

THE UNIVERSITY OF CALGARY

Exploring Social Interaction Between Robots and People

by

James E Young

A DISSERTATION

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE
DEGREE OF DOCTOR OF PHILOSOPHY

DEPARTMENT OF COMPUTER SCIENCE

CALGARY, ALBERTA

SEPTEMBER 2010

© James E Young, 2010

James E Young: *Exploring Social Interaction Between Robots and People*,
PhD Dissertation, © September 2010

THE 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 dissertation entitled “Exploring Social Interaction Between Robots and People” submitted by James E Young in partial fulfilment of the requirements for the degree of Doctor of Philosophy.

Supervisor, Dr. Ehud Sharlin
Department of Computer Science

Supervisor, Dr. Takeo Igarashi
Department of Computer Science, University of Tokyo

Dr. Saul Greenberg
Department of Computer Science

Dr. John Aycock
Department of Computer Science

Dr. Patrick Feng
Department of Communication and Culture

External Examiner, Dr. Sriram Subramanian
Department of Computer Science, University of Bristol

Date

ABSTRACT

Robots are rapidly advancing toward becoming autonomous and skilled entities in a wide range of environments, and it is likely that more and more people will soon be interacting with robots in their everyday lives. As this happens, we believe that it is crucial that robots are designed to be easy to use and understand, reducing the requirement for people and environments to adapt to the robot. Emerging research in Human-Robot Interaction (HRI) suggests that people have a strong natural tendency to treat robots as social entities, anthropomorphizing, zoomorphizing, and generally attributing them with social characteristics and roles. Our approach to HRI is to explicitly focus on the social layers of interaction, building robotic interfaces that use people's existing skill sets, and that explicitly attempt to integrate into familiar social structures. We refer to this approach of directly considering HRI in the context of the social human world as social HRI.

The field of social HRI is only just emerging: there is little general discussion which explains why social HRI is important or what exactly social HRI means, there is no methodology for approaching the specific consideration of designing and implementing social HRI interfaces, and there is no structured methodology for evaluating and studying social HRI. There are few social HRI interface designs and implementations — those that focus on social interaction between people and robots — and the scope of social HRI interaction possibilities is still relatively unexplored. In this dissertation we present what we believe is the first thorough exploration of the theory, design, implementation, and evaluation of social HRI.

We present a detailed analysis of social HRI, drawing particularly from selected works in social psychology and philosophy, and compose a social HRI-targeted theory that addresses why people tend toward social interaction with robots, how robots can leverage this, and what the implications are for both users and designers. We present a set of social HRI interfaces we designed, implemented and evaluated as a means to demonstrate and reflect on the practical and technical feasibility of applying social HRI principles to robot interface designs. We present the results from several extensive user studies, which we conducted as a means to learn from our interfaces, and to test, reflect on and further develop our social HRI theories.

Finally, we distill our overall efforts into a set of some of the very first social HRI-specific design heuristics.

Overall, our work highlights the importance of considering the broad landscape of social interaction between people and robots, and the usefulness of explicitly considering social aspects for robotic interaction design. We hope that by establishing the foundational social HRI groundwork, this dissertation will lead to and support continuing development of social HRI theory and new social HRI interface designs as this emerging domain evolves and grows.

PUBLICATIONS

Some ideas and figures included in this dissertation have appeared previously in the following publications:

JOURNAL PAPERS

James E Young, JaYoung Sung, Amy Volda, Ehud Sharlin, Takeo Igarashi, Henrik Christensen, and Rebecca Grinter. Evaluating human-robot interaction: Focusing on the holistic interaction experience. *International Journal on Social Robotics*, 2010. to appear, (Chapter 2, Chapter 3, Chapter 8).

James E. Young, Richard Hawkins, Ehud Sharlin, and Takeo Igarashi. Toward acceptable domestic robots: Applying insights from social psychology. *Inagural Issue of the International Journal on Social Robotics.*, 1(1):95–108, January 2009. doi: 10.1007/s12369-008-0006-y. (Chapter 3, Chapter 8).

REFEREED CONFERENCE FULL PAPERS

Cheng Guo, James E. Young, and Ehud Sharlin. Touch and toys: new techniques for interaction with a remote group of robots. In *ACM Conference on Human Factors in Computing Systems, 2009. CHI '09, Boston, USA, April 4–9, 2009*, pages 491–500, New York, NY, USA, 2009. ACM, ACM Press. doi: 10.1145/1518701.1518780. (Chapter 5).

James E. Young, Takeo Igarashi, and Ehud Sharlin. Puppet master: Designing reactive character behavior by demonstration. In *ACM SIGGRAPH / Eurographics Symposium on Computer Animation, 2008. SCA '08, Dublin, Ireland, July 7–9, 2008*, pages 183–191, Germany, 2008. Eurographics Association Press, Eurographics Association. (Chapter 7).

James E. Young, Min Xin, and Ehud Sharlin. Robot expressionism through cartooning. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 309–316, New York, NY, USA, 2007. ACM, ACM Press. doi: 10.1145/1228716.1228758. (Chapter 3, Chapter 6).

James E. Young, Ehud Sharlin, and Jeffrey E. Boyd. Implementing bubblegrams: The use of haar-like features for human-robot interaction. In *Proceedings of the second IEEE Conference on automation Science and Engineering, 2006. CASE '06, Shanghai, China, October 8–10, 2006*, pages 298–303, Los Alamitos, CA, USA, 2006. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/COASE.2006.326897. (Chapter 6).

James E. Young and Ehud Sharlin. Sharing spaces with robots: an integrated environment for human-robot interaction. In *Proceedings of the 1st International Symposium on Intelligent Environments, 2006. ISIE '06, Cambridge, UK, April 5-7, 2006*, pages 103–110, Cambridge, UK, April 2006. Microsoft Research Ltd., Microsoft Research Ltd. Press. (Chapter 6).

REFEREED CONFERENCE SHORT PAPERS AND EXTENDED ABSTRACTS

James E. Young, Kentaro Ishii, Takeo Igarashi, and Ehud Sharlin. Style-by-demonstration: Using broomsticks and tangibles to show robots how to follow people. In *adjunct proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (Late-Breaking Abstracts), 2010. HRI LBA '10, Osaka, Japan, March 2-5, 2010*, New York, NY, USA, 2010. ACM, ACM Press. (Chapter 7).

James E. Young, Daisuke Sakamoto, Takeo Igarashi, and Ehud Sharlin. Puppet master: A technique for defining the actions of interactive agents by demonstration (Japanese: Puppet master: 例示によるとインタラクティブなエージェントの動作作成手法). In *Proceedings of the Human-Agent Interaction Symposium, 2009. HAI Symposium '09, Tokyo, Japan, 2009*. (Chapter 7).

James E. Young, Ehud Sharlin, and Takeo Igarashi. The concept of a robot. In *Proceedings of the ACM/IEEE Human-Robot Interaction Pioneers Workshop, 2008. Pioneers '08, Amsterdam, The Netherlands, March 12, 2008*, New York, NY, USA, 2008. ACM, ACM Press. (Chapter 3).

BOOK CHAPTERS

James E. Young, Ehud Sharlin, and Takeo Igarashi. “what is Mixed Reality, anyway?” considering the boundaries of mixed reality in the context of robots. In Xiangyu Wang, editor, *Mixed Reality and Human-Robot Interaction*. Springer-Verlag, Berlin, New York, Heidelberg, 2010 (in press). (Chapter 6).

“If you do your best every day, your best gets better.”

— Hugh Young

ACKNOWLEDGMENTS

I owe a debt of gratitude to the many people who have supported, guided, and inspired me throughout my program over the last few years. I could not have done it alone. There are far too many of you to thank here, a task which would consume an entire chapter, and so I only focus on a select few. For the rest of you, you have my unending gratitude.

To Ehud Sharlin, it is difficult to overstate my appreciation for all the time, mentoring, and support which you provided through these last five years. Your supervision has gone well beyond expectations; you have become a mentor not only in academia, but have helped me through many stages of life as well. You truly are one of the kindest, most caring people I have ever met, and have become a dear friend.

To Takeo Igarashi, since my first internship in Japan you have been an integral driving force of my research, your ideas and feedback helped to shape every aspect of this dissertation. I thank you for all the time and support you have given me over the last few years and for all the doors you have opened for me, and I am looking forward to the many collaborations to come down the road.

To my supervisory committee members, Saul Greenberg and John Aycock, I thank you for all the thoughtful insight and direction you have provided over the last years, through my PhD roll-up proposal, candidacy, and finally my defence. To my defence committee members, Patrick Feng and Sriram Subramanian, thank you for all the careful questions and feedback to my dissertation and research, it has been very productive and educational.

To the Interactions Lab, I truly believe that this is one of the best labs in the world to do research. The communal support, amazing depth of knowledge, and family-like atmosphere has made grad school an amazing chapter of my life. In particular I would like to thank: Nicolai Marquardt for being a great friend and always smiling and laughing at my incessant complaining about something or other; Mark Hancock for entertaining my political rants and for offering a kindred spirit and understanding of the hell that is dissertation writing,

the morning routine was not for the coffee (although it was delicious, you should try it); Petra (π) for always being available, even when away, for a good chat to offload stress; Floh for always giving input on my writing and research (and for having an insatiable appetite for boiled eggs, sushi, and chicken); Chris for being particularly helpful through the job-hunt process; the lunch-table crowd for keeping it interesting; Cody and Cheng for keeping video games in my life; and absolutely everyone else, I wish I had the room to list you all.

To my colleagues and friends at the University of Tokyo and JST ERATO Design Interface Project, thank you for always welcoming me and being so incredibly helpful to a blundering foreigner in Japan. Thank you to: Daisuke for all the fun evenings at the Hub and great research discussions; Kenshi for being so supportive and positive, and always providing great perspectives; Kenta for the fun lunch chats and reflection on research; Alec for your love of cockroaches and eating other insects; Yotam for keeping it kosher; Greg for helping me understand the wonder of sheep; and all the rest for making my times in Japan unforgettable.

My many projects would not have been possible without the hands-on, involved support of many people. Thank you to: Cheng Guo for inviting me to work on his Touch and Toys project and for conducting a very lengthy evaluation; Min Xin for drawing awesome graphics that has made the Jeeves project successful; Amy and JaYoung for shaping my crazy social HRI ideas into a publishable paper; to Maria Bakardjieva and Richard Hawkins for giving me the opportunity to study new and exciting ideas outside the realm of computer science; Jeff Boyd for showing me how great computer vision is; Kentaro Ishii for helping with early prototypes of the Puppet Master to Robots project; Youichi Kamiyama for being an engineering genius with the dog-leash project; and Daisuke Sakamoto for great Japanese translation help.

To my parents Ellen and Hugh, my sisters Christina and Melissa, my Japanese parents Shigeko and Katsura, brother and sister Osamu and Asuka, and amazing nephews Tora and Taka, I thank you for being an amazing family that makes all of the difficult times worthwhile.

To my cat Neko, thank you for (usually) thinking that I'm awesome, or at least, warm to lay on, and for keeping my company during 3 am thesis writing sessions.

Finally, the generous financial support from the Natural Sciences and Engineering Research Council of Canada (NSERC), The University of Calgary, and the Japan Society for the Promotion of Science has made it possible to pursue my research with intent focus while maintaining a normal life.

To Aya

*Your unending love and encouragement supported me through the many
difficult and dark times. You brighten every day of my life,
I would never have made it without you.*

To my parents, Ellen and Hugh,

*Thank you for always believing in me, even when I did not
believe in myself. You are amazing parents
who are always there for me.*

CONTENTS

I	SOCIAL INTERACTION WITH ROBOTS	1
1	INTRODUCTION	3
1.1	The Emergence of Social Interaction With Robots	3
1.2	The Emergence of Human-Robot Interaction (HRI)	4
1.3	Understanding Social Interaction with Robots	4
1.3.1	How People Perceive Robots	5
1.3.2	Why does HRI Demand Explicit Consideration?	6
1.3.3	Robots that Leverage People’s Social Tendencies	6
1.3.4	Robot Emotions and Intelligence	7
1.3.5	Research Questions	8
1.4	Methodology and Approach	9
1.4.1	Fundamental Questions of Social HRI	10
1.4.2	Social HRI Interfaces	10
1.4.3	Observations and Studies	15
1.5	Significance	17
1.5.1	Contributions	17
1.6	Structural Overview	18
2	A SNAPSHOT OF THE SOCIAL HRI LANDSCAPE	19
2.1	Classic Human-Robot Interaction	20
2.1.1	Remote-Control Robots	20
2.1.2	Collocated Robot Operation	22
2.2	Social HRI Interface Designs and Observed Phenomena	24
2.2.1	People Respond Socially to Machines	24
2.2.2	Robot Design that Leverages Social Tendencies	29
2.2.3	Our New Social HRI Interfaces	34
2.3	Theoretical Foundations for Understanding Social HRI	34
2.3.1	Existing Social HRI Theories	35
2.3.2	Embodiment	37
2.3.3	Embodied Interaction	39
2.3.4	Sociology in Relation to Social HRI	41
2.3.5	Theory of Mind	42
2.3.6	Summary: Theoretical Foundations	44

2.4	Studying and Evaluating Human-Robot Interaction	44
2.4.1	Task Completion and Efficiency	44
2.4.2	Emotion and Affective Computing	45
2.4.3	Situated Personal Experience	47
2.4.4	Frameworks for Exploring Social Interaction with Robots	48
2.5	Summary: A Framing For Our Dissertation	49
3	DEVELOPING A THEORETICAL FRAMEWORK FOR SOCIAL HRI	51
3.1	Analyzing the Meaning of Social HRI	51
3.1.1	What is a <i>Robot</i> ?	52
3.1.2	What Does <i>Social</i> Imply for Social HRI?	53
3.1.3	Embodiment and Embodied Interaction with Robots	55
3.2	Why Robots Create Unique Interaction Experiences	57
3.2.1	Why Robots Encourage Social Interaction	57
3.2.2	From Anthropomorphism to Sense of Agency	58
3.2.3	Embodied Interaction Experience	59
3.3	Social Psychology Perspectives for Exploring Social HRI	61
3.3.1	The Socialization Process	62
3.3.2	Absorptive Capacity	63
3.3.3	Select Adoption-of-Technology Theories For Robots	64
3.4	Applying Selected Adoption-of-Technology Models to Robots	67
3.4.1	Two Domestic Robots for Targeted Analysis	67
3.4.2	Initial Exposure to Domestic Robots (TRA, TPB, MATH)	69
3.4.3	Control and Safety of Domestic Robots (TPB, TAM, MATH)	71
3.4.4	Hedonic Aspects of Social HRI (MATH)	73
3.4.5	An Additional Perspective: The Role of Social Design	76
3.5	Robot Expressionism — Ambiguity Can Be Helpful	78
3.6	A Theoretical Framework for Social HRI	80
3.6.1	Outlining Social HRI: Defining a Domain	80
3.6.2	Focusing on the Breadth: Three Perspectives for Social HRI	82
3.6.3	Factors Shaping Robot Acceptance	89
3.7	Summary: A Theoretical Social HRI Discussion	91
II	IMPLEMENTING AND EVALUATING SOCIAL ROBOTIC INTERFACES	93
4	A DOG-LEASH INTERFACE FOR LEADING A ROBOT	97
4.1	Past Robot-Leading Interfaces	98
4.2	Designing A Dog-Leash Robot Interface	99
4.2.1	Robot Following Behind the Person	101

4.2.2	Robot In-Front of the Person	102
4.3	Implementing our Dog-Leash Interface	103
4.4	Evaluating the Robotic Dog-Leash Interface	106
4.4.1	Design Critiques	106
4.4.2	Tasks	107
4.4.3	Study Procedure	108
4.4.4	Study Design	109
4.4.5	Results	109
4.4.6	Discussion	113
4.5	Dog-Leash Robot: Conclusions	115
5	EVALUATING “TOUCH AND TOYS”	117
5.1	Project Introduction	117
5.2	Evaluation Design	120
5.3	Results	122
5.3.1	Task-Completion Time	122
5.3.2	Usability	123
5.3.3	Preference	124
5.3.4	Touch	125
5.3.5	Toy	125
5.3.6	Robot Movements	126
5.3.7	Collisions	127
5.3.8	Two-Handed Interaction and Multitasking	127
5.3.9	Complexity	128
5.3.10	Real Robots	128
5.4	Discussion	129
5.5	Implications	131
5.6	Evaluating Touch and Toys: Conclusions	132
6	CARTOON ARTWORK ROBOTIC INTERFACES	135
6.1	Applying Cartoon Artwork to Robots	137
6.1.1	Interface Technology — Mixed Reality	137
6.1.2	Existing Use of Cartoon Artwork for HRI	138
6.1.3	<i>How</i> Robots can use Cartoon Artwork	140
6.1.4	Adapting Techniques from Cartoon Artwork	145
6.2	Implementing Bubblegrams	151
6.2.1	Haar-Like Features for Robot Tracking	153
6.3	Implementing The Jeeves Test Bed	157

6.3.1	Implementing Trash Jeeves	158
6.4	Informal Design Critiques	159
6.4.1	Bubblegrams	160
6.4.2	Jeeves	161
6.5	Conclusions on Cartoon Artwork	162
7	STYLISTIC LOCOMOTION AND PUPPET MASTER	165
7.1	Stylistic Locomotion: Robotic Style	166
7.1.1	An Overview of Style	168
7.1.2	Animation Interfaces for Stylistic Locomotion	170
7.1.3	Robot Locomotion: Robot-Based Stylistic Locomotion	174
7.1.4	Conclusions on Stylistic Locomotion	176
7.2	Puppet Master: Designing Interaction Style by Demonstration	177
7.2.1	An Overview of Programming by Demonstration	179
7.2.2	Puppet Master Interface Designs	180
7.2.3	Animation-Based Puppet Master	181
7.2.4	Robot-Based Puppet Master	186
7.2.5	Conclusions on Puppet Master	193
7.3	Animation Puppet Master Algorithm	195
7.3.1	Core Puppet Master Algorithm	196
7.3.2	Features and Dataset	196
7.3.3	Algorithm Loop	198
7.3.4	Similarity Metrics	199
7.3.5	Output Generation	202
7.3.6	Initial Profiling and Appending Training	205
7.4	Robotic Puppet Master Algorithm	206
7.4.1	Kinematic Model	206
7.4.2	Auxiliary Actions	209
7.5	Stylistic Locomotion and Puppet Master Studies, Animation	209
7.5.1	Study Procedure and Methodology	210
7.5.2	Observations	211
7.5.3	Discussion	215
7.5.4	Reflection	216
7.6	Stylistic Locomotion and Puppet Master Studies, Robots	217
7.6.1	Programmer Study	219
7.6.2	Designer Study	223
7.6.3	Observer Study	239

7.6.4	Stylistic Motion and Puppet Master Studies: Reflections	250
7.7	Stylistic Locomotion and Puppet Master: Conclusions	253
III	CONCLUSIONS	255
8	HEURISTICS FOR SOCIAL HRI	257
8.1	Heuristics	257
8.1.1	(H1) Employ Agency and Anthropomorphism	258
8.1.2	(H2) Improve Social Accessibility	258
8.1.3	(H3) Design for Specific Interpretations	259
8.1.4	(H4) Use Real Robots	260
8.1.5	(H5) Get People Involved	261
8.2	Evaluating our Interface Designs Using our Heuristics	266
8.2.1	A Dog-Leash Interface for Leading a Robot	266
8.2.2	Cartoon Artwork for Expression: Bubblegrams and Jeeves	267
8.2.3	Stylistic Locomotion and Puppet Master	267
8.3	Summary: Heuristics for Social Human-Robot Interaction	268
9	CONCLUSIONS AND FUTURE WORK	269
9.1	Research Questions	269
9.2	Research Contributions	269
9.2.1	Social HRI Fundamentals	270
9.2.2	Social HRI Heuristics	272
9.2.3	Social HRI Designs and Implementations	272
9.2.4	Social HRI Evaluations	274
9.3	Scope of Contribution	277
9.3.1	Application to Other Cultures	277
9.3.2	Application to Other Interfaces	277
9.4	Future Work	279
9.4.1	Social HRI Framework and Heuristics	279
9.4.2	Social HRI Evaluation	280
9.4.3	Social HRI Interface Designs and Implementations	281
9.5	Final Words	284

IV	APPENDIX	285
A	TOUCH AND TOYS STUDY MATERIALS	287
	A.1 Protocol and Questionnaires	287
B	DOG-LEASH ROBOT STUDY MATERIALS	311
	B.1 Materials and Questionnaires	311
	B.2 Task Instructions	325
C	STYLISTIC LOCOMOTION AND PUPPET MASTER STUDIES, ANIMATION	331
	C.1 Designer Study	331
	C.2 Observer Study	345
D	STYLISTIC LOCOMOTION AND PUPPET MASTER STUDIES, ROBOTS	359
	D.1 Programmer Study API	359
	D.2 Walking Paths	367
	D.3 Designer Study: Broomstick Condition	369
	D.4 Designer Study: Surface Puppet Master Condition	385
	D.5 Observer Study	393
	BIBLIOGRAPHY	405

LIST OF FIGURES

Figure 1.1	a dog leash for robots	12
Figure 1.2	screen shots of <i>bubblegrams</i> and <i>Jeeves</i>	13
Figure 1.3	a robot with an aggressive locomotion style	14
Figure 1.4	demonstrating style to a robot	15
Figure 2.1	iRobot Packbot and controller	21
Figure 2.2	collocated robot control	23
Figure 2.3	functionality-oriented robot designs	26
Figure 2.4	robots that people were asked to shut off, or “kill”	27
Figure 2.5	museum guides and roboceptionists in everyday environments	29
Figure 2.6	robots working within social norms	30
Figure 2.7	examples of androids	31
Figure 2.8	zoomorphic robots	32
Figure 2.9	Keupon the dancing robot	32
Figure 2.10	examples or robots that people find eerie	36
Figure 2.11	graph representing the “uncanny valley”	37
Figure 2.12	“theory of mind” experiment setup	43
Figure 2.13	the sensual evaluation instrument	46
Figure 3.1	embodied interaction with robots	56
Figure 3.2	situated holistic interaction context	60
Figure 3.3	two domestic robots used in our analysis	68
Figure 3.4	The Sony AIBO Robotic Dog	73
Figure 3.5	two sex robots produced by AndyDroid	74
Figure 3.6	Edvard Munch’s “The Scream”	79
Figure 3.7	interaction experience mutually shaped by people and robots	83
Figure 4.1	a person leading a robot with a leash	97
Figure 4.2	robotic dog-leash interface implementation	100
Figure 4.3	leash-robot following from behind a person	102
Figure 4.4	leash-robot being led in front of a person	103
Figure 4.5	dog-leash-robot leash mechanism	105
Figure 4.6	map, path, and way points for dog-leash study	107
Figure 4.7	result of NARS questionnaire on dog-leash study	112
Figure 4.8	comfort distances with the robot on and off a leash	113

Figure 5.1	touch and toys project interfaces	118
Figure 5.2	touch and toys interface details	119
Figure 5.3	Tangible User Interface (TUI) mapping to robots	119
Figure 5.4	touch and toys study setup	121
Figure 5.5	touch and toys ease of use questionnaire responses	124
Figure 5.6	conviction of ease-of-use-related responses	129
Figure 6.1	an example comic collage	135
Figure 6.2	example mixed-reality interfaces	138
Figure 6.3	cartoon-art related robot designs	139
Figure 6.4	our thoughtcrumbs interaction technique	142
Figure 6.5	the bubblegrams interaction technique	144
Figure 6.6	our decorations interaction technique	145
Figure 6.7	cartoon-artwork icons	146
Figure 6.8	cartoon-artwork text examples	147
Figure 6.9	cartoon-artwork facial expression examples	147
Figure 6.10	the rash Jeeves interface implementation	149
Figure 6.11	the recycle police Jeeves interface implementation	150
Figure 6.12	the clean-tracks Jeeves interface implementation	150
Figure 6.13	bubblegrams implementation	152
Figure 6.14	bubblegrams' hardware	153
Figure 6.15	Haar-Like templates for feature extraction	154
Figure 6.16	Haar-like templates applied to a Sony AIBO	155
Figure 6.17	AIBO poses targeted by individual classifiers	155
Figure 6.18	AIBO tracking using classifier voting	156
Figure 6.19	Jeeves test-bed interface	158
Figure 6.20	colour segmentation for robot tracking	159
Figure 7.1	concept sketch of stylistic locomotion path	168
Figure 7.2	Heider and Simmel's experiment on apparent behaviour	169
Figure 7.3	our generically-designed animation avatars	171
Figure 7.4	mouse GUI for interacting with stylized characters	172
Figure 7.5	animation table interface for interacting with stylistic characters	174
Figure 7.6	a robot following a person	176
Figure 7.7	a sketch of programming style by demonstration	181
Figure 7.8	using the mouse GUI for demonstrating style	183
Figure 7.9	using the animation table interface for demonstrating style	184
Figure 7.10	SMART Technologies tabletop with Vicon tracking	185

Figure 7.11	Surface puppet master interface for demonstrating to a robot	188
Figure 7.12	reflective markers and caster on the Surface puppet master TUI	189
Figure 7.13	a broomstick interface for demonstrating following to a robot	190
Figure 7.14	details of broomstick attachment to robot	191
Figure 7.15	the view from the Vicon camera system	192
Figure 7.16	iRobot Create with Bluetooth modem and reflectors	193
Figure 7.17	space for interacting with robots	194
Figure 7.18	features used by the puppet master algorithm	197
Figure 7.19	puppet master algorithm situational similarity	200
Figure 7.20	puppet master algorithm entity-path coherency	201
Figure 7.21	puppet master algorithm general-trajectory computation	204
Figure 7.22	kinematic model of the iRobot Roomba's movements	207
Figure 7.23	frequency-analysis for moving a robot toward a given target	208
Figure 7.24	the process of applying the puppet master algorithm to robots	208
Figure 7.25	histogram of puppet master uninformed behaviour matching	214
Figure 7.26	on-screen simulation used for programming condition	220
Figure 7.27	experience with robots versus average training time	230
Figure 7.28	observers classifying behaviours	244
Figure 7.29	participant responses to questions per each of twelve behaviours	245
Figure 8.1	process for building the interaction experiences map	264
Figure D.1	walk paths used for robots that follow	367
Figure D.2	walk paths used for robots that share their space	368
Figure D.3	walk paths used for observing randomized behaviours	368

LIST OF TABLES

Table 3.1	summary table of technology-adoption models	66
Table 4.1	participant rankings of the dog-leash interfaces	109
Table 4.2	changes in the SAM response post dog-leash behaviour	110
Table 4.3	results of rating the dog-leash robot as controllable	111
Table 4.4	results of rating impressions of the dog-leash robot	111
Table 4.5	results of the GODSPEED V dog-leash questionnaire	111
Table 5.1	average touch-and-toys task completion time	123

Table 5.2	touch and toys participant preferences	124
Table 5.3	touch and toys responses regarding using two hands	128
Table 7.1	behaviour matching to original keywords	214
Table 7.2	demonstration times for programmer condition	221
Table 7.3	answers to per-behaviour post-training questions	227
Table 7.4	breakdown tables for two per-behaviour questions	227
Table 7.5	results of participants matching own behaviours	228
Table 7.6	frequency table of post-test questionnaire answers	229

LISTINGS

Listing D.1	base class given to programmers	360
Listing D.2	skeleton class given to programmers	360
Listing D.3	EntityState listing	361
Listing D.4	RobotCommand listing	361
Listing D.5	RoombaSong listing	362
Listing D.6	utility functions included in API	362
Listing D.7	example provided to programmers	363

ACRONYMS

ANT	Actor Network Theory
API	Application Programming Interface
ANOVA	ANalysis of VAriance
GUI	Graphical User Interface
HCI	Human-Computer Interaction
HMD	Head-Mounted Display
HRI	Human-Robot Interaction
MATH	Model of Acceptance of Technology in Households

MR	Mixed Reality
MRIE	Mixed-Reality Integrated Environment
NARS	measurement of Negative Attitudes towards Robots Scale
PC	Personal Computer
RFID	Radio Frequency IDentification
RNT	Rotate 'N Translate
SBD	Style-by-Demonstration
SAM	the Self-Assessment Manikin
TAM	Technology Acceptance Model
TPB	Theory of Planned Behaviour
TRA	Theory of Reasoned Action
TUI	Tangible User Interface
WPF	Windows Presentation Framework
XAML	Extensible Application Markup Language

Part I

SOCIAL INTERACTION WITH ROBOTS

INTRODUCTION

1.1 THE EMERGENCE OF SOCIAL INTERACTION WITH ROBOTS

From robotic vacuum cleaners in over 2.5 million homes to robot receptionists in Japan, to autonomous robots carrying medicine around hospitals, robots are poised to become a part of life for much of the general public. Current research, some of which is outlined in this dissertation, strongly suggests that this is a trend and that robots will continue to permeate society. Similar to how we encounter computing in our daily lives it is likely that people will soon find themselves interacting with robots in a wide variety of contexts and scenarios. With this in mind, then, we believe it is important to consider how the general public will interact with, work with, and understand robots, and how the robots can integrate into people's social spaces.

When people work with other people, using communication mediums such as speech, gestures, the written word or art, we rely on complex levels of inter-personal common language and understanding as a base for communication. However, when we interact with robots, intrinsic common understanding is very limited, if it exists at all; robots think in the foreign language of bits and bytes, a language we humans cannot inherently understand. This poses a potentially serious communication problem, made eminent as autonomous and intelligent robots are already starting to enter our everyday environments.

Robots, by their very nature, are a unique technology in people's environments. People naturally perceive robots as active social players and often treat them as such, an inclination that plays a very important role in how people understand, react to, and ultimately interact with robots (Breazeal, 2003a; Kiesler and Hinds, 2004; Young, Hawkins, Sharlin, and Igarashi, 2009). One approach to Human-Robot Interaction (HRI), which we refer to as social HRI, is to embrace, leverage, and try to understand people's social tendencies toward robots. Understanding this social component of interaction, where people interact with robots as social players and not as mere mechanical devices, is critical for building a deeper understanding

how people interact with robots in general. In addition, social HRI designers can leverage these social inclinations in their robot interface designs. Doing so uses paradigms that people are familiar with, enabling them to utilize their existing social skill sets to easily understand, communicate and interact with robots. Developing an understanding of the social aspects of interaction between people and robots is a fundamental problem in HRI.

1.2 THE EMERGENCE OF HUMAN-ROBOT INTERACTION (HRI)

While many aspects of Human-Computer Interaction (HCI) research are indeed applicable to robots, robots have several unique properties that call for direct consideration. As such the field of HRI has emerged, as a sub-discipline of HCI, to explicitly study the ways people interact with robotic technologies (Kiesler and Hinds, 2004). Traditionally, robotics and early work in HRI generally focused on low-level robot control, dealing with basic robotic engineering issues such as task completion and efficiency. As robotic technology advanced, higher-level and interactive human control became practical, for example, remote-control interfaces that combine video and sensor data. From this has arisen the human-factors-centred HRI, further fuelled by increasing autonomy and practical utility of robots (Breazeal, 2003a; Kiesler and Hinds, 2004; Young et al., 2009). Still, many of the traditional and existing HRI tools do not explicitly consider or leverage the emergent social aspects related to interacting with a robot. Here enters social HRI, an emerging movement in HRI which we contribute to and attempt to crystallize in this dissertation.

1.3 UNDERSTANDING SOCIAL INTERACTION WITH ROBOTS

Understanding the social elements of interaction between people and robots is a complex and multi-faceted problem. Primarily, we believe this challenge revolves around understanding why robots are unique, and how robots impact the social aspects of interaction. Our exploration of this challenge surrounds the following overarching research questions: what does the tendency to treat robots as social entities mean for interaction between a person and a robot? How can robotic interfaces be designed to leverage this tendency? Which methodologies, structured techniques, taxonomies, and heuristics can be developed and

used for social HRI? In this section, we explore these questions and map them to our research goals and to this dissertation.

1.3.1 *How People Perceive Robots*

When people look at a laptop they do not simply see a plastic box with rubber buttons, filled with electronics and a battery. They will also see a portable device that facilitates their work, a social communication medium, an entertainment device, and so on; they generally see much more than a simple sum of the physical components, much more than the direct physical form, and much more than a collection of arbitrary tasks. People use this interpretation of an object to inform them on what they can expect from an object and how they should interact with it (Norman, 1988; Koffka, 1935). The fact that people naturally tend to interact socially with robots brings this consideration to the forefront: we can expect a person's *interpretation* of a robot to be rooted directly in how it is perceived within the context of the social human world. Thus, we argue that interaction designers who want (or need) to leverage the social layers of interaction must consider this flexibility of interpretation.

Interaction between people and robots can include complex layers of social interplay, similar to that which exists between two people or between a person and an animal. These layers are very important for shaping a person's experience of interaction (Norman, 2004), and acceptance of technology in general is highly dependent upon social factors (*factors of socialization*, McMeekin, Green, Tomlinson, and Walsh, 2002; Von Hippel, 2005). It follows that HRI designers should consider social factors in parallel with more traditional concerns such as goal-oriented utility, completion time, or accuracy. Designers must take a very broad, socially embedded and person-oriented perspective to understanding robotic interaction to help ensure their design is also socially valid.

This construction of perception, or interpretation, of a robot is highly subjective, dependent on many factors such as a person's background, culture, context of interaction, personality and preferences, and experiences with the robot. This perception can be particularly influenced by existing understanding of similar entities, previous experiences, and how they relate to the design of the robot itself. On the one hand, this complexity makes it difficult to precisely determine how a particular robot or interface will be perceived by a given individual. However, this flexibility means that designers can influence interpretation by designing robots to match

qualities and ideas that people already understand. Unfortunately, understanding the nuances of social communication to the point where they can be programmed into robots is a difficult problem (as also argued by Norman, 2007): ultimately, successful intelligent robots may be able to overcome these issues by observing and dynamically learning from the social environments they occupy, similar to how people do.

1.3.2 *Why does HRI Demand Explicit Consideration?*

At the core of the above discussion is the question of how robots are different than traditional technologies. It is still not well understood what it is about robots that encourages people to interact socially with them, or what exactly makes robots unique, although there is mounting evidence that robots do provide a truly unique experience (e. g. Forlizzi, 2007; Garreau, 2007, outlined in detail in Chapter 2). Building understanding of how robots are different is an important challenge for HRI that will help direct future research and explain interaction results, a question we explore in this dissertation.

The question remains, then, as to how HRI as a field is different from HCI. HRI is considered to be a subfield of HCI, in that a robot is a special case of a computer. However, as argued above and discussed in detail later in this dissertation, we argue that the inclusion of robots creates a unique, particularly complex, interaction experience. It is important to be aware of and to consider the consequences of this distinction, particularly in terms of which methods are transferable from the more-general HCI, under what circumstances, and which methods should be revised and re-evaluated.

1.3.3 *Robots that Leverage People's Social Tendencies*

We argue that robots should leverage people's tendencies toward treating them as social entities. Programming robots to both understand and communicate using social techniques, such as through externalized emotions, reinforces and leverages people's anthropomorphic inclinations, helping them to easily relate to and understand, as well as communicate with, the robots they interact with (Breazeal, 2003a; Kiesler and Hinds, 2004; Young, Sharlin, and Igarashi, 2008; Young et al., 2009). Of particular interest has been the role of this human-like robotic emotion in interaction with robots.

Hans Moravec, robotics pioneer, claims that as robots become increasingly complex we will be forced to program them with psychological-social models to help us understand their motives and communicate effectively with them. In his vision, they will have a “behavioural character” and he suggests that “many people will empathize and interact with [these robots] as they do with pets, and the robots will respond” (Moravec, 1998).

Cognitive scientist Donald Norman, HCI and design expert, also states that robots will need emotion to communicate effectively with people, and further argues that such emotion is important to provide people with a sense of satisfaction that they do not feel when dealing with emotionless machines (Norman, 2004). Norman further points to the practicality of such an approach, as people already have many natural mechanisms to deal with emotion, and interaction with other people depends on our ability to interpret their internal feelings. Robots, Norman argues, should take advantage of this by displaying their state and suggesting their intentions through expressing emotions similarly to how a person or pet does. The robot could show confidence to indicate that it understood a command or show frustration if it is having difficulty performing a task. These robotic emotions can provide a familiar layer of insight into a robot’s otherwise-alien technical state and algorithmic motivations. That is, we can expect these “social” robots to use emotions to broadcast and represent their state as much as humans do, and for people to intuitively understand them and readily accept them as they would other (living) entities. People already often do this with objects, giving human-like characteristics to anything that is vaguely lifelike: consider how people talk to their car, or how children treat teddy bears.

Emotion, however, is only one component of the greater social interaction picture. In this dissertation we highlight how people use many socially-rooted communication and interaction techniques — including emotion and beyond — that robotic design can leverage.

1.3.4 *Robot Emotions and Intelligence*

Emotions, intelligence, and other such human or animal-like characteristics are often used in an anthropomorphic or zoomorphic fashion to describe and discuss robots; this is standard practice in HRI. That is, expressions such as “robot emotions,” “robot intelligence,” “what a robot is thinking,” “motivations,” and so forth, are used to refer to components of robotic programming and a robot’s state. We adopt this style in our writing as it fits well within the

natural tendency to treat robots as social entities and also serves as a useful tool to build intuition surrounding difficult-to-understand robotic properties.

Some have slight resistance to this approach as they feel uncomfortable with using such human language when relating to machines and their properties. As interesting as the philosophical discussion on the nature of emotions and intelligence is, however, such questions are beyond the scope and focus of our work. Instead, we admit that the robotic emotions and properties we refer to are purely synthetic and fundamentally different from the human versions. We maintain, however, that this perspective is a useful one:

Thus, for the same reason that animals and people have emotions, I believe that machines will also need them. They won't be human emotions, mind you, but rather emotions that fit the needs of the machines themselves.

— Norman (2004)

In addition, we would like to make a clear distinction between robots using *externalized* human-like social techniques such as emotions that are perceivable from a person's perspective (or a robot itself interpreting human emotions), and robots that use emotions *internally* as part of their algorithmic decision making, use not perceivable by a person. While there is often an overlap between the two, we present this categorization as a method for describing the boundaries of our work. In our work we focus on *external* robot techniques that a person can directly observe and interact with, where a robot tries to both interpret and express information using human-like emotion or other social methods. The robot itself may not necessarily have an internal emotional state, but uses synthetic emotion to communicate and interact with people. In this dissertation, we do not further explore the *internal* use.

1.3.5 *Research Questions*

Our exploration into social interaction between robots and people has led to the following three research questions that we explore through this dissertation:

Q1 *What does the tendency to treat robots as social entities mean for interaction between a person and a robot?* While evidence points to the tendency for people to treat robots as social entities, the question remains as to how this impacts interaction. What is the

extent of this social interpretation, and, how does this impact how people perceive and interact with robots?

- Q2 *How can robots be designed to leverage this tendency in their interaction and interface designs?* The fact that people naturally tend toward social interaction with robots is something that robotic interaction designers can leverage for creating robots that people understand and are easy to work with. The problem then becomes exploring how robots and robotic interfaces can be designed to leverage social interaction. That is, which social scenarios, interaction and communication techniques can social HRI designs integrate in ways that will make sense to people and improve interaction quality?
- Q3 *Which methodologies, structured techniques, taxonomies, and heuristics can be developed and used for social HRI?* Given that robots are treated as social entities, how does a social HRI researcher take this into account for their robotic interface design, implementation, and evaluations? What sorts of tools or frameworks can be used or developed to aid researchers in accounting for and targeting the social aspects of human-robot interaction?

Throughout this dissertation we will detail how these research questions have taken us through various stages of exploration, finally arriving at heuristics for social HRI. We have examined various targeted social HRI problems, designed solutions, prototyped and built implementations, and studied people interacting with our systems, resulting in various research contributions discussed later in this chapter and throughout this dissertation.

1.4 METHODOLOGY AND APPROACH

We take an exploratory approach to examining our research questions: we develop theory surrounding social HRI, we explore how such ideas can be applied in practise through designing and constructing interface implementations that explicitly target robot use of human social techniques and working within social contexts, and we observe people working and interacting with our implementations. In this section we detail our approach of using these experiences to reflect on the fundamental research questions outlined above.

1.4.1 *Fundamental Questions of Social HRI*

Social HRI is in its infancy, and there is a general lack of established knowledge and practise in the field. Currently, we are lacking the analytical tools, vocabulary, and base understanding of how to approach social HRI design, implementation, or evaluation. Researchers from engineering, HCI, animation, psychology, and sociology, among other fields, are converging to create a new research domain, and much of the current work revolves around simply trying to understand what the field of social HRI is. Definitions and boundaries are actively being developed, where the preliminary task of articulating the proper questions to ask is still an important contribution.

In our work we explore the meaning of social HRI, tackling the questions of why a prominent social layer even exists for HRI and what this means for interaction. Drawing from previous work in social HRI, HCI theory such as Dourish's embodied interaction (Dourish, 2001b), theories from sociology and social psychology, as well as our own original research, we present the first attempt at clearly defining social HRI. We articulate the unique properties of robots and how they relate to interaction with people, and present an in-depth exploration into the wider social context of HRI (Chapter 3). We design, implement, and evaluate several social HRI interfaces to validate and further develop our theories. Finally, we extrapolate from our own experiences from researching, designing, implementing, and evaluating to provide heuristics for social HRI design and evaluation (Chapter 8).

We admit that the social HRI theory questions we are approaching in this dissertation are very broad and that our efforts of approaching them are still preliminary, and call for further long-term research. However, we argue that our social HRI theory contributions are fundamental and important for moving ahead in this emerging new research domain.

1.4.2 *Social HRI Interfaces*

We designed and developed eight robotic interfaces, and two involved original algorithms, that leverage social HRI principles, a set of contributions that concretely demonstrate (via proof-of-concept) social HRI motivations and concepts. All our interfaces present novel robot social interaction techniques that robots can incorporate into social settings, that is, example

solutions to social HRI tasks. These further provide platforms where we can get people involved, providing an opportunity to observe and study people’s social HRI experiences.

Our particular selection of interfaces reflects our exploration approach to social HRI, from both the interaction design and implementation engineering perspectives. We selected and implemented several different design approaches as a means of reflecting on the higher-level questions from varying viewpoints. In each case we target a particular social HRI concept, and attempt to remove or minimize other channels of interaction as much as possible. Below we outline our social HRI interfaces: robots that communicate through cartoon artwork, a dog-leash interface for robots, interactive, characteristic locomotion style (a project we call *stylistic locomotion*), and programming such style by demonstration (called *puppet master*).

1.4.2.1 *A Dog-Leash Interface for Leading a Robot*

There are many social skills that people use for interaction in daily life which, although learnt and not instinctual, are extremely well established in society. One example is walking and leading an animal such as a dog or a horse on a leash. Far from being a simple physical locomotion problem, leading an animal on a leash is a delicate interplay between the leader and the animal being led which involves various social communication cues, for example, lightly tugging on the leash, walk direction, hesitance, and so forth. Even for poorly-trained animals, either the animal quickly learns and (even roughly) cooperates, or the entire scenario does not work and the person gives up.

We apply this leash scenario to robots and build an interface where a person can lead a robot similar to how they may lead an animal (Figure 1.1). This interface not only leverages existing common knowledge, but it also fits within people’s perception of robots as social entities: they can walk the robot just as they would walk an animal such as their dog (Chapter 4).

The power of the dog-leash for robots interface lies in how we leverage existing social techniques to take a complex HRI task and turn it into an easy-to-understand casual one. This project serves as an exemplar of the benefits of the social HRI approach.

1.4.2.2 *Social HRI via Cartoon Art*

People regularly use indirect artistic or symbolic mediums such as the written word or drawings as a form of social and emotional communication. We believe that robots can also use these methods as part of their social interaction toolkit. In particular, modern cartoon



Figure 1.1: our dog leash interface for leading a robot

art (e. g., as shown in Figure 1.2), the simplified artistic and visual language found in comic books and animated cinema, can be a simple-yet-powerful social expression mechanism for robots (Young, Xin, and Sharlin, 2007).

We see cartoon art as a means to enrich a robot’s communication vocabulary in ways that are widely understood by the general public, as well as a demonstration of how subtle social cues can enhance the social HRI experience — we designed and implemented two cartoon artwork interfaces, *bubblegrams* and *Jeeves* (Chapter 6). Such robots will be able to utilize powerful cartoon annotations to express such things as human-like emotion. For example, a robot may show sweat drops on its forehead to express fatigue and the need to rest and recharge, or it may use simplified cartoon-like facial expressions to express happiness for completing a task or fear for not completing it on time.

From a human point of view, we believe that these expressions will foster a stronger understanding of the robot’s internal state, tasks, and goals. This improves the accessibility (ease of use) of interaction when compared to common low-level and limited expressions such as blinking lights, error messages, or beeps. In addition, we expect that cartoon art will compel people to draw from their understanding of cartoons when interacting with a robot. We expect that they will see robots that use cartoon artwork as fun, simple, and enjoyable, and this will have a direct impact on the robots’ acceptance as interactive social peers, within

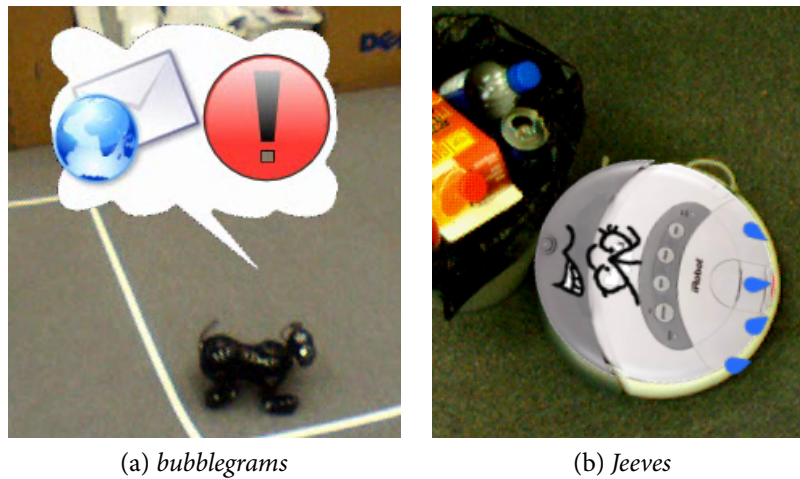


Figure 1.2: screen shots from our cartoon artwork interface implementations, *Jeeves* and *Bubblegrams*

a given task. As such, we focus our exploration on how the cartoon artwork can be integrated into a social HRI task and do not directly consider the science behind cartoon artwork itself.

1.4.2.3 *Stylistic Locomotion: interactive, characteristic locomotion style*

As is the case with people or animals, the way that a robot moves can influence how it is perceived by others. The combination of the style of its movement, including attributes such as gait, gestures, or locomotion movement patterns, can project a very strong social message. For example, many people find it easy to distinguish a happy dog from an aggressive dog simply by how it is moving, and one can often tell if a colleague is stressed simply by the way they are walking. We build here on a strong body of previous research in psychology and animation (see, e. g., the classic Heider experiments, Heider and Simmel, 1944) that show how people can construe (and construct) intricate stories, personalities, and emotions from basic motion patterns.

Our implementation deals with a very simple form of this subtle social communication: interactive robot locomotion, that is, the way that a robot moves around a space in real-time reaction to a counterpart entity (Chapter 7). We argue that through the style of locomotion, regardless of where said movement may take them, robots can communicate strong (but perhaps subtle) messages. For example, a robot can present itself as being either shy or aggressive, confident or unsure, by how it moves and how those movements interact with people. Such a situation is illustrated in Figure 1.3 where, for example, the robot is expressing aggression through its locomotion. This will be useful to inform a person on how they should



Figure 1.3: a robot expressing aggression toward the person through the locomotion style, from our *stylistic locomotion* project

react to a robot, for example, if they should shy away from an aggressive (and dangerous) lawnmower robot or if the robot they are interacting with is becoming impatient.

We believe that this work constitutes the first in-depth exploration of how robots can leverage the style of their interactive locomotion as a social component of interaction, and to design and build systems that enable such interaction to take place.

1.4.2.4 *Puppet Master: Programming Social Aspects of Interaction*

One ultimate design goal of robots is for them to understand, seamlessly fit into, and adjust to the dynamics of particular social environments such as domestic homes or (commercial or industrial) workplaces much the same as people do: robots must be able to actively participate in social communication in a meaningful way. Unfortunately, having a robot understand, analyze, and dynamically participate in a complex social scenario to the same extent as a person is a very difficult design and engineering problem. As an intermediate step, we propose that people can show a robot the style of how they want it to socially interact. That is, people can program a robot's social behaviour by directly demonstrating an exemplar to it, a technique we proposed and developed, call Style-by-Demonstration (SBD) (Chapter 7).

In everyday life, people regularly teach others how they want things done; teaching is a common social technique for working with others. We believe that robots can take advantage

of this skill, leveraging people’s existing understanding of how to interact with others, and people’s great ability to show others how to act, and how not to act, socially.

We developed a system which enables people to directly show a robot how to interact with a person, specifically, how to follow a person. While following by itself is a simple task, the focus of our work is to enable the person to specify *how* the robot is to follow the person with a social and emotional nuance, for example, aggressively or happily. Figure 1.4 illustrates how such demonstration may take place. In this case, a person is directly showing the robot how to interact with another person. Although programming-by-demonstration is not a new concept for robots, as far as we know we are the first to explicitly apply this technique to a socially-charged style, rather than a task-oriented goal, of interaction.

1.4.3 Observations and Studies

Evaluation and user studies are key components of research in the domain of HCI (Dix, Finlay, Abowd, and Beale, 1998; Sharp, Rogers, and Preece, 2007). We argue that involving people is particularly important for social HRI given the underlying principle of focusing on the



Figure 1.4: demonstrating to a robot using our *puppet master* interface the social style of how it should interact with a person

social aspects of interaction, and we apply these methods in our work to explore how people with very little (if any) prior robot experience perceive and react to our robots. Thus our design implementations serve as tools to learn about the underlying social HRI principles of interaction used to design and build the interface, in addition to the particular interface instances themselves. This closes the loop of our design cycle, where our theories about social HRI informed the development of interfaces, the interfaces enabled evaluation, and the evaluation reflects back on and informs our theories. In this dissertation, we present design critiques on our cartoon-artwork interfaces (*bubblegrams* and *Jeeves*), a formal evaluation of the dog-leash interface, five related formal evaluations of the stylistic locomotion and puppet master projects (combined), and a formal evaluation of an project external to this thesis (Chapter 5). Our studies are some of the first to specifically target the evaluation of social HRI.

Interaction with robots is different than interaction with other technologies. Traditional HCI evaluation tools do not explicitly consider these differences, and as such, may not be well-suited to social HRI. This is particularly the case when considering social interaction with robots, where highly-subjective details such as emotion are involved. While the evaluation of affective computing (e.g., Boehner, DePaula, Dourish, and Sengers, 2007; Desmet, 2005; Höök, 2005; Isbister, Höök, Sharp, and Laaksolahti, 2006) is an active research field which can sometimes be applied directly to robots, this does not explicitly consider the unique properties of robots. There is as of yet no explicit, solid foundation or targeted methodologies for the evaluation of social HRI and this is an open problem that we address in this dissertation.

Our approach to social HRI evaluation was to generally focus on qualitative-oriented exploration and describing participants' experiences, more than on measures of task completion or efficiency, using, for example, participant self-reflection, interviews, video, and long-answer questionnaires. Where numerical measures are used, such as time, we attempt to relate them to the participants' experiences and context of interaction. Much of our approach has emerged from our own experiences of conducting the evaluations, and throughout our chapters we highlight reflections on our application of these tools to robots.

In this dissertation we have explored the question of how to perform social HRI evaluations through our various studies listed above. These experiences have played an integral role in the development of our social HRI theoretical foundations, and we distill our experiences as an inclusion in our social HRI-targeted heuristics (Chapter 8).

1.5 SIGNIFICANCE

Our research constituted significant contributions to the relatively-new research field of HRI where, until recently, most robotics interaction efforts have focused on the engineering and direct control of the robot. Interaction generally followed direct remote-control or tele-presence paradigms with task efficiency and effectiveness being primary measures of success. The idea of people having a strong inclination to treat robots as social entities is only now emerging explicitly within the academic research fields of robotics and HRI, and questions of why this happens or how the existence of social components effects overall interaction are still relatively unexplored.

Our work is among the first to tackle the high-level questions of social HRI directly. We are the first to provide a comprehensive discussion on and definition of social HRI, including consideration of why robots are unique and why people treat them as social entities. We provide a new set of heuristics to serve as a straightforward social HRI design tool-box. We designed, implemented, and evaluated original HRI interfaces that leverage social interaction with people. In addition to the immediate contributions of our analytic discussions, robotic-platform implementations, and participant studies, we believe that our research exposes and describes interesting and useful social relationships between people and robots, information and understanding that is of interest to the greater HRI community.

