviour cues in an interruption context: proximity to person, gaze and head movement, body language, and speed of motion. Of these four, we found that just two (speed of motion and proximity) were actually enough to convey interruption context (i.e., urgency). However, some robots will have other behavioural capabilities, such as ambient

status indicators and arm / leg movement, at their disposal. We would like to explore these other behavioural cues within the context of minimal non-verbal interruption.

We believe there may be additional depth regarding interruption context to be gained from the use of additional behavioural cues. Our prototype did not communicate specific information about the interruption (e.g., type of message), as we focused on the stage of initiating and coordinating an interruption. While many participants suggested the robot was acting as a messenger delivering a message of varying importance, no specific information about the message itself was provided. Behavioural cues could be used by a robot to not only assist in coordinating the interruption itself, but to communicate a greater depth of context about the interruption as well.

7.3.2 Coordinated Interruption

In Chapter 2, we discussed McFarlane and Latorella's (2002) taxonomy which provides four design solutions to coordinate interruption: (a) immediate, (b) negotiated, (c) mediated (through some other entity), and (d) scheduled. Our robot interruption user study in Chapter 4 followed the *immediate* solution, where the robot requires the person to interact immediately, or ignore it entirely without any communication to defer the interruption until later. We would be interested to move beyond basic behavioural cues and explore robots that coordinate interruptions using McFarlane and Latorella's other three solutions as well. We discuss our thoughts on future explorations of these three solutions below, in their social HRI interruption context.

First, a *negotiated* solution would have the robot announce its desire to interrupt, and then allow the person to deal with the interruption or defer it until later. This type of solution requires some dialog, and a level of feedback from the person, which could be

communicated directly through verbal commands, or physical gestures. Another option is implicitly inferred intent, where the person's desire to defer an interruption is inferred implicitly without any direct interaction. In Chapter 3, our example of an emotion-sensing robot roughly relates to a negotiated interruption design, where the robot moves away from the person based on their current state or context. In an interruption context, a robot could infer the person's desire to defer an interruption when the person is busy. However, bypassing explicit intent may result in undesirable interruption coordination behaviour for some of the time. Indeed, there may be times when a person who is very busy would have not chosen to defer an interruption, particularly if it is important and relevant to their current task. We would be interested to study this type of negotiated situation across a variety of scenarios.

Second, a *mediated* solution would allow the robot to interrupt indirectly through a mediator object. McFarlane and Latorella suggest that an object such as a PDA or desktop computer notifier can decide when and how to perform an interruption. This approach could be useful for robot implementations that do not possess any decision-making capacity on their own regarding interruption coordination. For example, it may be feasible for a desktop computer to act as a mediator object if it can mediate the robot's need to interrupt with a person's level of interruptibility and current task type (e.g., web browsing, sending e-mail). This approach evolves robots from acting as mediator objects themselves to social entities that are expected to infer activities of other social entities (humans) from other mediator objects, such as PDAs or desktop computers, which themselves are humble computing entities that cannot move. We would be interested in studying the feasibility and intuitiveness of this design solution.

Third, a *scheduled* solution would restrict the robot's interruptions to a prearranged schedule, such as once every 30 minutes or every day at 9 o'clock and 1 o'clock. Although it may be most ideal to deliver any interruption when the person is the least busy, this approach could be used by a robot to deliver interruptions when information about the person's busy state is not available.

Clearly, an interruption can be coordinated by robots in many ways. McFarlane and Latorella's three methods of coordinating interruption explored above present many opportunities for future exploration. We are interested in studying these methods of coordination in more detail to learn their strengths and weaknesses when applied by robots in interruption contexts. We are also interested in comparing the different methods of coordination and their impact across a variety of settings and environments.

7.3.3 Interruption in Different Environments

We are interested in studying robots that are designed to interrupt people in different types of working environments. Our user study in Chapter 4 was situated in a simulated office environment. For all of the robot's interruption attempts, the participant was seated in the same location engaged in the same activity (i.e., conversing with the interviewer).

Of course, not all robots will operate in a similar environment. Specifically, we would be interested in exploration the use our minimal nonverbal interruption behavioural cues in noisy environments such as factories or shopping malls. In these environments, verbal and sound cues could be sometimes rendered inaudible and much less effective due to background noise.

Groups of people in meetings together are also prevalent in many environments. This adds difficulty to the design of a robot designed to interrupt one individual. If an

individual to be interrupted is within a group of people, the robot must be able to attract the attention of the person to be interrupted with minimal disruption to the people not being interrupted. The robot must be able to convey a suitable amount of information about the reason for interrupting while also respecting the privacy of the person being interrupted if the issue is sensitive.

7.3.4 Alternate Physical Forms

In Chapter 4, we briefly discussed our decision to use a minimal physical form for our interruption user study. In Appendix B, we explore our limited implementation of a teddy bear robot form. While the use of a teddy bear appearance may be desirable for many circumstances, we found that its use created many technological issues that hindered its ability to interrupt.

We feel that there is room for further exploration into the impact of a robot's physical form within the context of an interruption. Other robot forms could vary drastically in size and appearance. Exploration of whether these differences could have an impact on how a person perceives the robot's behaviour. These differences could be an important consideration when using a drastically different robot form (e.g., a robot which is very large in comparison to a person could be more intimidating than a much smaller robot, even if both are using the same behaviour).