1.5.1 *Contributions*

SOCIAL HRI FRAMEWORK — We contribute to the groundwork of social HRI theory, the body of ideas, principles and techniques distinct from particular implementations. We develop ideas of why people are inclined to treat robots as social entities and why this can be desirable, build the first definition of the meaning of social HRI, highlight (and partially map) the broad scope of social interaction with a robot, and present new vocabularies for describing and analyzing social HRI instances.

SOCIAL HRI HEURISTICS — We present a concrete set of heuristics for designing and evaluating social HRI, drawing from our entire dissertation: the theoretical contribution, and our experiences with designing, building, and evaluating social HRI interfaces.

SOCIAL HRI DESIGNS AND IMPLEMENTATIONS — We designed and implemented four novel social HRI interfaces that leverage people’s existing skill sets and their tendency to treat robots as social entities. Some of these designs have several vastly different implementations that illustrate different possible realizations of the same interaction approach, providing comparison points for analysis.

SOCIAL HRI EVALUATIONS — We detail two informal design critiques and seven formal studies, where we report on people’s experiences interacting with social HRI instances. Through these studies we investigate the question of how to conduct social HRI studies and explore various evaluation techniques and methods, summarizing our experiences both in each respective section as well as in our heuristics.

1.6 STRUCTURAL OVERVIEW

The remainder of this dissertation is organized as follows: in Chapter 2, we present an extensive treatment of the related literature relevant to our social HRI exploration. Chapter 3 is a thorough theoretical exploration into social interaction between people and robots, culminating in a theoretical framework for social HRI. Part II of this dissertation (from Chapter 4 to Chapter 7) details our various interface designs, implementations, and evaluations. We present a set of concrete design heuristics in Chapter 8 for the design and evaluation of social HRI instances, and conclude in Chapter 9.

A SNAPSHOT OF THE SOCIAL HRI LANDSCAPE

In this chapter we review current social Human-Robot Interaction (HRI)-related and relevant background, with a focus on outlining how our work fits within related research. We begin by highlighting current work in HRI, ranging from well-established task-oriented research to the more-recent, and more exploratory, social-oriented work. Following, we present existing theories which we later use to outline the fundamental reasons and motivations behind why people perceive robots in the way that they do and what this means for interacting with robots. We finish by outlining fundamental Human-Computer Interaction (HCI) work which we will later use to develop a specific framework for the study and evaluation of social HRI.

We first outline the large body of existing “classic” HRI research that revolves around the direct use of robots as tools for specific tasks. This research commonly considers how to enable people to accomplish these tasks effectively and efficiently, focusing on such issues as control and feasibility of communication for both collocated and remote robots. In addition to a wide range of implementations, work here includes theoretical frameworks for understanding what makes a control interface useful or successful.

Secondly, we discuss the emerging body of research in HRI that is explicitly focusing on social interaction, with current explorations increasingly showing that people do indeed have unique and strong social, emotional, and anthropomorphic reactions to robots. Some explore the boundaries of this reaction through carefully-designed studies, including questions of how robots integrate into social fabrics, and others look at how robots can (or should) explicitly leverage this social interaction in their designs.

Following, we briefly present fundamental related work from philosophy, sociology, and psychology that helps us understand the foundation behind the experience of interaction, and as such, will later help us lay the groundwork for understanding why people treat robots as social entities. As part of this, we also outline the (very limited) extent of existing work which directly attempts to explain various facets of social interaction with robots.

Finally, we present an overview of existing HCI methodology which we deem to be particularly relevant for our efforts to develop the study and evaluation of social HRI. In addition to

methods for exploring usability-oriented questions, we also discuss methods for evaluating questions of personal experience and context of interaction, emotion, and social norms.

Overall, we relate our work to these existing themes: we support the uniqueness of our overall approach in relation to established research, we outline the novelty of our particular interface designs in relation to current social HRI work, and highlight the current lack of methodology and theory for understanding and describing social HRI.

2.1 CLASSIC HUMAN-ROBOT INTERACTION

From the early days of robots in the 1970s and 1980s when only advanced engineers and programmers worked with them (Moravec, 1998), advances in robotic technology have pushed people to interaction roles beyond the more-traditional programmer or overseer to include a robot operator: thus the field of HRI has emerged (Kiesler and Hinds, 2004). Much of this “classic” HRI research has been on designing interfaces that enable a person to process and understand the state of the robot (or to gain HRI awareness, Drury, Scholtz, and Yanco, 2003) and to simultaneously provide the robot with appropriate movement and action commands.

In this section we outline classic HRI questions related to the remote control of robots (tele-robotics) as well as early work on collocated interaction with people. We do this as a means to highlight work that has generally fallen outside the social HRI focus and to show where social HRI enters.

2.1.1 *Remote-Control Robots*

One prominent focus for control-oriented HRI has been on remote-control robots (tele-robotics), for such applications as urban search and rescue, exploration (deep-sea, volcanic, space, etc.), or military reconnaissance; remote-control robots were used in the 9/11 rescue efforts and are now commonly used in search and rescue (Davids, 2002). There is a general problem of complexity surrounding remote control, as an operator must remotely monitor and have constant real-time awareness of the robot’s state and environment, perceived through its multi-dimensional and abstract sensor and state data. This data includes such variables as robot speed, cameras, or proximity sensors used to avoid collisions with the

environment, and can even include the robot's current morphology and configuration of moving parts such as arms. Imagine the common scenario of operating a robot, while under task-related pressure, via a camera mounted at the end of a multi-joint arm (see, e. g., the iRobot Packbot and controller shown in Figure 2.1); it can be quite easy to forget that the camera may not be pointing forward and, as such, to provide erroneous direction commands to the overall robotic platform based on the visual camera feed.

This general issue has been posited in classic HRI as a problem of *awareness* between the controller and robot, where the person needs *awareness* of the overall robot state and environment, and the robot needs *awareness* of the person's perspective in order to properly interpret commands. Drury et al. (2003) defined this HRI awareness as:

HRI awareness is the understanding that the human has of the location, activities, status, and surroundings, of the robot; and the knowledge that the robot has of the human's commands necessary to direct its activities and the constraints under which it must operate.

— Drury et al. (2003)

In fact, it was precisely a lack of HRI awareness that Drury et al. (2003) found to be the primary cause for mistakes made in urban search and rescue robot competitions and trials. (RoboCup Rescue, the urban search and rescue branch of the famous RoboCup competitions, has a very prolific and active community that meets for annual competitions.) This points to



Figure 2.1: the iRobot Packbot and remote control mechanism

the importance of HRI awareness to the wide acceptance of tele-operated robots in several (relatively narrow) domains such as military applications and space exploration.

There is a great deal of work in HRI that focuses on improving the awareness and control mechanisms provided by remote-control interfaces. One focus is the attempt at clever incorporation of sensor data into the interface (e. g., Baker, Casey, Keyes, and Yanco, 2004; Drury, Yanco, Howell, Minten, and Casper, 2006b; Kadous, Sheh, and Sammut, 2006; Sellner, Hiatt, Simmons, and Singh, 2006; Yanco, Baker, Casey, Chanler, Desai, Hestand, Keyes, and Thoren, 2005), with one project on improving tele-robotics effectiveness through inspiration from video game interfaces (Richer and Drury, 2003) and another coupling the input and output spaces via tangibles (Lapides, Sharlin, and Sousa, 2008). Other efforts are more specialized, for example, targeting the remote control of robot teams (Squire, Trafton, and Parasuraman, 2006), specifically unmanned aerial vehicles (Drury, Riek, and Rackliffe, 2006a), or attempts at developing metrics for evaluating the effectiveness of these interfaces (e. g., Jacoff, Messina, Weiss, Tadokoro, and Nakagawa, 2003).

Tele-robotics is an extensive research domain which we do not attempt to cover in a detailed manner. Rather, we present this to clearly differentiate this established area from our own work and focus on social HRI.

2.1.2 *Collocated Robot Operation*


The control of collocated robots demands an interaction paradigm very different from that of real-time remote control. Yanco and Drury (2004) addressed this as part of their updated awareness framework (original introduced above, Drury et al., 2003), adding components of human-robot proximity to the nature of control: the taxonomy added the five human roles of supervisor, operator teammate, mechanic / programmer, and bystander. This adds a focus on the overall human-robot control relationship in a much-less-mediated fashion than traditional interfaces.

One approach to collocated control is to interact via gestures presented directly to the robot, such as pre-decided hand gestures (Koenig, Chernova, Jones, Loper, and Jenkins, 2008) or pointing (Giesler, Salb, Steinhaus, and Dillmann, 2004; Sato and Sakane, 2000), sketching (Kemp, Anderson, Nguyen, Trevor, and Xu, 2008; Sakamoto, Honda, Inami, and Igarashi, 2009, see Figure 2.2a), with some work focusing on robot awareness of the person

and whether they are too close for safety (the robot knows it can injure, Sisbot, Clodic, Alami, and Ransan, 2008) or how to mix autonomy and direction (Wang and Lewis, 2007).

A theme with designing robotic interfaces for collocated robot operation has been to use graphics displayed onto the physical world (mixed reality) as a means to highlight the command given and to provide feedback of the command's progress (Dragone, Holtz, and O'Hare, 2006; Giesler et al., 2004; Sato and Sakane, 2000; Ishii, Zhao, Inami, Igarashi, and Imai, 2009), or to use Tangible User Interfaces (TUIs) as a means to map gestures to robot commands, for both robot steering and for robot pose definition (Guo and Sharlin, 2008, see Figure 2.2b). There has also been work on augmenting and scaffolding the environment directly, for example, using physical tags, to denote way-points or locations for commands (Marquardt, Young, Sharlin, and Greenberg, 2009; Zhao, Nakamura, Ishii, and Igarashi, 2009).

While research surrounding collocated robot operation is still often “classic” in how it approaches HRI (as a low-level interaction and efficiency task) it is a major step toward social HRI in comparison to the tele-robotics work presented above. Considering issues related to collocation starts to consider the person as a human being in their own body (not at a Personal Computer (PC) interface) and environment, complete with emotions and social structures; it looks at the person's embodiment (Dourish, 2001b). This is the direction we take for social HRI in this dissertation. We introduce existing research regarding these social, human-oriented aspects in the following section.



(a) Sketch and Run robot interface (Sakamoto et al., 2009), for collocated robot control (b) tangible user interfaces for HRI (Guo and Sharlin, 2008)

Figure 2.2: interfaces for the collocated control of robots

2.2 SOCIAL HRI INTERFACE DESIGNS AND OBSERVED PHENOMENA

Ongoing research on the kinds of classic HRI outlined in the previous section plays an important role in understanding how to design effective and efficient control interfaces for robots. However, as robots become a part of everyday lives for many people it becomes increasingly important to consider robots' roles and interactions beyond task effectiveness, particularly as an increasing body of research suggests that people naturally treat robots as social entities (as we outline below).

Overall, the work in this section is shifting away from viewing robots simply as mechanical tools and is exploring the idea of social and affective, and emotive interaction with robots. This sets the stage for our own social HRI work presented in this dissertation, and we use this opportunity to highlight the novelty of our particular social HRI designs.

This section has two components. First, we present work that both supports and highlights how people treat robots as social actors: that people tend to respond to them socially and that they have impacts on the social environments they occupy. Second, we discuss research that has explicitly attempted to leverage this tendency and design robots that directly use social communication mediums.

2.2.1 *People Respond Socially to Machines*

Most of us, at one time or another, have felt emotion toward machines, for example, frustration at them not working properly or sadness at a favourite device breaking (Picard, 1999). Reeves and Nass (1996) have designed and conducted various studies to explore the nature of these emotional reactions, considering how far people take the response, and how the emotional reactions compare to similar ones that may be felt toward a person. The results of their experiments are summarized in their book, *The Media Equation* (Reeves and Nass, 1996).

The series of studies outlined in *The Media Equation* show how people naturally tend to respond to machines, and computers, in much the same way as they respond to other people or living things, readily applying social rules and norms to technologies. Examples include people treating computers differently based on a male or female voice, how a large on-screen face can invade people's notions of personal space, and that people are generally polite to computers in similar ways to how they are polite to other people. The authors conclude

that the evolution of the human brain has not prepared it for modern technology, and so reacting to non-social technologies that portray life-like characteristics as social entities is a hard-wired instinctual response. As part of this work, the authors highlight how the design of a technology can influence and strengthen these reactions, and recommend that designers strategically leverage emotional response.

Follow-up work by Nass and Moon (2000) has also highlighted how people apply *learnt* social categorizations and generalizations to machines, including ideas of gender stereotypes and reciprocity. For example, in one study they found that people felt a television labelled as a *specialist* (in a given area) provided better content than one labelled as a generic television that provided a range of content, although the quality of the content did not change. This is an example of how following social norms can influence perceived usability and effectiveness.

2.2.1.1 *People Respond Socially to Robots*

As people respond socially to machines in general, we posit that it follows that they respond socially to robots: robots, as machines, also elicit social and emotional responses from people. However, we believe that robots are unique from other technologies in that the reactions they elicit are much more pronounced. That is, people's social and emotional reactions to robots are very salient and evident, and people themselves are acutely aware of these reactions.

We present research in this section that highlights how people's mental models of robots tend to be more anthropomorphic or zoomorphic than they are for other technologies or systems. While some of these models are generated undoubtedly by intentional design, for example, when people zoomorphise the dog-like Sony AIBO (Friedman, Kahn, and Hagman, 2003), similar reactions have also been shown to emerge when interacting with robots that have more mechanical, and arguably not zoomorphic, designs such as iRobot's Roomba vacuum cleaner or Packbot military robot (Figure 2.3), as discussed below.

Several studies have shown that some people treat the Roomba as a kind of pet (Forlizzi, 2007; Forlizzi and DiSalvo, 2006; Sung, Guo, Grinter, and Christensen, 2007), giving it human-like motivations and characteristics. Families have been found to regularly give names to their Roombas (consider how many families name their standard vacuum) and some talk to the Roomba while it works (Forlizzi, 2007; Forlizzi and DiSalvo, 2006; Sung et al., 2007; Sung, Christensen, and Grinter, 2009a), for example, saying "excuse me" when bumping into it, or calling it "dumb" or "pathetic" when it gets stuck (Forlizzi, 2007).




Figure 2.3: examples of functionality-oriented robot designs that people tend to anthropomorphise and zoomorphise

Sung et al. (2007) found that people attributed complex personalities to the robots as a means to explain the often quirky or inconsistent mechanical and operational characteristics of their particular Roomba. In the same study, Sung et al. (2007) found that in one particular case, a family expressed sadness at having to exchange their broken Roomba, named “Spot,” for another one rather than to fix Spot.

Similarly strong anthropomorphic inclinations have also been reported in military settings: soldiers have awarded robots “battlefield promotions” and “purple hearts” to their robots, and in at least one case demanded that a particular damaged iRobot Packbot (Figure 2.3) named “Scooby-Doo” be repaired instead of replaced at a fraction of the cost. In another case, an Army colonel cancelled a mine-sweeping-robot exercise as the robot was getting mutilated, stating that the exercise was “inhumane” (Garreau, 2007).

Some have developed targeted studies to try to understand the depths and relationships behind these reactions. Fussell, Kiesler, Setlock, and Yew (2008) considered how anthropomorphism tendencies can be dependent on experiences with robots. Sung et al. (2007) looked particularly at ideas of intimacy between a person and a robot, how this intimacy can be influenced by design and customization (Sung, Grinter, and Christensen, 2009b), and how this relates to acceptance and anthropomorphism. Others have looked at how hesitant people are to switch off a robot (Bartneck, van der Hoek, Mubin, and Mahmud, 2007a) or to “kill” a robot by smashing it with a hammer (Bartneck, Verbunt, Mubin, and Mahmud, 2007b), in relation to the level of intelligence portrayed by the robot (Figure 2.4). In both


experiments, people had increasing reservations about shutting off or “killing” the robot as the perceived intelligence level increased, even though in one case the robot had a very mechanical design (Figure 2.4b).

We have highlighted how people exhibit a tendency toward viewing robots as social entities, and that they apply social practises and rules to understand and interact with them. HRI research which directly targets these social questions is rare, and we have presented a bulk of the full scope here. In this dissertation we draw heavily on this social tendency toward robots with our own social HRI designs, and attempt to further the understanding of the phenomenon through our theoretical explorations and interface evaluations.

2.2.1.2 *Robots and Social Integration into Everyday Environments*

Robots not only provide cause for explicit social interaction, but also integrate into everyday environments as social actors. One theme in social HRI has been to look at how robots will integrate into people’s everyday lives. Several studies have looked at how the iRobot Roomba vacuum, arguably the most successful domestic robotic technology to date with millions sold, fits into social home environments (Forlizzi, 2007; Forlizzi and DiSalvo, 2006; Sung, Grinter, Christensen, and Guo, 2008), with some focus on long-term use (Sung et al., 2009a).

The Roomba has been found to have a “substantial and lasting impact” (Forlizzi, 2007) on people and social structures of homes. Much of this impact surrounds how people have been found to make changes to their work practises and environments as a means of ensuring the



(a) a Phillips iCat that begs to not be turned off (Bartneck et al., 2007a)



(b) participants were asked to “kill” the robot with a hammer (Bartneck et al., 2007b)

Figure 2.4: robots that people were asked to shut off, or “kill”

Roomba's success as a cleaning device (Forlizzi and DiSalvo, 2006; Forlizzi, 2007; Sung et al., 2007, 2009a). This includes the "roombarization" of environments (Sung et al., 2007), where people make significant changes to furniture layout or create special barriers to accommodate and support the Roomba. People were also found to make sweeping social changes to their cleaning and tidying habits and rituals to better suit the needs of the Roomba, for example, shifting from routine to more opportunistic and multi-tasked cleaning (Forlizzi and DiSalvo, 2006; Forlizzi, 2007). After introducing the Roombas, many cleaning duties shifted from a female-dominated to a male-dominated role, and even toward (older) children: they became more interested in cleaning due to the interaction with the Roomba (Forlizzi and DiSalvo, 2006; Forlizzi, 2007). For some, the Roomba became a pride point and social conversation piece, where they aggressively protected their own Roomba, both physically (against possible damage) and reputation-wise, and promoted the product to others (Sung et al., 2007).

There is also a body of work that considers Robots in the context of more public environments: people reacted quite positively to Valerie the "roboceptionist" (Gockley, Forlizzi, and Simmons, 2006, Figure 2.5a) who worked for nine weeks in an office environment. Robots have also been used in museums (Burgard, Cremers, Fox, Hähnel, Lakemeyer, Schulz, Steiner, and Thrun, 1999; Nourbakhsh, Kunz, and Willeke, 2003; Shiomi, Kanda, Ishiguro, and Hagita, 2006) to provide information and guide people around (Figure 2.5b). Some research suggests that the best way to attract interest may be for the robot to demonstrate awareness of a person's presence (Nourbakhsh et al., 2003), for example, the robots used by Shiomi et al. (2006) directly called people by name, made possible as museum visitors wore Radio Frequency Identification (RFID) tags for identification.

Other related research placed Robovie robots in train stations to give directions (Shiomi, Sakamoto, Takayuki, Ishi, Ishiguro, and Hagita, 2007) or to announce relevant train information (Hayashi, Sakamoto, Kanda, Shiomi, Koizumi, Ishiguro, Ogasawara, and Hagita, 2007). In the former example, Shiomi et al. (2007) found that people were very impatient with the robot's imperfect conversational abilities, and highlighted the fact that people are generally more patient with difficult-to-use computerized (and non-robotic) information kiosks, and Hayashi et al. (2007) also found that people were much less likely to ignore two robots that talked together about the announcements than a single robot that announced by itself: they hypothesized that people did not want to be engaged and felt that the robot pair would be less



(a) Valerie the “roboceptionist” (Gockley et al., 2006) (b) Robovie museum robots (Shiomi et al., 2006); the Robovie series were explicitly designed for effective social-level communication

Figure 2.5: examples of robots integrating into everyday environments

likely to engage them. This suggests that, at least in this case, people used their understanding of culturally-grounded social norms to decide how to interact with robots.

Here we have outlined how robots have impact on the encompassing social structures of environments that people occupy. In our dissertation work we formalize this level of social HRI and integrate it into both our theoretical discussions and interface evaluation techniques.

2.2.2 *Robot Design that Leverages Social Tendencies*

In this section we outline how researchers have been directly attempting to leverage the social and emotional tendencies outlined above, and explicitly use them in HRI design to facilitate natural and comfortable interaction between people and robots. We approach this discussion from two standpoints: first, we introduce how robots can leverage the more visceral-level and immediate forms of social interaction techniques, and second, we discuss how robots can try to fit into higher-level social structures. We relate these approaches to our own work in this dissertation, and highlight the limited scope and emerging nature of work in this area.


2.2.2.1 *Robots Leveraging Direct Social Interaction*

There is a body of research on robots that target socially-acceptable human-like practises. Some robots employ human-like speech practises of pauses, delays, and conversational fillers to keep from being annoying or confusing (Shiwa, Kanda, Imai, Ishiguro, and Hagita, 2008,

Figure 2.6a), or the use of deictic (context-specific) reference (Brooks and Breazeal, 2006), natural pronoun usage (Gold and Scassellati, 2006), or appropriately-timed head-nodding (Sidner, Lee, Morency, and Forlines, 2006) to structure conversations. Conversation can also be used to build common understanding (Stubbs, Hinds, and Wettergreen, 2007).

Robots can monitor a person's gaze to help decide how it should act (Atienza and Zelinsky, 2002), and can use gaze cues to influence interaction (Mutlu, Shiwa, Kanda, Ishiguro, and Hagita, 2009); people better comprehend a robot's speech when its gaze cues are human-like (Staudte and Crocker, 2009), and find it disturbing when a robot (purposefully, in this case) uses socially-inappropriate gaze cues such as suddenly looking at an irrelevant item (Muhl and Nagai, 2007). Robots can also use social meanings behind interpersonal distances and relative position, for example, for how it should *naturally* follow a person (Gockley, Forlizzi, and Simmons, 2007; Yamaoka, Kanda, Ishiguro, and Hagita, 2008, Figure 2.6b).

Some research has evaluated robots leveraging social techniques in real-world settings, for example, it was found that Valerie the “roboceptionist” (Figure 2.5a) could impact how people interacted with it based on the facial expressions it used (passive, happy, unhappy) (Gockley et al., 2006). Hayashi et al.'s (2007) train-station robots found they could change people's expectations and interactions simply by bowing. Similarly, a study on museum robots highlighted how robots could get people to pay more attention to them by social interaction, for example, via talking and hugging (Shiomi et al., 2006) or through interactive questions (Nourbakhsh et al., 2003) — people also expressed more enjoyment. This research also found



(a) a robot using conversation fillers to buy processing time (Shiwa et al., 2008) (original Japanese text replaced by English, right-most frame originally said ええと...)

(b) a robot attempting to follow naturally (Gockley et al., 2007)

Figure 2.6: examples of robots that try to work within established cultural and social norms

that people expressed little affection (treated robots coldly as machines) when curious about the robots' technical specifications.

Androids, humanoid robots that have the ultimate goal of passing for a human, are designed with the primary goal of interacting seamlessly using social techniques and appropriate behaviour. This is an active research area in HRI, with current examples including (Figure 2.7) the Repliee Q1 (MacDorman, Minato, Shimada, Itakura, Cowley, and Ishiguro, 2005), Geminoid (Sakamoto, Kanda, Ono, Ishiguro, and Hagita, 2007), and the CB² baby robot (Minato, Yoshikawa, Noda, Ikemoto, Ishiguro, and Asada, 2007). Androids are posited to work alongside people in the same way another person would, ultimately removing any robot-specific learning required to interact with them.

Rather than focus on humans, some research attempts to leverage people's familiarity with animals. Several projects embrace zoomorphism by making animal-like robots that use communication styles and cues that we are familiar with, such as the Sony AIBO (Figure 2.8a).



(a) Repliee Q1
(MacDorman et al., 2005)



(b) CB2 Baby
(Minato et al., 2007)

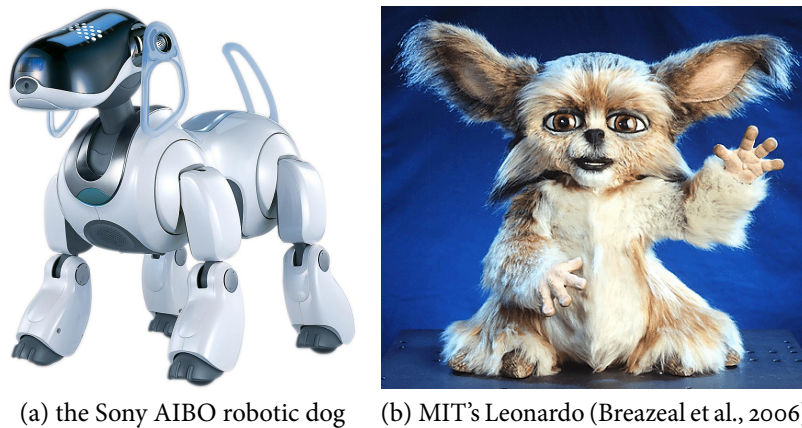


(c) The Geminoid and his creator (Sakamoto et al., 2007)

Figure 2.7: examples of androids

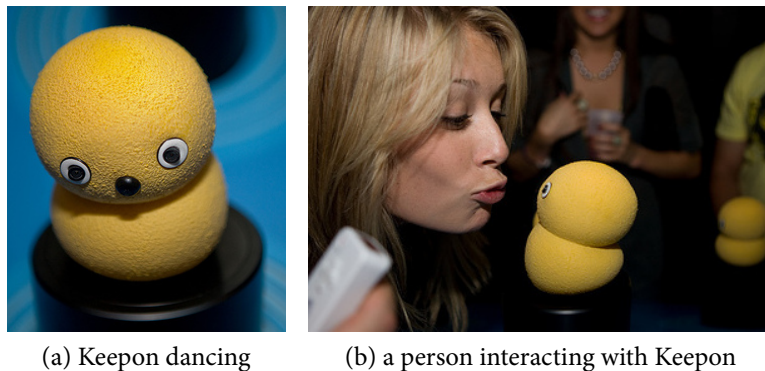
Similarly, Leonardo (Breazeal et al., 2006, Figure 2.8b) is a fantastical mammalian-like creature that uses facial expressions and full-bodied gestures (including its ears) to provide social feedback relating to its task, for example, showing whether it understands a command, or to communicate its state to people.

Some research looks to more subtle forms of interaction, for example, a robot named Keepon (Figure 2.9) looks at the use of simple rhythms to influence how people respond to it (Michalowski, Sabanovic, and Kozima, 2007). A core component of this robot is to use rhythmic movements, with music, to elicit strong emotional reactions of fun and pleasure; this goal is reinforced through its cute design. (We encourage the reader to view videos of Keepon on the Internet). Another project which focuses on movement style is a mechanically-designed search and rescue robot that uses its movement style to be more (or less) intimidating to the people they are rescuing (Bethel, Bringes, and Murphy, 2009).



(a) the Sony AIBO robotic dog (b) MIT's Leonardo (Breazeal et al., 2006)

Figure 2.8: examples of zoomorphic robots



(a) Keepon dancing (b) a person interacting with Keepon

Figure 2.9: the dancing Keepon robot, photos: Dave Bullock, eecue.com

The projects summarized in this section demonstrate how robots can leverage peoples' social interaction techniques to reduce the requirement for people to learn how to interact with them, and to communicate more seamlessly in ways people find natural and easy to understand. This approach is emerging and fairly new to HRI; our new interfaces for social HRI will add to this repertoire and expand the overall approach to new areas which have not yet been explored.

2.2.2.2 Robots Leveraging Higher-Level Interaction

In addition to more direct social interaction, robots can also explicitly work within higher-level social structures, such as those within groups of humans (teams). Human-robot collaborative teams are common in such applications as search-and-rescue, military operations, or even factories (Hoffman and Breazeal (2004) offer a comprehensive analysis of human-robot teams). Some research considers how robots can work within existing person-person team dynamics, for example, through exhibiting appropriate social and collaborative cues or by using other methods for sharing information and building common ground (Brooks and Breazeal, 2006; Fong, Thorpe, and Baur, 2002), for example anticipatory actions or turn taking (Argall, Gu, and Browning, 2006; Hoffman and Breazeal, 2007). Some have even looked at what happens if robots are on equal ground with, or hierarchically above, people (Xin and Sharlin, 2006, 2007). One ongoing project which explores many of these issues is NASA's Robonaut (Bluethmann, Ambrose, Diftler, Askew, Huber, Goza, Rehnmark, Lovchik, and Magruder, 2004), a robotic astronaut designed to work along with people as a team member.

Robots can conceivably work within existing social structures and interaction patterns in the domestic environment, although there is currently very little work in this area. We point to one idea from the study of intelligent environments, where Hamill and Harper (2006) (see also Hamill, 2006) suggested that intelligent technologies in general could learn from the interaction structures and patterns used with Victorian-era servants. We feel that robots could likewise leverage this well-developed, successful command structure to fit into this pre-established social knowledge.

People's natural (but sophisticated) teaching and demonstration abilities, and general familiarity with teaching, have been used for robot interaction. Primarily, these implementations are goal oriented where, for example, people teach robots particular paths (Saunders, Nehaniv, and Dautenhahn, 2006), or kinaesthetic motions such as for lifting an object (Gribovskaya

and Billard, 2008) or removing objects from a dishwasher (Dillmann, 2004). Breazeal et al.'s (2006) Leonardo (Figure 2.8b) learns button-pressing sequences via conversation with a person, using familiar social cues such as gestures and facial expressions to convey his understanding, confusion, and interest in learning (Breazeal, Brooks, Gray, Hoffman, Kidd, Lee, Lieberman, Lockerd, and Mulanda, 2004; Lockerd and Breazeal, 2004).

In our research we included components of the relatively-unexplored social HRI question of how robots can integrate into existing social structures, a question which is increasingly important as robots become more ubiquitous. We have attempted to integrate components into our interfaces, and have further addressed this level of interaction in our theory.

2.2.3 *Our New Social HRI Interfaces*

In this section we detailed the current state-of-the-art for social HRI designs and outlined how we believe our interfaces can meaningfully contribute. All our interfaces — cartoon artwork, dog-leash robot, *stylistic locomotion*, and *puppet master* — provide new ways for robots to leverage socially-understood techniques to help make complex human-robot communication and interaction accessible. Further, the dog leash and teaching interfaces in particular attempt to embed wider social structures into the robot interaction. An important element of our dog-leash robot is that we believe the robot-interface combination can integrate into public social structures (crowds) as the interaction concept is arguably familiar and easily understood by bystanders; they know how to react to it. Further, we explicitly consider the teaching question mentioned above for how a person can teach a robot, except that we take the novel angle of teaching of social *style* rather than task-oriented goal.

2.3 THEORETICAL FOUNDATIONS FOR UNDERSTANDING SOCIAL HRI

In this section we outline work that can help explain the reasons behind why people react to and treat robots differently than other technologies. We first present existing social HRI-related ideas which provide theory and understanding of social interaction between people and robots. We follow with various ideas from philosophy and sociology regarding interaction in general which we feel are particularly relevant for social HRI. We present the ideas of embodiment and embodied interaction to help understand the context within which people

exist and experience interaction. We also introduce select fundamental sociology concepts as a means to explain how social HRI happens within the greater context of society, and to help understand the impact of the word *social* in social HRI.

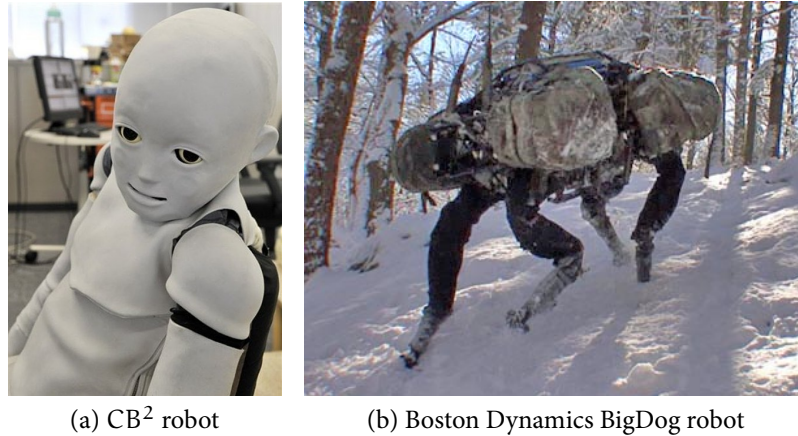
2.3.1 Existing Social HRI Theories

We have found very little work that directly addresses the core questions of understanding or describing social HRI. In this section we detail two such examples: a classification of robots which exhibit social properties, and a discussion of the problem of robot eeriness.

Breazeal (2003b) outlines a five-point classification of *social robots* for how robots can actively and intentionally leverage social interaction. There are *sociable robots*, those that use social models for their own internal purposes, and not necessarily for interaction with or external representation to people; as such this is not necessarily HRI. *Socially communicative* robots explicitly use human-like social cues and techniques to facilitate interaction with people. *Socially responsive* robots respond to social cues from people, either explicit (e. g., a smile) or implicit (e. g., subtle signs of impatience). Finally, *socially evocative* robots are designed to act in ways that promote anthropomorphism and encourage people to interact socially. Bartneck and Forlizzi (2004) also provide a robot-centric breakdown of how robots can be social, related to such categories as their form, knowledge of social norms, or modality.

These classifications serve as clear and useful tools to describe how robots leverage social interaction techniques themselves. For our work, however, we point out that the narrow scope of this classification is only a part of the wider social HRI picture. The need still remains to consider the broader social structures, both in terms of how they impact the perception of the robot and the robot's acceptance, as well how the robot itself impacts them. Further, this classification is robot-centric, and we argue that there is a need for a person-centric classification which outlines the person's reactions and perceptions, in addition to this framework which outlines the robot's socially-oriented design properties.

Another area of work in social HRI is the problem of robot eeriness, where some people find particular robots *eerie* or *creepy* even when they are not designed to be. Examples include the CB² robot of a human baby (Figure 2.7b, Figure 2.10a), or the Boston Dynamics BigDog robot (Figure 2.10b), which people particularly find creepy when the robot moves or, in the case of CB², makes baby-like noises; this is particularly evident in videos of these projects

(a) CB² robot

(b) Boston Dynamics BigDog robot

Figure 2.10: examples of robots that people find eerie

(in comparison to still images). Some researchers confront robot eeriness as a problem, as robots which elicit negative feelings may be received poorly by the people who work with it. We note, however, that eeriness can also be leveraged as part of a strategy for making robots unattractive, for example, for a night guard robot (Young et al., 2009).

One early theory which attempts to explain this is the “uncanny valley” (Mori, 1970, in Japanese with English translation, “bukimi no tani”, or, “不気味の谷”). Generally, this theory proposes that likeness to a human can be directly related to familiarity, where the more human-like a robot is, the more believable and comfortable people find it. However, Mori postulates, as likeness increases toward realistic, there is a breaking point beyond which comfort drops and the robot becomes eerie. Mori suggests the possibility of a conflict between what appears to be correct but, on some psychological level, what we can sense is wrong, a dissonance that Mori relates to working with prosthetic limbs or the deceased (Figure 2.11). Mori predicts that this “valley” of eeriness will not be overcome until robots mimic humans so well that we do not cue in on the fact that we are interacting with a robot.

Mori presented the uncanny valley in 1970 as a hypothesis, and it was not originally backed up with empirical evidence (Bartneck, Kanda, Ishiguro, and Hagita, 2009a; Mori, 1970). There has been a plethora of follow-up work seeking for evidence and developing a more comprehensive model (e. g., Geller, 2008), yet the theory itself is arguably difficult to test (Bartneck et al., 2009a) and some are now arguing that there may not be an uncanny valley at all (Bartneck et al., 2009a; Blow, Dautenhahn, Appleby, Nehaniv, and Lee, 2006). Rather, research has been pointing to a much more complex nature to the phenomena of eeriness, includes such variables as realistic appearance, behaviour, motor skills, quality of




Figure 2.11: a graph representing the “uncanny valley” of familiarity toward robots (Mori, 1970)

interaction, or facial expressions (Bartneck et al., 2009a; Brenton, Gilles, Ballin, and Chatting, 2005; Hanson, Olney, Pereira, and Zielke, 2005; MacDorman et al., 2005; Minato, Shimada, Ishiguro, and Itakura, 2004; Mori, 1970; Walters, Syrdal, Dautenhahn, te Boekhorst, and Koay, 2008). Further, eeriness is not limited to “human-like robots” as Mori suggested, as the uncanny valley ideas have been used extensively in animation (Bartneck et al., 2009a), and animal-like robots such as BigDog (Figure 2.10b) also have issues with eeriness.

Research on the problem of eeriness provides one of the only detailed explorations into a higher-level social HRI problem, as well as a glimpse at reasons behind why people may react to robots in certain ways. The eeriness problem itself demonstrates how important social HRI elements and concepts can emerge prominently via often subtle cues and traits, supporting the more-general idea of how subtle interface traits can have important social HRI implications. This has implications for all of our interfaces, particularly our project on conveying stylistic, interactive motion paths through often-subtle details (*stylistic locomotion*).

This section has introduced the primary extent of overarching social HRI methodologies and theory. Below we bring additional ideas to the discussion which we feel are relevant for social HRI, and will be particularly useful for our theoretical discussion (Chapter 3).

2.3.2 Embodiment

We discuss the concept of embodiment as a means to help describe the context and meaning of interaction between a person and a robot; we believe that the idea of embodiment can help

explain why people tend to treat robots as social entities. We first talk about embodiment in general, and follow with a discussion on how interaction itself is embodied (embodied interaction). We start by consulting the dictionary (WordWebOnline, 2010, “embodiment”):

“*embodiment.*” –

1. A new personification of a familiar idea. “*the embodiment of hope*”
2. A concrete representation of an otherwise nebulous concept. “*a circle was the embodiment of his concept of life*”
3. Giving concrete form to an abstract concept.

The term *embodiment* relates to the basic idea of encapsulation and attribution, where abstract ideas are given, conceptually, a more concrete form, and this form is then attributed with many of the properties and characteristics of the more abstract idea.

In artificial intelligence, embodiment is often used to attribute algorithms to particular entities or agents. For example, the famous Eliza chat bot was a personified embodiment of an underlying pattern matching algorithm (Weizenbaum, 1966). This sort of embodiment can represent algorithms spanning machines distributed around the globe, advanced sensor networks, cameras, or any other component available in the computer system, with the complex whole often being attributed to (embodied within) simple computer entities, animated characters, or even physical robots (Ziemke, 2001).

The form that the embodiment of these computer agents takes, whether it be a chat bot, animated character, or robot, is important in defining how the agent interacts with and exists within its environment. This “structural coupling” (Ziemke, 2001) between an agent and its environment defines the complete context under which the agent interacts with its environment. This context is very wide-sweeping, including the particular virtual or physical environment for the agent, whether people are involved, as well as the social structures and layers of that environment. Thus, for social HRI, the robot’s embodiment and *structural coupling* define how it must interact with the world within which it finds itself.

The embodiment of a robot has wide-reaching implications for social HRI. As an illustrative example, robots are often used to test models of human behaviour (e. g., Atkeson, Hale, Pollick, Riley, Kotosaka, Schaal, Shibata, Tevatia, Ude, Vijayakumar, and Kawato, 2000). For this application, the entity must arguably have a very human-like embodiment to ensure the

validity of any comparisons made between the results and real people (Ziemke, 2001). Even with advanced artificial intelligence programming (e. g., using neural networks), a machine's environment and *structural coupling* with it are fundamentally different than that of a person, and so the power to compare is inherently limited.

A similar argument can be made against the possibility of using machines to fully emulate human intelligence. For example, the *Chinese Room* argument contends that, despite acting in an intelligent way, a machine can never have understanding in the same sense as people do, due to fundamental differences in context: the machine's embodiment is fundamentally different and so is the derived meaning of interaction (Ziemke, 2001). For the same reasons, there have been questions raised as to the validity of using the HCI Wizard of Oz evaluation technique (Dahlbäck, Jönsson, and Ahrenberg, 1993; Maulsby, Greenberg, and Mander, 1993) to explore and test *social* interaction between people and robots; the discussion on the validity of Wizard of Oz for social interaction was an ongoing theme during Q&A at the HRI 2009 conference. The argument goes that, even though the *Wizard's* actions are filtered through the robot, the person's embodiment in the world gives them a fundamentally different standpoint from anything viable (arguably, possible) in a programmed robot behaviour, and so the end result may not be comparable to a similar behaviour performed by a robot: The difference is between testing a *social robot* and a robot that is *puppetteered to act socially*.

These examples help to illustrate the fundamental impact that embodiment has on the meaning behind interaction, and thus the importance of considering what the implications of embodiment are for social HRI. The key messages from this section are that robots' perception of the world is fundamentally different because of their particular *embodiment* (e. g., computer program running on a robotic platform) and *structural coupling* to the world (e. g., the particular set of sensors, actuators, and perhaps even network connections). We continue this discussion below.

2.3.3 Embodied Interaction

In addition to the concept of embodiment applying to robots, a seminal theme in HCI is the idea that *interaction* itself is embodied. We feel that this concept is particularly relevant to social HRI given the complexity and diversity of contexts within which interaction can

happen with a robot, and the unique role that robots can play in these contexts. In this section we introduce the concept of embodied interaction and how it relates to social HRI.

Dourish's (2001a) (see also Dourish, 2001b) idea of embodiment in general is very focused on what he describes as "real-time" and "real-space," perhaps a contrast to the more general ideas of embodiment presented in the previous section that includes such things as virtual actors. The reason for this difference becomes clear when one realizes that Dourish's stance completely surrounds questions of human experience (through his focus on phenomenology) and the meaning of interaction, all from a person's point of view. As such we feel Dourish's stance can be particularly informative for person-oriented social HRI.

People are embodied in the real world with a well-defined (but complex) structural coupling, including such things as our basic senses of touch, sight and hearing, as well as more abstract social layers. Similar to the general case of embodied entities as introduced above, then, the argument goes that humans must rely on their embodiment, and the details of their *structural coupling* to the real world, to build all understanding and meaning of the world. Dourish (2001a) describes this concept as *embodied interaction*.

"...by embodiment I mean a presence and participation in the world...So, physical objects are certainly embodied, but so are conversations and actions. They are things that unfold in the world, and whose fundamental nature depends on their properties as features of the world rather than as abstractions. So, for example, conversations are embodied phenomena because their structure and orderliness derives from the way in which they are enacted by participants in real-time and under the immediate constraints of the environment in which they unfold."

— Dourish (2001a)

Dourish's description of embodied interaction is an attempt at defining a person's structural coupling to their environment, that is, the interface through which people must interact with the world and construct meaning. Dourish follows Heidegger in discarding Cartesian dualism, the idea that mind and body are separate (Dourish, 2001a, p. 9), and pushes the idea that thinking and understanding itself happens inherently through our interactions with the world. All meaning and understanding (from a phenomenological perspective) derives from our experiences in the real world. This has powerful implications for robots — and our interactions with them — because through their physical manifestations, robots themselves are dynamic actors in the real world, with many similarities to living entities.

The power of the idea of embodied interaction is the perspective it offers, as it draws attention to the kinds of interactions which are more natural to people and also explains why they are natural: they emerge from the properties of our embodiment. This can be used to motivate the fields of tangible computing (see Ishii and Ullmer, 1997) and social computing (e. g., as with affective computing, see Picard, 1999) under the same explanation (Dourish, 2001a), as our combined physical abilities in the world and the history of our social experiences is the complete base for our understanding. Dourish's embodied interaction has particular relevancy for social interaction with robots (i. e., social HRI), and can be used to help explain how and why people interact with robots in the way that they do. We detail our interpretation of this relationship in Chapter 3.

2.3.4 *Sociology in Relation to Social HRI*

People are inherently social beings, and as such, our social interactions and our roles in society play a prominent part in every aspect of our lives, including our interactions with robots. Sociology, the study of society and human social activity, provides socially-oriented models and methods for how people interact with and understand the world and things in it (Giddens, 2009). As such, we believe that sociology can play an important role in describing and understanding social HRI. In this section we introduce the concept of how meaning and knowledge is socially constructed (in comparison to technological determinism) and a theory for articulating and describing the complex network behind this construction, and relate them to social HRI. We later use these as part of our discussion on the meaning of social HRI itself, in Chapter 3.

Social constructionism is the study of how meaning, knowledge, and reality is a construct of social forces (Berger and Luckmann, 1966). That is, all knowledge, and indeed the current state of society, is a construct of social relationships and interactions rather than being based on an ultimate truth or inevitable outcome of natural forces. Meaning is defined largely by the “social stock of knowledge” (Berger and Luckmann, 1966), including such things as customs, social structures, and shared knowledge.

Within the area of social constructionism are many theories surrounding the role of the “social” in the development and meaning of technology. These approaches generally reject the idea of *technological determinism*, where technology is seen as being neutral and

as developing along a natural, predictable and predetermined logical path outside social influence (Chandler, 1995). Rather, many argue that technology is intertwined within the social world, where technological development is driven by social forces, the very meaning of technology is socially constructed and realized (Bijker, 1993; Callon, 1980; Pinch and Bijker, 1984; Williams and Edge, 1996), and that technology itself is a social (political) actor in the world (Winner, 1987). This points to the idea of *interpretative flexibility* (Pinch and Bijker, 1984), which accepts that meaning, or interpretation, is flexible and dependent on the context and the people involved.

Of particular interest to us is Actor Network Theory (ANT), an approach to understanding and describing networks of interactions between *actors* and *networks*, including all material artifacts (such as automobiles) and non-material concepts (such as social forces) (Callon, 1987; Latour, 1987). ANT highlights how the meaning of entities can be a combination of both materials and concepts such as how a functioning school requires people, technologies, and ideas. Further, ANT defines a recursively-layered structure of the network, where all networks are entities themselves, and every entity is constructed from a network, for example, a school is simply an actor in a larger city network, while the school itself is a complex network. At the core of ANT is the idea that all entities in a network can have *agency* (the capacity to act in the world), and thus all entities can have meaningful influence and interaction with the network. We believe that ANT is useful for highlighting how people's perceptions of robots is the product of a complex network of technologies, social perceptions and expectations, media constructions, robot developers, industrial processes, and so forth, and that robots are active and influential actors in the networks (spaces, environments, contexts) they occupy.

The socially-constructed nature of the above theories highlights the sheer complexity and diversity behind meaning for a person and their interpretation of interaction with robots, and we see these as pointing to the wider context of social HRI. In our theoretical exploration (Chapter 3) we discuss what we believe this, and the idea of *interpretative flexibility* means for social HRI research.

2.3.5 *Theory of Mind*

Psychology tells us that people develop concepts and opinions of subjective “mental states such as ‘desiring,’ ‘knowing,’ and ‘believing’” (Whiten, 2006) as an important part of social

interaction with and understanding other people (Premack and Woodruff, 1978; Whiten, 2006). We believe that this *theory of mind* concept, the idea of entities (people or robots) having a model of another entity's mind, can be particularly relevant for social HRI.

Some work in this area considers how robots themselves develop theories of people's minds (Breazeal, Buchsbaum, Gray, Gatenby, and Blumberg, 2005; Scassellati, 2002). There are as of yet few social HRI projects or studies on people's theory of mind of robots, and results to date are self-described as being inconclusive, highlighting the need for future work in this area. In one particular study, Hegel, Krach, Kircher, Wrede, and Sagerer (2008) had set up a game (prisoner's dilemma) where a person would play against another person, an anthropomorphic-design robot, a functional-design robot, and a laptop (Figure 2.12). This study used brain scanning to examine how people perceived their opponent, where areas and types of brain activity were compared against existing understanding of brain activity patterns. This study found that people worked hard to understand the "intentionality" for all opponents, regardless of whether they were human, laptop, or robot, and preliminary results suggest that people empathized more with the anthropomorphic robot than the functional robot (Hegel et al., 2008). In a different study, researchers used descriptions of robots versus descriptions of people, combined with a game-like scenario, to judge how people perceive the intentions and capabilities of robots (Levin, Killingsworth, and Saylor, 2008). This project focused on predisposition rather than theory of mind that develops purely from interaction.

We believe that the concept of theory of mind, and how it may be applied to robots, can be very informative for social HRI and can be a useful perspective for exploring how people may react to and interact with a robot, a perspective we develop further in Chapter 3. The idea of



Figure 2.12: "theory of mind" experiment setup, left to right, anthropomorphic robot, computer, human, functional typing robot (not visible) (Hegel et al., 2008)

theory of mind also helps to explain why people anthropomorphise robots: they attribute a mind to perceived intelligence and autonomy.

2.3.6 *Summary: Theoretical Foundations*

In this section we have outlined, to the best of our knowledge, the limited scope of existing social HRI theory and methodology for describing and explaining social interaction between people and robots, and as such highlighted the need for further work in this area. Following, we have presented various ideas from philosophy and sociology which we believe can be relevant for understanding social HRI, and which to our knowledge have not previously been directly applied to social HRI, a connection we explore in full detail in Chapter 3.

2.4 STUDYING AND EVALUATING HUMAN-ROBOT INTERACTION

Observing people interacting with interfaces, and learning from these observations, is a core methodology in HCI; it explicitly extends the focus on the *human* component of HCI to beyond the interface designers themselves. This practise has followed to HRI and social HRI. Our goal in this section is to provide a selected summary of methodologies, techniques, and concepts from both HCI and HRI and explore how these existing methods apply to the unique social properties of interaction with robots. This includes an element of highlighting where we feel additional methodology is needed for social HRI.

We develop our discussion from the following categories of approaches to evaluation: task completion and efficiency, emotion, and situated personal experience, and conclude with a discussion on frameworks for exploring social interaction with robots. These perspectives serve as a mechanism to add structure to our discussion, as a means to highlight the different sorts of questions asked in social HRI studies, and as a way to point to the void in social HRI-targeted methodologies that we address in this dissertation.

2.4.1 *Task Completion and Efficiency*

Given the effectiveness-oriented nature of most classic computerized tasks and computer interfaces, traditional HCI evaluation has often taken a task completion and efficiency ap-

proach to usability evaluation, focusing directly on how an interface supports a user in their tasks, actions, and goals (Dix et al., 1998; Eberts, 1994; Norman, 1988; Sharp et al., 2007). This trend also exists in HRI where questions explored often centre around control-oriented issues, performance quality, the person's tactical awareness of the robots' environment, error rates and action mistakes, etc. (e. g., Drury et al., 2003; Guo and Sharlin, 2008; Richer and Drury, 2003; Yanco and Drury, 2004).

In addition to the direct utilitarian importance, concrete measures of task accomplishment and efficiency can be used as part of wider, interaction-experience oriented explorations. For example, these quantitative measures can support other data which highlights points related to engagement and interest (e. g., through task completion time or number of pauses), or whether and how much a person understands what the robot is trying to convey (e. g., through error rates). These techniques alone, however, can only provide limited insight on the social aspects of interaction, and so we argue that other techniques are needed for a more comprehensive view of the social HRI experience.

2.4.2 *Emotion and Affective Computing*

Some research in HCI specifically targets socially-situated interactions between people and computing technologies, with a particularly strong focus on human emotion. Much of the research in this area is performed under the title of affective computing, a domain which explores how interaction with an interface influences the emotional state, feelings, and satisfaction of the person (Picard, 1999), whether through deliberate design (e. g., Bates, 1994) or as an incidental artifact of interaction even without affective design (e. g., Isbister et al., 2006; Picard, 1999). This also includes computer awareness of human emotions (Picard, 1999). Given the socially-situated nature and tendency toward social interaction with robots, we feel that this body of work is particularly relevant.

Evaluation of affective aspects of interaction can be based on the monitoring of biological features such as heart rate, blood pressure or brain activity, or enumerating and measuring apparent behaviour data such as the number of laughs, or number and duration of facial expressions (Desmet, 2005). These methods can serve to quantify the difficult-to-quantify social-oriented aspects of interaction with robots such as types and amounts of emotion or the person's social involvement. However, evaluators should note the limitations incurred

when using such methods. Arguably, the ability to understand the rich and multi-faceted nature of social interaction will be limited and the validity of the gained insight reduced when emotions are simplified to a set of external quantities and discrete categories (Isbister et al., 2006; Strauss and Corbin, 1998).

Other affective-computing approaches focus on participant self-reflection, where people directly report on their experience with an interface and how it makes them feel (e. g., see Bates, 1994; Boehner et al., 2007; Höök, 2005; Höök, Sengers, and Andersson, 2003). Examples include think-aloud techniques, interviews, or surveys. This has the added benefit of accepting participants as expert judges of their own social interaction experience (e. g., with robots). Creative or artistic techniques have also been used to help people reflect on affective aspects of the interaction that are difficult to express with words. One such example is the sensual evaluation instrument, which asks people during interaction to handle a set of abstract, molded props (Isbister et al., 2006; Sanders, 1992, Figure 2.13). These objects are used to represent emotional states, and participants are later asked to use the props as physical memory aids and descriptive tools for their experience. Self-reporting, regardless of the media and mediators used, has the complication of often being done in retrospect (after, not during, an experience) and relies on people understanding their own emotions and being introspective and confident enough to openly express their thoughts.

Affective computing techniques can be very useful for exploring how people feel and think about robots, and how the robot affects their emotional state. We believe that social HRI, however, points to a wider picture that includes such things as social structures, and how



Figure 2.13: the *sensual evaluation instrument*, used to help people relate to their abstract affective experiences (Isbister et al., 2006)

all of these concerns relate to the multi-faceted physical, cultural, and social context within which interaction is taking place.

2.4.3 *Situated Personal Experience*

A person's interaction experience is situated within a wide, socially and physically-embedded context that includes such things as culture, social structures, and the particular physical environment they are interacting with. As such, we argue that the experience itself is a very complex and elusive concept that is difficult to explore with evaluation.

Existing evaluation approaches that focus on personal experience (and the context within which it happens) often aim to describe interaction experience rather than to explicitly measure it. Some argue that it is important to accept the complex, unique, and multi-faceted nature of experience (as perfect understanding is perhaps impossible, Sengers and Gaver, 2006; Strauss and Corbin, 1998), and evaluation should aim to find themes and in-depth description of the complexity (Bates, 1994; Höök, 2005; Isbister et al., 2006). This stance explicitly recognizes the complex and embodied nature of interaction with robots and as such many of the related data collection and analysis techniques can be used toward this goal. In fact, an emerging body of work in HRI considers interaction as a holistic and contextual experience that addresses such issues as how a robot meshes into existing social structures (exemplified in Forlizzi and DiSalvo, 2006; Lee, Kim, and Kim, 2007; Sung et al., 2008, 2007).

The approach of accepting complexity often uses qualitative-oriented techniques such as thick, detailed description based on participant feedback and interviews (e. g., Vaida, Grinter, Ducheneaut, Edwards, and Newman, 2005), collecting multiple viewpoints (e. g., across participants), or more structured approaches such as Grounded Theory (Strauss and Corbin, 1998), culture or technology probes (Gaver, Dunne, and Pacenti, 1999), or contextual design (Beyer and Holtzblatt, 1998). Longer-term interaction or interplay with social structures and practices can be targeted with in-situ, context-based ethnographic (e. g., Crabtree, Benford, Greenhalgh, Tennent, Chalmers, and Brown, 2006) or longitudinal field studies (e. g., Forlizzi and DiSalvo, 2006; Sung et al., 2007).

Another concept which is important to this view of evaluation is the idea that each person and their experiences are unique. This means that rather than trying to find an average, representative user, context-sensitive evaluation should perhaps value that individuals have

unique, different and personally-grounded experiences (based on culture, gender, education, language, etc.), and evaluators should take care when attempting to generalize any affective experience across people (Boehner et al., 2007; Strauss and Corbin, 1998). Further, the evaluators themselves will have personal biases toward the robots, participants, and the scenario. This bias, which some argue is unavoidable, should be explicitly considered and disclosed with the evaluation analysis (Strauss and Corbin, 1998).

The involvement of social structures in social HRI highlights that, since robots are often viewed as life-like entities, it is possible that person-person structures and norms may manifest between people and robots. For example, perhaps the observer effect (Landsberger, 1958) may be pronounced when interacting with robots: as interaction between a boss and a worker may change when they are being videotaped due to more pressure to act in a socially acceptable manner, the same change may happen between a person and a robot.

While these approaches consider many of the wider social and contextual components of the social HRI experience, they do not directly target the lower-level considerations of a person's emotions. Further, there is no explicit consideration of *how* these techniques can be applied to robots specifically, and it is up to the evaluator to devise appropriate methods. As such, we maintain that there is a need for structures and methodologies that aid evaluators in applying specific techniques such as the ones outlined above to the evaluation of social interaction experience with robots.

2.4.4 *Frameworks for Exploring Social Interaction with Robots*

So far in this section we discussed how existing HCI and HRI evaluation methods and techniques relate to the complex and contextual nature of social HRI. Complementary to this, evaluators can use frameworks as a means of dissecting this holistic, complex whole into more-targeted and focused units or perspectives, and use this as a means to direct evaluation. Frameworks can provide common vocabulary, provide means for comparison, and can serve as sensitizing tools to help evaluators focus on particular concepts. In terms of HRI, we argue for the need of frameworks to help evaluate and target such social HRI-related concepts as personal comfort, internal emotional experience, and social integration.

One particularly relevant example in HCI is Norman's three-level framework for analyzing how people interact with and understand everyday objects (or products, in this case), with

an explicit concern for emotion (Norman, 1988). Norman's framework highlights the stages a person may go through when dealing with a product over time: a) initial, visceral impact, b) behavioural impact, or how a person feels during use, and c) reflective impact, the thoughts one has after interacting with a product. The tendency to treat robots as social entities, however, suggests that the robot may not fall into the standard "product" category and as such this framework is perhaps limited in targeting social HRI.

Closer to HRI is Drury et al.'s (2003) *HRI awareness* conceptual framework, and specifically, the awareness (understanding) that both the people and robots have of the social structures and activities within a group. This work focuses on robots as team members in goal oriented tasks, and does not consider interaction outside this professional role. Perhaps the most explicit social interaction framework for robots is the classification of robots based on their socially-charged design characteristics and capabilities (Breazeal, 2003a), although this work is focused only on the robot design (and not a person's experience) and stops short of considering the wider context or the more-general social interaction that may occur.

To summarize, within the breadth of existing evaluation techniques and methods in HCI and HRI that we present above, there is no clear method that specifically targets breadth and depth of the holistic interaction experience for interacting with robots. Further, there is a lack of frameworks which can synthesize various existing methods together to target the socially-embedded nature of interacting with robots. In the following chapter, we present our initial take on classifying this rich interaction into a set of articulated concepts.

2.5 SUMMARY: A FRAMING FOR OUR DISSERTATION

In this chapter we presented a snapshot of the research and ideas which we feel best frames social HRI. The discussions presented here speak directly to our research questions (Section 1.3.5, page 8) and contributions (Section 1.5.1, page 17) outlined in Chapter 1:

We introduced *classic* HRI and discussed how social HRI poses new and unique questions. We outlined the existing (relatively-limited) extent of social HRI work and highlighted that the wide question of how robots can leverage social interaction in their designs is only minimally explored. The exploration further provided evidence that people are found to have particularly strong social and emotional responses to robots, motivating the importance of understanding social interaction with robots and the reasons behind why people interact

with robots in the way that they do. In this dissertation we continue to explore this landscape through our novel social HRI interfaces for how to leverage people's social tendencies to make interaction easy to do and easy to understand. These implementations also serve as interaction and technical proofs-of-concepts, showing how social interaction can be implemented in practise, and provide a means to study fundamental social HRI questions directly in our interface evaluations.

In this chapter we summarized existing social HRI theory and highlighted a void regarding social HRI-targeting methodologies. Further, we introduced a set of sociological and philosophical ideas which have not previously been applied to social HRI but we feel are crucial for understanding and describing it. These ideas help us analyze the core questions behind social HRI, such as what the tendency to treat robots as social entities means for interaction between a person and a robot.

Finally, we have detailed existing approaches to interface evaluation in HCI and related them to the kinds of questions we believe will be important to ask for social HRI. This has highlighted how there are as of yet no methods which specifically target the broad spectrum of social HRI, a void which we address in this dissertation through presenting new social HRI-targeted frameworks and methodology to help social HRI researchers account for social aspects in their designs.

In this dissertation we contribute to the state of the art in HRI by providing one of the first high-level methodological discussions on the fundamental questions of social HRI, novel interface designs and implementations that contribute to the expanding scope of social HRI, and extensive user evaluations which help explore core social HRI questions as well as the question of how to evaluate social HRI.

DEVELOPING A THEORETICAL FRAMEWORK FOR SOCIAL HRI

In this section we develop a theoretical framework for social Human-Robot Interaction (HRI), a set of definitions and concepts which can be used to describe, analyze, and discuss the details of social HRI in general or for a particular social HRI interface instance.

We start by analyzing the meaning of the term social HRI from various perspectives. We consider what *robot* adds to the meaning, as well as an exploration of the implications of the word *social*, and use the concept of *embodied interaction* (Dourish, 2001a) to outline the unique properties of interaction with robots. We summarize this analysis into a concise discussion on why social HRI is unique in comparison to other disciplines in how it is situated within everyday human social contexts.

As part of our exploration, and as a means of gaining insight into how people interact with and accept robots into their social structures we present tools from the field of social psychology, with particular focus on social factors of technology acceptance; social psychology deals directly with how people's behaviours and experiences are influenced by society. Following, we perform a detailed analysis using these tools on actual robots.

We finish this section with a summary of our multi-faceted exploration into a concise theoretical framework for social HRI, a primary contribution of this dissertation. This helps to address our core research questions of why people tend toward social interactions with robots and what implications this has for interaction. This further provides analytical tools that potential social HRI researchers can use to help understand how to directly account for these implications, or to develop robots that leverage social interaction in their designs.

3.1 ANALYZING THE MEANING OF SOCIAL HRI

In this section we analyze the term *Social HRI*, outline how we use it throughout this thesis, and how we believe it can be used by the wider field. The keyword *social* is emerging in the field of HRI as a means to distinguish classic HRI from emerging work that includes more

human-oriented social components of interaction. The term social HRI itself is rarely used, and even less commonly explained or analyzed, and we believe that our discussion in this dissertation is the very first thorough attempt at explaining and understanding it.

Throughout the remainder of this section, we develop and present an analysis of social HRI. We do this by first exploring the terms *robot* and *social*, and how ideas from *embodied interaction* help to highlight the unique nature of social HRI.

3.1.1 *What is a Robot?*

HRI is distinct from Human-Computer Interaction (HCI) in the explicit replacement of the word *computer* with *robot*; in this section we explore the implications of this change and what it means for our overall exploration in this dissertation.

The general population, arguably, has a practical understanding of what a robot is, but we argue that most people would have difficulty coming up with a clear definition. Robotists and HRI practitioners generally resort to domain-specific definitions or simply rely on *common sense* understanding, where robots are often described using loose criteria such as: machines that have intelligent behaviour, resemble (physically and behaviourally) a human or animal, are mobile, are able to physically interact with their environment, and so on.

The word *robot* originally emerged from a play to mean an artificial worker (Capek, 1970, book version, original play in 1921), although there were important elements of social impact to the story. Since then development in industrial applications and general automation, science fiction media, as well as science-fiction-inspired advanced research has muddled and diversified the meaning of *robot*. As such, *robot* as a term is currently subject to a large degree of *interpretative flexibility*, where its meaning is depending upon social context, the particular people interacting with the (particular) robot and the task at hand, rather than according to some universal meaning (Pinch and Bijker, 1984). The *social* understanding of a robot has not yet reached consensus (*closure*). For example, while a toy company may sell an electric, walking toy as a robot, others may argue that it is not a robot due to the lack of intelligence. Fleck (1984) predicted a movement away from ideas regarding a do-all *universal robot* toward application-specific robots, and argued that social understanding of *robot* will similarly move toward specific domains and usages. This is expected to eventually lead toward consensus, providing a clear distinction between robots based on categories

such as task, operation setting, and level of autonomy, for example, industrial, military, and domestic robots, although this consensus has not yet been reached.

The pure novelty of robots to the general public — that they have very little in terms of similar technology to compare to — means that *interpretative flexibility* is a particularly important consideration for HRI in general: how will people regard and perceive robots that enter their space? Will robots be seen as just another electronic gadget or appliance along with the cell phone, microwave or home theatre system? Which experiences and ideas will people draw from to build expectations of robots, for example, science-fiction, media, and general understanding of computers? Given people's predisposed tendencies to treat robots as social entities, the consideration of *interpretative flexibility* has particular significance for social HRI researchers who want to understand people's social interactions with robots, or attempt to design interaction for particular social interactions. *Interpretative flexibility* can both help to explain why a robot was received in a particular way, and means that designers can develop robot behaviours and designs to target specific interpretations.

Throughout the rest of this chapter we explore angles on the unique properties of social HRI, articulating specific properties and characteristics of how robots impact interaction. We believe this can provide insight on how to approach social HRI problems, how to design solutions, how to build tools, how to evaluate interaction, and so forth. Also, this can also prove as a means to compare robots to other related, perhaps non-robot, technologies, and to highlight which social HRI methods can be applicable to other fields, and vice versa, although this is not a task we explicitly undertake in this dissertation. We conclude our thorough exploration below with our own definition of robot as part of our theoretical framework, Section 3.6.

3.1.2 *What Does Social Imply for Social HRI?*

In this section we analyze the meaning of the word *social* in terms of what it implies for social HRI. Our analysis is admittedly simplistic from the perspectives of sociology or philosophy, as indeed we do not claim contribution to these fields. Rather, we propose that our interaction-oriented discussion serves the targeted purpose of clarifying what is meant by social HRI, why the addition of *social* (to HRI) is important, and what this addition implies for researchers.

We point out that the use of the word social is widespread in both everyday language and academia, particularly in the adjective form, for example, social networking, social media, social issues, social welfare, social change, social insects, social studies, and so on are examples of greatly disparate — but commonly encountered — terms. We note that *social* is generally used to refer to the interactions of people, an interpretation supported by the dictionary: “1. relating to human society and its members.” (WordWebOnline, 2010, “social”)

In particular, we are interested in *social* in the sense of being inherently people-oriented, relating to the interactions that people have with the world. Berger and Luckmann (1966) discuss this general knowledge that people have under the terminology umbrella of *the sociology of knowledge*, where “everyday life presents itself as a reality interpreted by men [people] and is subjectively meaningful to them as a coherent world.” It is this knowledge that we think is particularly relevant to the field of social HRI, including the full range from high-level social structures and “knowledge that guides conduct in everyday life,” to the detailed (and still subjective to social context) daily interaction practises such as conversation, negotiation, etc. The sociology of knowledge includes “everything that passes for ‘knowledge’ in society.” (Berger and Luckmann, 1966)

This agglomeration can be referred to as a *social stock of knowledge*, a person’s ensemble of understanding of how to interact in the world in everyday life (Berger and Luckmann, 1966). We use this as a key concept throughout this dissertation to represent the kinds of existing and familiar (to people) knowledge and abilities which robots can use in their interaction designs. Further, within a given group of people (be it society, association, or even family) there is a shared overlap of each individual’s *social stock of knowledge* that they rely on to interact in their daily lives. Throughout this dissertation we refer to this broad inter-personal, context-specific set of shared knowledge as *the social stock of knowledge* (in comparison to a person’s individual knowledge). This is useful to represent skills or social structures which a robot could be designed toward and which it would be reasonable to expect a member of the general public to understand.

For social HRI this points to people’s mastery of routine, everyday problems, their “recipe knowledge” (Berger and Luckmann, 1966), and how they naturally apply this to new interactions: we can expect this to surface when interacting with robots, particularly given how people tend to anthropomorphise robots. Thus, an important element of social HRI is the consideration of how people apply *the social stock of knowledge* to robots, and how robots

themselves (through design and behaviour) can adapt to this as a means of being easy to work with. Social HRI also deals with the social structures in everyday life. In addition to simply having knowledge of this (i. e., the robot itself having a *social stock of knowledge* it can use to interact with people), an important element of social HRI is how the robot itself fits into and impacts social structures, and how social structures adapt to the robot.

In addition to the more direct interaction mechanics described above, the term social also encompasses external representations of internal emotions, insofar as they can be interpreted by another person as a means of reflecting on the generating person's state. This is also related to *the social stock of knowledge*, as people have standard interpretations and understandings of these externalized expressions, and what they may mean for the other's internal state. Thus robots can elicit similar reactions to provide people with insight into their state, and robots can also interpret this in others. On a final note, as actors within human social structures, robots can also have impact on the internal state of people, on people's emotions.

Overall, we believe that the purpose and goal in using the term *social HRI* — instead of simply HRI — is to point to the communication and interaction techniques, as well as social structures, that have evolved and emerged between people and how they interact with the world. Thus, social HRI looks at the encompassing picture of how robots fit into this overall context, both in terms of how they explicitly interact with it, as well as how researchers can understand the impacts that robots have on this context.

In this section we have explained what we believe *social* implies for social HRI: it points to the *social stock of knowledge*, how the robot works within this and can use this to interact with people, and how the robot itself impacts people's social world and structures.

3.1.3 *Embodiment and Embodied Interaction with Robots*

In this section we discuss how the ideas of embodiment and embodied interaction can be used to explain people's tendencies toward treating robots as social entities, and to further highlight the fundamental importance of the social aspects of HRI.

A robot's *embodiment* and *structural coupling*, as discussed in detail in Section 2.3.2, has a foundational impact on how the robot interacts with the world: a robot's *structural coupling* completely defines the only path through which it can interact with its environment, and as such defines the only ways that it can interact with people. A person's *embodiment* and

structural coupling to the world likewise defines how they understand the world and the mechanisms available to them to interact with it. As such, interaction between a person and a robot can only happen through a series of abstractions, where all interaction (and thus the meaning of the interaction) happens completely within the shared environments accessible to their particular *embodiments*: this can be more simply described as *embodied interaction*. This relationship is outlined in Figure 3.1.

Perceptions of interaction are fundamentally linked to the *embodiment* and thus the environment of each respective entity. From a person's point of view, this is the physical, socially-embedded everyday world within which they spend their lives. This implies that a person fundamentally interprets interaction with a robot from this perspective.

While the overall *embodiment* relationship exists for any technology we interact with, we argue that the nature of this relationship with a robot is fundamentally different from most other technologies leading from the *embodiment* of the robot itself. Unlike the Personal Computer (PC), which stays where it is placed and must be actively engaged and enabled, a robot will physically interact with and alter its surroundings, and may not remain in a simply-defined allocated space. Unlike physically-safe PC-based virtual environments, the robot may move unexpectedly posing risk to person and property, monitoring the robot involves following its motions and physical state, and people may not have direct access to orthodox (and familiar) interfaces such as a keyboard or display panel. A robot's *structural coupling* with the real world involves interaction with people directly via people's physical capabilities. Most other computing technologies we encounter, despite having some sort of physical embodiment, require a person to engage the virtual interface through computer-oriented abstractions such as a mouse, on-screen widgets and virtual desktops, or arbitrarily-mapped

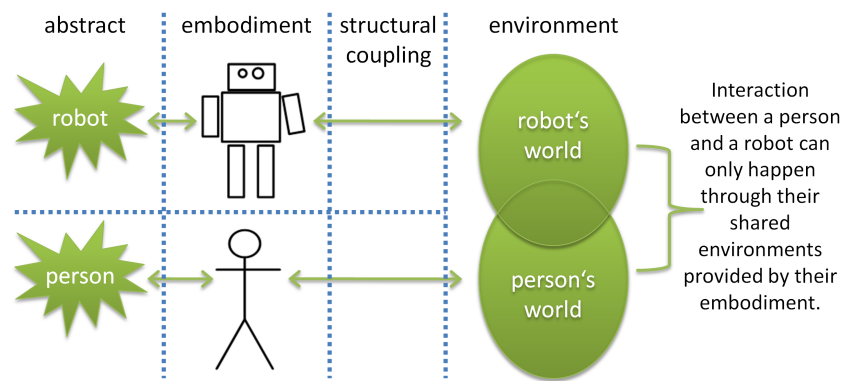


Figure 3.1: embodied interaction with robots

buttons. Even for more recent use-anywhere devices such as the Apple iPhone where interaction must integrate into everyday practises, and in this case attempts to leverage familiar touch interaction with physicality-based interaction metaphors, we argue that engaging the device (*the embodied interaction*) still requires the person to conceptually enter the virtual world of the phone and interact there. Finally, while we do have interactions (such as riding a bicycle or using a power tool) which we experience very directly within our own embodiment, these technologies generally do not elicit tendencies toward anthropomorphism or attribution of agency, a difference which we argue (in more detail below) has important impacts on the interaction experience.

In this section we have discussed what is meant by *robot*, *social*, and how *embodied interaction* points to the importance of the real human world, the environment context of interaction, for shaping how the interaction unfolds and how a person perceives it. This discussion helps explain why people tend to interact socially with robots (since interaction happens directly embedded within the social world), and it also explains why it is meaningful for robots to have physical *embodiment* (versus, e. g., on-screen simulated robots). Overall, *embodiment* and *embodied interaction* help to highlight how interaction with robots has fundamental differences from interaction with many other familiar technologies, and pushes us to explicitly consider the context of interaction. We explore this in detail below.

3.2 WHY ROBOTS CREATE UNIQUE INTERACTION EXPERIENCES

In this section we propose an explanation for why robots elicit unique, socially-charged interaction experiences. We propose that this is directly related to how robots elicit a unique sense of agency not usually experienced with other artifacts, in combination with their particular *embodiment*. Through this, we outline important properties of interaction with robots that social HRI researchers can use to understand social HRI in general, and how they can account for social aspects in both their robotic interaction designs and analyses.

3.2.1 *Why Robots Encourage Social Interaction*

In Section 2.2.1.1 we show that people naturally tend to respond socially and apply social rules to robots, and here we outline reasons why we believe this is the case.

Robots have well-defined physical manifestations, can perform physical movements and autonomously interact within peoples' personal spaces, properties that set them apart from other technologies such as a PC or microwave (see *embodiment* discussion above, also Norman, 2004). Further, robots' physical abilities to autonomously move and act in proximity of personal spaces (Dourish, 2001b; Harrison and Dourish, 1996), is considered to have a unique effect on the social structures surrounding interaction (Hornecker and Buur, 2006). As such, the way in which people apply social rules to robots, and the extent of this application, can be expected to be different than for other technologies.

Previous studies in non-robot human-computer interaction cases show that peoples' social tendencies toward technology can be deepened through socially-evocative technology designs (Reeves and Nass, 1996). Even for robots without explicit social designs, simple movements and abilities are often construed as lifelike (Forlizzi and DiSalvo, 2006; Sung et al., 2007), perhaps having this effect. Therefore, it is likely that robots that explicitly utilize such mediums as familiar human-like gestures or facial expressions in their designs will further encourage people to interact socially with them in a unique and fundamental way.

3.2.2 *From Anthropomorphism to Sense of Agency*

Section 2.2.1.1 highlights how people have been found to anthropomorphise robots more than other technologies and to attribute robots with qualities of living entities (e. g., animals or other humans) such as intentionality. We posit that perhaps it is this anthropomorphism embedded within social contexts (through the *embodied interaction*) that encourages people to readily attribute intentionality to a robot's actions regardless of its actual abilities. Perhaps this further relates to people applying their *social stock of knowledge* to make sense of the situation: an entity which moves around their space with some hint of intelligence is likely some sort of animal or living thing.

Intentionality helps give rise to a sense of agency in the robot — the word *agency* itself refers to the capacity to act and carries the notion of intentionality (Dewey, 1980). While people do attribute agency to various other technologies such as video game characters or movies (Reeves and Nass, 1996), we argue that the robot's physical-world embeddedness and socially-situated context of interaction creates a unique living-entity-like sense of agency similar that of living entities. We call this *active agency*. In a sense, then, for many people,

the robot becomes an active social player in their everyday world, and interacting with it is more like interacting with an animal or another person than with a technology.

Agency can help explain why people perceive robots to make autonomous, intelligent decisions based on a series of cognitive actions (e. g., Bartneck et al., 2007b; Nass and Moon, 2000; Reeves and Nass, 1996), and as such helps explain why people readily attribute lifelike qualities to robots. Further, agency contributes to the development of expectations of the robot's abilities (such as learning ability) or can create the expectation that the robot will be an active social agent, all in a much more prominent way than with more traditional technologies. In fact, it has been demonstrated that people tend to believe that even simple robots engage in reciprocal social interaction, and that people tend to develop strong emotional attachments to robots (e. g., Forlizzi and DiSalvo, 2006; Friedman et al., 2003; Marti, Pollini, Rullo, and Shibata, 2005; Sung et al., 2007). While people do sometimes exhibit emotional attachment to other artifacts, robots are unique in that they can legitimize and validate this relationship by actively responding to people's affection (Bartneck et al., 2007b).

Overall, this discussion suggests that robots have a presence in people's environments in a similar fashion to living entities, such that these robots naturally integrate into social worlds and encourage emotional involvement in ways not generally encountered with more traditional technologies, thus strengthening the sense of *active agency*.

3.2.3 *Embodied Interaction Experience*

As interaction is embodied within our social and physical worlds (Dourish, 2001b), a person's experience of interaction includes difficult-to-quantify thoughts, feelings, personal and cultural values, social standards, and so forth (Csikszentmihalyi, 1990; Dewey, 1980; Dourish, 2001b). Thus a person's experience cannot be fully or properly understood by reductive accounts or limited perspectives (Dewey, 1980), and the meaning of experience cannot be separated from the wider context.

We argue that robots' unique *active agency* and life-like presence makes this wider context a particularly prominent part of interaction experience. That is, the meaning of HRI often reaches well beyond the simple point of interaction (particular interface and particular actions) in a much stronger and deeper way than interaction with many traditional and more passive technologies and artifacts, making HRI a very unique instance of HCI. Following,

we expect that people will leverage the *social stock of knowledge* to inform how they should interact with robots similar to how they may for people (e. g., will people be too shy to undress in front of an advanced household robot they are not familiar with?).

This *holistic interaction context*, and how robots fit into this, is outlined in Figure 3.2. A person's experience of interaction, embedded within a broad context, is greatly influenced by the robot. The robot itself is a prominent and very active social and physical player within this context, with its social influence being similar in many respects to a living entity. The human and robot mutually shape the experience in a way not different from how two living agents may. Our discussion highlights how deeply interaction with robots is embedded in the social and physical worlds, and the uniqueness of this integration, compared to non-robotic HCI instances (such as interaction with a PC for example). This means that the range of factors contributing to the interaction experience with robots is immense, encompassing issues of culture, class, gender and age as well as social, political and economic structures and communication mechanisms (e. g., discussed generally in Dholakia, 2006; Rogers, 1995; Silverstone, 1991; Silverstone and Morley, 1990; van Zoonen, 2002).

In this section, we have detailed why we believe robots elicit unique interaction experiences from people: people's *embodiment* and the fact that they rely on a *social stock of knowledge*, combined with the insights gained from considering *embodied interaction* with robots, points to how interaction with robots is fundamentally unique. Robots are (without intentional design) anthropomorphic and generate a strong sense of *active agency* similar to a living creature. Overall this generates a very influential, socially- and physically-embedded *holistic*

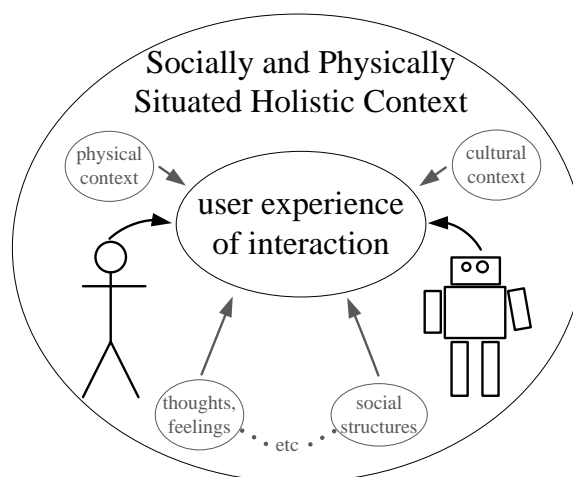


Figure 3.2: A person's experience of interacting with a robot is influenced by many real-world social and physical factors, where the robot plays an active role similar to that of a living entity.

interaction context within which a person experiences interaction. In the next section of this chapter, we introduce methods from the social sciences which we believe are useful for exploring the *holistic interaction context* in relation to how people perceive and accept robotic technologies.

3.3 SOCIAL PSYCHOLOGY PERSPECTIVES FOR EXPLORING SOCIAL HRI

In this section we take a look at social HRI from a broad perspective as a means to examine the *holistic interaction context* under which people interact with and form their understandings of robots: we ask, “what are the key social dynamics and factors that influence how people perceive, understand, and ultimately accept robots?” Our consideration of this question led us to an exploration into the social sciences, and in particular, to the general adoption of technology. We borrow concepts and analytical methods in order to help provide insight into the thought processes and influencing factors behind how people perceive robots, and use this as a means to re-conceptualize the robot design problem toward social considerations.

Here we explain our use of the terms *social psychology* and *sociology*. Overall, we generally default to the term *social psychology* as this work primarily deals with the more individual, personal perspectives and interactions with society (Myers, 1993). However, there are elements, particularly in this chapter, of looking toward interaction of society as a whole and so there are elements of crossover into more broad sociology (Comte, 2005).

Our acceptance-of-technology exploration deals primarily with domestic robots. One reason for the narrowing of scope is simply a practical reflection on the lack of existing relevant work in the area of technology adoption, as explained below. Further, the domestic context provides an excellent, easy-to-analyze exemplar of where social HRI will take a particularly prominent role: the social aspects of interaction are at the forefront in the home, and the home is also a popular target area for social HRI research.

We did not find any work which explicitly deals with the adoption of robotic technology in domestic environments. Most existing research into the adoption of robotic technology concerns the industrial application environment and generally focuses on financial, business, and economic concerns: for example, exploring specific tasks and goal-oriented problems in terms of *robotisability* (i. e., the ability to automate tasks with robots, Fleck, 1984), general industrial automation issues (e. g., Dosi, 1982; Fleck, 1984; Williams and Edge, 1996), or

macro and international-level industrial issues (e. g., Mansfield, 1989). These do not consider personal and social concerns (e. g., as with non-robot work, Steiner, 1995; Von Hippel, 2005).

Outside of consideration of robots, we did find a great deal of work that specifically addresses technology adoption in domestic contexts (such as, Ajzen and Fishbein, 1980; Davis, 1986; Mathieson, 1991; Venkatesh and Brown, 2001; Venkatesh, 2006, discussed below). We selected high-level perspectives on this process to give us more flexibility in considering the unique properties of robots, although social psychology perspectives on domestic environments have been explored in more formal and targeted frameworks: for example, Montalvo (2002) applies social psychology models analytically to argue that decisions to adopt are conditioned by subjectively-defined ‘willingness’ factors. While points covered by relevant existing work include personal satisfaction, status and other technology socialization concerns, none of the past efforts focuses on robotic technology, and none addresses the special socialization characteristics and problems presented by robots, and as such our application to robots is an original contribution.

In this section we use the idea that the meaning of interaction with robots is embedded in the *holistic interaction context* (a complex interrelated network of social, individual, and technological influence): this relationship goes beyond the definition of robot and applies to all aspects of the robot, including design and (perceived) utility. We present the general perspective of technology acceptance, several more specific adoption-of-technology models, and perform an analysis on specific robots, raising several key points which will likely be pivotal to robot acceptance by the general public. Following, as part of our theoretical framework we distill the robot-specific social psychology analysis into a set of factors that developers and designers can use for analyzing social HRI interfaces.

3.3.1 *The Socialization Process*

We believe that one of the most important and unique barriers to the widespread adoption of robotics is an especially-complex *socialization* process (Scitovsky, 1976), where people come to integrate a robot into their lives. We consider that the robotics environment is far more complex than for most already established domestic technologies, given the unique properties of robots outlined in the previous section. Further, the problems and implications of technology acceptance are far more significant in a domestic environment than in an

industrial one, where much of the work on robots has taken place: by design, it is intended that domestic robots will enter into our personal spaces, where their mere physical presence will have an effect on the spaces they occupy (Callon, 1987; Winner, 1987; Young et al., 2009).

Thus, the socialization of robots in the domestic context is far more than a conventional “human factors” design problem, in which barriers are overcome simply through the design of interfaces and routines. Neither is it merely a conventional “diffusion” problem whereby mass markets are created through positive feedback as more people experience and adopt a technology (Rogers, 1995; Stoneman, 1976). Instead, we argue that the socialization of robots is largely dependent upon *subjective personal perceptions* of what robots are, how they work, what exactly they are and are not capable of doing, and how they fit into social structures. Some even suggest that satisfaction with a technology is related more to the psychological expectation of acquiring it rather than its actual acquisition and use (Scitovsky, 1976). In the following sections we explore how these *subjective personal perceptions* are shaped for the general public, how this relates to robots that will enter everyday spaces, and how this discussion helps expose factors of the *holistic interaction context*.

3.3.2 *Absorptive Capacity*

In this section we introduce the idea of *absorptive capacity*, and relate it to social HRI. Core considerations underlying the study of technology adoption include, for example, how much and in what fashion a person or household is willing or able to adopt a technology, and how they are able to recognize the relevance or potential of new technologies as they appear. In an industrial or organizational setting, this *absorptive capacity* (Cohen and Levinthal, 1990) is generally seen to be generated by *related knowledge* — i. e. by existing capabilities upon which new capabilities can be built.

In trying to explain what determines the capacity, interest, or desire of individuals and households to adopt domestic robots, however, we have to consider what constitutes this *relevant knowledge*. From the perspective of an individual, the understanding of technology is typically the result of social rather than scientific, technological or industrial activity (Berger and Luckmann, 1966; Bijker, 1993; Callon, 1987; Clark, McLoughlin, Rose, and King, 1988; Williams and Edge, 1996). Thus, the *meaning* of a technology is not limited to the mechanisms, physical and technical properties, or actual capabilities of the technology. Meaning extends

also to how people think they must (or are supposed to) interact with technology and how it will (or should) integrate into and affect their lives.

While we do not re-visit the term *absorptive capacity* (Cohen and Levinthal, 1990) in the remainder of this chapter, we maintain the usefulness of the idea of considering adoption of a technology (and thus acceptance of robots) in terms of which factors create a predisposition toward acceptance. We believe that highlighting these factors can help improve understanding of the *holistic interaction context*, and thus which factors impact how people interact with robots. In the next section, we further delve into this adoption of technology approach.

3.3.3 *Select Adoption-of-Technology Theories For Robots*

In this section we present four social-psychology behavioural and decision-making models which we believe can help expose factors involved in the *holistic interaction context* and general perception of robots: they are, the Theory of Reasoned Action, the Theory of Planned Behaviour, the Technology Acceptance Model, and the Model of Acceptance of Technology in Households. We are the first to apply these models explicitly to robots. Rather than presenting each model in full detail, we distill each into simple representations that we feel focus on the robot-relevant components: we outline their primary focus, considerations, and perspectives as a way of bringing to light different ways to analyze robots in social contexts.

The Theory of Reasoned Action (TRA) (Ajzen and Fishbein, 1980) assumes that rather than being controlled by capricious subconscious forces people are generally rational and leverage information available to them. TRA bases this on observations of both *attitudinal* (i. e., personal) and *normative* (i. e., social) beliefs. Applied to the question of technology adoption, the attitudinal concerns include opinions of utility, efficiency gains, and how a technology fits into a given lifestyle. The normative beliefs include social views, pressures, expectations, and reactions to adopting a technology. Perceptions are more important than actual outcomes, and perceptions of outcomes can be more important than the perceptions of the technologies themselves. A person may acquire a technology simply because they *believe* it will have a positive impact (e. g., creating more free time), even if there is little or no actual evidence that it will do so (Ajzen, 1991). As key to shaping these beliefs, TRA points to lifetime experiences, and past actions and events. Sometimes beliefs are inferred from other knowledge, some beliefs being dynamic and others static (Ajzen and Fishbein, 1980).

The Theory of Planned Behaviour (TPB), an extension to TRA, adds an explicit focus on perceived behavioural control and points more to external factors (media, social acceptance, etc., Mathieson, 1991) than to “previous experience” as in the TRA model. This focus tries to accommodate the rapid change and perceived complexity of technology, where previous experience may be lacking and people are wary of difficulty of use.

A third model, the Technology Acceptance Model (TAM) (Davis, 1986), is specifically designed to explain and predict computer use, behaviour, and adoption. TAM lacks explicit consideration of social and normative variables and focuses on the perceived ease of use and usefulness of computers, based on external variables, as key to how people form attitudes. This emphasis represents a narrower (but focused) version of TPB’s perceived behavioural control (Davis, Bagozzi, and Warshaw, 1989).

These models take varying perspectives to the task of unveiling key characteristics of technology adoption (Mathieson, 1991). TRA may not handle problems associated with rapidly-changing technologies, while the focused nature of TAM may restrict the scope of its considerations, for example, if social pressure is part of a person’s evaluation of a technology’s ease of use. TPB would explicitly consider this in the framework from various viewpoints while TAM would simplify by integrating it with other ease of use concerns. However, the more thorough (and wide) nature of TPB makes it difficult to apply meaningfully across various contexts (it tailors criteria to each analysis).

The above models primarily take a personal perspective, and are less attentive to the household as an entity itself. The Model of Acceptance of Technology in Households (MATH) (Venkatesh and Brown, 2001; Venkatesh, 2006), a domestication-of-technology framework that focuses on the home as a macro-entity, was developed around an extensive longitudinal study of the adoption of PCs, primarily concerning the factors that people cited for or against adoption. Interestingly, the factors cited for adoption (primarily status and utility gains) did not line up with the factors cited for non-adoption (primarily fear of obsolescence and media influence), and only 45% of those who claimed they intended to adopt the PC did so six months later, suggesting that fears may overpower perceived gains.

MATH identifies that, in comparison to other contexts, household decisions have a more normative structure and are highly affected by social pressures, views of relevant others, and media (Burnkrant and Cousineau, 1975; Venkatesh and Brown, 2001). This also includes the perception of hedonic gains (entertainment, fun), family, friend and social network

influence, and perceived barriers or rules surrounding adoption, such as lack of knowledge (inability to properly use a product), prohibitive cost, or regulations requiring/restricting adoption of a technology (Compeau and Higgins, 1995; Venkatesh and Brown, 2001). Media influence from secondary sources such as TV and newspapers is particularly strong for early adopters (Rogers, 1995) where there are fewer informed friends and families to exert pressure, and the media often provides the first impressions. The hedonic value (pleasure) and social gains derived from a product, through both possession and use, have played a strong role in technology adoption in the past (Rogers, 1995), being the primary reason for adoption of such things as video games. Adopting a technology also has social gains including public recognition or being a knowledge reference within a social group (Venkatesh and Davis, 2000). From an attitudinal perspective the home has a strong focus on factors such as price, depreciation, maintenance, and space requirements: Venkatesh and Brown (2001) found that non-adopters (of the PC) often cited fears of technology obsolescence.

In the study behind the MATH model, status gains from having a new technology were cited as the primary reason for adoption, with social pressure from family members, hedonic gains, and personal utility cited as contributing reasons. For non-adopters (both “intenders” and “non-intenders”), the social influences and barriers were most significant, with negative influence from secondary sources being the largest factor, for example, where due to media representation parents fear for the safety of their children using the Internet. We provide a summary breakdown of the various theories in Table 3.1.

model	full name	brief summary
TRA	Theory of Reasoned Action	mix of attitudinal (e. g., utility, efficiency) and normative (e. g., social views, pressures, expectations) beliefs, based on previous experiences and beliefs inferred from those
TPB	Theory of Planned Behaviour	focus on perceived behavioural control based on external factors (e. g., media, social acceptance) more than experience, to accommodate rapid change of technology
TAM	Technology Acceptance Model	PC-specific, lacks explicit consideration of normative social considerations, primary focus on perceived ease-of-use
MATH	Model of Acceptance of Technology in Households	PC-specific, targets household rather than the individual, notes importance of hedonic gains and perceived barriers to adoption, and normative focus related to status gains and social pressures

Table 3.1: a summary of key points from the adoption-of-technology models presented in this section

In this section we introduced perspectives from social psychology which we believe are useful for both understanding and exploring the *holistic interaction context* between people and robots. We detailed how there is very little existing research relating to the broad picture of robot adoption, discussed the inherent socialization process behind the acceptance and use of robots, and introduced four acceptance-of-technology models which we believe are particularly relevant for robots. Overall, we believe that our approach as outlined in this section can help to provide insight into why a person may interact with a robot in a particular way and which factors impact this perception, and as such, provide an all-around better understanding of how social HRI researchers can design and build robots for particular interaction scenarios, and better-understand their evaluations. In the next section we apply the four models (in a detailed fashion) to real-world robotic instances.

3.4 APPLYING SELECTED ADOPTION-OF-TECHNOLOGY MODELS TO ROBOTS

The above models serve as a set of theoretical tools that enable us to analyze concerns and sources of influence for technology adoption, and thus, explore the *holistic interaction context*. In this section we utilize the perspectives of the above models, and the nuanced differences between them, to explore how such concerns relate to the unique properties of robots.

The four models above represent sometimes-conflicting perspectives on the domestication of technology. Rather than applying them in turn to robots, we organize our analysis around prominent themes that emerged, using each theory in place where appropriate. We first introduce two specific robots used for analysis, and then discuss them from these broad categories: initial exposure, control and safety, and hedonic aspects. We also introduce a new aspect, social design, currently not covered by the models. We primarily assume the social context of contemporary North American culture, an admitted limitation applied purely for practicality; inter-cultural studies remains an important part of future work.

3.4.1 *Two Domestic Robots for Targeted Analysis*

We analytically apply the acceptance-of-technology models to two particular domestic-robot cases, reflecting on factors that may impact how these robots are perceived. One domestic

robot, the iRobot Roomba (Figure 3.3a), represents a practical, successful commercial product, while the other, the RIKEN RI-MAN (Figure 3.3b), is a more futuristic research prototype.


The Roomba is an autonomous and mobile vacuum cleaner robot that is affordable, effective, and commercially successful. The Roomba has been introduced into existing home environments, with the overall product (design, implementation, etc.) being sensitive to existing in-home cultures and routines (Forlizzi and DiSalvo, 2006). The Roomba, however, is a utility robot which is meant to independently do its task while staying out of people's way. It is designed to allow a person to simply push a start button and walk away; thus it adheres to many characteristics of the traditional servant such as command-based interaction (as in Hamill, 2006).

The second robot is the RIKEN RI-MAN, a personal assistant robot currently under development that is designed to lift people who need assistance and to carry them around their homes. The RI-MAN can dramatically improve the quality of life for the people it is designed to help, and lower their dependence on other individuals. Unlike the Roomba, which works by itself (to clean floors) and stays out of the way, the RI-MAN is designed to work in direct contact with people, its users being the most crucial component of its design space. This introduces unique questions such as how robots like the RI-MAN will relate to personal space and privacy. There is also a trust concern; the RI-MAN can physically hurt people, or can cause problems by failing to perform as required, for example, by not carrying them properly from one location to another.

In the following section we analyze these robots using the above models to reveal considerations related to acceptance, and to map considerations of the *holistic interaction context*.



(a) iRobot Roomba vacuum



(b) RIKEN RI-MAN personal assistance robot

Figure 3.3: two domestic robots used in our analysis

3.4.2 *Initial Exposure to Domestic Robots (TRA, TPB, MATH)*

The core of the TRA model is that beliefs about a given technology are based on lifetime experience (Ajzen and Fishbein, 1980). This is supported by early studies suggesting that the way that robots are introduced to a home (or person) is crucial to the formation of lasting opinions of the technology (Forlizzi and DiSalvo, 2006). Since robots have not yet been adopted into the home on a large scale, perhaps experiences with other technologies will have a strong influence on beliefs thus positively or negatively shaping adoption. TRA suggests that which previous experiences people will draw on, however, is a function of how the robots themselves and the condition of owning a robot are perceived. It is possible that some robots will be seen as just another home appliance much like PCs, TVs, and personal music players, in which case people would draw upon their experiences with these devices in order to understand domestic robots. However, if robots are perceived as being fundamentally different from other domestic entities then it is not entirely clear which experiences people will draw upon. Given tendency to anthropomorphise, then perhaps people will relate to their experiences with children or animals, or robots may fall in between, with people building on past experiences and external sources, inferring new beliefs specific to robots.

The image of owning a robot is based on beliefs (not necessarily facts), and so, as MATH and TPB point out, media may have a strong influence on shaping these beliefs. This is particularly true for earlier adopters who have less experience to draw from, and may be amplified by the unique nature of robots. Perhaps the strong role of media and exposure to science-fiction has prepared people and has conditioned Pavlovian or Pavlovian-like responses (Pavlov, 1927) to domestic robots, such as fear of large robots or the attraction of cute, small robots.

TRA points to the utility, effectiveness, and price of robots. While we can expect the trend of utility gains from technology to be continued by robots, people must also *perceive* them as having a useful purpose. Recent findings (Forlizzi and DiSalvo, 2006) suggest that people without prior experience are not always ready to believe that robots are effective, hinting that other attitudinal factors may have to initially play a larger role in the way robots are perceived. However, utility may not be as key as it seems. Venkatesh (2006) found in their study that people who intended to adopt a PC cited utility as the motivation twice as often as adopters did in retrospect, suggesting that utility may be an excuse used as a rationalization when other factors (e. g., social status, being a knowledge reference) are closer to the real

motivation. Part of the Roomba's success could be that it is in the same price range as a regular vacuum cleaner, a cost-benefit ratio that people are familiar with. More advanced and expensive robots such as the RI-MAN will likewise need to create their own balance between price and quality of service, in consideration of the available non-robotic alternatives such as human nurses or care-takers.

MATH suggests that, following the technological trend, there will be social status gains or expectations associated with owning the newest technologies (including domestic robots) that may persuade people to adopt. Social pressures can also be manifested through concerned family members, such as children encouraging parents to adopt automatic vacuum cleaners (Forlizzi and DiSalvo, 2006), a point which may be very influential given current concerns surrounding ageing demographics in Western countries and Japan.

We know from our own personal experience that some people may be embarrassed by such automation technology, in that they are afraid to appear lazy to their peers. The Roomba is small enough to store in a closet and the nature of its work (i. e., it does the same task as a regular vacuum cleaner) makes it easy for an owner to conceal the fact that they have one, if they so wish. On the other hand, the Roomba has been designed and marketed as a stylish household appliance, which may help overcome some of these concerns. Conversely, the sheer size and mass of the RI-MAN, as well as the nature of its work, would make it very difficult to conceal if an owner wanted to. Such a problem of embarrassment, if it emerges, may disappear if adoption becomes more common and socially acceptable. In the RI-MAN case, the necessity of assistance may overcome such normative concerns, similar to canes and wheelchairs for people who experience a loss of mobility.

Venkatesh and Brown (2001) (through the development of MATH) found the concern of "obsolescence of technology" to be a large factor for PC adoption, although it is not clear how these concerns will map to robots. Conceivably, a robot is purchased for a particular purpose and will continue being useful until it stops working. This differs from the PC which, as software demands increase, can no longer execute software and perform the same basic tasks for which it was purchased (such as sharing documents, checking email, etc.) long before it physically breaks. Perhaps, then, robots will only be replaced when newer models offer a very large gain in capabilities and applications to new tasks or when it breaks, similar to an automobile. The hardware/software model of robots may lay between the PC and traditional

appliances that are generally not replaced until they stop working. Regardless, the resulting architecture will likely have a very large impact on the adoption of domestic robots.

MATH also points to a normative focus on perceived barriers to adoption, including possible legislation controlling the use of a robot or lack of facilities in the home to deal with a robot. Currently, as robots are not yet controlled by law and use standard household infrastructure (electrical outlets and Internet connections) this barrier does not yet exist. However, we can expect legislation to emerge with the proliferation of robots for such things as confining their use and controlling their collateral impact, for example, a lawnmower robot damaging the neighbours flowers, or an autonomous robot car that causes an accident.

In the following section, we continue our analysis of the acceptance of robots — and the *holistic interaction context* — through considerations of safety when interacting with robots.

3.4.3 Control and Safety of Domestic Robots (TPB, TAM, MATH)

The TPB model points to the importance of perceived behavioural control in forming opinions about technology (Mathieson, 1991). For example, people believing they can control when and how technology operates, or how adopting such a technology affects their social status. TAM narrows these criteria and places emphasis on the perceived ease of use.

We expect that both of these emphases, the broader behavioural control and more narrow ease-of-use, will be of importance based on the nature of the particular robot, its capabilities, and how it is designed to enter environments (e. g., its target tasks and how this relates to people). In either case, TPB points out the importance of the *perception* of the difficulty, i. e., the intersection of a person's skill set and the perception of the skills required to operate the robot. For example, given that early adopters of technology tend to be better educated (Rogers, 1995; Venkatesh and Brown, 2001), perhaps educated people have more confidence or skills around advanced technology. However, this may be less of a factor for domestic robots as there may not be many skills transferable from other technologies such as the PC, or from other contexts such as the workplace.

Although the general ability to control a robot is always important, we believe a key emphasis in the case of social HRI becomes one of personal safety. Despite safety tests and assurance by designers, the autonomous and physical presence and *active agency* gives the robot a “life of its own” that can override and hinder perceptions of control. Just as with

animals (or people), this fear will be a function of robot capability, size, and will be heavily influenced by experience. For example, similar to the Roomba, most people are not worried about a small kitten or a puppy as they feel they can control the animal if it gets out of hand. With larger animals, such as an untrained large dog, a cougar or a wild horse, this confidence is more difficult (or impossible) to achieve. Even with smaller animals (or robots), capabilities are key: approaching a wild and panicking adult cat is a very scary venture despite the small size as we know the cat has teeth and very sharp claws. The Roomba, however, has no claws, and is unable to hurt us as long as we keep our fingers away from the cleaning mechanism (a danger we are familiar with when using a regular vacuum cleaner), and so we feel safe around it. On the other hand, the RI-MAN can be viewed as a large trained animal: we can learn to trust it, but are still worried about what will happen if it breaks its training (programming).

The Roomba is marketed as a simple “clean with the touch of a button” device, a successful strategy where it only does a single task and only when commanded. Furthermore, its small size and harmless capabilities means it is easy to move or disable and the person can establish *virtual walls* which restrict the Roomba to a particular room or region. Regardless, people are sometimes worried about the Roomba bumping into furniture or knocking down breakables (Forlizzi and DiSalvo, 2006).

The RI-MAN may have more difficulty with control issues for several reasons. For one, it does complex tasks that involve performance ambiguity and so a person may not always be clear on what it is doing. Its physical size and weight make it impractical for an average person (let alone a needy person) to move or lift it in a dangerous situation. Further, the strength of the robot’s arms and its mobility makes the robot quite dangerous in a worst-case malfunction scenario. It is to be expected, then, that the damage-to-furniture type of concerns voiced regarding the Roomba will be dramatically amplified in the RI-MAN case, particularly given the contrast to the relatively small form and gentle movements of the Roomba. Until robots and artificial intelligence algorithms prove themselves to people, we expected that this doubtful and wary approach will be a strong factor in peoples’ considerations. On the same token, building an understanding of how a robot’s design can encourage this kind of worry or unease could be useful in some target applications, for example, a guard robot.

We directly address this control and safety concern, including the related social layers, with our dog-leash interface (Chapter 4). In the following section, we consider social HRI in light of the hedonic aspects of interacting with a robot.

3.4.4 *Hedonic Aspects of Social HRI (MATH)*

MATH suggests that hedonic gains, which have been shown to have shaped other technologies such as the PC in the past (Rogers, 1995; Venkatesh and Brown, 2001), will be a very important factor in the broader picture of social HRI. While the Roomba and RI-MAN, like the PC, do not directly address hedonic needs, they may do so indirectly: the Roomba saves time while the RI-MAN increases a person's mobility. Other robots may be designed for aesthetic purposes only, much like dynamic art. One such example is the SONY Rolly, which moves and dances while playing music.

The robotic toy is yet another recreational application. Devices like remote-controlled cars have long been marketed as “robots” even though they do not have many of the characteristics typically associated with robots. More recently, however, more genuine advanced robotics have been marketed to consumers in the form of toys. The prime example of this is the Sony AIBO robotic dog, a toy which can move around, sit, play with a ball or bone, take pictures and send them to the person's email, and even has a complex behavioural and artificial intelligence model to mimic a real puppy (Figure 3.4). Despite this, however, the AIBO was not commercially successful, and Sony stopped production. The exact reason for the toy's failure is not clear, but we believe it is related to the price and the dog's somewhat limited movement capabilities. It originally sold for \$2500 USD in 1999 (~\$3200 USD in 2009 dollars) and dropped to \$2000 USD by 2003, which is a steep price for a toy that has no direct utility or proven history, and it moved very slowly, got stuck easily, and could not traverse stairs.