The most appropriate type of physical form to use depends on its environment and the reasons it will have for interrupting. In a factory working environment, a robust, practical robotic appearance may be most appropriate, as the robot cooperates with people as a worker. In a home environment, a robot may operate more as a social

companion rather than a worker. Thus, a more social robot appearance, such as a dog or other pet, may be more inviting to people interacting with it.

In future work, we would be interested to study the different impact of a robot's interruption behaviour when using different physical form.

7.4 Autonomous Implementation

In our robot implementation that we describe in Chapter 4, we used remote controls, Wizard of Oz techniques, and pre-programmed behavioural sequences to move the robot. No sensors or communication with the participant were used at any point by the robot. This approach was effective for a research prototype implementation but is obviously not valid for final prototypes and for deployment.

We believe that an autonomous robotic which is interacting via non-verbal interruption cues and does not rely on humans through remote control or pre-programmed sequences, is feasible. For interruption, we believe this approach is the only feasible one for most applications. However, autonomous robots require an extensive engineering effort, far beyond our comparatively simple robot implementation. On top of the physical capabilities that our robot had, an autonomous robot would also require various safety and sensory features to work properly.

Safety is a primary concern with autonomous robots. Our participants were always in a seated position, so the risk of our robot colliding with people was minimal, even when moving at fast speed. Safety was designed into the procedures of our study, where the participants were always kept out of the robot's way by remaining seated, and the robot's motion path was specifically designed to avoid collisions with the walls and doorway in the study area. However, a robot that moves around an environment

autonomously would need to have a robust obstacle detection system to avoid collisions with walls, objects and most importantly, people. Even a single collision with a person causing injury could make use of the robot unfeasible in most environments.

Beyond safety and obstacle detection, an autonomous robot needs a robust navigation system. Navigation is necessary for the robot to know how to get from its current location to a chosen destination. The robot could navigate by using algorithms that analyze its current environment in real time, requiring the robot to have sufficient sensors to do this, or by using markers, which have been placed in the environment (e.g., RFID tags, etc.) to find the robot's location on a stored map.

An autonomous robot that interrupts also needs communication abilities to receive commands to initiate interruptions, and coordinate those interruptions with the participants. This communication could be accomplished in many ways, and multiple methods of supporting communication would be preferred. For example, a robot could recognize both voice recognition and hand gestures, which provides backup input method for people to use in case one does not work.

In future work, we would be interested to see robots that are designed to behave more autonomously in less controlled conditions to interrupt people.

7.5 Generalizing to Other Cultures and Environments

As we conducted our studies in a North American office environment, one may question how these results can be generalized to other working environments and other cultures. These questions can, of course, be addressed by performing more studies.

However, we believe our results, both quantitative and qualitative, illustrate a parallel to studies by Nass and Reeves (1996) describe in *The Media Equation*, where

people were found to treat media items (such as computers) as living things. Our quantitative results have shown that it is possible to use a minimal set of cues (i.e., speed of motion and proximity to a person) to convey urgency context about an interruption. Our qualitative comments by participants reveal an overwhelming social view of the robot with descriptions ranging from comparisons to animals and politeness, similar to how people might describe other living beings. We believe that these results seem to uphold the findings by Nass and Reeves.

Thus, and following the Media Equation parallel, we think that in order to predict how people might interpret robotic interruption behaviour in a specific context, it may be feasible to look at how people interpret human interruption behaviour, and expect that they will interpret robots using similar behaviour in the same way. We believe that to a large extent this lesson is general and will be applicable across cultures and various settings, but of course further studies will be needed to support this assumption.

7.6 Final Words

In this thesis, we have described the design and implementation of robotic interfaces that interrupt people using minimal nonverbal behavioural cues. The evaluations we conducted demonstrated that the use of minimal cues, such as proximity to the person and speed of approach, was valid in various interruption contexts. We hope that the research described in this thesis defines a starting point for designing robots that interrupt people using only minimal elements of their interruption behaviour. We also hope it will motivate others to build upon and extend our work to further explore how robots can interrupt people using only their physical behaviour.

APPENDIX A: STUDY MATERIALS

This appendix contains materials related to the user study described in Chapter 4 of this thesis. It includes:

- The informed consent form given to participants who participated in this study.
- The experimental protocol, which describes the actions taken by the experimenters while they administered the study.
- The interview questions used in the first phase of the experiment
- The general discussion topics used between interviews in the first phase of the experiment.

A.1. Informed Consent Form



Name of Researcher, Faculty, Department, Telephone & Email:

Paul Saulnier, M.Sc. Graduate Student, Faculty of Science, Department of Computer Science, University of Calgary
 Phone: 403.210.9499 Email: p.saulnier@ucalgary.ca

Ehud Sharlin, Faculty of Science, Department of Computer Science, University of Calgary
 Phone: 403.210.9404 Email: ehud@cpsc.ucalgary.ca

Saul Greenberg, Faculty of Science, Department of Computer Science, University of Calgary
 Phone: 403.220.6087 Email: saul.greenberg@ucalgary.ca

Title of Project: Investigating the effects of various robot behaviours in human task interruption scenarios

Sponsor: NSERC, iCORE

This consent form, a copy of which has been given to you, is only part of the process of informed consent. If you want more details about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information. The University of Calgary Conjoint Faculties Research Ethics Board has approved this research study.