Figure 3.4: The Sony AIBO Robotic Dog

A more successful example is the line of affordable (currently \$50–\$100 USD) robotic toys from Wowwee, including the humanoid Robosapien and a flying robot called Dragonfly. These examples, however, are very simplistic for robots: the Dragonfly is completely remote controlled, and the other models only have simple capabilities and weak ability to interpret their environment. Some of Wowwee's more advanced models, such as the Robopanda, are still extremely limited in their abilities. Because of this, perhaps these robots enter homes in much the same fashion as a remote-controlled car or battery-powered doll might. They

have a modern feeling of novelty that is perhaps induced by the buzzword “robot,” but to the average consumer they are still likely single-purpose toys that fit the existing *play* paradigms in the home. This question has yet to be explored via a formal study. If our assumptions are correct, however, this contrasts strongly with video games, the Internet, the PC, and television, which each provide a fundamentally different dimension to the world of domestic fun, each having a new range of interactive possibilities similar to how an advanced robot might.

As future robotic toys become more capable we may see a similar thing happen. For example, a humanoid robot that has just enough awareness of its surroundings to naively follow its owner and play simple games, help them fold the laundry, or even tell a few jokes, would be well beyond any toy available today. To many people, such a robot could become a kind of simple pet or companion, and would enable a whole new range of play possibilities not previously possible (as Isaac Asimov’s short story *Robbie*, Asimov, 1968, explores). With this in mind, presenting robots as toys may help to overcome understanding and acceptance barriers, and allow people to categorize these new entities effectively and easily.

One type of emerging robot is the personal sex-service robot, an application of particular interest to social HRI given the extremely human-focused and personal, and inherently social, nature of sexuality. Various producers around the world are working on such products (such as AndyDroid, Figure 3.5), and we expect that these will be successful, if the extremely successful adaptation of the PC (at least from a business, social integration perspective) or the existing markets for sexual devices including realistic dolls can be used as a model. The more interesting question for us, however, is what will happen when these sex robots become



(a) Andy, a female sex robot



(b) Nax, a male sex robot

Figure 3.5: two sex robots produced by AndyDroid

increasingly capable of interpreting, understanding, and intelligently interacting within their environment. How far will the human mind allow the anthropomorphism of machines to go? How much jealousy will people feel if their partner decides to have sex with a robot? Will people *fall in love* with their robots (Levy, 2007)?

If these robots become successful, even within a minority of people, such personal experience with a robot may be a key component of the acceptance of robotic technologies, as it would form a body of past experience and knowledge people could draw from (as highlighted in both MATH and TRA). Someone who feels they have an understanding and trust of robots through their experiences in the bedroom might be more willing to bring alternate models in to clean, cook, or play with their children, and thus hedonic considerations could have wide-sweeping implications for robot acceptance.

We believe the idea of robotic companionship and friendship to be particularly relevant to robot acceptance. The fact that people already anthropomorphise robots with human-like characteristics makes such a question particularly plausible. It will be no surprise if people start to feel an attachment to robots as already happens with material things such as sports cars, collectible items, teddy bears, or various other items that are important for personal reasons. Given the *active agency* of robots, however, these kinds of bonds may be closer to the kinds of bonds experienced between two people or a person and their pet.

Particularly for robots such as the RI-MAN, which has a human-like appearance, replaces a traditional human role, and provides a service that may result in a feeling of gratitude and perhaps emotional attachment from the owner, the development of a sense of companionship would be an almost-natural progression. This has happened, for example, in military settings (as outlined in detail in Section 2.2.1.1, see Garreau, 2007).

As intriguing as the idea of robotic companionship may seem, however, our social-psychology analysis leads us to believe that such factors are not likely to have a strong initial impact on intention to adopt a robot: there is little prior experience or a public body of knowledge to draw from. If anything, ideas of attachment may be assumed to relate to how people already get attached to toys or personal technologies such as cell phones, and it is unlikely that people will automatically consider the deeper reaches of the companionship factor until there is experience and a cultural understanding of such a phenomena. Initially, at least, companionship may just be a secondary product of purchasing a robot.

In this section we have discussed how hedonic considerations relate to social HRI and that perhaps, similar to the way games help drive the PC industry, entertainment may play a role in robotic acceptance, serving as a catalyst for the entire domain. This further highlights the importance of considering entertainment and other hedonic gains in relation to a social HRI design. In the following section, we highlight a hole in the existing social psychology models presented in this chapter, and introduce an additional perspective that we believe is important in relation to the acceptance of robots: explicit anthropomorphic design.

3.4.5 *An Additional Perspective: The Role of Social Design*

The social psychology models we present in this chapter point out how, in general, technologies must be designed to be sensitive to existing social structures and understanding. We note that these models do not consider the role of social design in the adoption of technology, that is, a technology's attempt to use familiar human communication techniques (e. g., emotion or spoken language) for interaction; leveraging *the social stock of knowledge* through its physical and behavioural design. We believe these considerations are of particular relevance for social HRI, as understanding how the use of social communication techniques impacts acceptance and perception is an important component of understanding how the robot itself impacts the broader *holistic interaction context*.

Robots are not the only technology which use social design in domestic settings. For example, emerging *intelligent home* technologies are often capable of sensing a home's physical environment and attempt to operate lights, temperature controls, music, and so on, in a socially appropriate manner (for the given household). As such, we see a large overlap between intelligent-home technologies and social HRI, where perhaps such a home could be viewed as a kind of robot, and our social HRI principles applied. However, such an analysis is beyond the scope of this dissertation.

Continuing with our analysis on the two domestic robots, the Roomba does not appear on the surface to be explicitly designed with social interaction in mind. It has a very-simple functional and mechanical appearance and utilizes simple blinking lights for status messages, although the sounds it makes can be construed as *happy* or *sad*, and the newest models have a synthetic-speech first-time use explanation mode. (The studies presented in Section 2.2.1.1, showing people's tendency toward anthropomorphism, primarily use the older Roombas

which did not use synthetic speech.) Despite this functional design, people anthropomorphise and zoomorphise the Roomba, giving it human and social characteristics, and find ways to fit it into their homes and social practises (see Section 2.2.1.1). In addition to being functionally useful, the Roomba as is can in a sense become a social participant in the family, perhaps in a similar fashion to, say, a pet hamster. This suggests that, just as people are willing to anthropomorphise robots with very little explicit design to that effect, people are perhaps willing to accept robots as social actors in their everyday lives even when the robot does not actively (programmatically or by design) support this.

The RIKEN RI-MAN, on the other hand, is explicitly designed to have a human-like appearance and interaction paradigm: it has a human head, face and arms, soft skin, ears that listen and a mouth that speaks, and social programming that allows it to follow communication protocols such as gaze during conversation. These characteristics by design attempt to encourage people to use their social skills to interact with the robot. It remains to be seen how this robot is received by the general public, but this level of human resemblance puts the robot at risk of eeriness problems, a potential concern in relation to social acceptance. Currently, the Roomba's approach shows promise for integration, but its target application does not entail much interaction with people. Perhaps successful domestic robotic interfaces will have to balance between the approaches of the Roomba and the RI-MAN, depending on the target application, where clever and simple, integrable design (Roomba) is balanced with explicit social techniques for advanced interaction with people (RI-MAN).

It is not yet understood how the existence of higher-level robot social savvy, such as the ability to fit into the existing social activity structures of the home, will affect the perception of the robot. Technologies such as the PC and the Roomba that have no explicit ability to interpret the social environment suggest that social understanding is not necessarily required for a technology to be successful in a domestic setting. The Roomba's case highlights how clever interaction design can be a substitute for actual ability, resulting in successful social integration; this idea has been posited as the idea of simple social "contracts" (Hamill and Harper, 2006) that the person and robot can easily adhere to. The problem with this argument, however, is the simplicity of how these socially-ignorant technologies physically interact within their environments. The goals of more advanced machines such as the RI-MAN require them to actively interact in spaces shared with people, and perhaps should have an understanding of what people are doing (such as a sleeping baby, a person using the

washroom, or a child doing homework) and alter its actions appropriately, with a calculated impact on the social fabric of the home.

Following from our analysis, we suggest that the importance of social abilities (to sense, intelligently interact, etc.) may be directly coupled with overall robot capability. It may be that the more capable a robot is, the more it builds expectations in people who interact with it. For example, it has been shown with robots that communicate with speech, people have high expectations due to the robots' apparent capabilities and are subsequently disappointed by the lack of actual depth in interaction (Hamill, 2006). This suggests that robot designers may sometimes want to lower the intelligence, or appearance of intelligence, of their robots in order to lower people's social expectations. Perhaps people will be forgiving and will accommodate robots' errors in much the same way they do with pets or children, finding it simply natural that the robot does not understand. For example, a dog is taught not to bite or bark excessively as people know dogs can learn this, but fish are not trained in the same fashion: rather, signs that say "do not touch" are affixed instead. Similarly, parents simply apologize when infants pull other people's hair, but when the infant becomes a toddler they are (usually) scolded and instructed not to do so. However, we expect this rationalization to break down for considerations of safety or security, as people may have zero tolerance for domestic robots that break plates or flood the floor while cleaning.

In this section, we applied established social-psychology perspectives, as well as our own social-design addition, to two existing robotic instances, and led a broad-but-thorough discussion through an analysis of the factors related to people's perceptions and acceptance of robots. Overall, this provides an in-depth exploration into the broader social structures embedded within the *holistic interaction context* of social HRI, and provides insight into the kinds of factors social HRI researchers may need to consider for their robotic designs. We summarize this discussion as part of formulating our theoretical framework, in Section 3.6.

3.5 ROBOT EXPRESSIONISM — AMBIGUITY CAN BE HELPFUL

In this brief section, we present *robot expressionism*, a concept which emerged from our cartoon-artwork exploration and has become a prominent concept throughout various phases of this dissertation research. *Robot expressionism* provides a perspective for both

designing and analyzing the abstractions between a robot’s technical state and what it conveys to people, and as such is useful for the overall social HRI approach.

When considering robots that represent their state using social language, we propose a parallel to expressionism in the arts, a movement during the early 20th century that emphasized subjective expression of the artist’s inner experiences and state of mind (WordWebOnline, 2010, “expressionism”). The expressionism movement was part of a shift in Western art where artists moved from motivations of direct representation toward more stylistic subjective representations of experience (e. g., see Figure 3.6).

Direct representation of robots’ digital state and information is arguably trivial to accomplish, but not necessarily intuitively and directly meaningful to people that interact with the robot. We propose the term *robot expressionism* (Young et al., 2007) to represent the idea that robots can distort and stylize their communication, perhaps using human social techniques, to provide a layer of insight into their internal state and algorithmic intentions, above and beyond simply offering raw data or direct representation.

Our use of the term *robot expressionism* emphasizes the subjective and perhaps artistic nature of interaction with people. Rather than striving for purely functional representations of robotic states and algorithms, robot expressionism motivates the abstraction of a robot’s state for the purpose of providing contextual insight into the robot’s state. Our approach also motivates how the ambiguity that




Figure 3.6: Edvard Munch’s expressionist painting, “The Scream”

can accompany abstraction can be desirable. We argue that sometimes people *clearly* understand *vague*, complex, non-concrete ideas such as confusion or frustration. For example, an expressionist robot may shrug if its artificial intelligence systems become overloaded and unable to cope with a complex question or interaction problem. While this offers no technical explanation, the shrug provides clear insight into the situation: the robot simply did not know the answer.

We use our *robot expressionism* idea throughout our dissertation as a means to explain the power of robots using abstractions for interacting with people. In the next section, we summarize this chapter into a theoretical framework for social HRI.

3.6 A THEORETICAL FRAMEWORK FOR SOCIAL HRI

In this section we present a theoretical framework for social HRI, based primarily on a summary of the various explorations presented in this chapter. This framework is further informed by our own experiences with designing, implementing, and evaluating social HRI instances, as detailed in the following chapters. We present this framework as the first such attempt to formalize, describe, and to add a structural means to inform the design, analysis, and evaluation of social HRI.

This framework has three key components: we outline the core concepts which we believe shape the domain of social HRI, we introduce a new set of three perspectives for classifying social interaction with a robot, and finish with broad considerations, which we believe will be particularly influential for how people form their perceptions of interacting with robots.

3.6.1 *Outlining Social HRI: Defining a Domain*

The three components of the term Human-Robot Interaction (HRI) explicitly identify two actors, the human and the robot, and the fact that they are interacting. Following, social HRI is HRI with an explicit focus on the encompassing social human-world within which interaction happens. For this discussion we consider robots to be machines that a) have a dynamic physical presence in the real world, and b) elicit a sense of agency and intentionality. This means that although machines such as elevators (a) have a dynamic physical presence in the real world, they are generally not considered a robot since people generally do not (b) attribute them with agency. However, if agency happens to arise in a given situation, then social HRI is applicable to understand, explain, and study interaction with the elevator.

The somewhat-vague meaning of the term *robot*, for example in relation to *interpretative flexibility*, raises questions of the boundaries of social HRI and the work we present in this thesis. We see social HRI itself as a perspective, a lens, on interaction between people and robots that focuses on social aspects. Further, our work and perspectives are particularly relevant for interaction with robots that fall under the above definition (have a dynamic physical presence and elicit a sense of agency). This does not suggest a hard-lined boundary to social HRI, however, as we see the applicable target area as a smooth drop-off function surrounding the core focus described here. That is, we expect that social HRI perspectives

are also relevant for explaining and understanding other interactions, for example, with the elevator as explained above, with virtual agents, or perhaps even the social aspects of interacting with a PC.

Social HRI provides a fundamentally different perspective due to robots' particular *embodiment* and capability to elicit a sense of *active agency*. This encourages people to treat robots as social entities, anthropomorphizing and attributing intentionality to their actions. As such, social HRI highlights the practical utility of designing robots that leverage *the social stock of knowledge* in their interaction designs, and encourages *robotic expressionism*: choosing abstract representations of robotic state that are easy for people to understand.

Robot's unique *embodiment* means that interaction unfolds embedded within the real-world *holistic interaction context*, where external factors such as the social and physical environment of the interaction take a particularly prominent role in shaping the interaction experience. The concept of *interpretative flexibility* is particularly informative here, highlighting how the meaning of robot is very fluid and dependent on the people involved and context of interaction. The *holistic interaction context* idea further means that the robot itself can have an impact on the human social structures. As such, we take the position that all interaction between people and robots has a social component, and therefore HRI is also social HRI. The difference between the terms is a matter of focus, as HRI researchers may not be particularly interested in social factors for a given project, although we maintain that social factors still exist even if not targeted in research. While the same argumentation can be used for the social aspects of HCI, it is how robots inherently encourage social interaction that brings this concern to the forefront.

Here we presented a theoretical definition of social HRI, outlining the intrinsic social component of interaction between people and robots. While this can be used to generally understand the scope and meaning of social HRI, there still remains a lack of explicit tools for analyzing and describing particular social HRI instances. In the remainder of this section we present two new such tools. First, we present a set of three perspectives which emphasize the breadth of interaction with robots, ranging from a person's visceral-level reactions to social-structure-level interactions. Second, we present a set of broad considerations which highlight factors that we believe will have a strong influence on shaping people's perceptions of robots and their acceptance of them.

3.6.2 *Focusing on the Breadth: Three Perspectives for Social HRI*

In this section we present three new perspectives that we believe can be leveraged for describing social HRI. In addition to being rooted in the theoretical discussion presented in this chapter, these three perspectives have emerged from our various experiences designing, implementing, and in particular evaluating social HRI instances.

We categorize social interaction with robots into three perspectives: *visceral* factors of interaction (e. g., the immediate, automatic human responses), *social mechanics* (e. g., the application of social languages and norms), and the more macro-level *social structures* related to interaction. Below we introduce the perspectives, and later followed by an in-depth discussion on the differences between them including an illustration of how they can be used to inspect and discuss interfaces and interaction.

Perspective One (P₁), visceral factors of interaction, focuses on a person's biological, instinctual, emotional involvement in interaction. This includes such things as instinctual frustration, fear, joy, or happiness, on a reactionary level where they are difficult to control.

Perspective Two (P₂), social mechanics, focuses on the explicit communication techniques used in interaction. This includes both the social mechanics that a person uses in communication as well as what they interpret from the robot throughout meaning-building during interaction. Examples range from gestures such as facial expressions and body language, to spoken language, to cultural norms such as personal space and eye-contact rules.

Perspective Three (P₃), social structures, covers the development of and changes in the social relationships and interaction between two entities, perhaps over a relatively long period of time (longer, relative to P₁ and P₂). P₃ considers the changes in or trajectory of P₁, P₂, as well as how a robot interacts with, understands, and even modifies social structures.

These three perspectives are not a hard-line categorization of the various components of interaction, or a linear progression of interaction over time. Rather, interaction happens simultaneously and continuously from all three perspectives, and there is crosstalk between the perspectives for any given interaction; these categorizations provide unique — but not mutually exclusive — perspectives on this complex relationship. Further, building from our idea of the *holistic interaction context* (Figure 3.2, page 60), we note that both the robot and the person play active roles, and the three perspectives can be used to analyze this. In Figure 3.7, the human-centric view considers how the person feels about, approaches, and

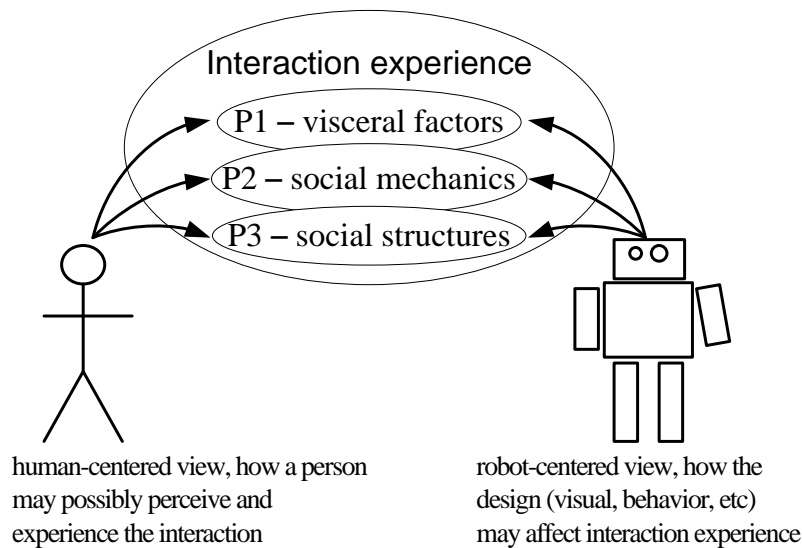


Figure 3.7: three perspectives on interaction experience, mutually shaped by two agents: human and robot

interprets the interaction experience, and the robot-centric view considers how the robot itself, including its design, behaviour and actions, influences the experience.

Given a particular robot, interface, or scenario, certain perspectives may be of greater interest than others for a given research question. However, we contend that components of all three perspectives exist in any interaction between a human and a robot. Following, deliberately not explicitly considering a particular perspective (or other such social perspectives) may limit the view and hinder potential understanding of a social interaction scenario.

Below we offer detailed descriptions of the three perspectives. Our approach revolves around using the perspectives to categorize and introduce existing literature and themes, serving as a simplistic case study highlighting the usability and applicability of the perspectives in describing and relating social HRI principles in a way not previously available.

3.6.2.1 *Perspective 1 (P1) — Visceral Factors of Interaction*

People have many visceral, perhaps largely instinctual, reactions to the world around them (Norman, 1988, 2004). These reactions are often difficult, if not impossible, to quell or restrict. Some of these reactions are nearly universal to all humans, such as smiling when happy, while others are cultural or individual-oriented, such as fear of insects or particular associations such as having a positive response to a Christmas theme. Many of these reactions are entirely internal, with very little or no outwardly noticeable effect, while others such as recoiling

from a spider are very externalized in their expression. Interaction continues to occur from this perspective (P₁) even for engaged or long-term interaction.

One example that highlights the importance of P₁ visceral interaction is the problem of innate robot eeriness (Ho, MacDorman, and Pramono, 2008; Mori, 1970). Another example is people's reluctance to interact with an anthropomorphic robot that appears taller than them (Lee, Forlizzi, Rybski, Crabbe, Chung, Finkle, Glaser, and Kiesler, 2009). Paro the rehabilitation robot was specifically chosen to take the form of a baby seal to elicit positive (P₁) emotional responses from people (Marti et al., 2005): people reported a great deal of emotional attachment toward the robot. Our own work (Chapter 6) uses familiar cartoon-artwork to explicitly anthropomorphise robots, and make them both familiar and fun, and give them a communication vocabulary of, for example, simplified and exaggerated facial expressions that people can naturally understand without thinking. All of these examples fall under our P₁ perspective.

Visceral (P₁-type) reaction is not limited to robots with explicit anthropomorphic designs. For example, the shape, speed, and patterns of a robot's movements also contribute to visceral reactions, as shown with our *stylistic locomotion project*. In one study, Roomba owners reported both excitement and enjoyment from watching how the robot moved around the space, even though the movements were random (Sung et al., 2007). A similar finding was reported with Bethel et al.'s (2009) search and rescue study where, based on the way it moved (e. g., with aggressive and sudden movements, or slower and softer movements), people reported feeling either more or less threatened by the robot, resulting in a deepening of the traumatic symptoms reported. The robot Keepon works largely on the principle of evoking P₁ reactions of fun and enjoyment from people simply through the way that it moves (Michalowski et al., 2007). Bartneck et al.'s (2007b) work showing how people were very hesitant to "kill" (or shut off, Bartneck et al., 2007a) a robot is perhaps an illustration of their P₁ reluctance to harm something which (somewhere inside) they feel appears to be living.

This perspective (P₁-type) of human reactions to the world is a very powerful and important part of the experience of interaction: fear, happiness, excitement, dread, and so forth, can have a large impact on the overall interaction experience. Robots make visceral reflection a particularly relevant component of interaction, as they elicit a sense of lifelike agency, and hence strong visceral responses that can play an important role in the reactions to the interface, to its acceptance or rejection. Thus in social HRI, visceral impressions form a crucial

component of the overall experience, and P_1 can be used to focus attention on these factors when assessing interaction with a robot.

3.6.2.2 *Perspective 2 (P₂) — Social Mechanics*

Many robots are designed to explicitly try to understand and communicate using social techniques such as those that are used between people (or perhaps between a person and an animal). This kind of communication consists of an extremely diverse set of social signals, responses, and other communication techniques, for example, such as the use of speech and voices, facial expressions, bodily gestures, and leveraging social norms such as appropriate appearance. We collectively refer to these communication techniques as the *social mechanics* of interaction, our second perspective (P_2).

People are very good at interpreting and understanding social mechanics, and are inclined to explain robot interaction using P_2 social communication techniques even where there is no communication intended. This is particularly true with robots, as their physical embodiment and active agency help make interaction with people inherently social. For example, despite having no internal social model, people attribute intentionality to the iRobot Roomba similar to how they may for another person or an animal and explain its floor-cleaning actions using P_2 . People have further been found to name their Roomba, have (one-sided) conversations with it, and dress it up to match their interpretation of its personality (Forlizzi and DiSalvo, 2006; Sung et al., 2007), all P_2 -type findings.

Examples of robots that attempt to explicitly use P_2 social mechanics are those that use such techniques as eye gaze cues, or head-nod recognition as an important part of interaction (Mutlu et al., 2009; Sidner et al., 2006; Staudte and Crocker, 2009), robots that appropriately yield to people while traversing a hallway (Pacchierotti, Christensen, and Jensfelt, 2006) or approaching seated people (Gockley et al., 2007), and those that convey an expression or mood to represent their state (Gockley et al., 2006). Robots' use of P_2 social mechanics extends beyond these more clear-cut examples, and includes subtle characteristics such as the tone and inflection of actions, components that can play a crucial role in overall interaction experience. One recent study identified that a subtle indication of team-play (i. e., by using the word "we") could largely increase the tolerance people have of robots' mistakes (Groom, Chen, Johnson, Kara, and Nass, 2010).

P₂ can also be more subtle, for example, it is conceivable that seemingly localized design decisions, such as a sporadic or rough (or *jerky*) arm movement, can taint the overall impression: one robot that debates using rough (perhaps aggressive) hand gestures may be received quite differently from another that uses smooth (perhaps docile) ones, or they would also be seen as different if the robots used a monotonous or bored versus excited voice in their statements. While there is a P₁ element to these interactions, P₂ explains how such actions can be explicit and intentional rather than visceral.

For much of P₂-targeting research the aim is to define how robots can comply with social practises and appear normal and acceptable in our lives. One approach to this has been to attempt to make robots that break what we accept as *normal* P₂ behaviour, as a means to both provide understanding of how people react when things go wrong, and to find the boundaries of what is seen as wrong. Notable research includes a robot that cheats while playing a game (Short, Hart, Vu, and Scassellati, 2010), one that purposefully talks in a disconnected manner (Takayama, Groom, and Nass, 2009), and one that uses inappropriate gaze cues to disrupt the flow of interaction (Muhl and Nagai, 2007).

For a robot to use P₂ social mechanics to the extent a human does poses difficult challenges such as sharp awareness of the context of interaction (e. g., interacting with a child versus an adult) or the culture they are interacting in (e. g., Asian versus European). One approach to this challenge is to program robots that can *learn* from their particular context on how to interact, mimicking the familiar P₂ methodology of teaching that people understand from the real world. Examples include our own *puppet master* projects, a separate project where people explicitly demonstrate to a robot how to push a sequence of buttons (Breazeal et al., 2004; Lockerd and Breazeal, 2004), or to observe and follow behaviours (Tanaka, Movellan, Fortenberry, and Aisaka, 2006); this last study showed that people perceived a robot that could learn as being more capable than the one that performed stock behaviours. P₂ does not imply complex social abilities, but the explicit and intentional use of social mechanics.

It appears that social mechanics (P₂) may be the most extensively studied area in social HRI, perhaps because it is often a clear part of the overall social interaction experience, and thus a clearer target for design. In this section, we have outlined what we feel are some of the current and active P₂ social-mechanic areas in social HRI, including both obvious and subtle components, robots that explicitly break communication practises, and robots that

manage complexity through learning from people. Exploring the vast landscape of P₂-type interaction is a rich area for future work.

3.6.2.3 *Perspective 3 (P₃) — Social Structures*

In addition to the arguably more salient P₁ (visceral) and P₂ (social mechanics) components of social HRI, interaction between people and robots extends into the holistic context of interaction (P₃). That is, the human environment and social structures are themselves components of interaction, where they both influence and are influenced in the process. One example of this kind of interaction is the relationship between a domestic robot and the social structures of the home: the existing home practises and contexts help define how people will perceive and interact with the robot, and the simple existence of the robot itself, and the fact that people interact with it, has an impact on the greater structures of the home.

Research in this area has shown that, for example, adopting cleaning-robot technology (a Roomba, in this case) in homes may shift who is responsible for the cleaning duties, from adults to young adults, and from women to men — a P₃ change (Forlizzi and DiSalvo, 2006). Other work has shown that robots can be attributed with moral rights and responsibilities of their own within the home and family (Friedman et al., 2003). In one case, a family expressed sadness at having to exchange their broken Roomba (named “Spot”) for another one, rather than having it fixed (Sung et al., 2007), evidence that the robot had P₃ interaction to become an important part of the household. The same phenomenon has been found in military contexts, where a bomb reconnaissance robot (named “Scooby Doo” by the soldiers) became a team member and was given medals by the team (Garreau, 2007).

Time can be a useful factor to consider in relation to how a robot fits into social acceptance and social structures; time can help highlight the extent of influence and a trajectory of how the social structures vary and evolve. For example, research has shown how a novelty factor can exist for robots, where they initially have an impact on structures, but are soon forgotten, with P₃ social structures returning toward their previous state. This has been demonstrated in research, where an office-assistant robot became forgotten after three months (Hüttenrauch and Eklundh, 2002), and a robot which was deployed into a classroom had much less interaction with children after two months (Kanda, Sato, Saiwaki, and Ishiguro, 2007b). Not all changes tend toward less use. Some studies have shown, for example, how people build emotional bonds with robots that strengthen over time, treating them as more than

mechanical beings, such as with people who treated their Roomba vacuum as a member of the household (Sung et al., 2009b).

Social HRI work that explicitly targets P₃ interaction is rare, perhaps due to the complexity and difficulty of exploring, explaining, or perhaps measuring social structures and the influence that HRI may have. This problem is exacerbated for longitudinal studies which may cover large and complex environments, such as homes and offices. However, P₃ can occur whether explicitly designed for or not, and some do study P₃ for robots, regardless of their explicit ability or intention to either interpret or interact on social structures (e. g., Forlizzi, 2007; Friedman et al., 2003; Sung et al., 2008; Takayama, Ju, and Nass, 2008).

In the above sections we have detailed P₁, P₂, P₃ and illustrated how they can be used to categorize and describe social HRI instances. This organizes the amorphous term *social* into a simpler social HRI-targeted toolkit. Below, we demonstrate how the perspectives can be used to analyze interaction experience.

3.6.2.4 *Describing Interaction Experience*

The perspectives can be used to clearly articulate complex, multi-level social interaction experience between a person and a robot. One could use them to define hypotheses, for example, that perhaps a particular robot will elicit pleasure (P₁: visceral reactions) and, when this occurs, the person will respond by using some social mechanics (P₂), such as waving at the robot or bobbing their head. As another example of use, in a hypothetical study, “people found the robot to be creepy, which they expressed both in P₁-type externalized reactions and P₂ gestures such as ‘keep away’ hand gestures, and this had very strong P₃-type interactions with the home.” In this example, the perspectives highlight the difference between perhaps involuntary P₁ and voluntary P₂ interactions, and the more individual P₁, P₂ in comparison to related P₃ social-structure impacts. We argue that the perspectives help simplify the communication of this concept.

In this section we have presented a new three-perspective categorization of social HRI for focusing on various aspects of the overall breadth of interaction. In the following section, we present a set of considerations that outline various aspects related to how people perceive robots within the *holistic interaction context*.

3.6.3 *Factors Shaping Robot Acceptance*

In this section we summarize our technology-adoption explorations into a set of factors which we believe are particularly influential to people's perceptions of robots. These have emerged primarily from our theoretical exploration (Section 3.4), although some points are further supported by from our own experiences studying our social HRI interfaces as highlighted through our several evaluations.

We present two layers of considerations which we believe can help designers and evaluators explicitly consider the social landscape of interaction within the *holistic interaction context*: first, we discuss social factors of concern to people regarding robots, and second we outline important sources of influence that shape how people perceive robots.

3.6.3.1 *Social Factors Affecting Acceptance*

The following factors are socially-rooted aspects which we believe will be very important for the acceptance of robots.

SOCIAL INTEGRATION — We believe that the question of how robots will integrate into existing social structures (such as the home or office) is particularly important; robots should attempt to act appropriately according to the wide context of interaction. This does not imply a requirement of complex social-situation-recognizing ability, however, as these goals can be achieved through clever and careful design.

SAFETY — Robots create a level of potential danger seldom experienced in the past with other domestic and everyday technologies: in a worst-case scenario, they can damage property or seriously injure and kill people. As such, we expect this concern to be disproportionately important.

ACCESSIBILITY AND USABILITY — The complexity of robots' capabilities raises serious accessibility concerns. We expect existing technology fears such as lack of knowledge and behavioural control, shown to have been a problem for PC adoption, will escalate given the physical presence and potential safety hazards of robots. Other barriers include facilities and space requirements within the home, financial practicality (affordability, maintenance and obsolescence) and legal barriers and regulations.

FUN AND ENJOYMENT — Fun and enjoyment are very important to people. In addition to entertainment robots, this includes secondary effects such as more free time due to utility gains. Further, we expect that robots will emerge to play a role in the basic human need of companionship and comfort, for example, in home environments as well as hospitals or possibly day-care centres.

SOCIAL PRESSURES AND STATUS GAINS — As robots become common we can expect the emergence of social pressures and status gains such as a family wanting to appear “modern.” The status gains associated with being perceived as a cutting-edge person, family, or establishment (e. g., store), or being recognized as a knowledgeable reference by neighbours or co-workers, has been important for technology adoption in the past.

3.6.3.2 *The Perception of Robots*

Here we present factors which people use to shape their perceptions of robots. Perceptions are often as or more meaningful than more objective facts about robots, and so robot designers should consider these factors in relation to how their robots will be perceived and accepted.

PREVIOUS EXPERIENCE — This includes personal lifetime experience as well as personally-inferred beliefs, with education and initial exposure being large factors. Given people’s limited exposure to robots, and how they attribute agency to them, perhaps previous experience with animals and children may be particularly influential.

MEDIA — Where previous experience is weak, media becomes an important source of information. This includes classic (perhaps science-fiction-like) literature, movies, and television, as well as various news sources, and can consist of both positive and luxurious as well as negative and dangerous portrayals. Designers can leverage these (negative or positive) media trends to affect perceptions of their robots, or with sufficient resources, media can be generated to attempt to shape perceptions.

PERSONAL SOCIAL NETWORK — We believe that opinions offered by friends, neighbours and family will have a large influence on how people perceive robots, despite the fact that robots are new and the social network itself will be less informed. Perhaps making an environment around the robot conducive to socializing may be helpful, for example, this could include particular robot designs that make it a conversation piece, or on-line support networks or integration into social networking sites.

In this section we provided a new set of influences that can be used to describe and break down the components of interaction between a person and a robot, and a clear set of criteria that outline the factors and influences we believe will be particularly important for how people perceive robots.

3.7 SUMMARY: A THEORETICAL SOCIAL HRI DISCUSSION

In this chapter we have provided a thorough theoretical exploration of social HRI that includes a discussion on how interaction with robots is unique, what the keyword *social* means for social HRI, and why people tend to treat robots differently than more traditional technologies. We further presented an exploration into social-psychology domestication of technology models, and a targeted robot analysis, as a means to help better understand the factors and concerns that are important to people in how they shape their perceptions of robots.

We synthesized our overall theoretical discussion into a social HRI framework presented at the end of this chapter. This framework contains a concise definition of social HRI that clearly provides the various useful vocabulary that emerged from our overall exploration. Further, we present a new vocabulary that spans the breadth of social HRI, from visceral-level reactions, to social mechanics, to interaction with social structures, and illustrated how this can be used to describe and discuss existing social HRI work. Finally, the framework contains a list of targeted factors that we believe shape how people perceive robots, and that we recommend social HRI researchers consider in respect to their robotic designs.

Our theoretical framework sets the tone for our work throughout the rest of this dissertation, and explains and motivates many of our design and evaluation-focus decisions. The framework itself, however, was constructed in parallel with (often informed by) our implementations and evaluations, and so we did not directly apply this exact framework to design and construct the interface designs.

Part II

IMPLEMENTING AND EVALUATING SOCIAL ROBOTIC INTERFACES

PART II INTRODUCTION

In Part II of this dissertation we introduce our original social HRI interface designs, implementations, and evaluations, work that represents an exploration of how robotic interfaces can be designed to integrate into and leverage *the social stock of knowledge*. Our selection of projects was led directly by our research questions as outlined in Section 1.3.5, page 8: (Q1) What does the tendency to treat robots as social entities mean for interaction between a person and a robot? (Q2) How can robotic interfaces be designed to leverage this tendency? and (Q3) Which methodologies, structured techniques, taxonomies, and heuristics can be developed and used for social HRI?

Our interaction designs illustrate new ways that robots can use familiar techniques and scenarios to make difficult robot interaction and control problems accessible (Q2), and our original implementations and algorithms serve as proofs-of-concepts and tools that researchers can use in their own social HRI implementations (Q3). Further, a core component of this work was to create scenarios with working robotic interfaces where we could observe the social aspects of people interacting with robots, helping to better understand the *holistic interaction context* in each case (Q1).

As such, our particular selection of interaction designs surrounded an exploration of which existing social techniques robots could be designed to use and which social scenarios they could integrate into, grounded in what can be practically achieved with modern robots. In retrospect, we notice that the selection of projects was somewhat arbitrary, as indeed we did not follow a carefully-structured framework in our selections, and this is perhaps a limitation of our research. in this dissertation. However, we maintain that our project selections serve our goal of exploring our research questions, enabling us to reflect on the breadth and depth of higher level social interactions between robots and people (i. e., the *holistic interaction experience*). Below we briefly discuss each project, and how it relates to our overarching theme.

A DOG-LEASH INTERFACE FOR LEADING A ROBOT (CHAPTER 4)

This chapter introduces a new interface design and implementation solution for controlling a robot that leverages the familiar scenario of leading an animal on a leash (Q2 and Q3). This includes a formal evaluation that explores social layers to people's interactions with this interface (Q1), including targeted consideration of how particular evaluation methodologies can be employed (Q3).

EVALUATING "TOUCH AND TOYS" (CHAPTER 5)

Here we present a detailed evaluation of an interface titled "Touch and Toys." This evaluation serves to highlight the importance of considering interaction with robots beyond questions of task efficiency or usability (Q1), and to further demonstrate how social aspects of interaction can be targeted with evaluation (Q3).

CARTOON ARTWORK ROBOTIC INTERFACES (CHAPTER 6)

This chapter demonstrates how robots can use familiar cartoon artwork to communicate robotic state in abstract (but easily-understood) ways (Q2). We present informal design critiques that improve our understanding of how social aspects impact overall interaction (Q1), and several original implementations and algorithms for realizing our prototypes (Q3).

STYLISTIC LOCOMOTION AND PUPPET MASTER (CHAPTER 7)

The *stylistic locomotion* and *puppet master* projects show how robots a) can use the style of their locomotion to communicate with people, and b) can learn directly from people's existing demonstration abilities how they should move (Q2). We present the results from extensive evaluations, outlining the various socially-oriented aspects of interaction that emerged (Q1), and reflect on our particular use of evaluation methodology to target social HRI (Q3). Further, we present several novel implementations and extensive original algorithms that show how these interfaces can be realized (Q3).

A DOG-LEASH INTERFACE FOR LEADING A ROBOT

In this chapter we propose a dog-leash interface for leading a robot, where a person holds the handle of a leash attached to a robot, and uses this to lead the robot to where they want it to go (Figure 4.1). This interface is practical from a utility perspective, for example, for a nurse who may bring along a medicine robot when traveling throughout a hospital, or an elderly person who may take a robot shopping with them to carry their groceries. This also speaks to our *accessibility and usability* factor of acceptance, as well as indirectly to the *fun and enjoyment* factor due to perhaps saved time or energy. We also believe that this method of leading a robot is inherently a social task, and we designed our project in an attempt to directly use *the social stock of knowledge*. Below we examine the dog-leash robot interaction concept reflecting on our social Human-Robot Interaction (HRI) theoretical framework presented earlier (Section 3.6), using our three social HRI perspectives: P₁ (visceral), P₂ (social mechanics), and P₃ (social structures).

The dog-leash robot interaction metaphor is fundamentally based on social interaction between a person (leader) and a dog (led), where ongoing communication between the two is required for the overall leading task to work. This includes such explicit *social mechanics* (P₂ interaction) communication as watching the other's movement direction and movement speed, pulling on the leash (by either entity), and adjusting actions accordingly. Also, we argue that there is a more-immediate visceral (P₁) level to this interaction: as soon as the leash is placed roles emerge, where the person knows they are to lead, and the dog (ideally) knows it is supposed to be led. Further, using a leash to lead something, be it an animal or toy, is ingrained in culture by the many cases where leashes are used to lead an animal, say, dog, horse, donkey, or cat, or in some cultures, a young child.



Figure 4.1: a participant leading a robot using our dog-leash interface

This technique is also conceptually supported by the simple physics of pulling, or pushing, something. Thus our robotic dog-leash interface draws from a person's *previous experience* and leverages their general understanding of how the world works. As such, we expect that our dog-leash interface will be naturally understood and easy to learn by members of the general public (related to the *accessibility and usability* factor of acceptance).

The dog-leash approach is also socially understood and fits into existing broader social structures (P₃ interaction): people not leading the robot, such as bystanders or a passerby in a crowd, likely understand from their *social stock of knowledge* of what is happening when they see a person leading a robot on a leash, and know what to expect. This this adheres to our *social integration* factor of social acceptance.

In the remainder of this chapter we first present an overview of our dog-leash interface design, including a discussion of related work, we follow with an implementation-details section, and finish with a formal evaluation of the dog-leash interface.

4.1 PAST ROBOT-LEADING INTERFACES

Tachi, Mann, and Rowell (1983) present a project where, rather than a person leading a robot, the robot leads the person as a guide-dog robot for the visually impaired. The robot has knowledge of the environment, and the person wears a stereo headset which notifies them (via coded aural feedback) if they are straying from the path. Communication happens via sonar, and no leash is used. The person in this scenario must learn the new aural-feedback code, and has no interaction whatsoever with the robot; the robot serves only as a kind of beacon that communicates with the person's headset.

Ootake, Fukaya, Syouzu, and Nagai (2008) present the only project we are aware of that uses a leash to lead a robot. In this project, force-sensors are used to detect in which direction a person is pulling so that the robot can follow, meaning that a fixed-length string must be kept taut at all times. In addition to placing constraints on how the person must walk, this solution means that the robot must always be behind the person, and cannot follow at the side or be itself led from behind (robot in front of the person). This project is primarily a technical contribution and these authors do not address the social aspects of interaction.

There is very little work that considers the social aspects of a person leading a robot. Gockley et al. (2007) present a study on how a robot can follow a person naturally, comparing

two different methods: copying the exact path taken by the person, or taking a shortest path (cutting corners). People reported the shortest path as feeling more natural. This work used a laser range-finder tracking techniques with no physical constraint or connection between the person and the robot. This work highlighted the fact that people do attribute notions of *natural* to robots' actions, although we believe the social elements of a robot following at its own pace are quite different from those for a robot being lead on a tethered leash.

The engineering problem of person-following robots has been approached by, for example, mounting laser range finders (e. g., Montemerlo, Thrun, and Whittaker, 2002; Kluge, Köhler, and Prassler, 2001) or cameras (e. g., Kleinhagenbrock, Lang, Fritsch, Lömer, Fink, and Sagerer, 2002) on robots to allow them to detect and follow the person. These sometimes require pre-calculated maps of the environment, and can be prone to failures when occlusions occur or busy environments are encountered. Another approach is to mount an active device such as a sonar on the person for the robot to detect (e. g., Bianco, Caretti, and Nolfi, 2003), although this can be heavy and is still prone to occlusions, noise, and reflections; robust person-following remains an open problem.

Our approach of having a physical leash between the person and the robot improves the person-detecting problem as only one person will hold the leash at once, and the robot can know roughly where that person is by closely monitoring the leash only. This can further improve the scalability of the system to crowded areas, and perhaps even to difficult interaction settings, such as rough and uneven terrains. Finally, our retractable leash can smoothly change length to match a person's walking pattern (or stride), and the robot can be in any relative position (in front, behind, to the side) while still tracking where the person is.


We present a dog-leash interface design and evaluation which was done specifically with social considerations in mind. Our interface enables a person to walk naturally, the robot to be technically capable of being at any relative position that the person may want it to be at, and the dog-leash metaphor is also familiar to onlookers. In the remainder of this chapter we detail our particular interface design, implementation, and evaluation.

4.2 DESIGNING A DOG-LEASH ROBOT INTERFACE


Here we present our dog-leash interface for leading a robot, with two distinct variations: the robot following behind a person and the robot in front of the person. In both instances our

focus is to create an interface that people can naturally understand, and quickly use with minimal (ideally zero) instruction. A person holding the leash is shown in Figure 4.2a.

Our leash interfaces are based on a spring-loaded retractable mechanism, where at rest the leash handle is at the robot, and can be pulled out to roughly four metres. While holding the leash there is slight tension from the spring, but it does not restrict the person from walking or moving their arms naturally with the leash smoothly extending and retracting while they walk. These kinds of leashes are popular and so we expect the mechanism to be familiar to many people. At the end of the leash is a handle for the person to hold (Figure 4.2b), with a red emergency-stop button mounted at a location easily pressed by the thumb. The exact same leash mechanism is used in all the interaction cases described below.



(a) a person holding the robot on a leash



(b) close-up of the leash handle, with the red emergency-stop button

Figure 4.2: our robotic dog-leash interface

One of our concerns with this interface is the danger associated with the robot — this relates to the *safety* factor of shaping acceptance (Section 3.6.3). A robot which is powerful enough to keep up with a person is likely also dangerous in a worst-case malfunction or accident scenario. While the general question of robot safety is an important research area beyond the scope of our dissertation, in this project we have included an emergency-stop button which the person can push at any time to cut the power to the robot. This speaks to the *safety* factor of acceptance in our theoretical framework.

Later in this chapter we present detailed discussions on our implementation and the study we conducted. First, however, we detail the different interaction cases.

4.2.1 *Robot Following Behind the Person*

In this scenario, the interface is designed for the person to walk normally holding the leash, and the robot to follow behind at an appropriate distance. The person does not need to concern themselves with how the robot will move, turn, etc., but can just expect that the robot will catch up. In the case where the person moves too quickly for the robot to keep pace, or if the robot somehow makes an error and becomes too distant, the leash will simply run out and the person will feel the tug on the leash, forcing them to slow down or stop.

In this design, the robot is set to keep the leash roughly 1.7 m long. As this distance increases the robot moves faster to keep up, and if this distance decreases the robot slows down or backs away from the person. The robot turns automatically to keep facing in the correct direction toward the person. Thus, as the person moves the robot follows, as the person stops the robot automatically stops, and as the person turns and changes direction the robot automatically adjusts its trajectory — no robot actions need to be specified by the person.

We developed two cases for this interface: the robot directly behind the person (Figure 4.3a) and the robot behind and to the side of the person, at an angle of roughly 45° (Figure 4.3b, the angle dynamically shifts between the left and right sides as explained below). We implemented both to have a visibility (at-angle is easier to see in periphery vision) versus space usage trade-off: much more space is required to lead the robot as it stands much further off to the side. For the behind-at-an-angle case, a right-handed person may interact differently than a left-handed one, for example, wanting the robot to stay behind on their left and not on the right. Our solution to this was to have the robot automatically detect which side it



(a) robot directly behind the person (b) robot behind and to the side of the person

Figure 4.3: the two ways that the robot can follow behind the person

should be on by where the person is in relation to it. Thus the person can change hands and walk naturally, and the robot will adapt by moving to the appropriate side. The interaction dynamics of the handedness question is important future work.

4.2.2 *Robot In-Front of the Person*

Rather than having the robot behind the person, and somewhat out of sight, putting the robot in front of the person enables it to be easily and constantly monitored during operation. As our robot was not quick (nor clever) enough to stay in front of the person as they walk on their path without additional input, we implemented a simple control scheme based on a push-stick metaphor. An image of a person directing a robot this way is given in Figure 4.4

Here, the person leads the robot from behind as if the robot was attached to a rigid stick, except that the spring-loaded leash makes this interaction less rigid than a stick would. As the person walks toward the robot and the leash gets shorter the robot moves away from the person, such that the person can walk at a comfortable pace and the robot stays in front of them. If the person backs away from the robot and the leash gets longer, the robot backs up to correct the leash length.



Figure 4.4: a robot being led on a leash from behind, where the person stays behind the robot

The robot's turning also follows the on-the-stick metaphor. The robot tries to keep the leash aligned with its vertical (front-back) axis much as how pushing a wheeled object on a stick tends to stay straight as you push. This means that, as they walk, the person does not have to manage small deviations in their path. For large turns, the person walks to the side of the robot as if to gain a better point from which to push the robot: to turn the robot left, they walk to the robot's right side and toward the robot as if they were pushing it with a stick. In this scenario the robot attempted to keep the leash at a length of roughly 1 m.

4.3 IMPLEMENTING OUR DOG-LEASH INTERFACE

We implemented our dog-leash interface for leading a robot using the Mobile Robots Inc. 3-AT, designed a custom-made retractable dog-leash mechanism, and used a standard Personal Computer (PC) for control (Figure 4.5). All software was written in C++ and Java.

Our dog-leash mechanism was designed using an off-the-shelf retractable dog leash, which we mounted atop the robot on an absolute (720 ticks / revolution) rotary encoder (Koyo Electric TRD-NA720NW). The assembly can rotate freely, so as a person pulls the leash and walks around the robot, the rotary encoder can sense in which direction the leash is directed, and following, where person is (Figure 4.5b). This information is sent to the controlling PC


over 802.11g using a Lantronix WiPort modem; the same modem connection is used by the controlling PC to send the robot movement commands.

Inside the leash assembly, the string is stored on a spring-loaded spool which can be pulled out and will automatically retract if released (Figure 4.5b). We attached a second, relative (spin-directional) rotary encoder (COPAL Electric 100-213-1, 64 ticks / revolution) to measure the rotation of the spool and sense when the string is being pulled or released. This information is sent back to the control PC over 802.11g using an additional Lantronix WiPort modem (Figure 4.5c) and is used to estimate the current length of the pulled string. Thus, the controlling PC senses both the angle to the person, and the distance (leash length), and can estimate where the person is in polar coordinates. Following, the PC generates locomotion commands to the robot to follow the person appropriately.


As the entire leash spool needs to turn freely around the robot (mounted on the rotary encoder) to facilitate natural walking, we did not want to run a cable from the top of the mechanism to the robot as it would snag as the leash turns around the robot. Our solution was to completely contain the top encoder, modem, and some batteries on top of the leash assembly as shown in Figure 4.5c.

To implement the on-leash emergency stop button (Figure 4.5b) we replaced the leash string with a two-strand wire with each wire connected to a terminal of the (normally-closed) button. These wires are inserted in series into the robot's existing emergency-stop mechanism based on a normally-closed circuit, such that if the button is pressed, the circuit is opened and the robot stops moving. We connected this wire from the leash to the base robot through two sets of slip-ring-and-brush assemblies — once from the spinning spool to the main leash assembly, and once from the leash assembly to the base robot (Figure 4.5b). Further, if the leash assembly breaks the emergency circuit goes open and the robot stops.


The software model for controlling the robot is based on closed-loop feedback, where the robot constantly monitors the person's position and fine-tunes its own behaviour in real time (at 15 Hz). For each follow style (robot behind, robot behind at an angle, robot in front), we have defined for the robot's locomotion algorithm a target leash length and target position zone in relation to the person, and have instructed the robot on how it should best try to reach that target. This is also dependent on the current position, for example, the robot may have to turn around before moving forward. The locomotion algorithm is designed



(a) dog-leash mechanism mounted on robot



(b) our custom-made dog-leash mechanism



(c) a top-view, showing the leash-length-sensor assembly

Figure 4.5: dog-leash-robot leash mechanism

so the closer the robot gets to its target leash length and position zone, the less drastic its movements are, leading to a smooth and stable result.

4.4 EVALUATING THE ROBOTIC DOG-LEASH INTERFACE

The dog-leash robotic interface study was an opportunity to investigate how to evaluate and explore core social HRI challenges, in addition to the more interface-specific questions. In this study we tried to explore a person's disposition toward the robot during the dog-leash interaction to target emotional responses and states, particularly focusing on using existing and previously validated questionnaires. Through these experiences we have improved our understanding of how existing methods apply to social HRI.

This study revolved around having participants complete simple navigation tasks with the robot where they picked up and dropped off items (carried by the robot) at designated locations, with the robot-following method as the independent variable. All study materials are included in Appendix B. The direct purpose of this evaluation was to test the basic usability of the interface, in terms of whether the robot could satisfactorily follow a person, and whether the leash interface makes sense to people. Looking more from a social HRI perspective, we also posed questions that target participants' emotional state, such as how they feel about the dog-leash robot when it is in front of them, behind them or to the side, and how interacting with our robot for a short time influences disposition toward robots. We also considered how far from the person the robot should be to help people feel comfortable.

4.4.1 *Design Critiques*

We performed several preliminary in-lab, informal design critiques which helped us fine-tune the protocol, robot behaviours, and the robot's follow distances, lessons we outline here. One such case is the idea of leash gestures: we implemented a system where the robot could detect a single or double tug, and act differently accordingly, for example, to pause or move more quickly. Initial testing revealed that this was confusing to people and that they did not use it, so we decided to omit this feature, and it was not included in our studies.

We also found that with the directly-*behind* condition people chronically kept turning to look behind them to see what the robot was doing, citing concerns over the robot colliding

with various objects in the environment. This led to the development and inclusion of the robot *behind angle* case. In the same series of tests, we found that a robot close behind was very uncomfortable and a little frightening, and so we designed the *behind* follow to be further away (1.7 m). On the other hand, having the *in-front* robot far away seemed to hinder sense of control, and particularly when turning the robot, having it closer made it feel easier to manipulate; the robot was closer with *front* condition (1 m).

4.4.2 Tasks

The participants' primary task was to follow a route with the robot and to pick up and deliver objects (carried by the robot) from and to designated locations. This path is illustrated in Figure 4.6, and detailed task instructions are included in Appendix B Section B.2; the route was designed to include both long and short passages, wide and narrow curves, and obstacles.