Purpose of the Study

This study is designed to measure how variances to a robot's unique abilities can affect how its behaviour is interpreted.

What Will I Be Asked To Do?

The robotic entity will perform a series of behaviours and with each demonstration you will be interviewed as to your interpretation of the entity's behavior and intent. Specific instructions will be provided to you at the beginning of the study and you will have the opportunity to ask questions throughout the study. The study and immediate follow-up will take approximately 1 hour. Following the active study participation, you will be interviewed about your experience.

Your participation is entirely voluntary. You may decline participation in parts of the study or withdraw from the study without penalty or loss of benefits.

What Type of Personal Information Will Be Collected?

Should you agree to participate, you will be asked to provide demographic information (e.g., gender, age, education). We will be observing your actions as well as videotaping you during the course of the session. Videotaping is mainly done because it is difficult for us to observe everything the first time. We often discover things by watching the videos later that we have overlooked during your session.

Additionally, a study administrator will be transcribing what you say during the study. The study administrator has signed a confidentiality agreement and will enjoy no access to any of the study data for their personal, academic or research purposes.

Are there Risks or Benefits if I Participate?

If you choose to participate in this study, you are not expected to encounter any foreseeable personal risks. You will receive \$15 remuneration if you choose to participate at the beginning before the study begins.

What Happens to the Information I Provide?

Your participation in this research is confidential. Only the researchers listed above will have access to your data. No

information that discloses your identity will be released or published without your specific consent to disclosure. We may need to cite your comments or your answers from the interviews; in this case we will cite you anonymously.

We might want to use clips or stills of the video for its use in presentations or other electronic media, but this can only happen with your consent. Please, indicate if you grant us permission to use video clips or still pictures of you. Please note that once photographed or videotaped images are displayed in any public forum, the researchers will have no control over any future use by others who may copy these images and repost them in other formats or contexts, including possibly on the internet.

We intend to use your data as part of ongoing MSc program research.

I grant permission to the researchers listed above to use video clips with sound and image of myself for public presentation or in other electronic media: Yes: ___ No: ___

All data obtained from this study will be stored in a locked cabinet and any electronic information will be stored on a computer only accessible through the use of a password. Collected data will be accessible only to the researchers named on this consent form. All data will be stored for a period of time no longer than three years. Information will be carefully disposed of (shredding for hard copies and deleting for electronic copies) when this investigation is complete.

You will be able to withdraw from this study at any point. If this occurs, any data collected up to that point about you will be discarded. You are also able to refuse to answer whatever questions you prefer to omit.

Signatures (written consent)

Your signature on this form indicates that you 1) understand to your satisfaction the information provided to you about your participation in this research project, and 2) agree to participate as a research subject. In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. You are free to withdraw from this research project at any time. You should feel free to ask for clarification or new information throughout your participation.

Participant's Name: (please print) _____

Participant's Signature _____ Date: _____

Researcher's Name: (please print) _____

Researcher's Signature: _____ Date: _____

Questions/Concerns

If you have any further questions or want clarification regarding this research and/or your participation, please contact:

Mr. Paul Saulnier
Department of Computer Science/Faculty of Science
P. Saulnier – Phone: 403.210.9499, Email: p.saulnier@ucalgary.ca

Dr. Ehud Sharlin,
Department of Computer Science/Faculty of Science,
Phone: 403.210.9404, Email: ehud@cpsc.ucalgary.ca

Dr. Saul Greenberg,
Department of Computer Science/Faculty of Science,
Phone: 403.220.6087, Email: saul.greenberg@ucalgary.ca

If you have any concerns about the way you've been treated as a participant, please contact the Senior Ethics Resource Officer, Research Services Office, University of Calgary at (403) 220-3782; email rburrows@ucalgary.ca.

A copy of this consent form has been given to you to keep for your records and reference. The investigator has kept a copy of the consent form.

A.2. Setting

The participant is seated in an office of a typical Computer Science laboratory on a couch with the open doorway in view. An interviewer (an actor) is in the room conversing with the participant. The robot is initially positioned outside the doorway. From there, it can pass by the doorway or enter the office. A video camera is positioned to record both the robot's behaviour in the doorway as well as the participant's reaction to this behaviour. Figure A.1 illustrates the setting used for the study.

The robot is stationed 100cm away from the doorway horizontally and up 145cm vertically. Inside the office are various obstacles as indicated in the diagram. The participant will be seated approximately 230cm away from the doorway. Once inside the office, the robot is free to move around in the space between the participant and the doorway as well as 170cm horizontally once past the initial obstacle. These measurements are arbitrary and based on the current environment available for study use.

During normal operation, the robot should return to its initial position in order for the automated macros to execute properly in subsequent scenarios. If this does not happen due to a malfunction, a marking on the floor indicates where the robot should be manually repositioned by the study administrator before executing any pre-programmed behaviour macros.

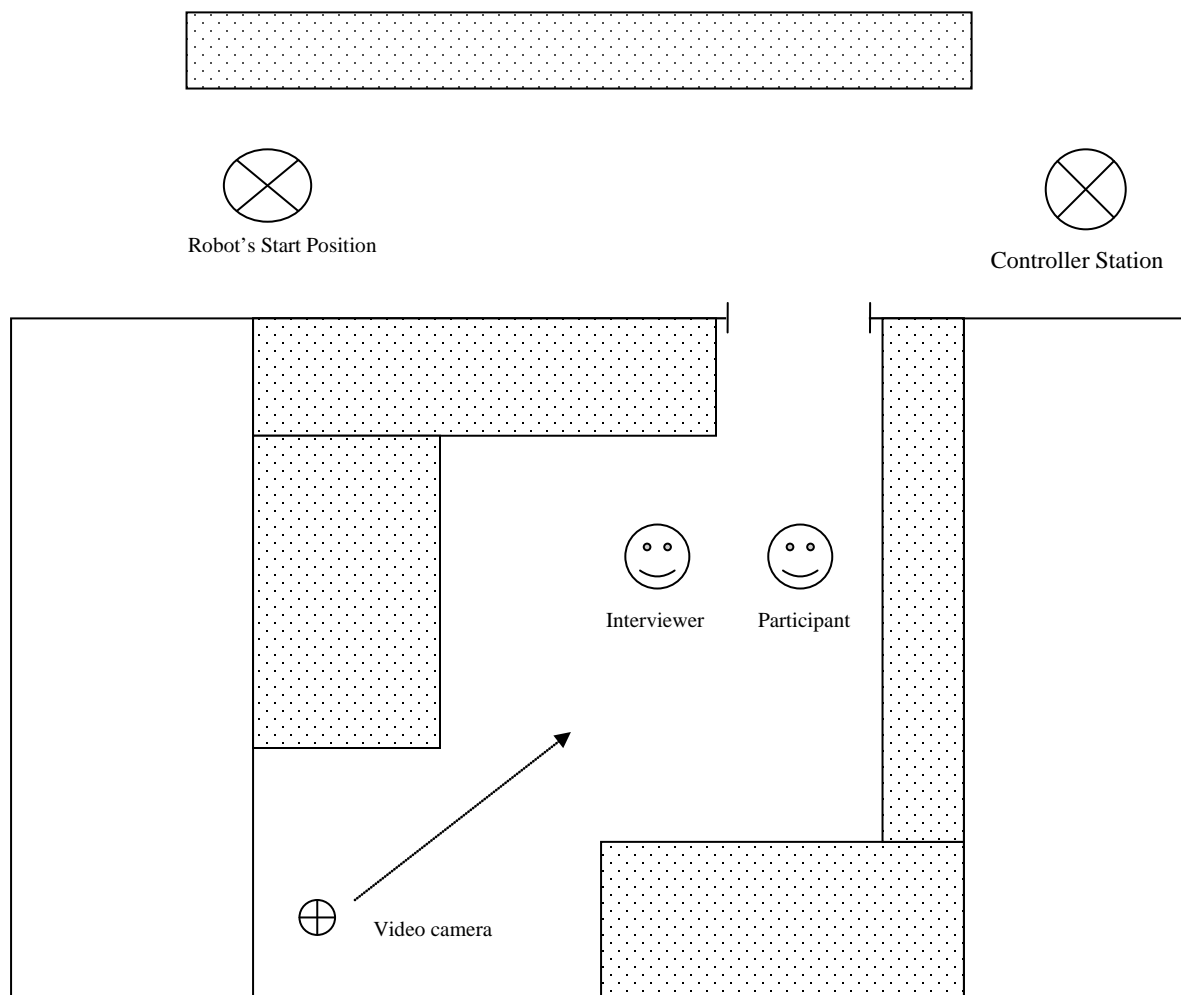


Figure A.1: Floor layout of user study.

A.3. Experimental Protocol

This section includes the script used by the interviewer as a guideline when communicating with the participant.

Introducing the Participant

“I am a researcher with the Interactions Lab in the Department of Computer Science. My research primarily deals with human-robot interaction, that is, how robots can interact with people. More specifically, I am interested in exploring how robots can interrupt people in a workplace atmosphere, such as an office. Shortly, we will move into the office, and have a casual discussion about different things. This will be the start of the first phase of the study. While we are doing this, the robot that you can see in the corner may try to approach you and interrupt our conversation in different ways, as if it was going to deliver you a message. For the purposes of the study, the robot will not talk, respond to you in any way, or actually deliver a message. While this is happening, I will invite you to talk about your experience and tell me what you think about what the robot is doing. The second phase of the study is very similar, with the robot still approaching you, but instead of interviewing you, I will ask you to rank the robot’s behaviour. I will tell you more about that when the time comes. Before we begin, I ask that you read this consent form and sign it. This study has been approved by the University of Calgary Ethics Board, so this is a formality to ensure you understand what you will be asked to do in this study. We will be recording you using a video camera for the duration of the study. Additionally, our study administrator will be transcribing what you are saying into the computer. Let me know if you have any questions.”

The consent form is signed.

Excellent. Let’s go into the office now.

Beginning of Phase 1

“Once again, this is a very casual atmosphere. Please feel free to think out loud and tell me what you think of what the robot is doing when it appears, particularly in the context of it trying to interrupt you. Tell me what circumstances might lead a robot to interrupt you in such a way. As the robot appears, I will ask you some questions before we return to our casual conversation. (*Video recording begins.*) Now the video camera is recording us and our study administrator is transcribing what you are saying.”