At the end of the study, we performed an auxiliary task to measure participants' comfort distance of the robot from the person. We investigated both the approach and withdrawal distance, where in both cases the person moved to approach or withdraw, and the robot stayed still. For approach distance we asked the participant to move toward the robot and stop as soon as they felt no longer comfortable with the distance (i. e., the robot is too close). For withdrawal we asked the participant to get as close as physically possible to the robot

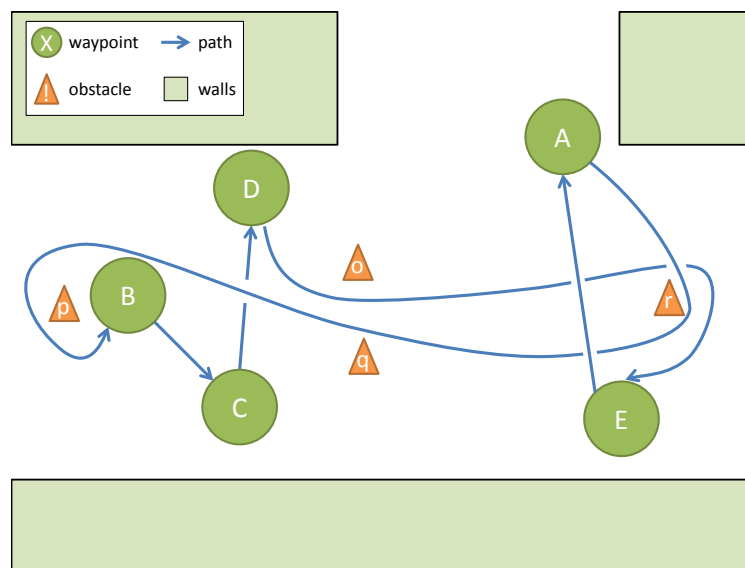


Figure 4.6: the map, path, and way points for the dog-leash study, participant starts at point E to go to first way point A

and move away until they felt comfortable. We did this procedure twice, once while the participant was holding the leash and once without the leash.

4.4.3 Study Procedure

We conducted this study using a structured protocol including the use of informed consent forms and the use of questionnaires. Participants first completed pre-test questionnaires designed to target previous experience with robots and related technology such as computers, video games, and driving or operating machinery such as automobiles or forklifts. We also enquired about their experience with pets (especially dogs). Following, participants were introduced to the robot and shown how to use the emergency-stop button.

The participants completed tasks for each behaviour, before each trial the participant could try the given behaviour, then performed two entire circuits around a pre-defined path (Figure 4.6) executing the pickup and delivery instructions (Section B.2). The obstacles were plastic cones which the participant had to walk around. After each behaviour we administered questionnaires to explore the participant reactions specific to the given behaviour, in part using the the Self-Assessment Manikin (SAM) technique (Morris, 1995) to enquire about participant emotional state, measuring *pleasure* and *arousal* on nine-point Likert-like scales. SAM was also used during the pre-test to serve as a baseline for comparison purposes. Also after each behaviour, we administered variants on the GODSPEED questionnaires (Bartneck, Kulić, Croft, and Zoghbi, 2009b), to measure perceived safety and likability.

After these tasks the participant performed the comfort-distance task, and post-test, we asked various free-form questions relating to participants' impression of the robot, feeling of safety, if they felt in control, and their overall preferences. Also, we used the measurement of Negative Attitudes towards Robots Scale (NARS) (Nomura, Suzuki, and Kanda, 2006) during both the pre- and post-test phases to explore participants' disposition toward robots and how it changed through participation. NARS assesses a person's general opinions of robots on three scales: negative attitudes toward situations and interactions with robots (*interaction*), negative attitudes toward social influence of robots (*social*), and negative attitudes toward emotions in interactions with robots (*emotion*). Lower scores mean more-positive responses.

4.4.4 Study Design

The main independent variable in this study was the robot behaviour type: robot following directly behind the person (referred to as *behind*), robot following behind at an angle (referred to as *behind angle*), and person leading the robot in front (*front*). We had a within-subjects design, such that each participant did the tasks with each behaviour, order counterbalanced between participants.

The study took place in Yokohama, Japan, in a model-home complex called HouseSquare Yokohama. Twelve male right-handed Japanese students ranging in age from 20 to 23 (M=21.1) participated in the study, for which they received ¥4500 (Japanese Yen, approximately \$55 2010 Canadian Dollars) for their participation.

4.4.5 Results

Participant rankings of the three behaviour types on the given questions are shown in Table 4.1. Friedman’s ANalysis of VAriance (ANOVA)s failed to expose a significant effect of behaviour type on how each behaviour was ranked ($\chi^2(2)=2.92, p=.232$), but an effect was found for the participant feeling the most in control ($\chi^2(2)=6.62, p=.037$), and the analysis suggests a trend toward the robot being rated as doing what the participant wanted it to do ($\chi^2(2)=5.69, p=.058$). Six participants stated (post-test questionnaire) that they would recommend a friend *front*, three would recommend *behind*, and three would recommend *behind angle*.

The change in emotional state after interacting with each behaviour as measured on *pleasure* and *arousal* via the SAM scale (Morris, 1995) is outlined in Table 4.2; only eleven responses

	behind	beh. angle	front	behind	beh. angle	front	behind	beh. angle	front
ranking 3	1	8	4	3	8	2	2	8	3
2	9	2	2	6	4	3	6	4	3
1	3	3	7	4	1	8	5	1	7
	I liked ... best			I felt most in control with...			I felt the robot did what I wanted it to do the most for...		

Table 4.1: The result table of how participants ranked the behaviours in relation to each other. Each number represents how many participants ranked that given behaviour as first, second, or third, for each question.

4	0	0	0	4	0	0	0
3	1	1	1	3	1	0	2
2	1	2	2	2	0	4	4
1	3	0	5	1	4	2	0
0	3	3	0	0	3	3	2
-1	0	0	1	-1	2	0	1
-2	2	4	2	-2	0	1	2
-3	1	0	0	-3	1	0	0
-4	0	1	0	-4	0	0	0
	behind	behind angle	front		behind	behind angle	front

(a) change in pleasure, where positive number indicates change toward positive end of scale

(b) change in arousal, where positive number indicates change toward excited / anxious end of scale

Table 4.2: table of number of participants who had particular changes in emotional state as measured by the SAM scale (Morris, 1995)

are included as one participant did not complete the questionnaires. The table shows how pleasure increased in comparison to pretest for the *front* condition, and decreased for the *behind angle* condition. However, Friedman's ANOVA failed to reveal an effect of behaviour type on pre-test versus post-behaviour pleasure ($\chi^2(3)=2.63, p=.452$). On the arousal scale, a Friedman's ANOVA suggests an effect of robot behaviour on arousal ($\chi^2(3)=23.67, p<0.001$). Post-hoc Wilcoxon Signed Ranks tests, with a Bonferroni correction (effects considered significant at $p=.008$, all six cases tested) failed to reveal further relationships ($p>.2$). No effect was found on behaviour type for how much the participant claimed they enjoyed the interaction after using each behaviour (Friedman's ANOVA, $\chi^2(2)=1.39, p=.499$).

Many participants stated that they felt in control of the robot (*back*: 8, *back angle*: 4, *front*: 8). On the other hand, many participants commented that the robot was unpredictable and it did not move as they wanted or expected (*back*: 7, *back angle*: 10, *front*: 5), with comments specifying difficulty controlling speed and turning. Note the disparity between responses to *behind angle* and the other two. Table 4.3 is a table of participant responses to questions relating to perceived control of the robot, and Table 4.4 shows responses to questions of impressions of the robot. Friedman's ANOVA tests did not reveal any effect of behaviour type on responses to any these questions. On the question of robot visibility, it was only mentioned as a problem for the *behind angle* behaviour, and was surprisingly not mentioned for walking directly *behind* (*behind*: 0, *back angle*: 5, *front*: 0).

	1	2	3	4	5	6	7	
not controllable	1	2	8	7	10	6	2	controllable
not predictable	1	1	7	6	8	9	4	predictable
not autonomous	0	3	10	6	11	5	1	autonomous

Table 4.3: cumulative result table of the questions targeting perceived control, value represents number of participants who gave that response

	1	2	3	4	5	6	7	
dislike	0	0	7	5	11	10	3	like
unfriendly	0	4	6	6	10	8	2	friendly
unkind	0	3	12	9	9	2	1	kind
unpleasant	0	3	5	9	10	8	1	pleasant
aweful	0	1	4	5	17	8	1	nice
aggressive	2	3	6	7	12	6	0	non-aggressive

Table 4.4: cumulative result table of the questions targeting participants’ impressions of the robot, value represents number of participants who gave that response

Table 4.5 shows the results of the GODSPEED V perceived-safety questionnaire (Bartneck et al., 2009b), which shows how participant responses generally indicated a safe or neutral disposition. We found a significant effect of behaviour on how participants rated their feeling on the *surprised to quiescent* scale (Friedman’s ANOVA, $\chi^2(2)=9.53, p=.009$), although post-hoc Wilcoxon Signed Ranks Tests (with Bonferonni correction for significance at $p=.017$, three cases tested) failed to reveal further significant relationships: *behind angle* versus *behind* ($Z=-1.85 p=.065$), *front* versus *behind* ($Z=-.36 p=.722$), *front* versus *behind angle* ($Z=-2.32 p=.020$). On the *agitated* versus *calm* scale participants tended to be more agitated with the robot behind at an angle ($\chi^2=5.25, p=.072$, average ranks: *behind*=2.33, *behind angle*=1.58, *front*=2.08). No effect was found for behaviour type on the other scales.

The comparison of general disposition toward robots before and after our study is shown in Figure 4.7 as the average results of each of the three NARS scales; these results are on a scale from 1 (not negative) to 7 (negative). Note that responses were generally low, meaning dispo-

	1	2	3	4	5	6	7	
anxious	0	3	9	5	13	6	0	relaxed
agitated	1	5	5	12	11	1	1	calm
quiescent	0	1	5	7	12	9	2	surprised
unpleasant	0	3	5	12	10	5	1	comfortable

GODSPEED V: Perceived Safety

Table 4.5: cumulative result table of the GODSPEED V questionnaire on perceived safety, where lower scores are seen as unsafe (Bartneck et al., 2009b); value represents number of participants who gave that response




Figure 4.7: the result of the NARS questionnaire on the dog-leash study

sition toward robots was fairly positive, particularly on the interaction scale. We conducted a two-way repeated-measures ANOVA on this data, with testing time (pre or post-test) and scale type (interaction, social, or emotion) being the factors. We found a main effect of scale ($F(2,22)=44.33, p<.001$), with visual analysis of Figure 4.7 matching the same relationships found by Bartneck et al. (2009b). Main effect of test time (pre versus post test) was not significant ($F(2,22)=.09$), and no significant interaction effects were found ($F(2,22)=.91$).

Participants used the emergency stop button for its planned purpose, i. e., for preventing accidents and in case of emergency. Observation and preliminary data analysis suggested no effect of behaviour condition on emergency button use. Participants also used the emergency stop button as a means to temporarily pause the robot, even though they could simply stop walking for the same effect; the robot would automatically stop. Participants mentioned that the lack of a pause button was problematic, as the emergency-stop button resulted in a complete system shutdown and, once released, the robot took roughly 10 s to resume. Several participants commented that the robot was hard to stop (*back: 2, back angle: 5, front 6*).

Figure 4.8 details the results of the comfort-distance measuring phase, both for the approach and withdrawal conditions. We present the differences between leash and no leash per participant to focus on the difference between the conditions. While the figure suggests that holding the leash makes participants require a further comfort distance, a two-way repeated-measures ANOVA failed to find a significant effect of with / without leash ($F(1,11)=1.23, p=.290$) or withdrawal versus approach ($F(1,11)=.017$) on comfort distance, and there was no factor interaction observed ($F(1,11)=.305$).




Figure 4.8: comfort-distance difference when the participant is holding the leash versus no leash. The more positive the value (>0), comfort distance is further while holding a leash than without a leash. The more negative a value (<0), comfort distance is closer. The vertical lines indicate the disparity of the comfort-distance relationship between the cases.

4.4.6 Discussion

The results of this study suggest that overall people found the robot relatively easy to use (fairly controllable and predictable), generally liked the robot, and was fairly relaxed and calm while operating it. The GODSPEED V questionnaire results (Bartneck et al., 2009b) suggest that people generally felt safe toward the robot, and at the least, did not generally feel unsafe, and the NARS tests indicated general positive attitudes toward the robot. Overall, this supports the idea that the general public is both capable and comfortable with using our dog-leash interface for robot control.

The *front* behaviour was generally preferred, with seven people explicitly ranking it as their first choice, and half saying they would recommend it as the best to a friend. There was a tendency toward this preference throughout the rest of the study data, although statistical tests failed to find significant numerical results, and so we believe this is a relationship worth further exploration.

General negative response to the *behind angle* was a prevailing theme throughout the study, particularly in terms of cited usability and control problems. For example, participants reported that with *behind angle* the robot was harder to see, it further gave them a worse

sense of control and tended to make them feel less safe and more agitated with the robot. Particularly surprising is how this feedback compares to the *behind* case: while some of the *behind-angle* problems can be justified given that the robot is not directly in sight of the person, the same complaints were not mirrored in the directly *behind* behaviour. This contradicts the original design intent of the *behind angle*, that is, allowing the robot to be at the side of the person, so it can be easier to see compared to the *behind* condition, requiring less effort from the person and reducing the need to turn and look. One musing related to this cause is that perhaps the *behind angle* had a wider footprint, i. e., the robot was to the side of the person and so as a team they required more width-space to move, resulted in increased difficulty of control and added a negative overall feeling.

The NARS breakdown supported previous work which showed how the participants are more positive toward general interaction with the robot than the idea of the interaction involving social and emotional elements. Perhaps this is an indicator that approaches such as the dog-leash, which use familiar techniques that do not directly involve emotion or explicit social characteristics in their designs may be more acceptable by people. While the dog-leash metaphor itself leverages existing social knowledge, this is very subtle, and the robot does not give the obvious impression to people of trying to be a social actor. We note, however, that participants did not convey a negative tone toward social elements of the interaction.

This study points to the need of an explicit pause mechanism, given that the robot's emergency button was primarily used for merely pausing the movement. While the robot does stop when the person stops, people reported that they felt uneasy about this and wanted a more-explicit mechanism. Perhaps this is related to trust in the robot, where an explicit mechanism could enable them to directly be in control.

Results from our comfort-distance task failed to reveal any effect of holding a leash, or approach versus withdrawal, on people's preference for robot distance. Regardless, one caveat with our comfort-distance study is that it was conducted with the person approaching the robot, whereas in real-life scenarios the robot would likely be approaching the person. We believe that this may have an impact on comfort-distance results as, for example, the person approaching the robot puts them in control, while they may feel a lack of control if the robot is approaching them.

The dog leash interface is a versatile platform and there remain various future-work questions which we hope to explore using it. From this current study, there still remains

further analysis of the data collected, in particular analysis of participant long-question answers and task-completion time data. However there is only a limited scope on the kinds of answers we can derive from our particular study and further studies would be needed to explore such questions as at what distance should the robot follow (during movement) and how does this distance relate to following position, or how does culture or gender relate to dog-leash interaction. One question of particular interest is the leash versus no leash variable, for example, how does being tethered to the robot relate to a person's feeling of responsibility, affect comfortable following distance, interaction style, or how people perceive the robot? What would have changed if the leash would be based on other, non-physical sensors of distance and direction? Another question is how does interaction change when other people are nearby, for example, does the operator get performance anxiety? Are they more sensitive to robot mistakes due to others being in the vicinity, very much like a person would often feel responsible for a misbehaving dog?

4.5 DOG-LEASH ROBOT: CONCLUSIONS

In this section we presented the idea of leading a robot on a leash as one may lead a dog. We discussed existing related approaches, detailed two new interfaces that we designed and implemented to enable our approach, and presented a formal evaluation of our interfaces.

Designing, implementing, and evaluating our dog-leash interface has helped us reflect on core social HRI questions, and made important contributions to our overarching research questions as outlined in Section 1.3.5, page 8. We demonstrated that robot design can use *the social stock of knowledge*, that is, people's familiarity with leading an animal on a leash, to make a difficult robot control problem accessible (question 2): the general public was able to complete complex robot-direction tasks using the dog-leash interface with very minimal to no training. We have provided insight into how to evaluate social HRI, presenting and tested methods for exploring and measuring a person's disposition toward robots, perceptions of safety, and even robot personality, in part through the application of existing and validated questionnaires (question 3). One reflection on this experience is that, when using standardized questionnaires, we felt the need to tailor and modify them to fit our particular questions of interest. Perhaps this speaks to the limited applicability of standardized approaches to some of the new social HRI challenges. Further, we addressed question 1 by designing and

conducting our study to target and analyze participants' subjective impressions, mental state and comfort, rather than performance ability (i. e., ability to control and complete tasks), thus shedding light on how people interact with robots.

While we argue for the simplicity of this interface, the greater *holistic interaction context* within which a person leads a robot makes this anything but simple. We believe that important remaining questions regarding this interface include, for example, questions relating to the robot always using one follow behaviour, or if there are conditions (e. g., person's mood) that dictate to the robot how it should follow, and how closely. We are also interested in *social integration*, for example, how other people feel about someone leading a robot around, particularly in public spaces where other people (and their children, or perhaps pets) could be injured by such a robot. How do these questions relate to culture, previous experience (e. g., with dogs), or perhaps even gender?


EVALUATING “TOUCH AND TOYS”

In this section we detail our evaluation of a project called *Touch and Toys: new techniques for interaction with a remote group of robots* (Guo, Young, and Sharlin, 2009). While the implementation and focus of the project itself falls outside the scope of this dissertation, and will be detailed merely to provide context and validity, designing and conducting the study for the *Touch and Toys* project served us as an important experience for evaluating social Human-Robot Interaction (HRI). In particular, this study highlighted how an apparently-cut-and-dry task-oriented efficiency experiment had critical elements of human concern and emotion. Even though the study was designed to chiefly target concrete measurements and quantitative comparisons of task completion time, emotive and social elements emerged. This social HRI evaluation experience helped inform us on how we can approach social HRI studies in general and is the main discussion point and goal of this chapter.


First, in order to help with evaluation context, we give a brief introduction to the project itself including the interaction design as well as the implementation. We follow with details on our study of participants using this interface to interact with a small group of remote robots in simple tasks, and summarize our findings as a set of specific design considerations.

5.1 PROJECT INTRODUCTION

This project proposes two new interfaces that enable high-level interaction with a remote group of robots by a single operator. Through the new tabletop-computer interfaces the person can configure and manipulate groups of robots directly by either using their fingers (touch, Figure 5.1a) or by manipulating a set of physical props (Tangible User Interface (TUI)s, Figure 5.1b). These interfaces — the large tabletop, and TUI and touch technologies — were specifically selected to leverage the physical and spatial nature of the robots, and to enable people to direct the robots easily using two hands simultaneously.



(a) the touch interface



(b) the toy (TUI) interface

Figure 5.1: the *touch and toys* project interfaces

The basic design of the interfaces enables the person to specify a target location and orientation for a given robot, with the digital tabletop reporting the actual current robot location. The target location is represented by an interactive icon in the touch case, or a physical toy object in the TUI case, and a line is drawn from the current location to the target to specify the robot’s planned movement trajectory. When the physical robot has reached the target location, the target icon or TUI is highlighted by a green halo. These details are shown in Figure 5.2. The *path-finding* algorithm employed is a simple three-step process: once a target is specified by the user, the robot first rotates itself to face toward the target location, it then attempts to walk straight toward the target with minor direction adjustments, and once it reaches the target location it finally rotates to the target orientation.

For the TUIs interface, plushie dogs were used, black and white, to respectively represent the Sony AIBOs they are to control, and a Frisbee to represent the white iRobot Roomba




Figure 5.2: *touch and toys* interface details: the black Sony AIBO dog is at its target location, shown by the green halo, and the white AIBO is on its way to its location, shown by the graphic and line

(Figure 5.3). Moving and rotating the TUIs is as natural to do as with any similar physical object, and the spatial mapping between the TUI and robots is direct. The plush-and-toy design of the TUIs makes the them familiar, a pleasure to touch and fun to use.

For the touch interface each robot is represented by a single icon. To move the icon, the participant could either translate it by touching the centre circle of the icon and moving it, or by selecting outside the circle and using Rotate 'N Translate (RNT) a technique that enables the simultaneous rotation and translation using only a single touch point (Kruger,



Figure 5.3: our TUIs and corresponding robots, two Sony AIBOs and an iRobot Roomba

Carpendale, Scott, and Tang, 2005; Hinrichs, Carpendale, and Scott, 2006). Figure 5.1a shows a participant simultaneously interacting with two robots.

The digital tabletop is a standard Personal Computer (PC) with four video outputs combined to form a high-resolution (2800 px × 2100 px) display projected onto a 146 cm × 110 cm 2-touch DVIT SMART touch-sensitive board. The TUI interface uses a Vicon object-tracking camera system to track the location and orientation of the TUIs on the tabletop surface. A second Vicon system tracks the robots and reports to the controlling PC, which commands the robots via 802.11 a wireless and Bluetooth.

5.2 EVALUATION DESIGN

The full evaluation design materials, including protocol, and questionnaires, are given in Appendix A.

Our approach to the evaluation of the *touch and toys* project did not involve any specific social HRI methodology or goals. We simply intended to ask participants to perform simple tasks with robots, with the independent variables being the TUI and touch interfaces, and the number of robots controlled simultaneously, and the dependent variable being task completion time. Throughout conducting the study and preliminary analysis, however, it became evident that there were many important social factors at play surrounding the interaction experience, resulting in them being a significant component of the results. Particularly with these later factors, our analysis and presentation approach emerged to become very qualitative-oriented, where we focus more on describing observations and interaction-experience dynamics than on exact measurements of observations.

Throughout the experiment, participants were positioned at the tabletop computer, separated from the real-robot space via a make-shift wall to avoid them observing the actual robots directly and to encourage them to use the interface (Figure 5.4a, Figure 5.4b). We presented the participants with a robot configuration using cut-out robot pictures on a white board and asked them to use the interface to position the robots as directed (Figure 5.4c). This was done in three stages, a one-robot, two-robot, and three-robot stage.


For each stage, the participants were asked to move the robots from a starting position to five configurations (in sequence) using both the touch and the toy interfaces in turn. The configurations were the same across interfaces, but changed with the number of robots. For



(a) the tabletop workspace with tracking cameras



(b) the robot area with tracking cameras, separate from the table area



(c) study administrator presenting a target robot configuration to the participant

Figure 5.4: touch and toys study setup

the one-robot case, the participant did the task for each the AIBO and the Roomba, for the two-robot case we used a single AIBO (white) and a Roomba, and for the three-robot case we used two AIBOs (one black, one white) and a Roomba. The exact robot configurations are given in Appendix A. The order that we presented the *touch and toy* interfaces, and the order that the robots were presented in the one-robot case were counterbalanced across participants, but all were presented with the one, two, and three-robot cases in order. The participants completed questionnaires before the study, after each stage and interface type, post-study, and had a final open interview which was taped (audio-only).

We recruited 23 participants, aged 19–47 ($M=25.5$, $SD=6.5$), 15 male and 8 female, from the university population to participate in our study. Each participant was paid \$10 per hour for their time (most took 1.5 hours and were paid \$15). There were 20 right handed, 1 left handed and 2 ambidextrous participants.

In addition to standard statistical tests on our measured data such as time, our qualitative analysis method focused around exploring our observations, participant long-answer comments, and the taped interviews, finding themes, and presenting them in summaries. We particularly focused on using participant quotes and comments in an attempt to capture and convey the often emotionally-charged and complex essence of what the participants were trying to say. The exact procedure conducted was to read through feedback, cluster quotes and comments into related themes, and then distill each theme into a paragraph to discuss. This clustering was done both in a word-processing application (via copy and paste) as well with paper cutouts of the comments.

5.3 RESULTS

Participants unanimously reported (100%) that the graphical feedback on the table was Hardware Companions?easy to understand and that it was *not* unnecessary, and we found no effect of the sex, age, handedness, or past experience of the participant on their reaction to the system. In the one-robot case, we found no statistically-significant effect of robot type (Roomba or AIBO) on how the participants used or reported on them. Finally, while there were some statistically-significant results related to time efficiency (as explained below), we found no consistent statistically-significant effect of interface type (touch or TUI) on task-completion time, only an effect of number of robots (Table 5.1).

5.3.1 *Task-Completion Time*

In the one-robot case a 2×2 ANalysis of VAriance (ANOVA) (technique × robot, or toy, touch × AIBO, Roomba) revealed no significant technique × robot interaction ($F(1,22)=.15$), which suggests that performance with a given techniques is not substantially influenced by the robot type. There was no main effect observed for technique ($F(1,22)=.54$). However, there was a main effect for robot ($F(1,22)=19.15$, $p<.01$), showing that the task completion time for

		AVG	SD
1-robot	touch	138.3 s	16.2 s
	toy	140.7 s	20.5 s
2-robot	touch	188.2 s	32.3 s
	toy	170.2 s	26.1 s
3-robot	touch	265.0 s	43.9 s
	toy	256.2 s	42.9 s

Table 5.1: average task completion time

the Roomba ($M=131.8$ s, $SD=10.34$ s) was different (11% faster on average) than the AIBO ($M=147.28$ s, $SD=21.43$ s).

In the two-robot case, a paired- t test revealed a significant effect of interface type on completion time ($t(22)=2.61$, $p=.02$). With the TUI interface, the participants completed the task ($M=170.26$ s, $SD=26.19$ s) 10% faster than with the touch interface ($M=188.22$ s, $SD=32.33$ s). In the three-robot case, a paired- t test failed to reveal a significant effect ($t(22)=1.24$, $p=.23$).

5.3.2 Usability

We asked four ease-of-use questions (via questionnaires) after each interface type and across all three robot cases (six times in total). The combined results are shown in Figure 5.5 which shows the percentage of positive responses (>4 on a 7 pt Likert) to each question respectively. On a finer granularity, when toy and touch received a similar amount of positive response, toy received a great deal more *strongly positive* responses than touch. For example, responses to the “precise control over robot movement” question in Figure 5.5 look similar across cases, but the strongly positive responses for toy/touch were 30%/7%, 30%/9%, 22%/9% for the one, two, and three-robot cases respectively.

Participants reported that (in comparison to touch) the toy interface gave more precise control over robot movement, and made it easier to move the robot to the target location and rotate the robot as required. Further, in the two-robot case participants said it was not confusing to monitor the two robots at the same time (70% toy, 61% touch) but easy to control the robots simultaneously (78% toy, 57% touch). With the three robot case, participants also said it was generally not confusing to monitor all three robots at once (70% toy, 52% touch)

and that they found it easy to form the group formations asked (83% toy, 57% touch). Further, Table 5.3, page 128 reports the percentage of participants that responded positively to questions about using both hands and controlling multiple robots simultaneously using the *touch and toy* interfaces. The table shows that they found it much easier to control two and three robots simultaneously with the toy interface than the touch interface.

5.3.3 Preference

For each of the one, two and three robot cases participants were asked how much they preferred each interface (one participant did not answer for the one and three-robot cases). The results, shown in Table 5.2, clearly show that people preferred the toy interface over the touch interface in the two and three robot case. This preference echoed in the written questionnaires and post-test interview as well. One participant explained that the toys gave

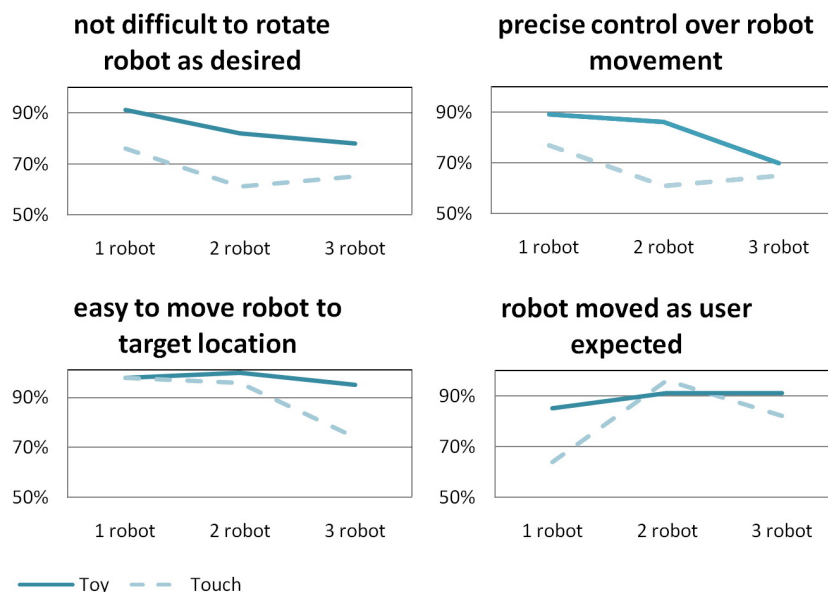


Figure 5.5: ease-of-use-related questionnaire responses

	1 robot	2 robot	3 robot
Toy	45%	83%	77%
Touch	45%	17%	14%
Neither	10%	0%	9%

Table 5.2: percentage of participants who preferred each robot case

them a “sense that [they were] in contact with the robot,” and seven participants wrote that they found it easier to visualize the robot position and orientation with the toy interface. One participant reasoned that the toys provide more visual cues about the orientation and organization than the flat images used in the touch interface.

5.3.4 *Touch*

Participants described the touch interface as being simpler due to less equipment and more precise and accurate due to the high resolution of the screen. Further, the touch was reported to be less intimidating because it was familiar and more similar to traditional PC interfaces. On the other hand, many people complained of the RNT scheme, with eleven people explicitly reporting that it was “unintuitive” to rotate the robot icon around the finger point. This is a property of RNT that participants liked for ballistic movements but which caused problems for precise rotation of the robot once it was at the target location (this matches previous findings regarding RNT, Kruger et al., 2005). RNT rotation moves the centre of the object, requiring a final corrective translation. Instead, participants recommended that it would be “more intuitive” for the robot icon to rotate around the centre, “spinning like a plate.”

Finally, with the three-robot case a few participants complained of visual clutter: three icons for the real robots, three icons for the robot-controlling widget, lines connecting them and the green halos crowd the interface. One participant complained that “for the touch interface, you have six pictures [displayed on the table]. It becomes confusing.”

5.3.5 *Toy*

Participants reported that the toys “were tactile and seemed more realistic” with their three-dimensional nature, with seven participants explicitly noting that with the toy it was “a lot easier to visualize what was happening [remotely]” and to visualize the robot configuration. Further, it helped make it “easier to understand the rotation” and other robot state, enabling them to “focus on collision avoidance.”

The primary complaint (mentioned by several participants) is that the reflective markers for the tracking system get in the way of grasp, where occluding the markers can make the system lose track of the toys and cause erroneous robot movements. They reported that the

marker areas become no-hands zones that distract them from the natural grasp-intuitiveness of the toy.

5.3.6 *Robot Movements*

Participants reported that the robots often moved unexpectedly, despite the contrary evidence shown in Figure 5.5, saying that it was often difficult to visualize the path that the robot would take and that the “robots seemed to take slightly different paths (than the one [they] planned).” The primary reason cited is that some participants expected the robots to copy or replay the movements given by the user, including sidesteps and exact paths, instead of moving directly toward a landmark target as the robots were programmed to do. This was explicitly described by ten of the users, and the problem was more prominent overall in the three-robot case and with the toy cases.

Another aspect of this was that the robots did not move consistently or in a straight line due to physical constraints and noise such as the robot mechanics and a somewhat uneven carpet. Because of this, robots sometimes had to correct their trajectory mid-movement. Participants further pointed out that our interfaces gave them no indication of the robot movement and rotation speed, or time to target location.

The robots have mechanical limitations and challenges with precise movements. As such, they sometimes had difficulties moving to the exact target location specified, and are sometimes off by as much as 10 cm. When this happened it was very obvious and visible to the participant and in the worst cases added considerable visual clutter.

With the toy interface, moving an object from one place to another was reported to be a trivial task by most participants. However, one participant said that “at times [they] forgot [they were] moving a robot and not only toys,” such that they would “pick up the first one and put it [at the target location] and then disregard” the robot, eventually resulting in collisions. However, with the touch interface, the same participant said that “if [the control] is on the screen, [they are] more likely to pay attention to where [the robots] are.”

5.3.7 *Collisions*

By far, the primary complaint overall was that the robots often collided in the multi-robot cases, with 15 participants bringing it up in their written comments as making them change their approach and increasing the effort required. Collisions were not dramatic (i. e., there were no loud noises or damaged robots), and we explicitly told the participants before and during the study that they could not damage the robots. Despite this, however, these participants showed a great deal of concern for the simple fact that the robots were colliding.

Part of this was related to how robots (i. e., two AIBOs) would occasionally push against each other and overlap their legs, taking special effort from the participant to remedy the situation. This annoyed a few participants, and several stated that they expected the robots to be smart enough to avoid each other. As five participants explicitly pointed out, they have to learn each robot's movement characteristics in order to make an efficient path plan and avoid collisions.

5.3.8 *Two-Handed Interaction and Multitasking*

One aspect we looked at is how participants utilize their hands in the experiment and if they use both at the same time. Table 5.3 summarizes our findings, which are echoed in the participant comments, showing how participants found toy easier than touch in general for simultaneous hand use, and for the two-robot case the toys were used to work with both robots simultaneously rather than one at a time as they did with touch. In the three-robot case, however, participants generally worked with one robot at a time for both the toy and touch interfaces.

Participants reported that it was easier to operate robots simultaneously when the movement paths were similar and parallel rather than different and crossing, and more specifically they resorted to sequential movements when they felt that collisions were likely. Conversely, referring to the touch interface one person said: “whenever I use both the hands there are strong chances of [sic] robots getting collide with each other.”

	question regarding robot use	toy	touch
2-robot	easy to control both simul.	78%	57%
	worked with both simul.	70%	43%
	worked with one at a time.	35%	74%
	used both hands simul.	61%	43%
3-robot	easy to control all three simul.	74%	48%
	worked with all three simul.	39%	26%
	worked with one at a time	61%	61%
	used both hands simul.	70%	52%

Table 5.3: percentage of participants that responded positively to questions about using both hands and controlling multiple robots simultaneously

5.3.9 Complexity

We found a correlation between the number of robots and certain properties of the participant responses. First, the conviction behind response (how strongly they agree or disagree) decreased as the number of robots increased. Figure 5.6 shows the breakdown of how strongly participants responded to four core questions asked throughout the experiment across the one, two, and three-robot cases (detailed in Figure 5.5, page 124), independent of the interface used, clearly outlining the trend to *weaken* their stance with the increasing number of robots. Further, the number of complaints (primarily regarding collisions) from the participants in both the written questionnaires and during the experiment greatly increased as the number of robots increased, although this can also perhaps be attributed, for example, to time and increased comfort to say their opinion. The trends of responses shown in Figure 5.6 suggests a general weakening of ease of use and control over the robot with the increased number.

5.3.10 Real Robots

In the post-test questionnaire participants were asked if the experiment should have been done with a simulation instead of real robots. Fifteen of the twenty-three participants felt that having real robots added value to the experiment. Reasons range from simple “the real thing is better” and “it is cool with real robots, more interesting than a simulation” to “real robots experience real problems. sims do not,” “I trust the results more with real robots,” “there

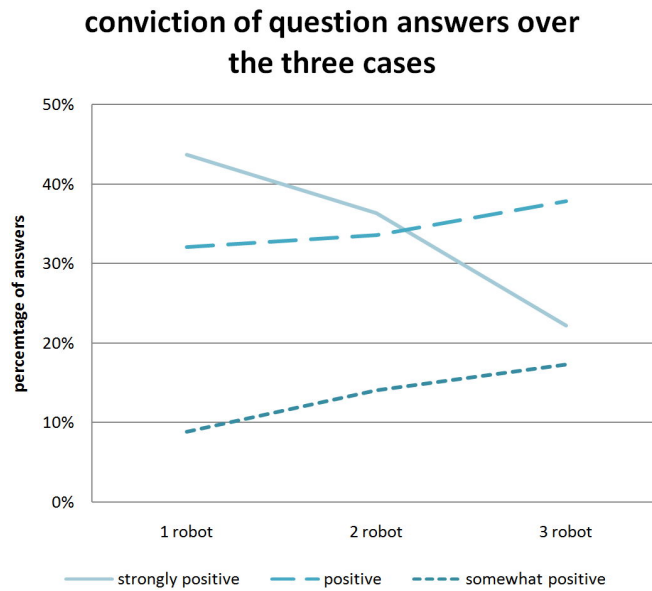


Figure 5.6: conviction of ease-of-use-related responses

was a real sense in knowing that real robots were colliding and that gave the situation more importance,” and “real robots and the monitoring provided me with a better understanding of speeds and limitations that a simulation would have a hard time to capture.”

5.4 DISCUSSION

COLLISIONS AND COGNITIVE LOAD — Participant concern over robot-robot collisions emerged as a much-more dominant issue than expected. That this was prominent despite the participants being informed and reassured that they could do no damage raises questions regarding feelings of worry over responsibility, concern for damaging robots, or perhaps some more-fundamental aversion to having the robots touch or collide. We believe that these factors emerge directly from people’s *social stock of knowledge*.

Following from the importance that participants gave this problem, in addition to our direct observations and analysis of written feedback and interviews, we are confident in drawing a direct link between increased collisions (e. g., as a product of more robots) and the observed drop in rating of ease-of-use and moving from two-handed back to one-handed interaction. Perhaps this is related to higher demands on the person as the number of robots increase and the task becomes more complex. As such, these problems — including the collision concern — may possibly be attributed to increased cognitive demand on the

participant due to more robots to deal with. This agrees with Drury et al.’s HRI awareness taxonomy (Drury et al., 2003) and supports their claims regarding how human-robot ratios affect interaction, and how awareness and control problems will grow with the number of robots. What we found particularly surprising is how discernible this effect was in our experiment, where only three robots are used with simple control mechanisms.

One participant’s observation that perhaps with the TUI interface they forget they are dealing with real robots suggests that while hiding low-level interface details can reduce cognitive load, it can at the same time hinder their HRI awareness.

INTERACTION EXPERIENCE AND EMOTION — Participants strongly favoured the toy interface in most respects. Our results link this success to core TUI concepts, as the people explicitly and repeatedly commented on the “intuitive” usability, the awareness gains, and the enjoyment they received from the interface. We expect that much of this was the familiarity with the attractive plush animals and Frisbee, as well as the task of arranging physical objects on a tabletop surface. That is, the direct mapping between the TUIs and robots, and the tabletop, increases comfort and enjoyment, and lowers cognitive load by exploiting people’s existing knowledge and natural understanding of the physical world.

Despite the favour, both interfaces were generally equally *efficient* in terms of the task-completion time. We believe this points to a deeper, but perhaps simple, dimension to our results. The participants found the TUIs “fun” and “felt” connected to the robots when using them, which had a direct effect on how they felt about the usability of the interface (helped them feel that they performed better, as in Norman 2004). There were also indicators that this led to a lower cognitive load. This is similar to how participants defended the use of real robots due to the *cool* and novelty factor. These findings directly correspond to recent arguments for incorporating emotion into design, and HRI specifically (e. g., Norman, 2004).

TWO HANDS OR ONE — The question of exactly when two-handed interaction is more effective is beyond the scope of our work, but in our experiments participants resorted to one-handed interaction as things got complex, confusing, or difficult. This can be seen as another indicator of mental load, and a benefit of simpler interfaces; they may promote multi-hand interaction and the versatility that comes with it.

INTERFACE DESIGN — Participant feedback directly outlined that both interfaces should be improved to afford robot limitations and our particular movement properties, for example, that they move in a straight line and do not replay input. We need to consider other interface styles, such as enabling the person to specify either a path or a target. Further, our interface could improve problems of visual clutter (e. g., when the robot did not line up perfectly with the input) which can impair a person’s ability to concentrate on their task. The disparity between the results for the *touch and toy* interfaces on many points, and the fact that it solidified with more robots, is a strong indicator that our TUI interface was better suited to the task than our touch interface. While our findings frame a TUI versus touch set of results, our results must be considered carefully. For example, our selection of the RNT technique (touch-only) had an overall effect on how *touch* was perceived. Further experimentation will be necessary before drawing strong TUI versus touch-type conclusions.

5.5 IMPLICATIONS

Here we distill our findings into a set of initial lessons and implications relevant for designing tabletop, touch, and TUIs for interaction with a remote group of robots.

people react to TUIs — TUIs have a strong impact on interaction experience, regardless of particular efficiency gains, that can change how an interface is approached, perceived, used, and evaluated.

robots change the experience — using actual robots (and letting the person know) instead of virtual ones has an impact on interaction experience.

complexity is related to how hands are used — people may utilize both hands when interacting with a group of robots through tabletop, touch and TUIs. However, they may resort to single-hand interaction when they are faced with increasing cognitive load.

indicate properties of movement — people should not be expected to extrapolate the robot path, speed, and task just from the robot motions, but instead the interface should clearly indicate these properties to aid them in planning and interaction and to improve their HRI awareness.

improved path-planning flexibility — enabling people to specify complex, multi-part paths and commands relating to macro-scale robotic actions may reduce their involvement and help them cope with more robots in complex interaction scenarios.

flexible granularity of control — users need to resort to lower-level control when the autonomy of the robot cannot solve a problem, such as a navigation complications or collisions. Good design should support detailed interaction as a backup option.

5.6 EVALUATING TOUCH AND TOYS: CONCLUSIONS

This evaluation effort points to the importance of considering HRI beyond questions of task-oriented efficiency, and highlights the wider picture of the person’s interaction experience. Despite the usability-oriented design, and lack of significant task-completion-time results, this study exposed various important findings related to a person’s comfort, familiarity, concern, cognitive load, and how they relate to robots and the given interfaces. Using our theoretical framework outlined in Section 3.6, these social HRI experience components fall under our *visceral* perspective (P1) on interaction (Section 3.6.2, page 82).

This study provides evidence that the fact of having real robots — instead of animated simulations — changed how people approached, used, and thought about the interface. While this emerged from comments only and was not methodologically controlled for, this points to importance of further exploring the question of why this phenomena happened and how we can control for it. Perhaps this points to inherent social layers to the task, for example, that perhaps the robots are seen as expensive or fragile. Further, although the two studies are quite different, the collision and concern aspect did not emerge in the least in our *animation table* animated *puppet master* study (Section 7.5), while a similar (human–robot) collision concern emerged in our *broomstick* robotic *puppet master* study (Section 7.6).

The *touch and toys* project addresses our research questions, presented in Section 1.3.5, page 8, in that it contributes to the understanding of how the particular evaluation methods used can address social HRI components (question 3). For example, this study points to the importance of leaving room for participants to reflect on their experience, and importance of being open to enabling alternate layers of the evaluation emerge. In this particular study, the social HRI-related findings surfaced through the open-ended questions and comment spaces on the questionnaires, as well as unstructured interviews, and through correlation

with task-oriented findings. Further, while it is not clear if these robot-oriented differences are related to robots being perceived as social actors, the results of this study help illustrate how the use of robots impact interaction (question 1).

Overall, this study illustrated how social human-oriented aspects emerged to dominate the study even though the original focus was on task efficiency. The experience of encountering this, analyzing the data and summarizing the themes helped us with important insight and informed on how to approach general social HRI evaluation questions. In particular, this study exposed how participant open-ended self-reflection questions can be a powerful method for social HRI, as in this case it gave the participants the opportunity to express their thoughts beyond the simple technical nature of the question that was asked.

CARTOON ARTWORK ROBOTIC INTERFACES

We believe that the simplified artistic and visual language found in comic books and animated cinema (e. g., as shown in Figure 6.1) can provide powerful expression mechanisms for robots that people can easily understand. Throughout this introduction we discuss our cartoon-artwork-for-robots approach reflecting on our social Human-Robot Interaction (HRI) theoretical framework (Section 3.6), particularly using our three social HRI perspectives: P₁ (visceral), P₂ (social mechanics), and P₃ (social structures).

A cartoon, traditionally, is a preparatory work or rough sketch (Merriam-Webster Online, 2010, “cartoon”). In modern times, the word cartoon has come to represent the stylistic art often found in comic strips and animated cinema. McCloud (1994) explores the language of comics and cartoon art and illustrates the power of comic artwork, and that it is rooted in basic human perception; such techniques found in cartoon art as exaggeration and simplification have been used throughout human history as far back as petroglyphs and cave writings.

Cartoon-art elements are encountered in everyday life, for example, in posters, magazines, advertising, even as traffic stop and warning signs. Much of this communication is visceral: that is, much of it speaks to basic human emotion (McCloud, 1994), transcends applications and often even culture (P₁ visceral interaction). People also regularly use cartoon-artwork, for example, in sketching or the use of emoticons with cell phones or on-line chat (P₂ social



Figure 6.1: example collage from the Calvin and Hobbes comic that highlight the versatile and powerful communication language of comics

mechanics interaction). Thus, a person's understanding of cartoon artwork is part of *the social stock of knowledge*.

The use of cartoon artwork for robots emerges from our idea of *robot expressionism* (Section 3.5), where we suggested using familiar representations (from *the social stock of knowledge*) to represent a robot's technical details in abstract but easy-to-understand ways. We believe that cartoon artwork enables the robot to communicate clearly across many language and cultural barriers, largely removing the requirement for a person to learn (*accessibility and usability* factor of acceptance) and further informing people how to interact with the robot and what to expect from it (related to *social integration*): cartoon-art can help foster a person's natural understanding of the robot's internal state, tasks, goals, and algorithmic intentions. For example, a robot which is faced with a situation it does not understand may use a stylized and simplified facial expression and place a question mark above its head to portray itself as a confused creature, it may use simplified cartoon-like facial expressions to convey happiness for completing a mission or fear for not completing a task on time, or when it has a low battery and needs to recharge it may display sweat drops on its forehead to express fatigue. Current approaches often use explicit representations of robot details such as a battery meter or *low-battery* light (e. g., as with the iRobot Roomba); we argue that these representations may not be immediately clear to people, and that people need to take the time to read, interpret, or understand what the indicators practically mean for them in the given task.

The versatility of cartoon artwork also gives it advantages over, for example, spoken natural language which has a stronger reliance on specific language, culture and temporal dependency than cartoon artwork: if a person misses parts of the spoken message it would have to be repeated, very much like in a conversation. Cartoon artwork may also help robots avoid eeriness problems (see Section 2.3.1, page 35). That is, we expect that people will apply their understanding of comics from *previous experience* and *media*, perhaps seeing the robot as a simplistic living-like (but not alive) entity, removing the imperfect link to real life which much research posits (including the *uncanny valley*) is behind the eeriness problem. Further, perhaps *previous experience* could further affect people's perceptions of *safety* (lack of danger) or even *fun and enjoyment*. Finally, cartoon artwork could be used by people as well, not only the robots (e. g., as in Ng and Sharlin, 2010).

In this chapter we explore the use of cartoon artwork for social HRI. We start with a theoretical exploration of how robots can use cartoon artwork, discussing both how the robot can integrate elements into interaction, and which cartoon artwork techniques we believe are particularly useful for social HRI. We further develop various designs for robotic interfaces that use cartoon artwork, and detail our two particular implementations: *bubblegrams* and *Jeeves*. We finish this chapter by presenting our informal design critiques on our interface implementations. Overall, this chapter highlights our exploration into how cartoon artwork can be used by social HRI, in particular, to leverage *the social stock of knowledge* to make complex robotic state information easy for people to understand.

6.1 APPLYING CARTOON ARTWORK TO ROBOTS

In this section we analyze the question of how robots can use cartoon artwork for interaction with people. We introduce existing social HRI-related work, outlining the novelty of our own approach. We explore *how* robots may use techniques from cartoon artwork for interaction with people, for example, how long cartoon elements may last for and which entity has control. Second, we explore *which* techniques from cartoon artwork robots can directly use in their interactions, for example facial expressions. We first, however, introduce the particular interface technology we use, as this has implications for our entire discussion.

6.1.1 Interface Technology — Mixed Reality

The use of cartoon artwork is an *approach* to interaction, and as such, could be implemented in any number of ways. For example, a robot could have customized on-board lights to show pre-designed cartoon artwork, large display panels for versatile (but localized) expression, or perhaps an on-board projector to project around the body. In our work we used a technique called Mixed Reality (MR) due to the versatility offered as a prototyping tool.

Mixed Reality (MR) is a concept which addresses how the physical and digital worlds can be combined into a single, integrated interaction space (Milgram and Kishino, 1994), which enables a person to seamlessly and simultaneously interact with both the physical and digital worlds (Young, Sharlin, and Igarashi, 2010). In practise, this is commonly accomplished through the use of visual, graphical overlays to augment objects in the real world such that


interactions with the physical object are reflected in the visual augmentation, or vice versa — an example is shown in Figure 6.2a where the computer model of the surgery area is projected onto the head of the patient, updated in real time as the patient moves or as cutting happens (Grimson, Ettinger, Kapur, Leventon, Wells, and Kikinis, 1998).

MR is often implemented using projectors that augment physical objects with graphics or using a window-to-the-world metaphor implemented using see-through displays which, when held up to the world, superimposes computer graphics on objects in its view. For example, Billinghamurst et al. (2001) present the Magic Book (Figure 6.2b), a physical real-world book which has animated computer graphics pop out of the page as a person peruses the book, where people can seamlessly move between the physical (book) and virtual (graphics); this is technically accomplished by having the person wear a head-mounted display. MR can also include other modalities, such as aural or haptic input and feedback.


We use MR as a base for all of our interface designs presented below. First, we discuss existing cartoon-art work for social HRI.

6.1.2 Existing Use of Cartoon Artwork for HRI

Cartoon artwork has been used extensively for non-robotic computerized entities. One clear example is the wide use of *emoticons*, the simplified comic-like facial expressions used in



(a) Grimson et al.'s (1998) MR surgery aide that displays a brain map on the patient's head.



(b) Billinghamurst et al.'s (2001) Magic Book, where virtual characters pop out of the physical book

Figure 6.2: example MR interfaces

email, chat (e. g., Kurlander, Skelly, and Salesin, 1996), and cell-phone texting applications, although these are static and do not represent interactive characters. Examples of interactive characters include avatars, such as interactive video-game characters which are drawn using cartoon artwork, and use cartoon techniques such as pop-out thought bubbles. In our work we considered how this approach transfers to robots (Young, Sharlin, and Boyd, 2006; Young and Sharlin, 2006; Young et al., 2007).


Dragone, Duffy, and O’Hare (2005a) (also Dragone, Holtz, Duffy, and O’Hare, 2005b) use MR technology to display cartoon-based avatars sitting atop robots. This work explored ideas of providing further agency to robots, but stop short of addressing the questions of the use of cartoon artwork. Further, this research meant for the cartoon character to be the sole point of interaction, where the person is expected to completely ignore the physical robot — the robot itself is intended only to provide a means for mobility within the real environment. In our work, we consider how techniques from cartoon and comic artwork can be used to augment and compliment the robot’s existing interactive presence.

We see some robots as being inspired by and in essence using principles of cartoon artwork, although this may initially not appear to be the case. For example, it is very common for anthropomorphic robots to employ cartoon-like simplified and exaggerated facial expressions, such as with Keepon (Michalowski et al., 2007, Figure 6.3a) and Kismet (Breazeal, 2002, Figure 6.3b) robots. MIT’s Leonardo robot (Breazeal et al., 2006) uses whole-body exaggerated human-like gestures to show such expressions as happiness, confusion, or surprise.

We believe that our discussion, techniques, and implementations as presented in this chapter are the first direct attempts at leveraging cartoon-artwork for interaction with robots. In the remainder of this section we explore *how* robots can use cartoon artwork for interaction



(a) several cartoon-like Keepon robots (Michalowski et al., 2007)



(b) Kismet with cartoon-like facial expressions (Breazeal, 2002)

Figure 6.3: examples of robot designs which we believe are using principles of cartoon artwork

and introduce particular elements from cartoon artwork which we believe are particularly useful. Following, in the latter half of this chapter we introduce our interface implementations and informal design critiques.

6.1.3 How *Robots can use Cartoon Artwork*

Here we present our theoretical exploration into how robots can use cartoon artwork for communication, including hypothetical interface designs (some implementations of these follow later in this chapter). Both the exploration and implementation of work was directed by the concept of MR: first we present what we call the Mixed-Reality Integrated Environment (MRIE) (Young and Sharlin, 2006), followed by a theoretical framework for designing interfaces (Young et al., 2006; Young and Sharlin, 2006; Young et al., 2010), and then detail several interface concept designs and implementations Young et al. (2006, 2007).