After each interaction, the interviewer asks questions as per section A.5 for a maximum of 2-3 minutes. Between interviews, the interviewer and participant engage in casual discussion starting from topics discussed in section A.4.

Transition from Phase 1 to Phase 2

“This concludes phase 1 of the study. Thank you for your assistance thus far. Your experience is very helpful for our study.”

Beginning of Phase 2

“We will begin phase 2 now, which is similar to phase 1. Instead of interviewing you with specific questions after each interaction with the robot, I will provide you with this Interruptedness Metre to rank how interrupted you feel by the robot’s behaviour. The space along this metre represents a spectrum of interruptedness ranging from barely interruptive on this side to extremely interruptive on this side. For each interaction with the robot, you will receive a marker which you will place on the metre in a spot you believe appropriately represents how interrupted you feel by the robot’s behaviour. There are no right or wrong answers.”

“In this phase, you will experience more interactions with the robot with less delay between each. As you observe the behaviours, please continue to talk about your experience, especially in the context of the robot using the interaction to interrupt you from some other task.”

“Do you have any questions before we proceed?”

Interviewer engages in casual discussion again with the participant while waiting for the first robot interaction to occur. After the interactions start, they will occur one after another with only a short delay in between to allow the study administrator to properly reposition the robot. While the participant is placing markers on the Interruptedness Metre, the interviewer can ask questions as per section A.5.

Conclusion of Phase 2

“This concludes the user study. I just have some questions for you to reflect on the whole study. (*Interviewer asks questions as per section A.5.*)”

Ending the Study -Full Disclosure

“Excellent. Now, I am going to share some of the study details with you. In addition to transcribing, our study administrator was also controlling

the timing of the robot's interruptions and ensuring it does not crash into anything. Each interaction was different in some way from the others, usually based on the position of the robot, speed of movement or type of head movement. We are interested in how changing these variables affects how people interpret the overall interruption. Your comments and participation in this study are very helpful in our research. Thank you for your involvement."

A.4. Conversational Topics

These questions are used to prime the conversation between the interviewer and participant while waiting for an interaction with the robot:

- What are you studying?
- What projects have you enjoyed in your degree? Not enjoyed?
- How long have you lived in Calgary?
- What do you like about Calgary?
- What do you do for fun in Calgary?

A.5. Interview Questions

These questions are used to prime the interviews occurring through the study. Other questions may be opportunistically asked by the interviewer to gain a better insight of the participant's experience. The questions below are grouped into sections corresponding to when they are used during the study.

Phase 1 - After Each Robot Interaction

- How would you describe your experience with the robot just now? Tell me as much as you can.
- What circumstances do you think may have caused the robot to behave this way?

During Phase 2 Marker Placement

- How would you rank this behaviour?
- Why are you placing the marker there?

At the End of Phase 2

- Can you offer your reflection of all the robot behaviour you have seen during the study?
- What aspects of the robot's behaviour most affected your experience with the robot?
- Do you any final remarks or comments regarding the study?

Other questions may be asked that relate to any observed anomalies, etc.

A.6. Robot Interaction Descriptions

This section describes the interaction behaviours and their variations used by the robot throughout the study. Each participant experiences these interactions in a scrambled order, which is kept the same for all participants.

Null Base Case

- This always occurs first before all other interactions
- Robot rolls past doorway from a distance at slow speed until it is out of view of participant
- No stopping, no head movement, no looking into office

Base Case 1

- Robot rolls toward proximity position at speed
- Looks toward participant.
- Waits 15 seconds
- Looks forward again, moves away from participant
- Continues rolling out of view of participant

Fixed Conditions

- Gaze (toward participant)
- Rotation (none)

Variables

- Speed of motion (slow, fast)
- Proximity position (far from doorway, in doorway, next to participant)

Variations of Base Case 1

- Speed: slow, Proximity position: far from doorway
- Speed: fast, Proximity position: far from doorway
- Speed: slow, Proximity position: in doorway
- Speed: fast, Proximity position: in doorway
- Speed: slow, Proximity position: next to participant
- Speed: fast, Proximity position: next to participant

Base Case 2

- Robot approaches participant in office at speed
- Head movement is defined by specified gaze
- Robot rotates back and forth in place at speed with gaze. This emphasizes whether the robot is looking at the participant or not.
- Wait 15 seconds until stopping the robot in place
- Study administrator recalls robot after interview is done

Fixed Conditions

- Proximity position (next to participant)
- Rotation (rotating in place)

Variables

- Speed of motion (slow, fast) for both robot and head
- Gaze (locked on participant, various directions)

Variations of Base Case 2

- Speed: slow, Gaze: locked on participant
- Speed: fast, Gaze: locked on participant
- Speed: slow, Gaze: various directions
- Speed: fast, Gaze: various directions

A.7. Instructions for Study Administrator Tasks

Main Goals

- Execute robot behaviour macros at the desired time
- Ensure the robot returns to correct spot marked on the floor of the study space
- Transcribe the participant's remarks during the interviews
- Bringing the Interruptedness Metre to the interviewer after phase 1 is complete
- Generally helping to administer the study.