6.1.3.1 *Mixed Reality Integrated Environment*

When considering *how* a robot can use ideas from cartoon artwork, one is quickly steered by the realities of the technology. Size and mobility constraints of common robots such as the Sony AIBO robotic dogs (Figure 3.4, page 73) often result in limited interfaces with few buttons and a limited display, and the context of use means that people may not be fixedly seated in front of these interfaces.

Many techniques from cartoon art, such as thought bubbles, can enable a robot to break free from the limitations of its physical body and gesture capabilities. This pushed us to consider the interaction possibilities if robots were not limited by their immediate graphical displays and could use the entire environment, including the air and space around them, freely adding colour, animation, and cartoon-like annotations to any location on or around their bodies, or in the surrounding environment. This suggests an integrated interaction environment, where digital (cartoon-like) and physical components, and both robots and people, could seamlessly intertwine. Given that MR is a technology which can enable these sorts of interactions, and for us served as a concrete tool for our exploration, we call this the Mixed-Reality Integrated Environment (MRIE) (pronounced *merry*).

The MRIE is based on the assumption that, provided that technical and practical implementation challenges are addressed, virtual information can be integrated directly within

the entire three-dimensional, multi-modal real world. One could imagine a parallel virtual world superimposed on the real world, where digital content, information, graphics, sounds, and so forth, can be integrated at any place and at any time, in any fashion. We present this MRIE as a conceptual tool for exploring how robots and people can interact using MR, where both people and robots can create, modify, destroy and interact with MRIE elements.

6.1.3.2 *Our MRIE Taxonomy*

Here we present our taxonomy for interaction within the MRIE. Our taxonomy maps MRIE interaction possibilities using four key variables: *lifespan*, *ownership*, *activity*, and *virtuality*. The development of this taxonomy emerged from our own explorations into how the MRIE could be used to interact with a robot.

LIFESPAN – the *lifespan* variable determines how long instances of a MRIE interaction technique last. For example, a robot may place a permanent element into the environment (separate from itself) for information purposes, resulting in an arbitrarily long or permanent *lifespan*. On the other hand, a robot may display a surprise mark which is designed to disappear immediately, resulting in a very short *lifespan*.

OWNERSHIP – the *ownership* variable determines which robot or person, if any, owns a technique instance; *ownership* reflects who has the control of the element and how control may be distributed among people and robots. This could be used to avoid others modifying elements, to add a level of trust, to know who left the message, and so forth. Elements of course may also be completely public, with no declared ownership.

ACTIVITY – the *activity* variable determines how active an interaction technique instance is. This is, in general, a sense of how dynamic (visually, aurally, etc.) an element is, through attracting attention or being interactive. An example of a technique with very low *activity* is an element which displays a static decoration on a wall; this technique does not actively invite attention, and does not react to interaction attempts. A variation on this technique which uses animation or other methods to gain attention would have a higher *activity* level. An example of a technique with high *activity* can be a MR interactive menu system which incorporates three dimensional animation and sounds for interaction purposes. For example, upon creation, this menu could make a *popping* noise to notify the user of its creation, and could react richly to a person's interaction.

VIRTUALITY – the *virtuality* variable is based on Milgram and Kishino’s (1994) *virtuality continuum*; it categorizes the representation technique as somewhere between purely physical and purely virtual. A purely physical technique could be, for example, physically touching a robot and getting a physical action response, while a purely virtual technique could be the use of virtual reality to interact with a robot’s parameters. Most techniques lay somewhere in between. Note that *virtuality* in our taxonomy includes all forms and modalities of virtual information, including graphics and sound.

We have not developed the idea of the MRIE or the MRIE taxonomy beyond what is discussed here (Young and Sharlin, 2006; Young et al., 2010). We see further exploration of the MRIE, the taxonomy, and how the taxonomy can be mapped to the MRIE and interaction instances, to be an important direction for future work. In the following sections we present our existing explorations into interaction possibilities, where we directly use the taxonomy to define three instances: *thoughtcrumbs*, *decoration*, and *bubblegrams*.

6.1.3.3 *Thoughtcrumbs*

Inspired by breadcrumbs from the Brothers Grimm’s Hansel and Gretel, *thoughtcrumbs* are bits of digital information that are attached to a physical, real-world location (Figure 6.4). A robot can use these to represent thoughts or observations, or a person could also leave these for a robot to use. These can also perhaps be interactive, offering dynamic digital information, or enabling a person or robot to modify the *thoughtcrumb*. For example, search and rescue robots may use *thoughtcrumbs* to leave information such as air quality, temperature levels, and potential risks at particular locations for the human teams that are following them.



Figure 6.4: A concept sketch of our *thoughtcrumbs* interaction technique. A robot can leave a *thoughtcrumb* behind (left pane), to be later used by a person (right pane).

Thoughtcrumbs can have any length of *lifespan*, depending on how long the information is deemed relevant, possibly automatically expiring after a time or remaining part of the environment until someone explicitly erases it. A short-lived *thoughtcrumb* may be a note left by a cleaning robot to say that the floor is wet; this *thoughtcrumb* would expire once the floor is dry. A long term *thoughtcrumb* could be a set of directional arrows left by a robot to direct a flow of traffic, which would be left until explicitly destroyed, possibly weeks later. We can also interesting to consider how *lifespan*, or *age*, can be conveyed in visual representations, for example, such as an old *thoughtcrumb* being wrinkled, rusty and faded.

In terms of *ownership*, *thoughtcrumbs*, once placed, may be public elements within the shared environment such that others can remove or modify them. This fits many of the examples already presented, as a cleaning robot may destroy *thoughtcrumbs* which a person placed asking it to clean, or a human may remove *thoughtcrumb* notes left behind by a cleaning robot. *Thoughtcrumbs* may range in activity, from low activity for more ambient information and communication, to high activity when the information becomes more critical or important. The *virtuality* of the *thoughtcrumb* can be digital (as in Figure 6.4) or physical, or simultaneously both. For example, a *thoughtcrumb* may use a physical token to show its location, but virtual graphics floating beside the token (when, e. g., using a MR device) to convey the information.

A Radio Frequency IDentification (RFID)-based robotic *thoughtcrumb* interface implementation, an effort which is out of the scope of this dissertation and was led by Nicolai Marquardt, illustrates how the *thoughtcrumb* idea can be used to both motivate and explain a system and implementation (Marquardt et al., 2009).

6.1.3.4 *Bubblegrams*

Based on comic-style thought and speech bubbles, bubblegrams are designed to represent a robot's internal state and expressions. These elements are overlaid onto a physical interaction scene, floating in proximity to the robot that generated it (see Figure 6.5a), enabling people to interact with the robot simultaneously in the digital and physical realms. Bubblegrams can be used by the robot to show information to a person, and can perhaps be interactive, resembling an interactive physical display directly within the task space, allowing a person to interact with elements within the bubble (Figure 6.5b).

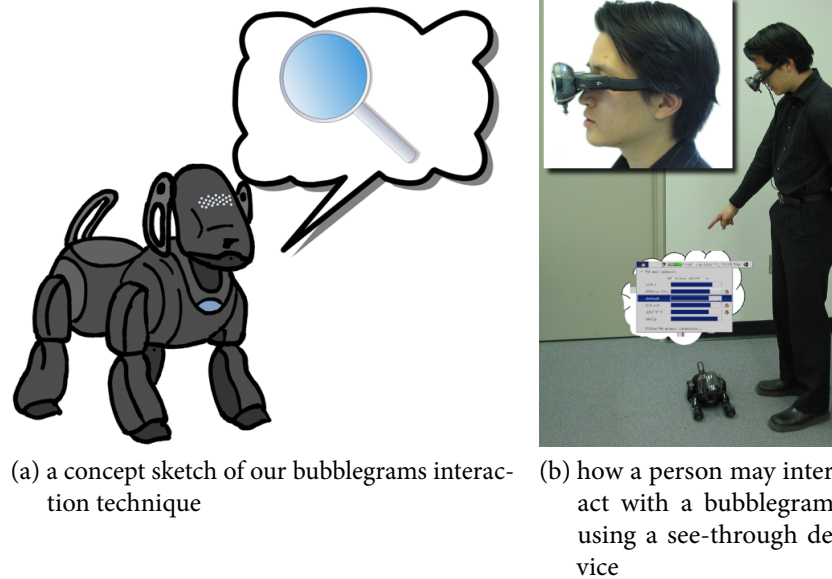


Figure 6.5: concept sketches of the bubblegrams interaction technique

Bubblegrams are designed for specific short-term interaction (*lifespan*), and are generally not used for long-term tasks; they are designed to convey current information or for immediate interaction. For example, a surprise bubblegram floating over a robot's head may last for five seconds, and a system-menu bubblegram will be destroyed as soon as the interaction is complete. Following the comic-style bubble motivation, and given the fact that bubblegrams are used to represent a particular robot's communication, *ownership* is directly attributed to the entity that created it. Bubblegrams can range from low to high *activity*, ranging from a static graphic with no interactivity to a full-fledged animated and interactive menu, and has medium *virtuality*, since they bring complex digital data directly into the physical interaction space. We have implemented a version of *bubblegrams* detailed later in this chapter.

6.1.3.5 Decorations

Another interactive example within the MRIE is *decorations*: a robot can use the entire physical environment as an area which it can *decorate* with its particular interests, synthetic emotions, fears, curiosities, and so forth (Figure 6.6). These decorations are intended to be personal to the robot, and serve as a non-technical way to provide insight into the robot's workings, algorithms, and reasons for why it may do things in particular ways. For example, a robot may place its favourite snapshots on a wall and decorate a room based on some observations that it found interesting. A person could then view this space, getting insight into the state




Figure 6.6: A concept sketch of our decorations interaction technique. A robot can freely decorate an environment with items of interest.

and personality of the robot, for example, that it is afraid of stairs or enjoys (perhaps can see better in) the sunlight.

We see decorations as generally persistent elements with a long, unspecified *lifespan* which means they exist until explicitly destroyed by the owner — by nature they are intended to be fairly static. Perhaps rapidly changing decorations could be a sign of instability in the robot itself. Likewise, given that they are personal to the robot, decorations are *owned* solely by the robot that created them. Decorations may be generally designed to have low *activity* given that they have no direct, immediate interaction purpose. Decorations have medium *virtuality*, decorating the real environment with virtual decorations.

In this section we have so far introduced methods for *how* robots can use techniques from cartoon artwork as a part of an interface, using the conceptual MRIE as a base. We provided a taxonomy on various points of interest related to interaction in the MRIE, and used it to present, define, and explore three interface designs which highlight how the MRIE can be used for interaction. In the remainder of this section we focus instead on proposing *which* techniques from cartoon artwork we believe can be useful for social HRI.

6.1.4 Adapting Techniques from Cartoon Artwork

Here we discuss three techniques from cartoon artwork which we believe are particularly useful for social HRI and the kinds of interactions we envision between people and robots: *icons*, varying *text styles*, and simplified and exaggerated *facial expressions* and gestures.

6.1.4.1 *Cartoon-Art Icons*

We use *icons* to refer to annotations such as movement lines, dust marks, and dizzy stars or heart symbols. Our use of the term *icon* is taken from McCloud (1994)'s *Understanding Comics*, and includes all signs, symbols, icons, and indexes as commonly defined in semiotics (Chandler, 2002). We propose that *icons* will provide robots with a method to show movement, emphasis, or emotion such as happiness or surprise, all using a very simple technique. The power and simplicity of this technique is illustrated in Figure 6.7.

6.1.4.2 *Cartoon-Art Text Styles*

Text styles are used in comics and cartoons for emphasis and variance in emotion or intent, both in dialogue and as added situational information. *Text styles* may be varied from typeset to hand printed or handwritten text, and can have varying colours, fonts, weight, shapes, distortions, and decorations. A robot can use these variations to add subtle or obvious overtones and meaning to text and letters, in similar ways as shown in Figure 6.8. In addition, *text* is often combined with thought or speech bubbles (*bubblegrams*) which can also be stylized in their own way.

6.1.4.3 *Cartoon-Art Facial Expressions*

Facial expressions and gestures in cartoon and comic art, often simplified and exaggerated, are used as a way of conveying an emotional state. These expressions usually focus on the expressive parts of the face, using such features as the eyes, wrinkle lines, eyebrows, and mouth, and are often used to make otherwise-inanimate objects anthropomorphic. Further, exaggeration enables much more emphasis than in real life, such as over-sized eyes

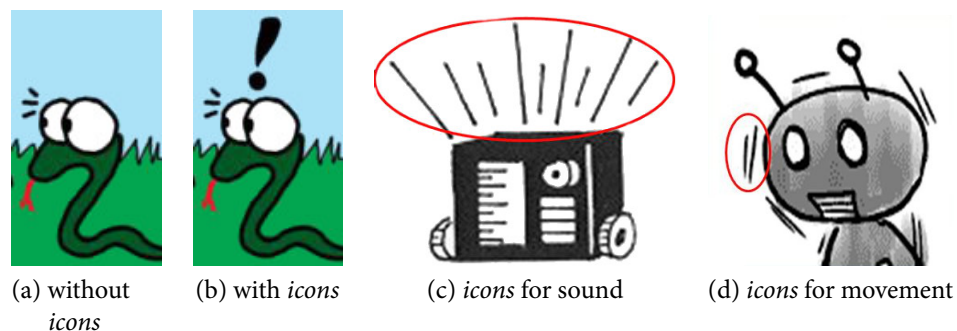


Figure 6.7: cartoon-artwork *icons*



Figure 6.8: examples of cartoon-artwork *text styles*

or mouth, for targeted emotional state communication (see Figure 6.9). Robots can use *facial-expressions* elements to representing technical information in an easy-to-understand way: for example, an iRobot Roomba robotic vacuum cleaner could show a depressed face to represent that its brushes are dirty and it is not cleaning as well as it should be.

6.1.4.4 Introducing Jeeves: Cartoon-Art Scenario Test-Bed

We present *Jeeves*, a scenario test-bed for exploring the use of cartoon artwork for robots. The premise of *Jeeves* is as a domestic robot (an iRobot Roomba in our implementation, detailed below) with various tasks around the household. We use *Jeeves* develop proof-of-concept cartoon-art interaction scenarios that utilize and highlight the versatility of cartoon art.

Jeeves can simultaneously use a combination of *icons*, *text*, and *facial expressions*. For example, a face may have heart-shaped or swirly eyes, or *text* may be decorated with various symbols. In addition, we add to our MRIE taxonomy and identify three primary ways in which we envision a robot could place visual elements. Cartoon elements can augment the robot directly, for example, to add a face or *icon* tattoos. Elements can augment the immediate area around the robot, for example, to place *icons* such as motion lines, stars showing dizziness,



Figure 6.9: cartoon-artwork *facial expressions* and gestures

or expressive words such as “VROOM” to give the impression of speed. Or, elements can augment other physical objects and spaces within the environment, for example, by using sound *icons* to exaggerate the noise from a sound system it deems to be too loud.

Below we present three example *Jeeves* interaction scenarios: *trash Jeeves*, *the recycle police*, and *clean-tracks*. We selected these three as *Jeeves* is designed around a domestic design scenario (and robot). Throughout our explanations we relate back to the MRIE taxonomy (Section 6.1.3.1); all cases have medium *virtuality* as they integrate both physical and virtual elements into one entity (as explained below).

Trash Jeeves (Figure 6.10) is a simple scenario which involves our robot butler, *Jeeves*, being obstructed by a garbage can while cleaning the floor. Noticing a near-by person, *Jeeves* makes efforts using its physical presence and (high *activity*) cartoon art (robot *ownership*) to get a person to assist it: it physically bumps the garbage can and tries to push it while expressing fatigue and annoyance using cartoon art (short *lifespan*, changing as the robot’s state changes). The cartoon art annotations both augment the robot directly, as well as its vicinity, and utilizes simplified *facial expressions* and cartoon-art *icons*. We have this scenario fully implemented and functional, as presented later in this chapter.

The recycle police (Figure 6.11) is an environmentally-friendly robot which roams a room looking for recyclables. The robot tags found items (with low *activity* tags) and continues searching, with the hopes that a person will notice the tag and recycle the item (public *ownership* of cartoon elements, *lifespan* until item is recycled). This scenario augments the robot officer, the direct vicinity and leaves MR *thoughtcrumbs* in the environment (Section 6.1.3.3, page 142), and uses cartoon-art *facial expressions*, *icons*, and *text* to communicate with people. Screen shots of actual content are shown in Figure 6.11, although this scenario is only partially implemented — tagging is static and the robot does not actually search for recycled goods.

Clean-tracks (Figure 6.12) *Jeeves* leaves *thoughtcrumb* tire tracks (low *activity*, robot *ownership*) behind on the floor to show where it has cleaned, tracks which persist until the robot is done cleaning the room (task-length *lifespan*). This can provide a person with a sense of progress of the robot’s work and which areas have been cleaned. *Clean-tracks* uses cartoon-art *icons* to realize the *thoughtcrumbs* concept. The current implementation is a mock-up only: while the tracks are drawn live on the person’s display, they are static and do not change and track as the robot moves. The static live image is shown in Figure 6.12.




Figure 6.10: Key-frames of the Trash *Jeeves* interface implementation (live images). Hard at work, *Jeeves* needs assistance to move a trash can.




Figure 6.11: Key-frames of the *recycle police Jeeves* interface implementation (live images)



Figure 6.12: a live image of the *Clean-Tracks Jeeves* mock-up, showing *Jeeves* leaving tracks behind in the environment as it cleans the floor


In this section we have discussed *which* kinds of elements robots can use from cartoon artwork, and presented a taxonomy on *how* they can use them. Overall, this section was a detailed exploration of how cartoon artwork can be used for social HRI interface design. In the next two sections, we present details of our *bubblegrams* and *Jeeves* implementations.

6.2 IMPLEMENTING BUBBLEGRAMS

We built a working *bubblegrams* interface using MR to float interactive thought-bubble interface above a Sony AIBO robotic dog as shown in Figure 6.13. To see the *bubblegrams*, the person wears video see-through goggles (Figure 6.13b) — they can see the real world *through* the goggles in real time, but the goggles can draw additional computer graphics (e. g., *bubblegrams*) on top of the view (Figure 6.13c) (Giesler et al., 2004 gives a comparison of video see-through versus optical see-through, motivating video see-through). The person holds a tablet PC which also displays the *bubblegrams* and enables them to use a stylus to interact with the *bubblegram* (Figure 6.13d). Our proof-of-concept implementation enables the person to use this to navigate through a menu system (Young et al., 2006).

Our robot is a standard Sony AIBO ERS-7, running a custom *random-walk* behaviour developed using the open-source C++ Tekkotsu robot-programming toolkit. We selected the AIBO as it was the only robot available to us at that time. We acknowledge that the form of the robot (a puppy) can add distractions to the cartoon-artwork bubblegram communication channel. However, we took steps to reduce the anthropomorphic nature of the robot, avoiding the use of the robot's lights, motorized tail and ears, keeping the neck fixed, and using a very rigid and constant walk algorithm; this further helped to simplify the tracking problem, outlined below. We developed a 802.11g wireless link between the tablet PC and the AIBO, but this was only used as an initialization mechanism in this particular application.

We use a portable tablet PC carried in-hand to drive the system and enable the person to interact with the *bubblegrams*. This is a Toshiba Portege M200, 1.5 GHz Intel Pentium M, 1 GB RAM. Our Head-Mounted Display (HMD) is the light 100 g Icuiti DV920 HMD (640 px×480 px resolution); to obtain video see-through capability we attached a modified Creative Web-cam to the front centre of the goggles (Figure 6.14). The video feed from the camera is displayed on the HMD directly to give a see-through glasses illusion, with the MR graphics drawn on the video before it is displayed. The tablet PC software was implemented in C++, using Intel's Open Computer Vision (Open CV) toolkit to interface with the video data from the camera. To detect the AIBO in the live video feed, thus enabling *bubblegrams* to work, we developed an original computer-vision algorithm, below (Young et al., 2006).




(a) a person interacting with a bubblegram using our MR implementation



(b) goggles worn to see *bubblegrams*



(c) *bubblegrams* seen through the goggles



(d) interaction with the *bubblegram* through a tablet PC

Figure 6.13: our bubblegrams see-through-device MR implementation Young et al. (2006)



(a) our HMD with Toshiba Portege M200 tablet PC



(b) a camera attached to an HMD for see-through capability

Figure 6.14: *bubblegrams*' hardware implementation

6.2.1 Haar-Like Features for Robot Tracking

A primary challenge with our MR implementation of *bubblegrams* is deciding where, in the person's vision through the head-mounted display, the bubblegram annotation should be drawn. For this, the computer needs to know where the robot is in the video feed, to be able to properly draw the bubblegram in an appropriate location in respect to the robot. Robots, particularly small ones like the AIBO, pose specific challenges for computer vision: they move around, get closer and further away, change orientation and shape, such as moving the head and legs, all meaning that they can look very different at different times. Further, the AIBO is shiny and seemingly-random specular reflections on the robot's surface add a great deal of variation to the robot's appearance.

While there are many practical vision solutions to this problem that would enable the interaction we are looking for (e. g., robot augmentation with markers), we elected for a marker-less approach to improve practical simplicity of the system by not having to annotate

the robot. As robots start to enter everyday environments on a large scale, then we believe that there are wider implications to the general problem of detecting and tracking robots, and in the long run, we see this as a promising, scalable approach to interacting with multiple robots in complex environments

Given the complexity and difficulty of this tracking problem, we restrict our system to the black ERS-7 AIBO on our in-lab grey carpets, and restrict the AIBO to a set of poses that are practical for our particular task: the AIBO will always be standing or walking in the same pose, and it will always have its display (lights) off. Its legs will still move during walking, the head can move dramatically up and down and look to the side. Further, the different poses, such as facing the person, sideways, or from behind, are quite different.

Our approach is based on a technique which uses what are called *Haar-like features* in combination with machine-learning to detect human faces in a picture (Viola and Jones, 2001). The *Haar-like features* implementation we employed in our algorithm is freely available and packaged with the Intel Open CV Library. In the remainder of this section we introduce Haar-like features and introduce our original method for leveraging them for robot detection. We finish by discussing how we used the existing Intel Open CV implementation and machine-learning system for our purposes.

6.2.1.1 *Haar-like Features*

Haar-like features are an adaption of Haar wavelets (Haar, 1910) to computer images, where two-dimensional templates are used to extract images' properties. These templates are rectangular, varying in size, and subdivided into white and black regions which roughly correspond to the frequency-sensitive form of Haar wavelets (Figure 6.15). The intuition is that a particular template can be used at a particular location to match frequency-specific intensity distributions about the region, for example, soft or hard edges, where the size of the template corresponds to the frequency being targeted. Figure 6.16 shows how these templates may be




Figure 6.15: example templates used for extracting *Haar-like features* from an image




Figure 6.16: an example of how *Haar templates* can be applied to a Sony AIBO (webcam images)

applied to our Sony AIBO, where a template may identify the darker body over the lighter legs, or the darker legs with lighter background in between, etc.


We noticed that many of these features are quite robust to certain changes in the image. For example, as our AIBO in Figure 6.16 turns many features, particularly lower-frequency ones, still apply. Thus given a detector (based on a particular set of features) for an AIBO pose, the AIBO can move, change shape, and rotate to a significant degree while still holding many of the light/dark relationships and still being detected. We also found that the range of change allowed before breaking can be increased by lowering the strictness of the classifier, although this increases the number of false positives, a problem which we discuss in the next section. However, this robustness can only be taken so far and we were unable to create a single detector that would reliably detect the AIBO from various angles.

6.2.1.2 AIBO Poses

Our solution to the problem of real-time, pose-independent detection of an AIBO is to break the problem into discrete AIBO poses to be detected independently, using an separate classifier for each pose. The robustness of the classifiers as described above enables some overlap between each classifier, such that all poses are covered by a carefully-planned subset. We selected the poses: AIBO from the top, side, front, and back (Figure 6.17). The problem remains as to how to combine the results from these separate classifiers into a single all-encompassing one. We present our solution below.



Figure 6.17: AIBO poses targeted by individual classifiers (webcam)

(a) how the different classifiers *vote* on the AIBO's location

(b) samples of AIBO detection

Figure 6.18: AIBO tracking using classifier voting

6.2.1.3 Classifier Combination

Ideally, our pose-targeted classifiers would be mutually exclusive in that one and only one classifier would detect an AIBO in any given frame; the classifier best suited to the AIBO's current pose. However, due to the large overlap between the classifiers, exaggerated by low classifier quality settings, it is the general case to get multiple (often conflicting) results, both from any given classifier and between classifiers.

Our solution is to use a voting scheme, where the positive hits from the various classifiers *vote* on the most likely positive hit. The Image region with the most votes wins, and is selected as the most likely positive hit (Figure 6.18). Thus we leverage the overlap and multiple-hits properties of the weak classifiers. We also used temporal information from the video feed to improve tracking quality. We made assumptions such as how fast the robot can move through the scene or change size between video frames to implement basic filtering and tracking smoothing, resulting in a more-robust result.

Finally, our tablet PC (Toshiba Portege M200, 1.5 GHz Intel Pentium M, 1 GB RAM) was not powerful enough to run the classifiers simultaneously, and so we scaled the image data to half-resolution (to 320 px×240 px) to improve speed while maintaining satisfactory results.

6.2.1.4 *Training and Image Library*

The face-detection algorithm we are using (Viola and Jones, 2001) requires a lengthy training phase using included software. We provide a sample image library of what is to be detected (faces, originally), and a classifier is produced. As part of this we need to select parameters such as the acceptable false-positive rate — this is how we relaxed the strictness of the classifiers to have more matches for voting (for more details please see Young et al., 2006).

The image library requires both positive (with AIBO) and negative (without AIBO) image samples, where the sample images should match the general properties and environment of the target detection scenario. To create this database we recorded video sequences of the interaction environment, both with and without the AIBO, using the same camera as will be used for detection during runtime. Using this technique we extracted, manually cropped and sorted more than 1300 positive and negative images. The classifier we created using this technique is very sensitive to the image library. For example, placing the AIBO on a different floor such as our white hallway linoleum results in very poor results.

In this section we detailed our *bubblegrams* implementation and original computer vision algorithm. In the following section, we present our *Jeeves* test-bed implementation.

6.3 IMPLEMENTING THE JEEVES TEST BED

In our *Jeeves* implementation the person holds a tablet PC, and can look *through* it via a live video feed to observe the cartoon artwork (Figure 6.19). This is a window-to-the-world paradigm, where the tablet displays the video feed on its screen such that it appears as if you are looking through the tablet as a transparent window to the real world. The tablet can then freely draw the appropriate cartoon artwork anywhere in the scene. Currently, this implementation does not enable input to the system and only focuses on enabling the robot to communicate using cartoon artwork. We see our MR interface as an enabling and test-bed technology only, enabling rapid full-colour prototyping of animated entities anywhere on or around the robot. We do not suggest that people necessarily should rely on see-through devices to see a robot's cartoon expressions: an implementation could conceivably use any methodology deemed appropriate, such as on-robot displays.

We achieved video see-through using the Toshiba Portege M200 tablet PC (1.5 GHz) Intel Pentium M, 1 GB RAM), and mounted a Creative web-cam on the tablet itself: the camera

video feed is displayed on the tablet's screen to provide the see-through (transparent-tablet) illusion. To achieve the MR, we affixed a large fiducial marker to the Roomba (Figure 6.20a). We use ARToolkit (Billinghurst et al., 2001) to track the full six-dimensional location and position of the robot, information which enables us to correctly draw the cartoon artwork properly rotated, scaled, and positioned in 3D over the robot. We also used an iRobot Roomba (discovery model) as our robot (Figure 6.19). We selected this robot over the Sony AIBO used for *bubblegrams* for its simplicity and non-anthropomorphic appearance, helping to keep an interacting-person's focus on the cartoon artwork.

Using this setup we achieved the preliminary prototypes of both the *recycle police* and *clean-tracks* interaction scenarios (Section 6.1.4.4). The details of *trash Jeeves* are given below.

6.3.1 Implementing Trash Jeeves

To implement *trash Jeeves* (Section 6.1.4.4) we needed to track the Roomba to determine where it was, and specifically, in relation to the trash can. To accomplish this we installed a ceiling camera which looks down on the environment, and augmented both the Roomba and the garbage can with conspicuous brightly-coloured markers. This camera is attached to a




Figure 6.19: *Jeeves* interface, showing the view through the window-to-the-world interface (live image)

host PC which we have programmed to use colour segmentation (via Intel's OpenCV library) to track both the robot's and the garbage can's locations (Figure 6.20). The host PC connects to the Roomba using Bluetooth to control its movement, and also uses this connection to listen to the Roomba's sensors. The host PC then uses an 802.11g wireless connection to the tablet PC to report the robot's current state and inform on which cartoon artwork sequence should be shown. The entire interface is implemented using C++.

In this section we detailed our working implementation of the various *Jeeves* interaction scenarios, including the specific solution for the *trash Jeeves* scenario. Following we detail our various informal design critiques, for both the *bubblegrams* and *Trash jeeves* scenarios.

6.4 INFORMAL DESIGN CRITIQUES

We performed informal design critiques on both the *bubblegrams* and *Jeeves* interface implementations, where we asked several lab members to use the interfaces and to generally comment on their experience with them. In both cases participants were graduate students from our lab, and had no prior experience or familiarity with our work. The simple results from these early experiences were paramount in early directing of our overall social HRI approach, ultimately shaping the direction of much of this dissertation.

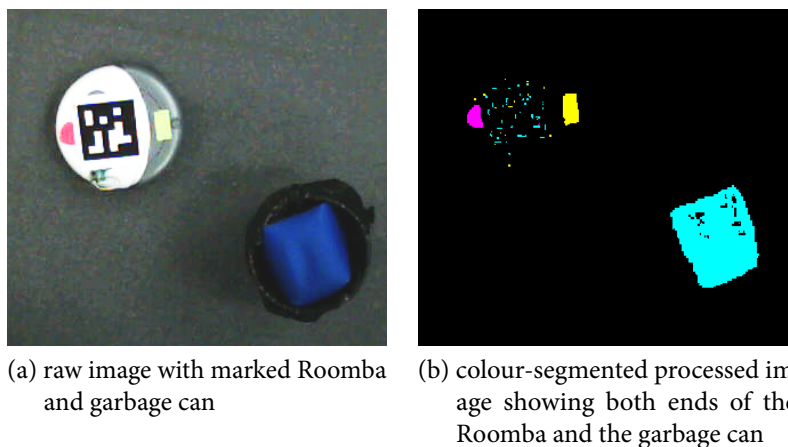


Figure 6.20: colour segmentation for robot tracking in *Jeeves*

6.4.1 *Bubblegrams*

We recruited three participants for the *bubblegrams* informal design critique. In addition to showing that they understood the cartoon artwork, they unanimously expressed excitement about the interface and that they enjoyed using it. The interest surrounding *bubblegrams* was not generally focused on the fact that the participants were interacting with new technologies such as robots, an HMD, and MR. Rather, participants expressed excitement about the bubblegram itself; the thought bubble floating above the robot. Their comments pointed to the familiarity of the graphics, how they already understood the meaning behind the bubble, and how this familiarity helped make the interface less intimidating.

This finding early in our social HRI exploration was initially very surprising for us, that strong familiarity with components of an interface can have such an impact on the interaction experience. This led to further exploration of how cartoon artwork can be applied and ultimately to the development of the *Jeeves* interface. Further, this led to the consideration of how common familiarity, the *social stock of knowledge*, is integral to interaction with robots, a consideration which has become a backbone of this dissertation.

We decided to not perform a formal study on the *bubblegrams* interface. The primary reason was what we saw as a difficulty of creating a valid (believable to the participant) scenario which would have the participant engage the interface for long enough to conduct a study. We see the development, design, and conducting of a formal *bubblegrams* study to be important future work. Below we detail our preliminary evaluation of our *bubblegrams* vision algorithm.

6.4.1.1 *Evaluation of our Bubblegrams Vision Algorithm*

To serve as a benchmark during algorithm development we recorded a two-minute video sequence of the AIBO walking around from various angles and distances with busy backdrops; we informally found the results to be acceptable as the tracking was only temporarily lost with dramatic camera movements resulting in blurry images.

During mock bubblegram interaction sessions with lab members we documented the detection success rate. When movement was minimal and the person was interacting with the bubblegram, detection rate was nearly 100%. Overall through the pilots, including during

rapid movements, our algorithm correctly detected the AIBO 70% of the time, with false positives 14% of the time, and false negatives 7% of the time.

In this section we have detailed our informal design critique for the *bubblegrams* interface, where we found that people reacted very positively to our device and understood the intended communication of the cartoon artwork. Below we detail our design critique for *Jeeves*.

6.4.2 *Jeeves*

We recruited four participants and asked them to use the *trash Jeeves* scenario. We approached this critique as a feasibility exercise with the purpose of observing initial reactions and comments regarding all aspects of the system, including use of the implementation hardware, the overall idea and our particular use of cartoon art. Participants were asked to comment about their experience using a think-aloud technique.

All participants enjoyed the interface, saying that it was, for example, “cute,” “fun,” and “interesting.” All participants were observed to anthropomorphise the Roomba, and one person mentioned at one point that the Roomba seemed “really angry.” While everyone mentioned that they noticed the Roomba was stuck, only one person attempted to move the trash can out of the robot’s way (see Figure 6.11, page 150). One person mentioned that the augmented Roomba seemed “more personal than most electronics,” and that this is an improvement over their existing home appliances.

The participants also had some suggestions about the work. One person voiced concern that the cartoon art may be “too flashy” and annoying, and that a robot should be careful as to how it uses these techniques. Another participant complained about the tablet MR interface, stating that they put more energy into working with the MR interface (holding it, aiming it) than viewing the cartoon art. Finally, one person raised concerns about the quality of the animation and noted that the interaction may be improved with the addition of sound.

This preliminary evaluation demonstrates some of the potential of *Jeeves*: everyone understood the Roomba’s state, no training was required, they all anthropomorphised the Roomba, and at least one person related to how personal the robot became as a result of the cartoon art. However, participants also raised some important concerns which highlight directions for future work, such as that we have to be careful when creating cartoon content to fit the level and types of cartoon art to the task and scenario at hand. Too subtle or too distracting art may

not only be ineffective but can be frustrating for people. Also, when creating cartoon art the visual quality of the content is important, including the quality of the art and the animation, as problems in the content can disrupt the communication with the person. Finally, while we used MR for test-bed purposes, real-world implementations should consider using a more light-weight means to communicate the cartoon artwork.

While we thoroughly explored the possibility of conducting a full-scale formal study of *Jeeves*, for various reasons we ultimately decided to focus on our other projects; some of our evaluation ideas are outlined in our future work section (Section 9.4.3.3). One reason is the difficulty we had with creating sufficient task validity, where the participant would be doing something engaging and meaningful where the cartoon-artwork communication could play a role. Further, many questions we developed related more to understanding the various kinds and potential applications of cartoon artwork, and how they could map to robot states. While these questions are meaningful, we decided instead to focus more on core social HRI questions through the development of other interfaces as outlined in this dissertation.

In this section we have presented informal design critiques for both the *bubblegrams* and *trash Jeeves* interfaces. These studies showed how our participants understood the cartoon artwork without requiring explanation, and in general had a positive response to our robots that communicated through cartoon artwork.

6.5 CONCLUSIONS ON CARTOON ARTWORK

In this chapter, we have explored how robots can use cartoon artwork as part of their communication and interaction repertoire. We have looked at the ways that such techniques can be used by robots for interaction (through our discussions on the Mixed-Reality Integrated Environment (MRIE), Section 6.1.3), and explored which techniques from cartoon artwork robots can use (through the *Jeeves* project, Section 6.1.4.4).

Our exploration into cartoon artwork helped us to better understand social HRI, and has directly addressed our research questions (Section 1.3.5, page 8): we demonstrated that robots *can* use cartoon artwork for indirect communication of robotic state through our interface designs and implementations (*robot expressionism*). Our design critiques demonstrate that people understand without explanation, using *the social stock of knowledge* (question 2). We have provided theoretical tools which social HRI designers can use to develop their own

cartoon-artwork interfaces, as well as provided technical solutions (including an original computer vision algorithm) to realize cartoon artwork implementations (question 3). Finally, this work has improved our knowledge of how anthropomorphism impacts interaction (question 1), as people reported that the explicit emotive cartoon artwork made the robot more fun and personal.

There is still a great deal of research to be done regarding using techniques from cartoon artwork for robots. The theory which we present in this section was developed as a means to build and understand our own interfaces, and should be further developed to consider which additional cartoon artwork techniques we can use, and in which additional ways. Further, it is important to develop a better understanding of the meaning behind the techniques and ways they are used, for example, in relation to being disturbing, encouraging anthropomorphism, and so forth. One other area we are particularly interested in is the expansion of our *Jeeves* test bed (beyond MR) to enable smoother and more natural forms of interaction and prototyping. There also remains the question of how a robot's use of cartoon artwork relates to the *holistic interaction context* (Section 3.2.3), for example, how such a robot would impact social structures, or which cartoon-related sources of influence people draw from when deciding how to interact with a robot wielding cartoon artwork. All of these questions remain important future work for the use of cartoon artwork for social HRI.

STYLISTIC LOCOMOTION AND PUPPET MASTER

We believe that the style of a robot's behaviours and actions will be particularly important for interaction with people, and that people should be able to demonstrate their preferred styles to a robot the same as they would to another person. This chapter presents both our *stylistic locomotion* and *puppet master* social Human-Robot Interaction (HRI) projects as they both surround the notion of style, and they share components of interface designs and implementations. Due to the broad scope and duality of this chapter, we present below an outline of what is contained, and introduce the individual projects themselves in their respective sections.

The link between these projects is that we developed a single underlying algorithm (the *puppet master* algorithm) to realize both designs. As such, many of the interface implementations, as well as the evaluations, addressed both projects simultaneously. Further, given the complexity of the systems involved, we first developed animation-only (non robot) prototype solutions before moving to the final robotic implementations. The non-robot, animation versions were created as a means to initially explore both the concept and viability of the *stylistic locomotion* and *puppet master* approaches. Following, we transferred the system to robots to enable people to directly experience and interact with robots that convey a style through locomotion paths. This explains the chapter outline as presented below.

We first introduce the *stylistic locomotion* project (Section 7.1), and then the *puppet master* project (Section 7.2), both complete with related work and original interface designs. Shared implementation details are presented in the (latter) *puppet master* section.

We give a full account of the underlying *puppet master* algorithm which we developed to implement both projects. We do this first for the animation *puppet master* algorithm (Section 7.3), followed by the extension to robots, the robotic *puppet master* algorithm (Section 7.4). Finally, we detail our formal multi-part evaluations for both *stylistic locomotion* and *puppet master*. First we present the animation study (Section 7.5) and follow with the robot study (Section 7.6).

These systems serve as proofs of concept for the general idea of portraying style and emotion through interactive robotic movement, and the bigger idea of robots communicating through style. We believe that we are among the first to develop interactive robotic characters that elicit stylistic reactions, and the very first to do so explicitly for locomotion paths. These interfaces serve as test beds for exploration into the how people interact with and respond to entities that elicit stylistic motion.

7.1 STYLISTIC LOCOMOTION: ROBOTIC STYLE

In this section we introduce the *stylistic locomotion* project, relevant related work, and present several social HRI interface designs for realizing interactive, stylistic locomotion. We introduce this project while reflecting on our social HRI theoretical framework (Section 3.6) in the text below, including our three perspectives on social HRI: P1 (visceral), P2 (social mechanics), and P3 (social structures). For this project we use style to refer to “a distinctive manner of expression” or “a distinctive manner or custom of behaviour of conducting oneself” (Merriam-Webster Online, 2010, “style”).

As people generally care a great deal about the style of objects and technologies that they possess, we expect they will want (P1 visceral) attractive, pleasing, and (P3 social structures) appropriately-designed robots the same as they want an attractive table, wristwatch, or car — design is directly related to user experience and satisfaction, and as such, *social integration* (Norman, 2004; Young et al., 2009). Research further has shown how, in general, style is important for strengthening communication, creating believable personalities and encouraging anthropomorphism, and ultimately has been linked to creating a more rewarding and engaging experience (Bates, 1994; Breazeal, 2002; Reeves and Nass, 1996; Thomas and Johnston, 1981). Thus, people have predispositions about style from *the social stock of knowledge* that robots need to be designed to accommodate.

Robots can integrate style into their actions (into *how* they perform their tasks) to convey socially-charged and easy-to-understand meaning, leveraging *the social stock of knowledge*. Given *active agency* and people’s tendencies to treat robots socially, style becomes a particularly important part of a robot’s emergent (P1 visceral) character such as a robot being perceived as acting *aggressively* or talking *shyly*. Other examples include: a lawnmower robot may move *sluggishly* to indicate that it requires maintenance, a cleaning robot may do its

tasks *meekly* and *shyly* to avoid interrupting people and to let them know it is trying to stay out of the way, and a guard robot may do its patrol *aggressively* to warn would-be intruders (both P₁ visceral and P₂ social mechanics). We further believe that these characteristics lend themselves to integration into robotic interfaces without changing the goals of the actual task the robots are doing.

The perceived *style* of a robot's actions is dependent on the *holistic interaction context*, the person(s) and the robot(s) involved in the interaction, among many other factors — what is perceived as a particular style by one person or in one context may be perceived entirely differently by someone else or in a different context, making the problem of robots acting in particular styles to be particularly challenging. However, wide-sweeping generalizations of style do exist, exemplified in the world of art. Actors, animators, and other artists have an incredible ability to develop and deliver believable characters and personalities that reach across personality types, demographics, and often even international borders. In fact, the power of art exemplifies the importance, universality, and power of style, as illustrated in this quote in specific regard to animation:

Conveying a certain feeling [sic.] is the essence of communication in any art form. The response of the viewer is an emotional one, because art speaks to the heart. This gives animation an almost magical ability to reach inside any audience and communicate with all peoples everywhere, regardless of language barriers.

— Thomas and Johnston (1981, p. 15)

While we avoid the question of whether social HRI itself is an art form, the above discussion highlights the element of art and feeling involved in the style of robotic communication, particularly when actions are intended to be interpreted by people (this is *robot expressionism*). Thus we believe that there is a large overlap between social HRI and the areas of acting, animation, and that social HRI can benefit from these approaches.

As an initial exploration into how robots can communicate through style we limit our focus to the locomotion path of robots. That is, we consider the style of how a robot moves around a space and how that style can be easily and readily recognized and understood by a person. In particular, we consider how such elements of style can be embedded in locomotion. For example, Figure 7.1 is a hand-drawn concept sketch of examples that demonstrate how path alone can be used to communicate emotion.

In our work we look at how a robot's movements relate to a person's movements. For example, a robot that is aggressive toward a person, showing it through its locomotion path, needs to consider how and where the person is moving to maintain their *aggression*, lest it ends up being aggressive toward another person, object, or an empty area of space. These characteristic, *interactive* locomotion paths are much more complex than static paths (e. g., as show in Figure 7.1), as they must take into account the other path in their own movements. Further, when dealing with people, this becomes a real-time challenge, where decisions about how to move and react to the person's actions must be made instantly, with little ability to look ahead.

Next we present an overview of relevant work that helps us understand the power of the style of motion. We follow with detailed discussions on both our animation and robot-based interface designs, before moving on to the next project, *puppet master*.

7.1.1 An Overview of Style

There has been extensive work in both psychology and animation surrounding style and emotion from motion paths. In psychology, Heider and Simmel (1944) show how simple geometric shapes (triangles and circles) moving around a screen (with a setup as shown in Figure 7.2) is enough to generate a social response. With only the movements and interactions between the shapes (no sounds, etc.), people were shown to construct stories that attributed personalities and gender to the shapes, and there was a surprising amount of similarity across participants. For example, the big triangle was described as an aggressive male trying to

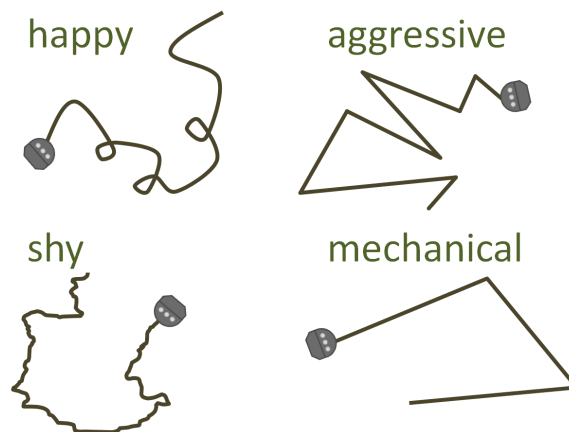


Figure 7.1: concept-sketch examples of a robot's stylistic locomotion path




Figure 7.2: a mock-up sketch of a scene as used by Heider and Simmel (1944) in an experiment that uses motion to show apparent behaviour

corner the female circle (mother, or just female), and the small triangle was trying to outwit the large triangle to help the circle. This finding supports the generalizability of the emotion and social communication that can be conveyed through motion. Kassir (1982) directly related this work to computer animation, and others did similar work that showed that even simpler situations using only a single point-like object, can push people to perceive animacy based on how the object moved (Scholl and Tremoulet, 2000; Tremoulet and Feldman, 2000).

Thomas and Johnston (1981) highlight that making the connection between motion, style, and emotion was one of the important innovations of Disney. In particular relation to our work, that the realization that *how* a character moves between locations in a scene is a very important part of the overall character experience (Thomas and Johnston, 1981, p. 30, 346–357). The importance of movement is also evident for interactive characters in computer animation (e. g., as integral in Bates, 1994; Dontcheva, Yngve, and Popović, 2003). In one particularly related project, Amaya, Bruderlin, and Calvert (1996) showed how scripted actions such as “pick up a glass” or “knock on a door three times” can be made to be “neutral,” “shy” or “angry” based on transforms applied to the motion data. This supports the idea that, *how* something is done can be separated from the actual goal task. However, as far as we know none of the previous work considers how such techniques can translate to robots. For example, much of the work in animation does not consider the real-time challenges of generating motions, and as such are not directly transferable to robots that must interact in real-time with people.

Perhaps more related to our work is research that has looked at proximity behaviour of a robot near a human. Much of this has focused on the practical, engineering mobility and vision challenges of a robot interacting near a person (e. g., Byers and Jenkins, 2008; Liem,

Visser, and Groen, 2008) or the development of goal-centric predictive models (e. g., Chueh, Joshi, Au Yeung, and Lei, 2006). Others have looked at social aspects such as how close a robot should be to a person (Yamaoka et al., 2008) and how to make natural following behaviour (Gockley et al., 2007); in this case, copying a path versus shortest route. We look at the broader case of a robot conveying style through its locomotion path, where questions of comfort distance and natural movements are included in the *holistic interaction experience*.

In our work we posit that style can be embedded within the locomotion path of the robot itself similar to how style is embedded in animation. Further, we focus on how this style is realized in part through direct real-time interaction with a person's movements. The question we explore in our work is: do these characteristic, interactive locomotion paths apply to social HRI as they do for on-screen animated entities? How does the robot (unlike an animated character) sharing the same physical environment with people affect the interpretation of their movement style? Do people interpret the motions as having social meaning, and construct emotion and characters, for the robots the same as they do for animation? We explore this by developing interfaces to enable such social interactions (detailed below) and performing evaluations to determine how people interpret them (Section 7.5, Section 7.6).

Here we have outlined the relevant work which supports our approach of communicating through locomotion style, and highlighted how we believe our work is unique. Below we detail particular *stylistic locomotion* interface designs.

7.1.2 *Animation Interfaces for Stylistic Locomotion*

Here we present two original interfaces that enable a person to interact with animated entities that communicate using stylistic locomotion paths. We first present a mouse-based Graphical User Interface (GUI) (which we call the *mouse GUI*), and then our tabletop Tangible User Interface (TUI) interface (which we call the *animation table*).

In our interface designs we attempted to isolate the communication to the style of the locomotion path only. We did this by simplifying the entities used in the interfaces as much as we felt possible, and so not choosing complex character designs was a deliberate choice.

In both cases, we utilize an example set of character styles to provide a wide cover of interaction style possibilities: these are *playful friend*, *lover*, *bully*, *stalker* and *afraid*. For each of these, the computer character acts in the style given in respect to the person's character,

using only its motion path and how this path reacts to the motion path of the person's character. We made these six style selections based on our own judgement of emotion cover. Methodologically approaching and analyzing the breadth and types of style, using frameworks such as Plutchik's (1980) wheel of emotions, remains important future work.

7.1.2.1 *The Mouse GUI for Stylistic Locomotion*

The goal of the mouse-based GUI design is to leverage the familiar mouse paradigm to enable people to interact with our characters. In this design, a person can control the locomotion path of an on-screen character (called the *main* character), to which a second, computer-controlled character (called the *reacting* character) will react to in real time to exhibit a particular style. The style of the reaction is embedded in *how* the reacting character moves around the screen in relation to the person's character. We have attempted to design the look of our primary characters to avoid bias toward how people interpret the characters' movements, although they have eyes which make them anthropomorphic in nature (Figure 7.3).

Figure 7.4 shows the interaction canvas where people control their character. The top image uses the generic characters described above (using a generic blue background), with the yellow character being the person-controlled *main* entity and the green being the computer-controlled *reacting* entity (Figure 7.4a). In addition, we have developed two other scenarios, one with a kitten (*reacting*) playing with string (*main*, person-controlled, Figure 7.4b), and one with a baby duckling (*reacting*) interacting with a mother duck (*main*, Figure 7.4c).






Figure 7.3: our generically-designed animation avatars


The person uses a standard Personal Computer (PC) mouse and clicks on the screen to drag the *main* character around, where the character is moved along the path of the mouse. Since the standard mouse lacks a rotation sensor, we mapped the character's looking direction to its movement direction. That is, with this interface the character cannot move sideways or backwards, and is always looking forward into the direction it is moving. We chose this solution rather than the use of other means such as keyboard input for look direction to keep the interface simple to use. Following, the computer-generated interactive character reacts and responds to the person's character's movements in real time, with the generated character acting in a particular stylized fashion. There is no collision detection algorithm



(a) generic interface



(b) kitten and yarn character interface



(c) mother duck and duckling interface

Figure 7.4: mouse GUI interaction canvas, with generic, cat, and duck variants, for interacting with stylized characters

employed, such that the entities can overlap with each other and with any components of the environment without impact on interaction.

In the generic entity case (Figure 7.4a) the characters have a top-down view and so the character is simply rotated to look in the appropriate direction. In the cat and duck cases, however, the characters are given in a sideways view (Figure 7.4b, Figure 7.4c). This means that they cannot simply be rotated as moving in an opposite left/right direction would have the character upside-down. This is solved by flipping the character to stay upright while moving in the appropriate direction.

Although serving as an important prototype step, we found the mouse-based interface to be limited in terms of interaction, particularly in how the person cannot specify the look direction of the character. This relates to a larger concern with the input–output mapping between the computer mouse and an on-screen entity, where in addition to the omission of input rotation there is a strong indirection between the movements of the mouse on the table and the on-screen entity (Sharlin, Watson, Kitamura, Kishino, and Itoh, 2004). Our *animation table for stylistic locomotion* presented below addresses some of these concerns.

In this section we introduced the *mouse GUI* interface design and detailed how it realizes *stylistic locomotion*. The implementation details are grouped with the *puppet master* project and presented in Section 7.2.3.1.

7.1.2.2 *The Animation Table for Stylistic Locomotion*

We developed a TUI-based system to extend our *stylistic locomotion* system beyond the capabilities of the *mouse GUI*. Here, rather than using a mouse, the *main* character is controlled using a physical TUI puck. This puck is placed on a horizontal tabletop computer where it is tracked in real time and computer graphics of the interacting characters are drawn on the screen. This is illustrated in Figure 7.5, which shows the yellow *main* character underneath the person's puck and the green *reacting* character near by.

With our TUI, a person can grab the physical puck and move the *main* character as they wish, with the *reacting* character responding in real time with a given style. Puck manipulations on the table are directly translated to system input used to control the *main* character, and this control is mirrored back to the person through the *main* character's entity being drawn under the puck (the yellow avatar in Figure 7.5). This tightly couples the input and output spaces, an approach which has been shown to improve the ease-of-use of the




Figure 7.5: our *animation table* interface for interacting with stylistic characters

interface (Ishii and Ullmer, 1997; Sharlin et al., 2004). In addition, the puck enables the person to explicitly specify the character direction to tell it in which orientation it should look, by simply rotating the puck as they move it. This enables further flexibility of interactions, for example, the character moving sideways, circling while maintaining a look direction, backing up, etc.

Our animation implementations serve as proofs-of-concept for enabling a person to interact with interactive characters that portray style through their locomotion path. These interfaces (implementations detailed with the *puppet master* project in Section 7.2.3) were important stepping-stones that both enabled us to explore the core ideas, and fine-tune the implementation of our algorithms with flexible animated characters, before the development of our robot-based system described below.