Phase 1

This phase, lasting about 20 minutes, includes 4 interactions with the robot that are each followed by interviews.

1. Ensure the robot is correctly positioned on the floor and facing forward.
2. When the interviewer/participant discussion becomes casual as per the conversational topics in section 10.4, page 52, select the appropriate macro following the order on your scenario order sheet.
3. Transcribe as much of what the participant says as possible. Do this until the interview returns to casual discussion.
4. Correct the robot's position if necessary.
5. Wait 2-3 minutes before beginning the next robot interaction. Try to select a time when the interviewer and participant are actively talking.
6. Repeat these steps until all 4 interactions with the robot are complete.

Phase 2

This phase, lasting about 20 minutes, includes 11 interactions with the robot which are each ranked by the participant using the Interruptedness Metre.

1. Ensure the robot is correctly positioned on the floor and facing forward.
2. Wait until the interviewer has explained phase 2 to the participant.
3. Run the desired macro according to your scenario order sheet.
4. Transcribe as much as possible of what the participant says about the robot's behaviour and their experience until the robot returns to its starting position.
5. After the robot returns to its starting position, correct its position, and immediately begin the next robot interaction, even if the participant and interviewer are talking about the study.

End of Study

After phase 2, the interviewer will conduct a final interview with the participant. Please transcribe as much of this as possible.

Malfunction and Recovery

Robot Collision

If a malfunction occurs (i.e. robot runs into an obstacle), click Macros > Stop Running All Macros and manually replace the robot to its starting position using the laptop's arrow keys or your hands. Re-run the scenario again or follow the instructions of the interviewer.

Incorrect Macro Executed

If you accidentally run the wrong macro, make a note of this by indicating which macro was run, and resume the correct scenario order after the incorrect scenario is complete.

Restarting Robot

If the robot fails to respond to arrow key or other commands, do this:

1. Close the Controller on both the laptop and VNC window
2. Open Controller again on the VNC window and click Robot > Connect to Robot
3. Open Controller on the laptop and click Robot > Connect to TCP Server
4. If this still fails, restart the robot using its side power switch and repeat steps 1-3

If these steps do not work or something else happens, enlist the help of the interviewer.

APPENDIX B: TEDDY BEAR ROBOT FORM

In Chapter 4, we explained our decision to use a minimalist physical robot form for our user studies. In this appendix, we explore a teddy bear physical robot form prototype as a side project that was begun as an attempt to consider form, but was abandoned for many reasons. It is included here not so much as a definitive statement about the use of form, but to illustrate some of the difficulties encountered.

Our choice of this form was based on the ease of obtaining a large teddy bear, and the ease of modifying it for use in the implementation of a robot. The teddy bear (Figure B.1, left) measures approximately 120 cm tall and 70 cm wide. We implemented a set of sensors and physical abilities to turn the teddy bear into a robot that could be used to interrupt a person.

B.1. Implementation

The implementation consists of four main components: a custom motor assembly which moves the teddy bear's head, a set of touch sensors placed at different parts of the bear's body, a speaker for producing audio feedback, and custom controller software that makes it all work together.

Before any modifications could be made to the teddy bear, we had to first open it along its seam in the back, and remove some of the stuffing material. We replaced the back seam with a zipper to enable easier access to the bear's interior later. About half of the teddy bear's stuffing was removed from its body and head to make room for electronics to control the sensors and head movement.

The robot head consists of a custom motor assembly (Figure B.1, right) that is capable of moving with two degrees of freedom (left/right and up/down). The assembly

consists of motors that are driven by a controller circuit board as well as a framework of short pipes, metal arms, and coat hangers that move the head as the motors turn. Our first attempt at the motor assembly was not powerful enough to actually lift the weight of the robot's head upwards. The motion used was also quite jerky at times, reducing its realism. We revised this part of the implementation to use more powerful motors that could move with more fluid movement with sufficient strength to move the head with ease.

To potentially enable some level of interaction, the robot was equipped with Phidgets Inc. touch sensors in its two "hands" (i.e., paws of the teddy bear) and nose. When its hands are touched, the robot turns its head toward the hand that was touched. When its nose is squeezed, the robot emits a honking noise through the audio speaker.



Figure B.1: Teddy Bear robot form.

Left: Exterior view

Right: Motor assembly removed from head (normally obscured from view inside head)

Although these particular mappings may not be practical for interrupting people or interacting with them in all contexts, they demonstrated the kind of responses that are possible with this particular form.

At this point, we abandoned development of this robot form as part of our current work due to numerous challenges that inhibited its usefulness in interruption. We will discuss these challenges below.

B.2. Discussion

Although we did not perform a user study using the teddy bear prototype, its implementation does reveal that it is much more difficult to implement a working and viable robot form actor, such as a teddy bear, vs. putting a few rudimentary features on an existing robot, as we did with our minimalist robot form in Chapter 4. In particular, we encountered three technical challenges that prevented our teddy bear form from being viable.

First and foremost, the complexity of animatronics engineering is a big problem. While many large institutions with considerable resources (e.g., Disney, Sony with the AIBO robot dog, etc.) are able to successfully work in this area, it is more difficult for amateurs or researchers without specific engineering skill sets. We believed these circumstances contribute to the following two problems that we encountered.

The second problem is that physical form implies certain functionality. For example, using a legged object, such as a teddy bear, implies mobility such as walking. This is much harder and more expensive than wheeled mobility.

Finally, side effects are introduced if animatronics are not done properly. In our case, the motor noise was distracting, possibly due to the kinds of motors we chose. This

noise intensified as the robot moved its head faster, potentially drawing all attention away from the head movement toward the noise only. Another side effect was a possible uncanny valley effect due to unrealistic motion of the robot's head. The motion used by the motors was difficult to animate in a precise way, and failing to do so diminished the robot's realism.

Overall, the challenges we encountered were not unusual or particularly insurmountable. We felt use of the teddy bear robot form did benefit from a novelty factor, even for many of our lab colleagues who are already familiar with seeing robots. Due to its size, the teddy bear robot was very eye catching and hard to miss.

As mentioned in Chapter 7, we would be interested to see future work that explores other robot physical forms within the context of interruption.

References

1. C. Breazeal. Social Interactions in HRI: The Robot View. *IEEE Transactions On Systems, Man, And Cybernetics—Part C: Applications And Reviews*, 34(2):181, 2004.
2. A. Chapanis and CM Overbey. Studies in Interactive Communication: III. Effects of Similar and Dissimilar Communication Channels and Two Interchange Options on Team Problem Solving. *Perceptual & Motor Skills*, 38(2):343–374, 1974.
3. S. Cohen. After Effects of Stress on Human Performance and Social Behavior: A Review of Research and Theory. *Psychological Bulletin*, (88):82–108, 1980.
4. N. Dahlback, A. Jonsson, and L. Ahrenberg. Wizard of Oz studies—why and how. *Knowledge-Based Systems*, 6(4):258–266, 1993.
5. K. Dautenhahn, M. Walters, S. Woods, K. L. Koay, C. L. Nehaniv, A. Sisbot, R. Alami, and T. Siméon. How may I serve you?: A robot companion approaching a seated person in a helping context. In *HRI '06: Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pages 172–179, New York, NY, USA, 2006. ACM.
6. S. Fallon. "Modified Roomba Detects Stress, Runs Away When It Thinks You Might Abuse It." *Gizmodo*. 16 02 2009. Web. 18 Dec 2010.
<<http://gizmodo.com/5170892/modified-roomba-detects-stress-runs-away-when-it-thinks-you-might-abuse-it>>.
7. J. Fogarty, S.E. Hudson, C.G. Atkeson, D. Avrahami, J. Forlizzi, S. Kiesler, J.C. Lee, and J. Yang. Predicting human interruptibility with sensors. *ACM Transactions on Computer-Human Interaction*, 12(1):119–146, 2005.
8. J. Fogarty. Constructing and evaluating sensor-based statistical models of human interruptibility. PhD thesis, Pittsburgh, PA, USA, 2006. Adviser-Hudson, Scott E.
9. J. Forlizzi and C. DiSalvo. Service robots in the domestic environment: A study of the Roomba vacuum in the home. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pages 258–265. ACM, 2006.
10. T. Gillie and D. Broadbent. What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research*, 50(4):243–250, 1989.
11. K. Grifantini. "A Robot That Knows When to Back Off." *Technology Review*. MIT, 13 03 2009. Web. 18 Dec 2010.
<<http://www.technologyreview.com/blog/editors/23104/>>.
12. K. Hayashi, D. Sakamoto, T. Kanda, M. Shiomi, S. Koizumi, H. Ishiguro, T. Ogasawara, and N. Hagita. Humanoid robots as a passive-social medium: a field experiment at a train station. In *HRI '07: Proceedings of the ACM/IEEE international conference on Human-robot interaction*, pages 137–144, New York, NY, USA, 2007. ACM.
13. S.M. Hess and M.C. Detweiler. Training to reduce the disruptive effects of interruptions. In *Human Factors and Ergonomics Society Annual Meeting Proceedings*, volume 38, pages 1173–1177. Human Factors and Ergonomics Society, 1994.
14. E. Horvitz, J. Apacible, and M. Subramani. Balancing awareness and interruption: Investigation of notification deferral policies. *UM*, 2005:433–437, 2005.