7.1.3 *Robot Locomotion: Robot-Based Stylistic Locomotion*

Rather than interacting with the virtual entities through input mechanisms and displays such as outlined with the interfaces presented above, transferring the animation systems to real robots enables a person to interact with the entity in their own physical space. In addition to robots serving as an end-application motivation, we believe that this transfer is justified by the physical and social nature of the robot adding social layers of depth and meaning to the stylistic interaction, relating to the unique agency properties of robots.

One important difference of our *stylistic locomotion* animation systems and robot ones is that we extend the expression language of the robot beyond locomotion path to include

emotive noises. The robots can make either a happy or unhappy noise while it moves. This as an initial exploration into how interactive style can be extended beyond locomotion path, a goal we analyze in our studies below (Section 7.6). We selected sound because it is a completely different modality than locomotion path, and is a simple (in terms of interaction) addition in comparison to, for example, more complex motion gestures or facial expressions.

Another important difference is that for the robot implementation we added a focus on the broader task of the style, for example, a robot *following* a person in a given style. The reason for this is practicality: a) freedom of movement was limited given the small physical space we could reasonably use for our robots, and b) it aids our evaluation of this system (explained in Section 7.6) to have a structured task. For the robot implementation our set of style scenarios are: *following politely*, *attacking a burglar*, *happy to see the person*, and *stalking the person*. As with the animated personalities above, our choices of the four styles are based on our own sense for achieving a wide cover and are not grounded in any particular theory on emotion.

In this robot implementation, a person moves (walks) about a space and the robot reacts to the person's movements, and moves (on wheels) to follow accordingly based on a particular style (Figure 7.6). That is, using the vocabulary introduced for the animation cases above, the person is the *main* character themselves, and the robot is the *reacting* character. With this setup the person becomes an integrated part of the interactive environment. The robot will immediately respond, enabling the person to experience the robot interacting with them (or observe a robot interacting with another person) within their own physical world. Perhaps a conceptual parallel in the animation case (e. g., our *mouse GUI* or animation table) would be the person being able to move into the screen and become an animated character.

One core limitation of this interface is the assumption that the robot only needs to monitor the motion path of the person in order to interact appropriately; all of the person's other body language, gestures, facial expressions, or sounds are ignored. We maintain, however, that these constraints are important within the context of an initial exploration for limiting the complexity of the problem to a practical, solvable subset. We believe that the resulting behaviour serves as an important proof-of-concept that can scale to more complex cases.

In this section we detailed the *robot locomotion* interface for interacting with *stylistic locomotion* robots. Interface implementations are detailed with the *puppet master* project in Section 7.2.3.2 below. Next we conclude the *stylistic interaction* section.



Figure 7.6: a robot following a person in a given style

7.1.4 Conclusions on Stylistic Locomotion

In this section we motivated the importance of style for interaction with robots, detailed how robots may use style as a component of their interaction for explicit *robot expressionism* communication, and explained how we believe that this leverages the existing stock of social knowledge — perhaps this can further help the robot to fit into existing social structures. We presented three interfaces which enable a person to interact with entities that communicate using style: the *mouse GUI*, the *animation table*, and the *robot locomotion* interface designs.

We see our interface designs as a first step toward the design of robots that communicate using *stylistic locomotion*. There is still a great deal of work to be done to expand the style of robot's actions beyond our straightforward locomotion paths to other forms of communication, such as full-body gestures, use of gaze, or even voice. We believe that future work on interactive robotic style to also consider how robots can accurately interpret the style of people's interactions such that robots can react appropriately to them.

Our current implementations serve as proofs-of-concept for the idea of robots communicating through the style of their actions. Further, these explorations were integral in developing our understanding of social HRI, particularly through our formal evaluations

(Section 7.5, Section 7.6) which helped to highlight both the breadth and the sheer complexity involved with interaction between people and robots.

7.2 PUPPET MASTER: DESIGNING INTERACTION STYLE BY DEMONSTRATION

In this section we introduce the *puppet master* Style-by-Demonstration (SBD) project, an overview of existing programming-by-demonstration work, and present several interfaces for realizing the creation of *stylistic locomotion* through demonstration. In the introduction below we relate this project to our social HRI theoretical framework (Section 3.6), including using our three social HRI perspectives for discussion: P₁ (visceral), P₂ (social mechanics), and P₃ (social structures).

Ideally, robots that integrate into people's lives would be able to independently observe, learn from, and adjust to everyday situations much the same as how people do. This would match *the social stock of knowledge* and enable the robot to seamlessly fit into existing social structures (*social integration*). Unfortunately, solving this artificial intelligence problem is still prohibitively difficult, and modern-day systems only manage to implement a small portion or subset of this greater goal. In our work, we target the focused problem of robots learning the *style* of interaction from people, enabling people to teach a robot a given style by demonstrating it to the robot directly. We call this approach Style-by-Demonstration (SBD).

In our everyday lives we often show others *how* we want things done, using our familiar P₂ social skills (from *the social stock of knowledge*). For example, a person may instruct a friend who is helping them move by explaining which boxes go where, and how something should be packed. In addition to teaching everyday tasks, people are also adept at demonstrating appropriate social behaviour, much of which includes the style of the particular actions. For example, how to perform a business handshake, how to hold one's hands when defending a soccer goal, or how to properly use a paint brush. We argue that, given both the importance of and range of personal tastes and preferences of individuals (Norman, 2004), it makes sense to enable people to directly instruct a robot on how they want it to act (the style), much as how people can show their employees or servants how they want them to act; this relates to the *accessibility and usability* factor of robot acceptance. Robots that can learn style directly would have an edge on *social integration*, as they could easily be trained to match the whims and preferences of any given person or context.

SBD can also be useful for robot manufacturers, as many robots may not be designed for end-user customization and will be pre-programmed from the manufacturer where the creation of the personality, style, and emotion of interaction can be left to professional engineers (or perhaps designers). Creating such behaviours through programming is a common approach for (similar) interactive entities in animation (e. g., Blumberg and Galyean, 1995; Maes, 1995; Reynolds, 1987), where low-level stimulus-response-type behaviour of entities (Wolber, 1997) or task-oriented goal planning and collision avoidance (Dinerstein and Egbert, 2005; Dinerstein, Egbert, and Ventura, 2007) are explicitly described using logical event sequences and conditionals. The resulting style of these goal- and efficiency-oriented actions tends to be very *mechanical*, predictable and boring. Further, advanced programming techniques are generally required to get robots to do even simple tasks in the real world (e. g., Breazeal and Scassellati, 1999; Kanda, Kamashima, Imai, Ono, Sakamoto, Ishiguro, and Anzai, 2007a), where the programmer explicitly defines how it should act in all given situations, sometimes using complex psychological models (Breazeal, 2002). Adding style to these robot actions, for example, to program robots that pick up objects, shake your hand, or follow you, to do these actions in an aggressive, timid, or careful fashion, all interactively (to unpredictable human input) and in real time, further complicates the programming problem. Conventional programming tools have no explicit mechanism for expressing style or a character's personality and emotion and are therefore not well-suited to the task. To be implemented, styles must be defined in algorithmic and logical terms, a non-trivial task even for expert programmers.

We propose that SBD takes advantage of people's existing skills (*social stock of knowledge*) to enable them to naturally and easily teach the style of actions to a robot, in a similar manner to how they may teach another person. SBD can simplify the creation of style-oriented content (of behaviours), providing a means for anyone, be it technical programmer, and non-technical designer, or someone with no technical robot skills, to program a robot's style. In the remainder of this section we provide an overview of programming-by-demonstration, and detail our particular proof-of-concept SBD interfaces.

7.2.1 *An Overview of Programming by Demonstration*

Programming by demonstration is a very active and mature field of research, spanning such domains as animation, robotics, and GUI programming. The problem of creating interactive robotic behaviours is very similar to that of creating interactive behaviours for animated entities: there has been a great deal of work in animation, computer graphics, and video games that aims for life-like, convincing interactive behaviour. As such, much of the related work described here is emerging from the field of computer graphics and animation. We believe that our work is also of significance to the field of programming-by-demonstration, in addition to the field of social HRI, a relationship we outline in this section.

Programming by demonstration emerged in the mid 1980s (Halbert, 1984) and was soon commonly used to automate GUI operations (e. g., Cypher, 1991; Maulsby, Witten, and Kittlitz, 1989; Perlin and Goldberg, 1996), with no consideration for agents or style. There are several systems which enable animation design by performance demonstration, where a person can directly show the computer by performing actions (Dontcheva et al., 2003; Hertzmann, Oliver, Curless, and Seitz, 2002; Igarashi, Moscovich, and Hughes, 2005a,b; Thorne, Burke, and van de Panne, 2004), and this approach has also been applied to robotic motion (Frei, Su, Mikhak, and Ishii, 2000; Raffle, Parkes, and Ishii, 2004) — some of these enable the creation of very expressive, stylistic motions. However, the results of these systems are generally static (not interactive) and do not respond appropriately to interaction from a person, often focusing on the playback of the demonstration.

Robot-targeting work is often goal oriented, including teaching navigation routes (Kanda et al., 2007a) or how to perform specific physical tasks (Gribovskaya and Billard, 2008; Lockerd and Breazeal, 2004; Otero, Alissandrakis, Dautenhahn, Nehaniv, Syrdal, and Koay, 2008). Matsui, Minato, MacDorman, and Ishiguro (2005) created an android that moves naturally based on observing a person's movements, but this is real-time performance (direct mapping of input to output) and the robot did not *learn* anything — it could not act on its own. Our work is unique in that we focus on the robot *learning* by demonstration, not being remote-controlled, and emphasizes the *style of interaction* rather than a particular task goal or movement.

A related approach is to synthesize new motion in real time from an example database (Lee and Lee, 2004; Lerner, Chrysanthou, and Lischinski, 2007; Wiley and Hahn, 1997). These

approaches generally require fairly large example databases as well large amounts of detailed pre-processing (often programmer-assisted). While some systems interactively respond to user input (joystick control for moving obstacles, other characters, etc.), the mapping from the user input to the output is explicitly (and often tediously) defined by the programmer and extremely sensitive and specific to particular behaviours within particular contexts (Halbert, 1984). Furthermore, the target of these systems is primarily physical accuracy and plausibility (punch, jump, walk, collision avoidance, etc.), not the explicit design of style *per se*, or personality emerging from interactive motion related to another entity.

Others incorporate pre-scripted style and emotion into their robotic interfaces, for example, Lockerd and Breazeal (2004) use pre-scripted responses as an important part of a person teaching a robot. In these instances, however, the pre-scripted emotional sequences are used as representations of algorithmic states such as an indication of not understanding a command, or having low confidence. This is an entirely different usage and problem from the more general case of teaching style and emotion in interaction.

While programming by demonstration is a popular approach for both animation and robotics, work generally focuses on a task-specific goal. For systems that do target style, the results are static and not interactive. As such, we believe that our work on interactive Style-by-Demonstration (SBD), the use of programming-by-demonstration to teach a robot the style of interaction, is unique and the first to develop this question. In the remainder of this section we detail our *puppet master* SBD interfaces.

7.2.2 *Puppet Master Interface Designs*

As with the *stylistic locomotion* interface designs presented earlier in this chapter, for practical and initial proof-of-concept reasons we focus our SBD discussion on the more-narrow — but well focused — case of how people can create characteristic, interactive robot locomotion paths by demonstration.

Following the *stylistic locomotion* discussion, here we define the robot as the *reacting* entity which interacts with the person (*main* entity) in a particular style. Our SBD interface designs further have two distinct phases: demonstration, where a person shows the robot the desired style of interaction, and the generation phase, where the robot attempts to interact using the style shown in the demonstration phase.

In the demonstration phase, two entities are required: while the *main* entity moves around a space, a person shows the *reacting* entity the style in which they want it to interact with the person. This is outlined in Figure 7.7 where in the left pane (demonstration) a person is showing a robot how to follow another person, and in the right pane (generation) the robot is autonomously following the person in the style demonstrated. It is important to note here that both the *main* and *reacting* entities are required for the demonstration phase: as we are creating *interactive* behaviour, the reactor entity must have an entity to react to. Also, the movements of the *main* entity during demonstration must be realistic and representative of how the it will move in practise (during generation). If this is not the case, then the *reacting* entity may be trained to react to perhaps unrealistic main-entity movements. This is why demonstration of *paired* motion of the two entities is needed.

The generation phase uses *stylistic locomotion* interaction (Section 7.1): the person controls the *main* entity and the *reacting* entity is automatically generated in real time.

Next we outline how we have implemented our proofs-of-concept by first developing an animation-only system, and then scaling our system to real robots. Animation enabled for an easier first-step proof-of-concept and a flexible development environment for prototyping our algorithms since animated characters are much simpler to control and manipulate.

7.2.3 Animation-Based Puppet Master

The interfaces that we used in our animation-based *puppet master* are based on those presented for *stylistic locomotion* (Section 7.1), where they provide the generation phase. Here we describe our expansion of those interfaces to enable the demonstration phase.




Figure 7.7: a conceptual sketch of our programming style by demonstration robotic interface

7.2.3.1 *Demonstration with our Mouse GUI*

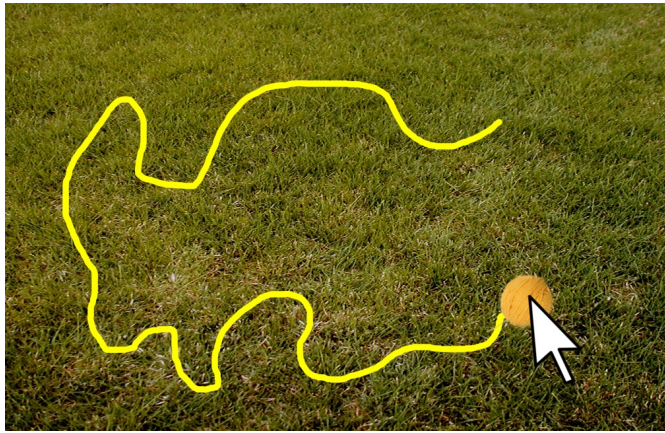
Our mouse GUI initial-prototype concept system used a single mouse and on-screen animated entities for interaction (Section 7.1.2.1), and we have the *main* and *reacting* demonstrations performed sequentially. With sequential training, a person first demonstrates an example input path to represent the *main* character. Demonstration is done using the mouse, where the person clicks on the screen and drags the mouse around to move the on-screen character, and the character's look direction (orientation) is locked to the movement direction. Once the demonstration of the *main* character's movement path is complete, it is replayed to the person who then demonstrates how the *reacting* character should respond to and interact with the main character. Finally, once both demonstration phases are completed the person can control the *main* character and interact with the now-generated *reacting* character. This process is illustrated in Figure 7.8.

This prototype served as an important milestone in the *puppet master* to robots interface design, particularly, as a flexible test bed to fine tune the underlying *puppet master* algorithm. There were several limitations, however. In addition to the indirection and limited expressibility with the mouse GUI interface, we believe that the inability to simultaneously demonstrate the paired interaction greatly hinders the ease of use of the interface: the person has to plan ahead, and demonstrate the full *main* path completely, before they can demonstrate the reaction. We believe that simultaneous demonstration more-accurately supports the puppeteering-like manner of demonstrating style, better enabling spontaneous creativity and natural teaching of stylistic behaviour, which leads us to our next interface implementations.


The mouse-based GUI was developed in its entirety using the Java programming language, version 6. The mouse input and graphical rendering were accomplished using the built-in Java Swing GUI toolkit. The core interactive and stylistic interactive behaviours, as well as the SBD logic, were realized using the *puppet master* algorithm discussed below (Section 7.3). First, we detail our other *puppet master* interfaces.

7.2.3.2 *The Animation Table for Puppet Master*


We extended our *stylistic locomotion animation table* interface (Section 7.1.2.2) to support the demonstration phase by adding a second TUI puck. The addition of the second puck (Figure 7.9a) enables the paired demonstration for the *main* and *reacting* characters, such that example *main* character input and example *reacting* character paths can be simultaneously



(a) First, a person demonstrates how the *main* entity may move.



(b) Then, as the demonstrated *main* entity's path is replayed, the person demonstrates how the reacting entity should interact with it.



(c) Finally, the person controls the *main* character and the *reacting* character is generated and presented to match the style of the demonstration.


Figure 7.8: Using the mouse GUI interaction canvas to sequentially program style by demonstration

performed. This can either be done two-handed by a single person (Figure 7.9b) or by two people working together (Figure 7.9c), where one person controls the *main* and the other controls the *reacting* character. For the generation phase only one puck is used to control the *main* entity, and the *reacting* entity is automatically generated in real time. In this section we detail our *animation table puppet master* implementation.

We used a rear-projected SMART Technologies 1.473 m × 1.090 m, 2800 px × 2100 px custom high-resolution digital tabletop (Figure 7.10a). The rear-projection means that the person's hands, arms, and the TUIs do not cast obstructive shadows. The *mouse GUI* software (on which this implementation is based) was initially too slow to run full speed at this high resolution, a problem which was solved by using Java Swing's two-dimensional graphics-acceleration architecture. We used two plastic paint-cups, turned upside-down with the lids taped to their bottoms, as pucks for a person to grab, and move and manipulate as TUI input to the system (Figure 7.9a).



(a) the two pucks on the tabletop



(b) a person using the two pucks to demonstrate movement style



(c) two people demonstrating movement style together using the pucks

Figure 7.9: people using the *animation table* interface for programming style by demonstration

The pucks are tracked using a Vicon camera-based motion-tracking system. In this implementation, eight Vicon MX-F40 cameras (Figure 7.10b) were mounted on tripods and distributed regularly around the tabletop, with camera lenses chosen appropriately to the distance from the table. The cameras sense only within the infrared spectral range, and are equipped with infrared emitters as seen in Figure 7.10b as the red ring around the lens. Reflective markers were placed on the cups to be seen by the cameras, with each having a unique spatial configuration of markers to enable the Vicon system to differentiate between them (Figure 7.9a). We found the TUI cup material to be particularly reflective and often confused the Vicon system, so we scuffed them with sandpaper and for the particularly troublesome yellow cup, added a non-reflective sticker to the top (Figure 7.9a).

The Vicon cameras are connected to a central Vicon Ultranet unit, which reports the camera information through an ethernet link to a host PC for processing. We use a software package called Nexus, provided by Vicon, for initial position processing, and relay this to our software on the same PC via a local-socket connection and serial protocol. We have the Vicon configured to track and report the pucks' locations to our software at 100 Hz.

A challenge with this configuration is that the Vicon-coordinate six-degrees-of-freedom position of the pucks need to be correctly and consistently translated to the coordinates of the tabletop; this enables the pucks to serve as mapped input and enables the accurate drawing of graphics below the pucks. The origin and orientation of the Vicon space is dependent on a daily-performed manual calibration, and so the position and orientation of the table (in



(a) SMART Technologies Tabletop with surrounding Vicon motion-tracking cameras



(b) Vicon MX-F40 camera with infrared emitter

Figure 7.10: SMART Technologies tabletop with Vicon tracking

Vicon space) changes regularly, complicating the integration of puck data to our software. Our solution is to follow the standard Vicon calibration procedure with an additional step where we define the corners of the table in the newly obtained Vicon coordinate system. We do this by placing the puck in three decided corners of the table, once to define the origin, and twice more to define the x and y axes of the table. This data is used to construct a linear-algebra basis which transforms a Vicon-space vector to the table space. The resulting (standardized) coordinates of the pucks can be directly used in the mouse-GUI software.

The sub-millimetre accuracy of the Vicon tracking enables us to use the height of the puck to determine if the puck is on the table (we use 3 mm off the table as the threshold), a mechanism we use to mark the beginning and end of the demonstration session. This works in lieu of the mouse-button clicks used in the *mouse GUI* system.

In addition to serving as proofs-of-concept for enabling a person to demonstrate style of interaction to entities, these animation-based interfaces provided an important test-bed where we rapidly prototyped and evaluated various interface and algorithmic possibilities. The puck TUI in particular enabled us to develop the added capabilities of simultaneous demonstration and dynamic entity orientation (i. e., not fixed to movement direction), non-trivial additions to the *puppet master* algorithm (Section 7.3). Finally, the animation prototypes enabled us to perform an evaluation on core *puppet master* ideas, validating the approach and algorithm before moving on to the robotic implementations (Section 7.5). In the following sections, we move from our animated-entity renditions of *puppet master* to implementations with robots.

7.2.4 Robot-Based Puppet Master

Here we present two robotic interfaces for *puppet master*, for creating characteristic, interactive locomotion for robots — this is in contrast to the above interfaces which focused on animated entities. For both of these, the generation phase uses the *robot locomotion* interface (Section 7.1.3): the person (as the *main* entity) can walk around a space freely, while the robot (the *reacting* entity) monitors the person's movements and reacts appropriately and autonomously in real time, based on the style demonstration given.

We further use the same behaviours as in the *robot locomotion* interface design: *following politely*, *attacking a burglar*, *happy to see the person*, and *stalking the person*, selected as

behaviours that can reasonably work within the limited size of the physical robot interaction space, and to add a degree of structure (weak task) for our evaluations.

The particular robot we are using in our interfaces (an iRobot Roomba) has a two-wheel-plus-caster configuration, meaning that it has unique movement properties — it cannot move sideways but can turn on the spot. As such, one of the challenges of adapting the SBD from animation to robots is that we need to ensure that the motions demonstrated are actually reproducible by the robot. For example, we do not want a person demonstrating to a robot to move sideways to circle around a person, when the real robot is not capable of the exact action. Further, our robot implementations can use *happy* or *sad* noises as part of their interaction repertoire in addition to movements (discussed in Section 7.1.3). As such, robot-based *puppet master* interfaces need to both enable the person to demonstrate sounds, and need to incorporate the particular movement capabilities and constraints of our robot.

Below we present these two new interface designs: one is a Microsoft Surface tabletop computer with a movement-constrained TUI (we call the *Surface puppet master*) and the other is a broomstick attached to a real robot (we call the *broomstick* interface). The *broomstick* interface was an attempt to remove as many layers of indirection as possible from the interaction with the robot, and the *Surface puppet master* was a proof-of-concept for constrained TUIs on tables as well as an explicit comparison point for evaluation (Section 7.6).


7.2.4.1 *The Surface Puppet Master — Demonstration for Robot using a Tabletop and TUI*

We built a small, motion-constrained wheeled puck (TUI) placed on a Microsoft Surface tabletop computer for *puppet master* demonstration. The TUI puck (Figure 7.11a) is small enough to grab comfortably with one hand and move around, and its physical form is a familiar mouse which is mounted on the top. Unlike our more-generic for-animation *animation table* TUI, this TUI is designed to enforce the actual movement properties of our robot (e. g., not allowing sideways movement) through the structure of its wheeled undercarriage, and enables the person to specify when the robot should produce sounds through use of the mouse buttons. Further, in contrast to the *animation table*'s much-larger SMART Table tabletop system, this puck is placed on a smaller Microsoft Surface, enabling a person to comfortably demonstrate style over the entire surface without walking around the table (Figure 7.11b). Two-person interaction is not possible, but this was not part of this phase's design goals.

The puck is used to control the *reacting* entity for purposes of demonstration, while the *main entity* (representing a person) is realized by an animated (pre-scripted) happy-face icon moving over the table (Figure 7.11). We selected a pre-scripted *main* character during demonstration as an attempt to simplify the demonstration process, a limitation we felt was reasonable given our narrowed interaction scenarios. In addition, this was used to add regularity during our evaluations, where people simply focused on showing a robot how to follow the pre-scripted person.

The TUI puck consists of a wireless mouse mounted on a mechanical base (Figure 7.11a). The base was built using a mixture of pieces from various Tamiya robot kits, and has a two-wheel-plus-caster design to constrain movements to match those possible by the robot (cannot move sideways). The TUI's location on the table is tracked using the Microsoft Surface's built-in diffused-infrared touch detection mechanism. We attached reflective tape (visible to the Surface) to the TUI's bottom such that it hovers above the table without touching (Figure 7.12), reflecting light back into the Surface and registering as touches. The arrangement of the tape forms a rough isosceles triangle: the centre point between all three corners is used as the TUI's location, and the point furthest from the centre is used as the front to deduce the TUI's look direction. The mouse buttons are used to trigger robot actions (sounds, in this case).

The Microsoft Surface software was implemented using Microsoft C# and Windows Presentation Framework (WPF), and we used WPF's integrated Extensible Application Markup Language (XAML) to build the TUI puck tracking event system and to implement the animated smiley face. We programmed the Surface's computer to track the TUI's location, and report it



(a) *surface puppet master* TUI for demonstration



(b) *surface puppet master* with Microsoft Surface and acTUI

Figure 7.11: our *surface puppet master* interface with TUI for demonstrating to a robot




Figure 7.12: bottom view of the base showing the reflective tape, and the caster and wheels providing movement constraints

(with the simulated human's positions) over an ethernet link to our primary PC running the *puppet master* algorithm and the *robot locomotion* interface.

This section detailed our *Surface puppet master* interface and design, an important component of our overall *puppet master* project. Following, we present the *broomstick puppet master* interface.

7.2.4.2 Broomstick Puppet Master — Demonstration using the Broomstick Interface

Here we present our *broomstick puppet master* interface for in-situ, immersed demonstration of characteristic, interactive robot locomotion style. Our *broomstick* interface enables a person to directly control a robot using a broomstick (which is attached to a robot). Manipulating a broomstick is a natural and familiar mechanism that directly builds on *the social stock of knowledge*, and so we expect this interface to be easy to understand.

For demonstration, one person (representing the *main* entity) walks naturally in the space while another person uses the *broomstick* robot to demonstrate an interactive style to the robot (representing the *reactor* entity), as illustrated in Figure 7.13. When demonstration is finished a real (non-*broomstick*) robot enters the space and follows the person in the manner that was demonstrated (the *robot locomotion* interface). As both the demonstration and generation happen in the robot's space, this enables the person to become a part of the interaction rather than to externally control demonstration.

This robot-on-a-broomstick interface is a regular aluminium broomstick attached to our robot (Figure 7.13a) via a two-axis swivel, allowing the robot to be freely moved forward,

backward, left and right. The robot is the same as used during generation, enforcing all demonstrated movements to match those reproducible by the generating robot. Pushing or pulling on the broomstick moves the robot forward or backward, and the robot can be directly turned by twisting the broomstick, or by tilting the broomstick while moving the robot forward and backward. We have further installed two soft-press buttons (Figure 7.13b) on the broomstick as a means to enable the person to demonstrate to the robot when it should produce the pre-programmed sounds.

Our *broomstick* interface for demonstrating to a robot was constructed using a standard household broomstick. We trimmed the base, a small push-broom flat bottom type, to fit the holding bay of the iRobot Roomba Create, and fastened it with two screws (Figure 7.14a). We cut two holes in the broomstick and installed the two snap-in panel-mount soft pushbutton switches to enable the person to perform robot sounds (Figure 7.14c). To communicate these button presses to the controlling PC, we modified a Phidget remote wireless clicker by soldering our buttons to the clicker's terminals, and added a power switch to avoid battery drain when not in use (Figure 7.14d, Figure 7.14b). On our main PC, we used a standard Phidget interface kit to receive the signal. The robot was initially heavy to push, as movement



(a) a person is pushing a robot using the attached broomstick



(b) buttons on the broomstick to make robot sounds

Figure 7.13: a broomstick interface for directly demonstrating to a robot the style in which it should follow a person

torqued the Roomba's (un-powered) drive motors. To solve this issue we disconnected the motors from the wheels by removing the connecting belts.

We attached reflective markers on the robot and the person to allow the movements to be tracked using the Vicon camera motion-tracking system. The pattern of these markers was carefully chosen to make the four objects, the broomstick robot, autonomous robot, and the two slippers the person will wear, easy to differentiate. Figure 7.15 shows a view of this scene from the Vicon Nexus software, where the camera positions are shown in relation to the floor, and the four tracked objects are highlighted. The two smaller ones are the tracked slippers, while the two larger ones are the autonomous and broomstick robots, respectively.

Here we presented the interaction design and implementation for the *broomstick puppet master* interface. Below, we present the implementation details of the *robot locomotion* inter-




Figure 7.14: details of broomstick attachment to robot




Figure 7.15: a screen shot from the Vicon Nexus software, showing the camera positions and real-time tracking of marked objects

face for *stylistic locomotion*. We place the discussion here for structural reasons, to include it in the same section as the other implementation details, although this is not a *puppet master* (but rather is a *stylistic locomotion*) interface.

7.2.4.3 Robot Locomotion Project Implementation

Here we present the implementation details for the *robot locomotion* interface, where a person can interact with a robot that has an interactive, characteristic locomotion style. The style output itself is achieved using the robotic *puppet master* algorithm, detailed in Section 7.4.

For robot output we use an iRobot Roomba Create (Figure 7.16a). The robot is tracked using a Vicon camera motion-tracking system (discussed in detail in Section 7.2.4.2, page 189) by attaching reflective markers to the robot for the Vicon system to track (Figure 7.16b, Figure 7.17c). We used nine tripod and ceiling-mounted Vicon MX-F40 cameras to track robots in this interaction space (Figure 7.17b).

As the robot must interact with a person, the person who interacts with the robot also needs to be tracked. This is accomplished by having the person wear marker-augmented slippers to be tracked by the Vicon system (Figure 7.17d). The person's location is taken as the midpoint between their two feet, and their look direction is likewise the interior angle bisector of the look directions of their individual feet. The resulting path of the person's movement




Figure 7.16: iRobot Create robot with Bluetooth modem and reflectors


is smoothed to remove artifacts inherent from human gait such as apparent stopping and starting as they shift their weight between feet. The tracked locations are reported back to the controlling PC at 100 Hz for processing and planning in order generate the next robot action. This action is sent to the robot using a wireless Bluetooth connection (Figure 7.16b).

In this section we have detailed our various robot-based *puppet master* interfaces and implementations, as well as the implementation for the *robot locomotion* interface. In addition to leveraging the familiar concept of demonstrating to others, these interfaces further constrained the demonstration motion path to the particular movement properties of our actual robot — this was accomplished without requiring training to use the interfaces.

7.2.5 Conclusions on Puppet Master

In this section, we presented several interfaces which we designed or extended to enable a person to demonstrate to an entity (either an on-screen animated character or a real robot) how to interact with another entity (a second animated character or a real person) in a particular style: the *puppet master mouse GUI*, the *puppet master animation table*, the *Surface puppet master*, and the *broomstick* interface. We believe that these interfaces are the first that directly address the specific problem of demonstrating interactive locomotion style, both for animated entities and robots.

The core idea — teaching through direct demonstration — leverages *the social stock of knowledge* in that people are very familiar with teaching others by demonstration in everyday life. This further takes advantage of how robots elicit agency and have a social presence,



(a) space for interaction between a person and a robot denoted by the yellow tape



(b) ceiling-mounted vicon cameras



(c) reflective markers on person's slippers and the robot (markers artificially highlighted by camera flash)



(d) the slippers worn by the person with unique marker arrangements

Figure 7.17: the *locomotion robot* interaction space where a person can interact with a robot's characteristic locomotion paths

such that people can demonstrate to robots as they may demonstrate to another person, making the very idea of teaching to the robot familiar and easy to understand. Our work is still (intentionally) limited and transferring it beyond locomotion will pose new challenges. Could a robot learn an appropriate handshake from people? Eventually, could a family simply show a robot how it should act when at home, when at school, when at the mall, and so forth? Further, it will also be important to consider how a robot can determine *when* to learn the same as a person does, such that the teacher would not need to explicitly set a learning mode. Hopefully, these questions will be answered by the vast, active research area of programming by demonstration.

The interfaces we designed have been important for showing the feasibility of our concept of programming style by demonstration, as well as for how considering robots as social entities can facilitate and guide interface design. This exercise included the development of an original learning algorithm, the *puppet master* algorithm (Section 7.3, Section 7.4), and the reflection on social HRI questions via formal evaluations (Section 7.5, Section 7.6).

7.3 ANIMATION PUPPET MASTER ALGORITHM

In this section we describe the implementation of the *puppet master* algorithm, our original Style-by-Demonstration (SBD) algorithm for designing interactive, characteristic locomotion paths by demonstration. This algorithm also provides a means for enabling robots to communicate using *stylistic locomotion*, as it could use a behaviour previously created using *puppet master*. The *puppet master* algorithm is the core which made-possible all the interfaces presented in this chapter.

The benefits of the *puppet master* algorithm align with the motivations for developing it: it focuses on the style of an interactive locomotion path rather than task-oriented goal, it does not require a large or pre-processed training database, an average of 33 s training only in our study (Section 7.3), and it runs in real-time to handle unpredictable input.

The *puppet master* algorithm was first created for our animation-only interfaces, as a preliminary low-constraint arena to develop the details of the algorithm. Following, we adapted the algorithm to the real-world constraints and challenges of working with robots, as detailed later in Section 7.4.

7.3.1 Core Puppet Master Algorithm

Here we detail the core *puppet master* algorithm. Throughout this section we use the terms *main entity* and *reacting entity* to refer to the two roles respectively, consistent with our discussion earlier in this chapter.

The *Puppet master* algorithm is a pattern-matching algorithm that focuses on the relationship between the two entities, and how this relationship changes over time, for example, one entity may approach another, circle around, etc. The core idea behind the *puppet master* algorithm is motion style by analogy. That is, if first given a demonstration of the desired interaction style, then given new input the *reacting* entity should automatically respond to the *main* entity in a way that is *analogous* to the demonstration (as illustrated in Figure 7.7, page 181). Our approach and algorithm is based on the image analogies technique (Hertzmann, Jacobs, Oliver, Curless, and Salesin, 2001), which learns static image filters from example image pairs. However, it is only conceptually related to the Hertzmann et al. (2001) technique due to the differences of our target application, features of our data set, and the real-time needs of our application.

The *puppet master* algorithm has two discrete phases which we detail here: first it requires a demonstration of the desired interactive style, and next it attempts to generate a similar interactive style in real time. The key problems to be solved through the algorithm design are to a) develop techniques for extracting what a person would see as the behavioural *style* from the training, and b) develop a robust generation algorithm that can use the training data to direct output in a way that maintains the stylistic characteristics of the training. All of this, extraction and generation, must happen in real time.

Below we outline the features which we extract from training to represent style, and detail our algorithm for applying the appropriate training data to real-time interaction.

7.3.2 Features and Dataset

The *puppet master* algorithm focuses on two broad components of locomotion path to represent style: the relationship between the *main* and *reacting* entities, and the isolated detailed movement texture of the *reacting* entity. The first component is important for such situations as one entity following another, circling around, or aggressively blocking their way.

The second component is important for representing *how* the *reacting* entity does the actions, for example, moving slowly and smoothly, fast and aggressively, or shakily. To capture these components we extracted the following features from our training data (Figure 7.18).

VELOCITY — the magnitude of the vector between an entity's position and its previous one. This captures speed and acceleration-related aspects of behaviour such as different entity reactions for stopped, accelerating, or slow input.

RELATIVE POSITION — position of the *reactor* in relation to the *main* entity's position and look direction (coordinate space). This captures relational behaviour such as following, circling, and approaching. One scalar value per axis: representing how much the reacting character is behind or in front of, and to the left or right of the main character.

NORMALIZED LOOK DIRECTION — look direction normalized to movement direction, with 0 (radians) pointing forward. This is the relationship between where a character is looking and moving, e. g., if it is backing up or moving sideways.

RELATIVE LOOK DIRECTION — the difference between the entities' look directions. This captures turning away shyly when observed or aggressively facing an opponent.

Δ DIRECTION — change in direction from one step to the next, represents the shape of the locomotion path (not in relation to the other entity). This feature helps to identify similar movement shapes and styles such as shaky or smooth.

Throughout the development of the *puppet master* algorithm we explored many other features not listed here, for example, direct path data, distance and Δ distance between characters, absolute screen location, etc. For some of these, we found that the properties were redundant from our interaction goal perspective, and were already captured in existing




Figure 7.18: the features used by the *puppet master* algorithm

features, for example, the various distance properties are captured in the *relative position*. For others such as the absolute location, we found that the system generally works better for the more generic case, such that a trained behaviour would work regardless of orientation or position on screen.

As interaction happens over a period of time in the *puppet master* algorithm these features are stored in a time-dependent array, where at each point in time the above features can be calculated for the two entities. For demonstration, the result is a time-series array of features, and a second lookup-table array is created during generation that is constantly being appended to as real-time interaction happens.

7.3.3 *Algorithm Loop*

Here we formalize the main components of the *puppet master* algorithm. During the demonstration phase, no processing takes place and the movement paths of the two entities are simply stored, with the basic features (Section 7.3.2) calculated on the fly. The generation phase is much more involved.

Below we outline the general stages of the main system loop: at each stage, we observe the current situation, search the training data for the best match, and use the result to generate the next *reacting* entity output.

Loop

- * observe current real-time entity states
- * search training for most-similar situation
- * generate new movement from best match

In our implementation we have this loop running at 40 Hz. This rate was selected through trial and error for various reasons. We found that increasing beyond 40 Hz resulted in little perceptual gain in terms of the behaviour while harming performance, and dropping below this number quickly had decremental effects on the result: entities' movements became rougher, and we had less data available for smoothing (discussed below) hindering the overall quality of the result. However, 40 Hz did allow us to provide real-time reaction from the user perspective, and maintain high generated movement quality.

The intuition is that the *puppet master* algorithm continuously compares the current real-time entity situation with the training data, and the training most relevant to the current

situation is used to direct the generation of the next entity output. The result is a meshed mix of training data from various source locations. This method results in the *reacting* entity being able to react immediately to changing *main* entity actions. While observing the current entity states is trivial from the algorithm's perspective, both of the remaining phases are quite involved and are discussed separately below: first we discuss the similarity metric used to determine the most-similar situation, followed by our method for generating new output based on the training.

7.3.4 Similarity Metrics

The *puppet master* algorithm's similarity metric attempts to find the best-suited piece of training data to be used in output generation (Section 7.3.5). A core concept behind our similarity matching is that it has two parallel and competing metrics: *situational similarity* and *entity-path coherency*.

Our similarity metrics are calculated over a neighbourhood, or history, to encapsulate trends over time. That is, rather than simply looking at the entity relationships and configurations at a particular point in time, the recent history (1 s, or 40 samples in our implementations) is used as well. Looking at a neighbourhood enables the algorithm to capture tendencies over time, such as the reactor approaching or turning around the main entity, and so forth.

For both metrics, the numerical similarity between the window of current data and a given window of training data is calculated using the cumulative Euclidean distance squared between each corresponding pair of features. That is, since the windows are all of the same size they can be aligned and feature data at each point in one window is compared to the corresponding point in the other window.

Each feature set (Section 7.3.2) can be viewed as a multi-dimensional point, with each scalar-valued feature being a dimension. Then, the Euclidean distance (squared) can be taken between corresponding points as if they were plotted in multi-dimensional space; the smaller this value is the more similar the two features are (this can also be viewed as applying an L2 norm). These values are summed over the window to yield the overall similarity value. We use squared results simply as a mechanism to avoid expensive square-root calculations.

Below we outline the two competing metrics, followed by a description of our mechanism for balancing them and deciding which one wins the competition.

7.3.4.1 *Situational Similarity*

The *situational-similarity* metric is an attempt to capture relational actions and movements such as how the *reactor* is interacting with the *main* character, for example, how far behind or in front the reactor is, and whether it is following, circling, approaching, or turning away shyly, etc. As such it primarily uses relative features (*relative position, relative look direction*) and ignores global and absolute position and look direction. The exception is that the absolute velocities of the two characters are also used to consider speed- and acceleration-related aspects of behaviour, such as whether a character is moving slowly or quickly, stopped, or moving backward.

This metric searches the entire training dataset via a moving window over the fixed neighbourhood size, comparing at each point to the current real-time window on the features mentioned above, and the closest match is selected (Figure 7.19).

7.3.4.2 *Entity-path coherency*

The *entity-path coherency* metric focuses on the fidelity of the generated output in relation to the demonstrated movements of the *reactor* during training, for example, if the output path itself is like the trained path; this largely ignores the *main* character and the relation between the two. The intuition behind this metric is to try and ensure that the movements performed by generated output match the kinds of movements given in training, irrespective of how they relate to interaction with the *main* entity, for example, if the *reactor* generally moves in a smooth or jittered fashion, slowly or rapidly, etc.




Figure 7.19: *puppet master* algorithm situational-similarity matching compares recent real-time data to the entire dataset.

To represent the *reacting* entity's movements, this metric considers the relationship between where the entity is looking and moving (*normalized look direction*), for example, if it is backing up or moving sideways, and the general shape of the locomotion path (Δ *direction* of actual movement and *velocity*), whether it is turning, and how quickly it is moving. In addition to this some elements of the relationship to the *main* entity are used (*relative position* and *relative look direction*). This is important as some aspects of path-coherency are dependent on the other entity, such as when the *reacting* entity wants to finish a circle around the *main*, even though the main character is moving quickly.

Rather than searching the *reacting* entity's entire training data set (as with *situational similarity*) to find similar paths, this metric focuses on regions of training data recently used in generation. This considers how accurately the algorithm has been continuing a given patch of training, and if this similarity is good, then we can be confident that it is reasonable to continue that patch. This is a particularly good fall-back plan when the current *situational similarity* is weak. This is illustrated in Figure 7.20.

7.3.4.3 Similarity Metric Balancing

At each iteration of the *puppet master* algorithm, we need to have a method for combining the results from the above-two similarity metrics. As done in image analogies (Hertzmann et al., 2001), at a given step we choose one of the two metrics, alternating between them using a particular methodology. Image analogies statically weights the resulting similarity scores with a coefficient to add bias; the metric with the best weighted score is selected for that step.

This initial approach did not work with our application and resulted in a problem we call coherence loops: when *entity-path coherency* is used to generate output for several consecutive steps then the result of generation, by design, will be increasingly similar to the training data, and *entity-path coherency* returns increasingly strong scores. *Situational similarity* is




Figure 7.20: *puppet master* algorithm entity-path coherency matching examines the regions of the training data recently used in generation.

eventually ignored, and the *reacting* entity starts to loop through sections of training with no regard for the *main* entity. This issue does not occur in image analogies as all data is given at the beginning (the entire picture), allowing the use of look-ahead and multi-resolution approaches. Multi-resolution is difficult in our system, however, as we are generating in real time and cannot look ahead in our input data.

Our solution is to make the previously-static weighting coefficient dynamic, such that it follows a target-use ratio between the two metrics. The coefficient is automatically and continuously tuned (linearly) each step of the *puppet master* algorithm to bias the metric-use ratio toward the target (a similar algorithm is used in texture synthesis systems to match overall colour histogram Kopf, Fu, Cohen-Or, Deussen, Lischinski, and Wong, 2007). Using a dynamic bias (rather than a static bias) is important to still give the *puppet master* algorithm the flexibility to use the metrics based on the quality of their match, but tends toward an attempted balance between the two. In our implementation we use a 1:1 target ratio.

Rather than selecting one metric exclusively in each stage as above, we attempted to mesh the results from both metrics in several ways. For example, averaging the results from each for a unified output, or using *situational similarity* for inter-entity relational data and *entity-path coherency* for movement texture. We did not get satisfactory results with these approaches.

Once the best-matching similarity metric is selected, the training data immediately following the source region (Figure 7.19, Figure 7.20) is passed on to the generation system: we call this the target.

7.3.5 Output Generation

Once an appropriate target entity state is chosen from the demonstration data, we use it to generate the *reacting* entity's next movement. The naïve approach is to simply copy this data directly to the output (as in texture synthesis and image analogies Hertzmann et al., 2001). The problem for us, however, is that due to the particular features we use, it is often not possible to solve for a movement that matches all features. There is usually no single movement for the *reacting* entity that would satisfy the appropriate *relative position*, *relative look direction*, *movement velocity*, and Δ *direction* as shown in the demonstration. For example, moving the entity to the appropriate *relative position* may result in a *velocity* and Δ *direction* that is very different from the demonstration data.

Another challenge of output generation is noise. There is a great deal of noise inherent in the unpredictable input, which exacerbates problems of noise within the similarity metrics themselves as results can rapidly jump between grossly disjoint regions of training. This is further worsened by the similarity-balancing technique, where rapid changes between the two metrics can add extra noise. The generation algorithm must intelligently deal with noise while maintaining the style that was demonstrated to the system. In image analogies (Hertzmann et al., 2001), noise was largely dealt with by using multi-resolution approaches and look-ahead, where localized noise was anchored to lower-frequency trends in the image — this was not possible in our application.

Our original solution to this overall problem has two parts, a) we use various techniques for motion smoothing and b) we apply frequency analysis (Fourier analysis). We decompose the motion (and thus the style) into its low-frequency (intentional move to certain relational position) part and high-frequency (texture of the motion) part and treat them separately. While Fourier analysis has been done on motion paths before (e. g., Unuma, Anjyo, and Takeuchi, 1995), we believe we are the first to use it in terms of interactive style, emotion and personality. We present these phases below.

7.3.5.1 *General Trajectory Generation (low-frequency component)*

The aim of our general trajectory generation phase is to reproduce the low-frequency component of the demonstrated motion, i. e., the general, overarching movements and inter-entity relations. Our base technique is to construct a vector to move the *reacting* entity from its current location to the target relative position, and then scale the vector to match the target velocity (Figure 7.21), with the normalized look direction (orientation relative to movement direction) being copied directly to the output — this is a form of basic smoothing of the output path. Although this makes the entity move toward the target position rather than be at that position, the real velocity of the entity is a very important component of the output path and this approach creates very convincing results. Further, the high rate of generation (40 Hz) helps to avoid the entity lagging behind where it intends to be. This simplification also helps remove much of the noise from the movement, as the entity is constrained to reasonable (based on training) movement speeds.

After this process there still remains considerable noise in the output, which we filter by applying a simple linear smooth (average) over the world-coordinate movements of the entity




Figure 7.21: how the *puppet master* algorithm computes movement for general-trajectory

(three samples in our current implementation). This results in a more convincing, stable, and consistent generation.

We hit a problem with generating satisfactory output for the normalized look direction, a problem which we have related back to the nature of normalized look direction itself. That is, since this variable is tied to the movement direction then as the entity moves, rapid positional movements can result in rapid changes in normalized look direction — even if the person felt and intended to keep the entity’s look direction smooth and static. In particular, an entity moving rapidly forward and then backward has a normalized look direction which rapidly alternates between (in radians) 0 and π . Our solution to this is to limit rate of change of the actual world-coordinate look direction. This lowers the amount of noise in the resulting look direction, but some jitter remains. This is an important problem for future work.

The glaring problem with this solution so far is that by removing the high-frequency noise we also remove the high-frequency movement detail and texture. This means that intentional jitter and rapid movements given in the training data will not be produced in the output. Our solution to this problem is presented below.

7.3.5.2 Detail Incorporation

To restore detail removed by the above technique we first extract the original detail from the training data and apply it to the generation output in real time — the high-frequency data from the demonstration is used to perturb the generated (smooth) trajectory. The idea is to remove the high-frequency noise but re-introduce the lost high-frequency data.

We do this frequency analysis by using Haar wavelets (Haar, 1910), where we extract the high-frequency detail from the movement path of the demonstration data. We extract the path texture from the raw motion-direction feature, and incorporate the detail directly into the output using Haar wavelet recomposition, performed in real-time with no look-ahead.

A single application of the discrete Haar decomposition scales our path data to half resolution and stores the removed high-frequency detail separately. This gives a frequency cut at $f_s/2$ where f_s is the sampling rate (40 Hz in our implementation): given k cumulative decompositions, this cut is at $f_s/2^k$. The resulting k high-frequency datasets (one per decomposition) can be re-composed to form a single high-frequency-only signal that we use in our generation. Our system uses four-level Haar decompositions, a frequency cut of $f_s/16$, or about 2.5 samples/s. We found this to capture sufficient detail without having enough low-frequency components to affect general trajectory.

While the smoothing above compensated for the various instabilities in the matching system, the source demonstration region used to pull the high-frequency detail still suffers from the noise problem — the exact source training data used changes as dramatically as dictated by the change in the similarity matching from one step to the next, and interweaving detail from rapidly alternating training locations resulted in noisy output not coherent to the training. Our solution is to apply the extracted detail in patches (16 samples): a patch is used in subsequent steps until the end when a new patch is selected. These are 0.4 s long so any delay between changed behaviour and choosing a new appropriate and matching detail patch is minimal. The effect here is to remove noise in the source region which is used for detail, rather than to smooth the detail itself.

7.3.6 Initial Profiling and Appending Training

While extensive profiling was not performed on the *puppet master* algorithm, we found the performance adequate for our usage. We used the *puppet master* algorithm during an experiment where it was run on a Pentium 4, 3.0 GHz PC, and it maintained the target of 40 Hz for up to about 80 s of training data. A behaviour generally requires less, as outlined in our evaluation results (Section 7.5).

We added to the *puppet master* algorithm the ability to append training data to an existing training set. This worked simply by recording an additional training dataset, and during similarity matches the multiple training datasets were searched. While this worked fine in practise, we found in our pilot studies that people generally did not want to add to training — they wanted to start over. As such, we did not explore or evaluate this capability further.

In this section we explained in detail the workings of the *puppet master* algorithm, and how it enabled both the animation *stylistic locomotion* and *puppet master* SBD. In the following section, we detail how we applied this algorithm to robotic output.

7.4 ROBOTIC PUPPET MASTER ALGORITHM

Robots are imperfect physical machines that work on irregular surfaces, often cannot move as expected, and must adhere to real constraints such as movement speed or physical design. They cannot be moved directly as with animated characters, and we found that it is generally difficult for a robot to reach a target output provided by animation *puppet master* algorithm within a reasonable amount of time — the result of directly applying the animation algorithm was that the robot lagged far behind and the resulting style was not recognizable.

Given the effectiveness of the existing animation *puppet master* algorithm and the fine-tuning which was required to obtain it in the animation context, our approach was to use the *puppet master* algorithm as a canned unit and to integrate it into a broader robot-specific framework, rather than changing the *puppet master* algorithm itself. One change to the algorithm for the robotic implementation was that we reduced the generation frequency to 20 Hz, as the robot's slow response and movement times removed advantages gained from higher rates.

In the robotic framework we used a multi-component translation layer between the *puppet master* algorithm and the robot based on the use of a kinematic model of the robot's movement capabilities. We use this to perform a simple form of frequency analysis, where we directly drive the robot with movement texture (high frequency), and modify the movements to tend toward proper localization with the person (low frequency), components we describe below. For our discussion in this section we replace the animation *puppet master* algorithm's *main* and *reacting* entities with the robot and the person, respectively.

7.4.1 *Kinematic Model*

A kinematic model of a given robot's movements (an iRobot Roomba, in this case) enables us to solve how a given robot command, such as drive velocity and turning radius in the case of the Roomba, will translate into a real-world movement, for example, Δ position

and $\Delta\text{orientation}$. Inversely, given both a current and target robot location and orientation, this model can solve for the robot command (velocity, turning radius) that can best try to reach that target within a given time step. The simple mathematical model of the Roomba's movement is shown in Figure 7.22, where in addition to the direct labels, θ is the $\Delta\text{orientation}$, p is the starting position, and p' is the end position.

We use this model to perform frequency analysis as outlined in Figure 7.23. First, to accomplish the low-frequency component we pass the target from the *puppet master* algorithm to the model to receive the robot command to reach the target state as soon as possible (Figure 7.23b). However, with this the fine texture details of the motion (high-frequency component) are lost.

For detail, we noted that the target location received from the *puppet master* algorithm, over time, forms a textured path of how the robot should move (Figure 7.23c). If we calculate the delta of the target state at each time step and pass it through the robot kinematic model, we can get direct-drive robot commands that produce the delta (via robot velocity and turning radius) and as such the movement texture. This attempts to reproduce high-frequency detail and texture components of the target path. With this alone, however, robot localization (position relative to the person) is completely lost.