15. J. G. Kreifeldt and M. E. McCarthy. Interruption as a test of the user-computer interface. In *Proceedings of the 17th Annual Conference on Manual Control*, pages 655–667, 1981.
16. A. Matsuzaki, H. Kobayashi, H. Tomita, K. Kikuchi, K. Inoue, Y. Kondo, M. Hirayasu, and A. Ichise. Message delivery system operable in an override mode upon reception of a command signal, December 12 1989. US Patent 4,887,152.
17. D. Maulsby, S. Greenberg, and R. Mander. Prototyping an intelligent agent through Wizard of Oz. In *CHI '93: Proceedings of the INTERACT '93 and CHI '93 conference on Human factors in computing systems*, pages 277–284, New York, NY, USA, 1993. ACM.
18. D.J. McFarland and J.R. Wolpaw. Brain-Computer Interfaces for the Operation of Robotic and Prosthetic Devices. *Advances in Computers*, pages 169–187, 2010.
19. D. C. McFarlane and K. A. Latorella. The scope and importance of human interruption in human-computer interaction design. *Hum.-Comput. Interact.*, 17(1):1–61, 2002.
20. A.N. Meltzoff, R. Brooks, A.P. Shon, and R.P.N. Rao. “Social” robots are psychological agents for infants: A test of gaze following. *Neural Networks*, 23(8-9):966–972, 2010.
21. J. Millan, F. Renkens, J. Mourino, and W. Gerstner. Noninvasive brain-actuated control of a mobile robot by human EEG. *Biomedical Engineering, IEEE Transactions on*, 51(6):1026–1033, 2004.
22. M. Mori. Bukimi no tani [the uncanny valley]. *Energy*, 7(4):33–35, 1970.
23. B. Mutlu and J. Forlizzi. Robots in organizations: the role of workflow, social, and environmental factors in human-robot interaction. In *HRI '08: Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*, pages 287–294, New York, NY, USA, 2008. ACM.
24. C. Nass, Y. Moon, B.J. Fogg, B. Reeves, and C. Dryer. Can computer personalities be human personalities? In *Conference companion on Human factors in computing systems*, pages 228–229. ACM, 1995.
25. S. Nishio, H. Ishiguro, and N. Hagita. Geminoid: Teleoperated android of an existing person. *Humanoid robots-new developments. I-Tech*, 2007.
26. I.R. Nourbakhsh, J. Bobenage, S. Grange, R. Lutz, R. Meyer, and A. Soto. An affective mobile robot educator with a full-time job. *Artificial Intelligence*, 114(1-2):95–124, 1999.
27. R.W. Obermayer and W.A. Nugent. Human-computer interaction for alert warning and attention allocation systems of the multimodal watchstation. In *Proceedings of SPIE*, volume 4126, page 14, 2000.
28. OCZ Peripherals. Neural Impulse Actuator. Retrieved December, 2008. <http://www.ocztechnology.com/products/ocz_peripherals/nia-neural_impulse_actuator>
29. E. Paulos and J. Canny. PRoP: personal roving presence. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, page 303. ACM Press/Addison-Wesley Publishing Co., 1998.
30. C. Nass and B. Reeves. *Media Equation: How people treat computers, television, and new media like real people and places*. CSLI Publications, Center for the Study of

- Language and Information Leland Stanford Junior University, Cambridge, UK, first paperback edition, 1996
31. P. Riley. Discursive and communicative functions of non-verbal communication. ED143217/FL008785, ERIC (Educational Resources Information Center), Sewell, NJ, 1976.
 32. S. Satake, T. Kanda, D.F. Glas, M. Imai, H. Ishiguro, and N. Hagita. How to approach humans?: Strategies for social robots to initiate interaction. In HRI '09: Proceedings of the 4th ACM/IEEE international conference on Human robot interaction, pages 109–116, New York, NY, USA, 2009. ACM.
 33. P. Saulnier, E. Sharlin, and S. Greenberg. Using Bio-electrical Signals to Influence the Social Behaviours of Domesticated Robots. In Adjunct Proc. Human Robot Interaction (Late Breaking Abstracts) - HRI'09, page 2 pages plus poster, San Diego, California, March 11-13 2009. Late-breaking abstract. Includes poster shown at conference.
 34. P. Saulnier, E. Sharlin, and S. Greenberg. Exploring interruption in HRI using wizard of oz. In DVD Proc. 5th ACM/IEEE Int'l Conf on Human-Robot Interaction - HRI'2010, page 2 pages, Osaka, Japan, March 2-5 2010. IEEE/ACM. Late Breaking Report.
 35. N.A. Storch. Does the user interface make interruptions disruptive? A study of interface style and form of interruption. In Posters and short talks of the 1992 SIGCHI conference on Human factors in computing systems, page 14. ACM, 1992.
 36. A.L. Strauss and J.M. Corbin. Basics of qualitative research: Techniques and procedures for developing grounded theory. Sage Publications, Inc, 1998.
 37. T. Stevens. "University of Calgary researchers teach little robots to be scared of angry humans." Engadget. AOL, 17 02 2009. Web. 18 Dec 2010.
<<http://www.engadget.com/2009/03/17/university-of-calgary-researchers-teach-little-robots-to-be-scar/>>.
 38. R. M. Taylor. Integrating voice, visual and manual transactions: some practical issues from aircrew station design. The structure of multimodal dialogue, pages 259–265, 1989.
 39. TechRadar.com. OCZ Neural Impulse Actuator Review. Reviewed on July 3rd, (2008) and retrieved December, 2008.
<<http://www.techradar.com/products/computing/peripherals/input-devices/other/ocz-neural-impulse-actuator--269721/review>>
 40. K. Wada, T. Shibata, T. Saito, K. Sakamoto, and K. Tanie. Psychological and social effects of one year robot assisted activity on elderly people at a health service facility for the aged. In Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on, pages 2785–2790, 2005.
 41. F. Yamaoka, T. Kanda, H. Ishiguro, and N. Hagita. How close?: model of proximity control for information-presenting robots. In HRI '08: Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction, pages 137–144, New York, NY, USA, 2008. ACM.
 42. B. Zeigarnik. Das behalten erledigter und unerledigter handlungen. Psychologische Forschung, (9):1–85, 1927.