We combine the low-frequency (move to the target location / orientation) and high-frequency (detailed texture of motion) components (Figure 7.23d) by taking a weighted average on the two components, with a heavier weight on the detail. This is done directly on the robot-action level, on the raw velocity and turning radius robot commands. The data-flow process is illustrated in Figure 7.24. The intuition behind this balancing is that detailed movements toward the target state are generally unchanged while movements




Figure 7.22: kinematic model of the iRobot Roomba's movements

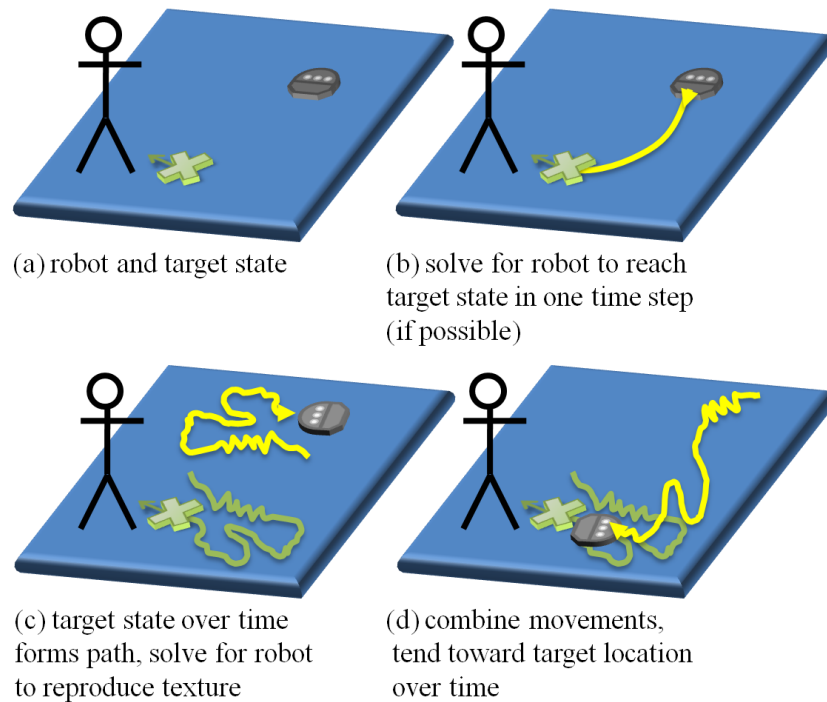


Figure 7.23: frequency-analysis algorithm for bringing a robot toward a given target state

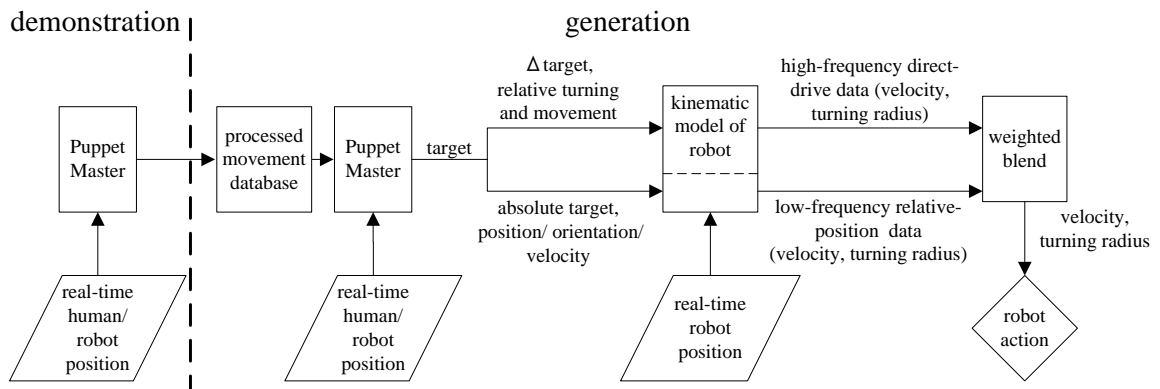


Figure 7.24: the overall process of how the *puppet master* algorithm is adapted to the robot application

which tend away from the target are dampened and altered to tend toward the target, while still maintaining much of the detail. Forward and backwards moving are handled inversely (turning toward or away from person, for example), and special-case handling helps to ensure that the bias does not destroy the texture, for example, to maintain appropriate forward or backward movement correctly.

We believe that our approach of using a robot-specific kinematic model, combined with a simple frequency analysis approach, is a unique method for programming a robot to reproduce an open-ended general motion path.

7.4.2 *Auxiliary Actions*

As part of moving the *puppet master* algorithm to robots we extended the original algorithm to enable the demonstration of discrete actions (i. e., robot sounds), which can be specified by the trainer in-situ during training. Technically, this was accomplished by adding two binary variables to the *puppet master* algorithm's feature set which signal when the actions were triggered. During generation, as particular training data is used in output construction, the associated action triggers are included in the target output and then performed by the robot.

Our current application, two discrete robot sound effects demonstrated through button presses, serves as a proof of concept that the *puppet master* algorithm is scalable to and can mesh with actions that are not derivative of motion paths.

In the last two sections we have provided a thorough account of the *puppet master* algorithm and its application to robots. This algorithm was a crucial enabling component of every interface in both the *stylistic locomotion* and *puppet master* projects, both major contributions of this dissertation. In addition, the *puppet master* algorithm, and its integration to the various interfaces, serve as important proofs-of-concepts for SBD. As such, we hope that this work will be extended to other SBD instances, for example, we see the simple inclusion of sounds in the *puppet master* algorithm as an important example regarding the scalability of the algorithm. The robot sounds could be replaced by any discrete pre-programmed robot action, for example, taking a picture or picking up an object.

In the remainder of this chapter we present our formal evaluations on both the animation and robot *stylistic locomotion* and *puppet master* systems.

7.5 STYLISTIC LOCOMOTION AND PUPPET MASTER STUDIES, ANIMATION

In this section we detail our evaluations on both the creation of an entity's characteristic, interactive locomotion paths (through SBD) and on how people perceive and interact with such entities. As both the *stylistic locomotion* and *puppet master* projects were developed in phases, we first present a study which focused only on the animated version, using the *animation table* interface (Section 7.1.2.2, Section 7.2.3.2). We discuss our study on the robotic variants in Section 7.6.

Our approach to the evaluation of our animation-targeted interfaces is two-fold. First, we evaluate our systems from a technical standpoint, including questions of feasibility and usability. For example, whether the system interface is usable by everyday people, whether our algorithms and implementations work as we expect them to, where the weaknesses of our algorithm and interface are, and how much (and what sorts) of style, emotions, and personality are captured by the *puppet master* algorithm. We also saw this as an opportunity to get people to interact with our entities that communicate using characteristic, interactive locomotion paths, and to ask questions such as: do people accept the idea of these entities communicating through locomotion path only? Do people *play along* with the idea of characters expressing personalities through motion path?

Related to these two perspectives, we performed the evaluation in two parts. In the first part, which we call the *designer* study, we asked participants to design new behaviours using our system. In the second part, which we call the *observer* study, participants only interacted with entity behaviours created in the first study; they did not create behaviours.

We initially conducted a pilot study to evaluate the study protocol and procedure. Five participants (two female, three male) joined the *designer* pilot and two participants (one male, one female) joined the *observer* pilot. These pilots exposed language and questionnaire wording that was confusing or strongly biased users toward particular responses, which we attempted to remedy for the full study.

7.5.1 Study Procedure and Methodology

Both studies consisted of a structured, pre-scripted protocol and various pre- and post-test questionnaires which assessed various aspects of the participant's background, ideas concerning robots, and their overall experiences during the study. All study documents and materials are attached in Appendix C.

The artist study explored how members of our general university population can use our system to create interactive behaviours. In one-hour sessions, participants were first asked to design five particular interactive character behaviours given the following keywords: *lover*, *bully*, *playful friend*, *stalker*, and *afraid*. Participants completed a short written survey about the result and experience after each behaviour. Following, we evaluated the internal validity of the design by loading the five behaviours each participant created in a scrambled order

(fixed across participants) and asking them to interact with, and recognize, each behaviour. Participants were not notified ahead of time that they would be revisiting their own designed behaviours.

The *observer* study considered how general people react to the behaviours created using our system, and whether a sense of character emotion and personality emerged. We subjectively selected five behaviours created by participants in the *designer* study (one per each of the five behaviour types), and participants were asked to “interact with and explore the characters” for each behaviour presented in a fixed order, and to “describe the character” in a questionnaire. Care was given to avoid anthropomorphic language when presenting the task to the participants, avoiding words such as “personality,” “behaviour,” and “emotion.” Following this process, participants were asked to interact with a set of “other” behaviours which were in fact a scrambled ordering of the same behaviours. This time the participants were asked to match each of the behaviours to the list of “correct” behaviours as given in the *designer* study.

The analysis of this study focused around exploration of participant comments, written answers, and verbal think-aloud results in addition to the measurements of behaviour matching success rates and time data. As such, we designed the questionnaires to use emotion and social-oriented Likert-like scales — with added room for comments — in addition to various long-answer written questions. Our analysis methodology consisted of exploring the result data for themes and points of interest, all of which emerged from the data and were not pre-determined. These analysis points were extracted from the data primarily via actual participant comments and Likert-like scale results, and included in the below results as directly as possible.

7.5.2 Observations

Twenty students (ten per study) from varying disciplines were selected from our university population and paid \$15 for participation. All people reported some to extensive programming experience and strong confidence with computers. In the *designer* study (two female, eight male), four participants reported artistic experience with three having formal training and one identifying themselves as an artist, and three participants reported basic animation experience. Ages ranged from 19 to 32 years ($M=22.8$, $SD=3.8$). In the *observer* study (four

female, six male), nine participants reported artistic experience with five identifying themselves as artists, and four people reported animation experience (two extensive). Ages ranged from 19 to 27 years ($M=23.7$, $SD=2.71$). All participants had no prior exposure to the system and no participants from the *designer* study took part in the *observer* study.

7.5.2.1 *Designer-Study Results*

Eight of the ten participants in the *designer* study identified 100% of their own behaviours. Further, in 74% of the cases participants agreed or somewhat agreed (using five-point Likert-like scales) they were satisfied with the resulting behaviour, and in 22% of the cases they neither agreed nor disagreed. The mean training time of accepted behaviours was 32.5 s ($SD=18.0$ s, $min=9$ s, $max=85$ s). The average number of trials required before accepting a behaviour was 1.7 ($SD=0.9$, $mode=1$ at $freq.=56\%$, $max=4$ trials). The average amount of time a participant spent testing a result before accepting it was 70.0 s ($SD=68.2$ s). In 46% of the cases participants disagreed that the generated behaviour felt mechanical with 26% neither agreeing nor disagreeing. In 48% of the cases participants agreed that the behaviour felt human-controlled (42% somewhat) with 26% neither agreeing nor disagreeing.

In the post-test questionnaire, on seven-point Likert-like scales, all 10 *designer* participants agreed (five strongly) that they enjoyed using the system, while seven disagreed that the system was frustrating to use (one strongly, two somewhat), all reported that the characters were fun to play with (six strongly, two somewhat) and six participants reported that movement jitter in the animation of the resulting entity behaviour was distracting. The two participants who failed to recognize their own designed behaviours were also the only two who did not use puck orientation during behaviour training, resulting in poor quality behaviours. This last result lead to us understanding a problem with our algorithm, as we have discussed in Section 7.3.5.1, page 203, third paragraph.

Four designer participants were notably and particularly immersed in the interface. Some made exaggerated faces, noises, and spoke to the characters while training. One artist used the “jaws” theme while training the “afraid” behaviour, and another commented “what a jerk!” when observing their “bully” character. Participants generally expressed excitement about and satisfaction with the capabilities of the system: “the system responded accurately and behaviour was smooth, human-like, with a human touch,” “it’s even a better stalker than I am!”, “it almost looks as if someone is controlling it,” “it did exactly as I wanted! Very

entertaining! (maybe it's just me?),” “nailed it!”, “I like it! I can see its bright future,” “the playful friend is a hoot!” Several participants commented on the robustness of the system, and one participant was excited that it “even reacted consistently with what [he] thought of after the fact.” Also, participants enjoyed the *animation table* system, finding it “super easy and intuitive to operate. Instant results.”

Several participants reported issues with the system, commenting on the resulting generation as well as the overall behavioural simplicity: “it felt a bit mechanical with some movements,” “as complexity of behaviour rises it feels more mechanical,” “if you pause to catch your breath, the system takes it as deliberate behaviour,” “I need to try more complicated behaviours,” “this setup cannot interpret smaller actions that well,” “he doesn't have hands so I can't punch,” “difficult to imagine what one pretty slime does to bully another pretty slime.” Further, six of the ten people had issues with occluding the Vicon markers on the controller puck — see Figure 7.9a, page 184. Several participants also verbally commented on the difficulty of concentrating on controlling and demonstrating to both the *main* and *reacting* characters at once, although this was not reported in the questionnaires.

7.5.2.2 *Observer-Study Results*

In the first part of the observer study participants were simply asked to interact with and describe prototype characters, and were not prodded to look at *behaviours* or *emotions*. At this point the participants completed a questionnaire which asked them to reflect on their experience. On a six point scale titled “the character felt...” ranging from “extremely mechanical” (1) to “somewhat mechanical” (3, 4) to “not mechanical at all” (6) the average across all behaviours was 4.04 (SD=1.19, Mode=4 at 36% frequency). On another scale ranging from “a human is controlling it” (1) to “somewhat lifelike” (3, 4) to “not lifelike at all” (6), the average response was 3.4 with a mode of 5 at 24%.

We analyzed the open questions to look for identification of each behaviour in respect to what it was originally trained for. Out of the fifty behaviours (five each across ten participants), behaviours were identified using the exact keywords used in the artist study nine times, and another ten times using very similar words (for example, “girlfriend” instead of “lover,” “naughty, trying to bug me” instead of “bully”). Out of the ten participants, who were not asked to look for behaviours or given information about them, we found that two did not identify any behaviours, two identified one behaviour, three identified two behaviours, one identified




Figure 7.25: a histogram showing how many behaviours were matched by how many people

three behaviours, and two participants identified four (Figure 7.25). Furthermore, in the open-ended questionnaires 52% of all behaviour descriptions were using social and behavioural descriptions (28% purely social), 34% of all the descriptions were using mechanical language (18% purely mechanical), with 14% being roughly a half-half mix.

For the second part of the *observer* study, participants matched the five behaviours against the original keywords used as shown in Table 7.1. Here, the diagonals show the number of participants (out of the 10) who matched correctly.

On the final questionnaires, four participants agreed that the characters were sometimes confusing (one somewhat), one neither agreed nor disagreed, and five disagreed (one strongly, one somewhat). One strong observation throughout the study was that participants tended to see social characteristics and used anthropomorphic language. For example, the observers mentioned: “the guy who kept sucker-punching,” “each one could bring to mind some real-life analogy,” “he needs more confidence,” “I liked the part when it came close to my character ... kind of like a dog who is happy to see you,” “He keeps trying to either hit you or kiss you,” “like an annoying kid brother in my face,” “he [the stalker] seemed like he wanted to

		Actual Trained Behavior				
		Lover	Bully	Playful Friend	Stalker	Afraid
Matched to	Lover	6	1	3	0	0
	Bully	0	5	4	0	1
	Playful Friend	4	3	2	1	0
	Stalker	0	0	1	6	3
	Afraid	0	1	0	3	6

Table 7.1: how people matched the behaviours to the original keywords used to train them

approach me, but he was too shy,” “facing it and watching it panic like it had been discovered somewhere where it shouldn’t be was fun,” “she [playful friend] is like a little sister who wants to talk to me.”

The participants were asked on the final questionnaire to describe the things they liked and disliked about each character. While some of these comments were analysis oriented, such as “actions were vague, subject to interpretation,” many of the comments referred to the participant’s opinion of the character’s personality. For example, for the afraid character (which stayed away from the participant’s character) one participant wrote “I didn’t really like anything, didn’t even give me a chance to get to know him,” and others complained that the character “tries to invade my personal space. I like a nice personal space bubble,” or “it doesn’t feel friendly!”

Similar to the *designer* study, some participants commented that the characters felt a bit fake when the jitter was too noticeable and several participants complained that the personalities were too simple: “the personalities were very blunt, they were easy to see,” “I wish they could touch each other.” All *observer* participants enjoyed the experiment (six strongly agreeing). Seven of the ten participants reported the pucks frustrating to use (all of these commented on how easy it was to occlude the Vicon markers), with the remaining 30% disagreeing or strongly disagreeing. However, several people commented that the *animation table* was “easy to use” and “intuitive.”

7.5.3 Discussion

The fact that 80% of the *designer* participants recognized 100% of their behaviours and were satisfied with the results suggests that the *puppet master* algorithm successfully supports some level of expression, and captures sufficient personality-related characteristics for recognition by the designer. That this was accomplished without prior experience or training at on average 32.5 s shows, even for outliers (e. g., the 85 s case), our algorithm enables people to create recognizable, characteristic, interactive behaviours very quickly. These results point out that people understood the core idea, and that the *puppet master* approach leverages people’s existing understanding of teaching, emotional style, and movement, to make a complex idea accessible.

Finally, that this was accomplished in on-average 1.7 attempts by novices supports the quality of the result and that people can satisfactorily and easily create behaviours. Further, this points to the importance of enabling rapid prototyping, as people did not mind to throw away their result and try again or try something new. How SBD supports rapid prototyping is an important question for future work.

The *observer* part of our study demonstrated that in 38% of the cases behaviours not only emerged but closely matched the artist keywords, based on motion only (Table 7.1). We believe that this supports our claim that our algorithm captures the personality and style of the demonstrated behaviour. Further, the results in Table 7.1 seem to hint at crosstalk between similar behaviours: for example, afraid and stalker are often mistaken for each other while lover, bully, and friend are rarely mistaken for stalker or afraid. This shows that, even in the cases where behaviours are not matched properly, there is still a strong component of feeling and style captured from the demonstrated data. To further understand this result it will be important to explore our choice of behaviours, and how they fit into peoples' understanding of emotion and personality.

Both studies suggest a strong sense of engagement. The explicitly-positive study results and the verbal excitement suggests that the participants were interested and mentally involved with the design process. Further, the extensive use of social and anthropomorphic language supports the idea that these characters, through locomotion path only, can create engaging and interesting personalities and characters.

7.5.4 Reflection

This study was important for our work for several reasons. These results were very encouraging in relation to the technical feasibility of teaching robots style by demonstrating to them, and for our particular approach taken with the *puppet master* algorithm. On the same token we learnt about the importance of certain issues which we thought were not important, such as jitter in the *reacting* entity's movement, and spent considerable effort to improve this for the robot implementation.

The feedback has also helped us to define a long-term agenda for SBD, where the participants illustrated the importance of eventually handling over-time evolving and more-complex behaviours, and also for moving beyond movement texture into other dimensions and action

(e. g., such as “punching”). This last point led us to explore the addition of discrete actions to the *puppet master* algorithm, as presented in Section 7.4.2.

While this interface design, implementation, and evaluation did not directly use robots, this work was particularly important for us in terms of designing and performing evaluations that target emotional, social, and personality-oriented layers of interaction: how engaged were the participants? How did they *feel* about the entities and did they attribute them with personality? We also confirmed that the very idea of teaching an entity stylistic, interactive behaviour made sense to people, and that it is something they can easily understand: training participants to create behaviours took less than a minute, after which they were able to create behaviours in on average 32.5 s. The SBD concept is much easier, quicker, and clearer to show and do than it is to describe using text in this dissertation.

Finally, this study served as our first landmark formal evaluation in relation to this dissertation work. Important gains for us here were in learning about the power of qualitative-oriented description and observation, and this study helped us understand the importance and the scale of what can be learnt simply by getting people to interact with our system and observing them. In the next section, we detail our evaluation of the robot versions of both *stylistic locomotion* and *puppet master*.

7.6 STYLISTIC LOCOMOTION AND PUPPET MASTER STUDIES, ROBOTS

In this section we detail our major multi-part study on the robot versions of the *stylistic locomotion* and *puppet master* interfaces. All study documents and materials are attached in Appendix D for reference. This effort includes three related studies: a programmer study, a designer study, and an observer study.

For the programmer study we recruited four experienced programmers to create robot behaviours using the Java programming language, and then to create the same behaviours using our *broomstick* SBD interface, and to reflect on their experiences with both. For the designer study 12 participants created behaviours using the *Surface puppet master* interface and 12 different participants used the *broomstick* interface, and reviewed the results through the *robot locomotion* interface. Finally, for the observer study 12 participants were recruited to observe several *robot locomotion* stylistic behaviours and comment on their experience.

Here we enumerate the overarching goals of these studies:

1. A practical evaluation of both the SBD approach for robots and the robotic *puppet master* algorithm.
2. To test the practical usability of both the *Surface puppet master* and *broomstick* demonstration interfaces, as well as how well they support SBD.
3. A comparison between programmed behaviours and those created using the *puppet master* SBD interfaces.
4. An expert-programmer design critique of our *broomstick* interface and SBD in comparison to traditional programming.
5. To create a situation where people are interacting with and observing robots. This gives us the opportunity to observe people's reactions, consider such questions as how they react to robots with interactive personalities and if they understand the personalities, as well as their attribution of agency, intentionality, and anthropomorphism.

We based these evaluations on four robot behaviours: a *polite follow*, a robot *stalking* a person, a robot that is *happy* to see the person, and a robot that is *attacking a burglar*; throughout the remainder of this section we refer to these shorthand as *polite*, *burglar*, *happy*, and *stalker*. These behaviours are a variant of those used in the animated study (Section 7.5), based on feedback from the animation study regarding task validity (e. g., participants asked “why am I doing this?”). We reduced the number of behaviours from six in the animation study to four to shorten the study and to accommodate the longer times potentially required to both train and observe the robots in comparison to animated entities, and renamed some (i. e., *playful friend* → *happy to see you*, *bully* → *attacking a burglar*) to better match a believable scenario with a robot.

We first performed statistical analysis of peoples' training time and their ability to match behaviours generated by the *puppet master* algorithm, tested across the two interface types. However, our main analysis approach in this study was to give detailed description of the participant interaction experience. We aimed to find themes in participant responses and represent the themes using direct quotes and questionnaire responses wherever possible. We explored applying statistical methods to understand the distributions and nuances of participant responses to socially-oriented questions, for example, through using non-parametric tests on Likert-like scale responses.

7.6.1 *Programmer Study*

To serve as a design critique of our *puppet master* SBD approach, we recruited four experienced programmers from our graduate-student lab at the University of Calgary to a) program robot behaviours using an Application Programming Interface (API) which we provided (detailed in Appendix D Section D.1, page 359) and b) create the same behaviours using our *broomstick* interface. We then interviewed the programmers regarding their experiences (unstructured, video-taped). We also video-taped the programmers' *broomstick* demonstrations, and later used this to extract task completion and time information (one participant requested not to be video taped).

The design of the API was intended to be simple and immediately usable, minimizing the learning required to program a robot behaviour. As part of the API we considered providing the transform of the robot's location in terms of the person's coordinate space (i. e., behind/in front, left/right, relative angle) as used by the *puppet master* algorithm itself, but decided against this as it may bias how the programmers approach the behaviour creation.

The programmers used an on-screen simulation of the real robot as a way to enable them to rapidly test and debug their ideas: Figure 7.26 shows the robot and a happy-face icon denoting the person; the X and red arrow can be used to manually tell the robot where to move in real time using the mouse as a means to test motions. We gave the programmers two hours to create the four behaviours.

When conducting the pilot study of the programmer condition, we considered fixing the length of the follow-up *broomstick* demonstration phase (e. g., to 60 s). However, immediate participant complaints regarding this resulted in us taking a more free-form approach for all subsequent demonstration studies, where participants could demonstrate for as little or as long as they wished.

7.6.1.1 *Programmer Study Results*

All programmers took the full two hours to create their behaviours, and all were able to create their behaviours in the time given, although one programmer stated that they would require a great deal more time to implement proper "nuanced behaviours." When programming by demonstration, all programmers were observed to "act the characters," making faces, laughing, etc. Table 7.2 shows how many attempts at demonstrating each behaviour each




Figure 7.26: on-screen simulation used for programming condition

programmer took and the length (in seconds) of the demonstration. The “total time” listed is the duration of creating all behaviours, from start to finish, including observation, thinking, brief discussion, and retraining time.

When asked which they preferred, direct programming or programming by demonstration, all programmers articulated a set of trade-offs rather than preference, for example, “the programming spoke to the scientist in me, and the other, the broomstick demonstration, spoke to the non-scientific part of me.”

The programming approach was touted as being more accurate and kept the person “in control” in comparison with the demonstration. Because of this, one person stated that they felt like they had “a lot more power to do something creative.” However, control is not easy or complete; one programmer noted that “when you’re programming something you have to anticipate ... what kind of situations can come up and how [the robot] should react ... that’s not a natural way of doing things.” The programmers made statements highlighting the difficulty of direct programming. For example: “hard to debug the program even though I have the simulated environment,” “even when I program I don’t know exactly what is going to happen”, and “when I see problems, I still don’t know why it happens.” Programmers mentioned that by focusing on style the “types of things [they were] trying to express were more nuanced, more complicated behaviours” than the “easily expressed things like sine

behaviour	programmer		
	1	2	3
polite follow			
tries	2	1	1
time	44 s	31 s	24 s
stalker			
tries	1	1	1
time	65 s	46 s	31 s
attacking a burglar			
tries	1	1	1
time	51 s	44 s	25 s
happy to see you			
tries	2	1	1
time	40 s	37 s	24 s
total time	14 m 49 s	8 m 52 s	7 m 40 s

Table 7.2: demonstration times for programmer condition; only three are shown (of four) as one participant requested not to be video taped, and we did not record the data directly

waves” that they are used to creating in interactive characters. One programmer noted that programming these behaviours first helped to highlight the sheer difficulty of the real-time problem, and helped them to appreciate the demonstration system.

Programmers noted that the *broomstick* is much faster and easier than direct coding. By using the demonstration system they “did not have to think technically or analytically,” and could more-easily program styles. As such “there is a huge time-saving potential here.” One reason cited for the *broomstick*’s success is that people are very skilled at “understanding changing situations on an instant-to-instant basis and [can] essentially make up [their] own behaviours on the fly.” However, several programmers pointed out that the learning by demonstration cannot be perfect as they are “at the mercy of the system” and their “demonstration is just a small part of the bigger thing.” They are “relying on its interpretations of [their] intentions, rather than on [their] actual intentions. There is no way to directly convey intentions.” As part of this they had no way to specify hard constraints such as “stay away from the corner.”

We asked the programmers to give an informal design critique of our *broomstick* interface. One programmer mentioned that the robot can be difficult to turn quickly; however, this

programmer tried to train the robot to turn much-more quickly than the real robot can perform. That is, although the interface is limited, the movement capabilities of the real robot are as well. Another, who has professional game development experience, said that although they would not consider using this as-is for something at the forefront of the game (such as a main character), they feel there is real potential for the approach for side-line characters. Another programmer pointed out that the inherent inaccuracy of the demonstration system is not necessarily a problem, as perhaps the *broomstick* can be used to capture basics and serve as a prototyping method for behaviours later programmed. This is related to suggestions on how to mix the pure demonstration approach with more logical components, for example, enable demonstrators to explicitly specify which components are important, or give them easy-to-understand parameters to tweak when observing the result.

Two of the programmers noted the amount of physical energy required to demonstrate using the *broomstick* (in the large space), that it was more physically exhausting than the act of programming. They mused about the use of a remote controller to reduce the fatigue problem, although they both admitted the result would likely be more difficult to use and control than the *broomstick*. One participant suggested a tabletop system as way to keep direct movement while lowering the effort required (our *Surface puppet master* system was not implemented at this point).

We evaluated the resulting programmed behaviours by subjectively selecting one of each behaviour type from the entire experiment, and including it in the *observation* study.

7.6.1.2 Programmer Study Discussion

The results of this study help support the idea and approach of SBD, and that it is applicable even for expert and experienced programmers. As outlined in Table 7.2, all participants managed to complete the creation and evaluation process in less than 15 m, substantially less than the two hours taken for programming.

What we found interesting beyond the time-efficiency results is how readily the programmers acted and *got into the characters*, laughing and making facial expressions to match what they were demonstrating. Even for scientists and engineers who *do* have an understanding of the technical nature of the robot, our results help support the idea that programming style to robots via demonstration leverages their innate social understanding and skills, and as such, they readily accept and embrace the approach.

One of the benefits of this study was that the participants have both a technical understanding of the problem, and a *social stock of knowledge* that makes the demonstration familiar and comfortable; therefore, their comparisons between the two techniques are particularly informative. The programmers feedback added depth to our understanding of the accuracy / control versus time / ease trade-off. This includes such observations that programming enables them to be creative in ways that demonstrating does not, and that the complexity of the programming approach means there is still a layer of uncertainty and mystery despite the extra control. Further, many ideas were proposed on how to combine the programming (more control) and demonstration (easier to do) approaches, for example, by using the demonstration as a rapid prototyping tool, or by including easy-to-understand parameters or conditions which the demonstrator could specify or tweak.

Regarding the *broomstick* interface, the results brought the fatigue issue to the forefront, and the tabletop suggestion in particular reinforced our development of the remote tabletop version of this system; chronologically, the *broomstick* was developed and this study conducted simultaneously with development of the *Surface puppet master*.

Many of the participants' observations help us to better-understand the limitations of demonstration, for example, that there is potentially no optimal solution as machines cannot understand a person's intentions, only their actions, and this interpretation is subjective to the demonstration-learning algorithm used. We point out, however, that people suffer from the same problem: people cannot know others' intentions, only what can be deduce from interaction. Regardless, this suggests that we should aim to better understand the particular biases introduced by any given algorithm, and how this relates to target applications and usage scenarios.

7.6.2 *Designer Study*

The designer study had participants use our *puppet master* interfaces to demonstrate interactive, stylistic robotic behaviours. The entire experiment protocol, summarized below, is given in full detail in Appendix D, Section D.3 and Section D.4.

7.6.2.1 *Designer Study Design*

There were two independent variables in this study. One was the interface used, *Surface puppet master* or *broomstick*, manipulated between subjects (half did each). The second was the behaviour which the participant created (within subjects), *polite*, *burglar*, *happy*, and *stalker*, such that each participant created all behaviours. Participants were randomly assigned to an interface condition, and behaviour order was fixed across participants.

We recruited twenty-four participants from the general university population as behaviour designers, split evenly across the between-subjects condition, and were paid \$20 CAD for their time (the experiment took roughly one hour). In the *Surface puppet master* condition, one of the twelve participants was excluded from the results due to a time complication. We did not limit the time of the phases in this experiment, and this particular participant arrived late and took a lengthy time to complete the initial questionnaires: they did not complete even one behaviour before we had to move to the next participant. One participant's data was also excluded from the *broomstick* condition. In this case, it was clear that due to a language barrier the participant did not understand the instructions, despite their (repeated) verbal affirmations. This was clear in how the participant approached the experiment and was also echoed through the form and content of their written responses in the questionnaires. Of the remaining 22 participants, 11 were female (11 male), aged 19–34 ($M=26.9$).

7.6.2.2 *Designer Study Tasks*

There were two tasks in this study: behaviour demonstration, and matching of created behaviours. Each participant was asked to demonstrate the given behaviours with their given interface. After each demonstration, the participant would observe the resulting generated behaviour on the *robot locomotion* interface (Section 7.1.3), where the experimenter walked around the space and the robot interacted with the experimenter based on the demonstration. The participant could choose to continue on to demonstrate the next behaviour or re-train the current behaviour if they were not satisfied.

The second task involved presenting the participant with the behaviours they just created (shuffled-order, shuffling fixed across participants and conditions). Participants could watch each behaviour for as long as they wanted, but could not return to a previous behaviour once they moved forward, and were asked to match the observed behaviour to the ones they created (via a sheet with the behaviour descriptions).

The paths that the experimenter throughout the phases were pre-designed to incorporate both short and long segments and turns. Further, we used generation paths different enough from the demonstration path so that the robot was not expected to simply replay the trained behaviour verbatim. We used one path for demonstrating the *polite* and *stalker* behaviours, those based on the robot following the person, and a different path for the generation phase of these behaviours. Yet another path was used for demonstrating the *burglar* and *happy* behaviours, as they are based on more-general (not necessarily following) interaction, with a fourth walk path used for the generation phase of these. Finally, a fifth path was used for the final random observe stage. These paths are given in detail in Appendix D, Section D.2. The demonstration paths were simulated on the table using the happy face, and the observer paths were walked by an experimenter with the real robot.

7.6.2.3 Designer Study Procedure

We conducted this study using a structured protocol including the use of informed consent forms and questionnaires. Participants first completed pre-test questionnaires which enquired about demographics and predisposition toward robots, artistic experience, and general technical ability. Before starting the study we gave an example of how demonstration works, and allowed the participant to try the physical interface (*broomstick* or *Surface puppet master*).

The participants completed the demonstration task for each behaviour, followed by a questionnaire that enquired about their experience and satisfaction with the interface in relation to that behaviour. After all behaviours were created the participants performed the matching task. Finally, the post-test questionnaire enquired about the overall experience.

7.6.2.4 Designer Study Quantitative Results

TRAINING TIME — The grand mean of training time across all cases and both conditions was 50 s (SD=37 s, min=4 s, max=261 s). We applied a mixed-design ANalysis of VAriance (ANOVA) (within-participants behaviour type [polite, burglar, happy, stalker] × between-participants interface type [*Surface puppet master*, *broomstick*]), applying a logarithmic transform to the time data to improve the normality of the distribution. Mauchly's test indicated that the assumption of sphericity had likely been violated for the main effect of behaviour, $\chi^2(5)=14.43$, $p=.013$, therefore degrees of freedom were corrected using

Greenhouse-Geisser estimates of sphericity ($\epsilon=.66$). No significant main effect of behaviour was found ($F(2, 39.5)=3.02, p=.061$), findings echoed by non-significant repeated-measures contrasts ($p>.2$). There was further no effect found of interface type ($F(1,20)=2.42, p=.136$) on training time, and no interaction effect found between behaviour and interface type ($F(3, 60)=1.42, p=.247$).

OBSERVATION TIME — The grand mean of the observation time, how long the participant observed the result of their creation before moving on, was 115 s (SD=68 s, min=24 s, max=450 s). We applied a mixed-design ANOVA of the same form as above, with logarithmic transform on the time data. The results show a main effect for behaviour ($F(3,60)=6.29, p=.001$), with repeated-measures contrasts indicating an effect between the first and second behaviours trained (*polite* and *stalker*, $F(1,20)=8.00, p=.010$). As the conditions' orders were not counter-balanced, this is perhaps a learning effect. The results further suggest no main effect on time for interface type ($F(1,20)=1.08, p=.311$), and no interaction effect on completion time between behaviour and interface type ($F(3,60)=.13, p=.945$).

PER-BEHAVIOUR QUESTIONNAIRE — We present the overall summary results of the per-behaviour (post-training) questionnaire as a frequency table in Table 7.3. Applying Friedman's ANOVA (as we do not assume normality of these distributions) failed to expose significant effect of behaviour type on participants' opinions on whether it "makes sense to teach robots this behaviour by demonstration" ($\chi^2(3)=6.26, p=.100$), or for whether participants were "satisfied with the result" (and $\chi^2(3)=7.10, p=.069$).

Table 7.4a presents a per-behaviour breakdown of responses to "the resulting behaviour felt overly mechanical." Friedman's ANOVA shows a significant effect of behaviour type on the scores ($\chi^2(3)=16.43, p<0.001$). *Post hoc* pair-wise Wilcoxon Signed Ranks tests, with a Bonferroni correction (effects considered significant at $p=.008$), indicated that the average ranks (relative score per behaviour) were significantly different: participants ranked their polite as more mechanical than both their burglar ($Z=-3.09, p=.002, r=-0.66$) and their happy ($Z=-3.21, p=0.001, r=-0.68$). No significantly-consistent relationships were found for the rankings of the remaining relationships: stalker-polite ($Z=-1.43, p=.153, r=-0.30$), burglar-stalker ($Z=-1.58, p=.114, r=-.34$), happy-stalker ($Z=-1.10, p=.272, r=-.23$) and happy-burglar ($Z=-0.877, p=.380, r=-0.19$).

	strongly disagree	disagree	neither agree nor disagree	agree	strongly agree
	1	2	3	4	5
you were satisfied with how well the system captured the behavior you were trying to demonstrate	0	8	10	54	16
the resulting robot behavior felt overly mechanical	4	17	36	27	4
the resulting robot behavior felt natural, organic, possibly human controlled	2	20	31	26	9
I think it makes sense to teach robots this behavior by demonstration	0	6	15	43	24

Table 7.3: table of answers to per-behaviour post-training questions

	strongly disagree	disagree	neither agree nor disagree	agree	strongly agree
	1	2	3	4	5
polite follow	0	3	6	12	1
stalker	0	4	10	7	1
burglar	4	3	10	4	1
happy to see you	0	7	10	4	1

(a) the resulting robot behaviour felt overly mechanical

	strongly disagree	disagree	neither agree nor disagree	agree	strongly agree
	1	2	3	4	5
polite follow	2	7	6	5	2
stalker	0	7	7	7	1
burglar	0	4	7	8	3
happy to see you	0	2	11	6	3

(b) the resulting robot behavior felt natural, organic, possibly human-controlled

Table 7.4: breakdown tables for two per-behaviour questions

Table 7.4b shows a per-behaviour breakdown of responses to “the resulting robot behaviour felt natural...”. Friedman’s ANOVA reports a significant effect of behaviour type on the scores ($\chi^2(3)=9.51, p=.023$). Pair-wise *post hoc* Wilcoxon Signed Ranks tests (using Bonferroni correction for significance at $p=.008$) did not reveal any further details regarding this difference: stalker-polite ($Z=-0.88, p=.382, r=-0.19$), burglar-polite ($Z=-2.29, p=.022, r=-0.49$),

happy-polite ($Z=-2.27, p=.023, r=-.48$), burglar-stalker ($Z=-1.27, p=.123, r=-0.27$), happy-stalker ($Z=-1.54, p=.123, r=-0.33$), or happy-burglar ($Z=-0.07, p=.943, r=-.02$).

BEHAVIOUR MATCHING — Participants' attempts to match their own behaviours (post-demonstration, scrambled order as explained above) are shown in Table 7.5: overall average matching-success rate is 67%, $SD=.37$, $max=100\%$, $min=0\%$. The mode at 11 (half of participants) is 100% correct. We applied a Friedman's ANOVA on the binary result of whether a given behaviour was matched correctly, and did not find any significant effect of behaviour on correct identifications ($\chi^2(3)=2.61, p=.455$). Mann-Whitney Tests (non-parametric test on two independent samples) did not find a significant effect of interface type ($U=48, Z=-.90, p=.438, r=-.19$) on a participant's ability to successfully identify their own behaviours.

POST-TEST QUESTIONNAIRE — We present a summary frequency table of the post-test questionnaire results in Table 7.6. Mann-Whitney tests did not reveal an effect of interface type on the participants' overall enjoyment ($U=61, Z=0, p=1$), disappointment ($U=56.5, Z=-.27, p=.797, r=-.06$), frustration ($U=51, Z=-.72, p=.519, r=-.15$) or that the results were as intended ($U=58, Z=-.22, p=.847, r=-.05$). These results were also echoed by visual analysis of scatter plots.

7.6.2.5 *Characterization of Designer Study Participants*

Here we summarize the characterization of our participants, based on our enquiries regarding their background. While we did not control for the below factors in our experiment design (technical ability, programming experience, prior robot experience, artistic background, disposition toward robots, sex), we used this experiment as an opportunity to explore potential effects which may exist — this could be used to direct future, formal, studies.

	polite follower	stalker	attacking a burglar	happy to see you
polite follower	14	5	1	2
stalker	6	14	0	2
attacking a burglar	1	0	17	4
happy to see you	1	3	4	14

Table 7.5: frequency table of participants matching own behaviours, the diagonal is the correct match

	strongly disagree 1	disagree 2	somewhat disagree 3	no opinion 4	somewhat agree 5	agree 6	strongly agree 7
I enjoyed the overall demonstration and observation of behaviors	0	0	0	0	0	10	12
I was often disappointed by the resulting robot behaviors	2	6	6	2	5	1	0
The process of demonstrating to a robot was often frustrating	4	12	5	1	0	0	0
The resulting robot behaviors conveyed the style and personality I intended	0	0	0	0	10	11	1

Table 7.6: frequency table of post-test questionnaire answers

TECHNICAL COMPUTER ABILITY AND PROGRAMMING EXPERIENCE — Participants self-reported their “technical computer ability” on a scale from one (“absolutely none”) to four (“I can install new software”) to seven (“I have a technical degree”). The results are: rating of 2 (1 participant, 4%), rating of 3 (1, 4%), rating of 4 (3, 14%), rating of 5 (3, 14%), rating of 6 (7, 32%), rating of 7 (7, 32%). Further, written responses to questions of programming experience were coded into none (5, 23%), some (4, 18%), and extensive (13, 59%). Experimenters noted that participants with an engineering background appeared to train for longer. However, informal analysis of the data did not reveal any correlation between self-reported technical ability and our observations.

PRIOR EXPERIENCE WITH ROBOTS — We coded participants’ self-reported prior experience with robots into none (15, 68%), experience interacting with (1, 5%), and experience creating or programming (6, 27%). Figure 7.27 highlights a potential relationship between prior experience and mean training time. With informal analysis we did not find any other correlation of prior experience on our measurements.

ARTISTIC BACKGROUND — We coded participants’ self-reported artistic background into none (15, 68%), amateur or hobby (5, 23%), and professional experience (2, 9%). Informal analysis did not reveal any correlation between artistic background and our observations.

DISPOSITION TOWARD ROBOTS — We coded participants’ disposition toward robots and their role in our futures into negative, neutral, and positive, for both the pre- and post-test

questionnaires. Pre-test: negative (4, 18%), neutral (8, 36%), positive (10, 46%). Post-test responses on disposition were: negative (2, 9%), neutral (4, 18%), positive (16, 73%). Visual inspection via frequency tables suggests a small movement toward more-positive disposition after working with our robots.

SEX — Informal analysis did not reveal any correlation between sex and our observations.

7.6.2.6 Designer Study Qualitative Results

ATTITUDES TOWARD SOCIAL INTERACTION WITH ROBOTS — Regarding the overall idea behind our system, one participant expressed wishes that “more studies like this were conducted, because some of existing robots really lack in the human interface quality.” Others said that this “is what the future is looking for. Robots need to be trained quickly and easily if we are to implement them in every home,” “if properly executed it makes life easier to human in many respects,” and as “each person’s interpretation of aggressive would be different, it wouldn’t make sense to pre-program the behaviour.” One participant said that, in particular, “when the instructions are ambiguous (e. g., what is ‘excited’ anyway?) it’s a good idea.” However, some participants had general doubts about the approach such as being “not sure ‘social’ robots will be very important,” and one participant said that they are “not looking




Figure 7.27: experience with robots versus average training time

for a pal. I am looking for something to do the things I don't want to do," or "I don't think I would worry about 'personality,' ... I don't see it as a person." Some expressed uneasiness, for example, "it would be scary when its too natural and gets a 'soul,'" and "I'm kind of scared that robots can control humans someday." When asked directly if they would care about the personality and style of the robot, we received positive responses from eight of the eleven respondents in the *broomstick* case and seven of the eleven in the *Surface puppet master* case. Surprisingly, we did not receive any direct negative responses to the idea of training a robot.

POSITIVE RESPONSE AND ENGAGEMENT — Participant feedback given both spoken and in written answers was generally positive. Comments ranged from basic descriptives such as "robot copies demonstrated behaviour well," or "the robot followed the demonstration pretty accurately," to more-enthusiastic "I am amazed!", "the robot did a good job in this case, better than I tried to do," "I think it was even more aggressive than me!" Some expressed pride and joy at being able to create robot behaviours: "when I saw all of them at the end, I was quite happy," and one person particularly liked their *happy* behaviour, exclaiming that "it looked like a little dog and his owner!", and another thought that "it [the happy behaviour] can welcome to the visitors while entering to home." We generally found that this kind of engagement increased throughout the experiment (perhaps due to comfort and learning); some participants believed that the robot improved throughout the experiment, although technically nothing changed in the algorithm. We noticed that several people who expressed negativity toward robots, still used anthropomorphic language when discussing the robot.

BEHAVIOURS — Comments surfaced regarding the selection and range of behaviours. Some participants noted that "it was easy to identify distinct tasks but hard to identify similar tasks, i. e., excited versus aggressive." A few participants expressed reservations about having, for example, *happy* robots: "I don't see a point in teaching this behaviour, but if you must, this method works." One person noted that a "simultaneous sorrow effect should work because people are not always happy."

Some people related to the length of training time, saying that the given behaviours "feel too complex to learn in a short period of time," and felt that "the robots can learn [the] behaviour with more coherent training," although no training-time-versus-quality trade-

off emerged from the actual study. Some participants mused that we may be using neural networks, and therefore suggested that a longer training time is better.

GENERAL LIMITATIONS OF SYSTEM — One theme of comments surrounded the limitations of the robot's capabilities, for example, that the robot itself was too slow. Some felt too restricted by both the robot's limited movement capabilities and the focus on locomotion path as, for example, "a dog would run around, jump, move its tail and follow its owner," or by the fact that the robot has no knowledge of the space boundaries; participants often asked about this. Others noted that the robot "can only do what is shown to it, no creativity," a concern which manifested in practical issues such as "the robot reproduced too many details and not just the general idea," or "at the beginning, I had a hard time with moving my robot, the robot reproduced that as well." One participant concluded that this "requires the teacher to be a *good* teacher," and another gave the example of "my definition of 'excited' resulted in high velocities which the robot was unable to reproduce."

ROBOT-HUMAN COLLISIONS — The issue of robot-human collisions was very prominent, and participants hinted at questions of responsibility, i. e., "what if the robot doesn't learn? Will it hit/harm the human/itself?" In response to this several participants stated that "the robot needs some underlying assumptions like 'don't drive over person' [that] should take priority over demonstrations." Some suggested that there are safer ways to be aggressive, for example, "maybe aggressive motions would be too much of a threat/danger, but the angry sounds would be good."

PROBLEM OF MOVEMENT JITTER — The most prolific complaint voiced is that the robot's movements "seemed really jerky" and were often "too abrupt." Some explained this using technical descriptive language such as "the robot made too many turns back and forth instead of just turning the right amount." Most participants, however, used behaviour-oriented or emotive language to explain the jerkiness: some felt that the "robot seemed a bit confused," and noted that "[the robot] stops frequently and becomes confused about which way to go next," or "the robot looked like [sic] thinking and deciding." One person stated: "if I didn't know the behaviours the robot is mimicking, I would say it's trying to act confused." Others thought the jerkiness showed that "the robot was being repeatedly confused, meaning it

did not fully grasp the behaviour” and it “did not fully understand what to do,” relating the problem back to learning ability. Some, however, tried to rationalize the utility of the jerkiness, for example, “I guess an angry robot can be jerky.” One participant suggested that “[the robot] might need more sensors to stop seeming confused.”

The jitter problem was found to be tied to participants’ interpretations of whether they found the robot’s actions as *mechanical* or *natural, organic*. Some noted that “the behaviour felt relatively mechanical mainly because of the rapid changes of direction and speed,” while others highlighted that this may be a localized problem: “only the jittering part was not natural,” and “although jittery, it followed in a very human way.” Further on the *mechanical* point, several participants reflected beyond the interactive movement path: “I felt more mechanical characters, just the movement was natural,” and “if I didn’t hear the noise from the robot, it would more natural.” This had a direct impact on how the characters were received, for example, “the robot is loud, so the stalker effect is less noticeable.”

TAKING PERSONAL RESPONSIBILITY — Many took personal responsibility for generation problems. Example comments include: “I think I could do it better. The robot learn perfectly what I did,” “maybe it was my fault as a demonstrator,” or “not very satisfied. I think reason was not very efficient demonstration.” This also applied to the jitter problem: “I was a little too shaky when I showed the robot behaviour,” “it was jerky, but then again, I was moving quickly during training.”

ROBOTIC SOUND — Several participants expressed regret that they “totally forgot about sound!” and expressed the desire to use it. Other participants claimed that “the robot sounds were necessary to understand the behaviours,” and that, for example, “it gives the users a much better chance of distinguishing polite follow and happy to see you,” with one person stating that “it’s the only way to tell the difference between happy or angry.” Several participants commented on the importance of sound to the overall character, saying that “it shows emotive responses,” “it brought a human dimension to the experience,” and that “if we can see the expressions it’s important to hear them.” One participant pointed out how sounds can have well-established meaning to people, such as “danger’ or ‘alarm.” Several asked for additional sounds, for example, “for the burglar an angry sound would have been useful.”

REFLECTIONS ON THE *BROOMSTICK* — Most participants expressed positive responses toward the *broomstick* interface, for example, “it was handy and helpful in movements.” One noted that “a body suit for the performer” may be better, but conceded that “regarding the type of student robot, [this] was the best way of control.” Some participants expressed difficulty, such as “it was a little bit hard to demonstrate using a broomstick,” for example, “it would be hard for me to make the robot rotate on the spot.” Two participants in particular raised concerns of how the properties of the *broomstick* influence training and thus results, as “the broomstick causes the movement to be a certain way,” “which will add some inconsistency to the robot’s behaviour, making it indistinguishable whether the robot has learnt the poorly-performed behaviour or did not learn appropriately.” One person expressed confusion over the imagery of the robot (an iRobot Roomba vacuum) and the *broomstick*: both have a strong image of cleaning, which was not a part of the experiment.

REFLECTIONS ON THE *SURFACE PUPPET MASTER* — The primary complaint regarding the *Surface puppet master* interface was the physical size of the table: people thought that “the task space should be larger / puck be smaller,” and that “simulator icons are very big compared to the area size. Hard to move ‘robot’ smoothly without driving into ‘person.’” One participant raised a concern that “a puck on a 2D screen does not represent very well the space that a person takes.” Given the task of training a robot, one participant felt that the table is “probably a good idea because it ... gives the robot some time to grasp the technique,” suggesting that the robot needed time to learn (there was a time delay between when the person demonstrated on the tabletop and observed the results, as they had to change locations). One participant (who was not aware of the *broomstick* interface) asked if they could “maybe demonstrate to the robot visually? I feel that demonstrating or teaching the robot directly may be more effective.” The small-space concern mentioned above manifested in the *locomotion robot* case, for example, people felt that the “rectangle is a bit small to see if the robot did follow the person politely.” Further, several requested to “include barriers and items inside the space,” particularly as “the stalker effect isn’t noticeable without barriers and obstructions.”

In this section we presented the results from our robotic *puppet master* designer study. Below we discuss the implications of these results, first starting with general discussion, and brief discussions on the results pertaining to particular interfaces, and our study methodology.

7.6.2.7 *Designer Study Discussion — Overall Reflections*

GENERAL SUCCESS OF STYLE BY DEMONSTRATION — Our results support the idea that the core SBD idea makes sense to participants, that they understand and accept the idea of teaching stylistic behaviours to robots. No participants were observed having or reported problems demonstrating the behaviours or understanding what to do, there were no explicit complaints regarding the teaching, and responses to direct questions of whether the idea makes sense were positive. Some directly applauded the idea of customization, and expressed that they were acutely aware of the importance of individual interpretations and meanings behind stylistic behaviours. Further, many related the teaching action to what they understand from teaching living things, for example, that more time teaching is better, or that confusion is a product of not understanding the learning. That we informally found no correlation between technical ability, programming experience, and artistic ability and any factor we tested for, including training time and success rate, suggests that not only is the system accessible to non-experts, but that perhaps being better-trained in a relevant area may not influence the use or results of our system; our SBD implementation leverages the more generally-accessible *social stock of knowledge*.

Our results also reflect on the *puppet master* algorithm, that it enables people to quickly ($M=50$ s) create stylistic, interactive robot behaviours that are reasonably identifiable, with half of the participants perfectly matching their behaviours (67% overall match success). Participants also generally expressed satisfaction with the results, both in their written answers and Likert-like questionnaires.

PEOPLE UNDERSTAND TEACHING — It became clear that people are adept at casual teaching and understand the intricacies and complexities behind it. Some demonstrated that they were aware of core weaknesses: they knew that the robot cannot distinguish between intention and actions (without detailed verbal explanation) and has “no creativity,” the properties of the training system (*broomstick* or *Surface puppet master*) has an impact on the training style, there are implications to the limitation of locomotion path and two simple sounds, and the approach requires the person to be a good teacher that fits the robot’s needs. These points raise important questions regarding if SBD is for everyone, i. e., can anyone be or does everyone want to be a teacher?

PEOPLE ANTHROPOMORPHISED — Although some participants expressed resistance to the idea of robots having emotions or personalities, the same participants were observed to readily attribute the robot with anthropomorphic qualities and used intentionality and agency to explain the robot's actions; it appears that they naturally tended toward this despite their predispositions. This illustrated potential parallels to Reeves and Nass's (1996)'s Media Equation, where people were found to treat media as living things despite their often explicit and adamant opposition to the idea.

JITTER PERCEIVED AS CONFUSION — The degree to which people interpreted the jitter as a robot personality trait was both surprising and not obvious, for example, being interpreted as the robot changing its mind or being confused about what to do. We find it serendipitous that underlying uncertainty in the algorithm, manifested as robot jitter, was accurately and naturally interpreted as confusion. This finding supports the idea that people apply *the social stock of knowledge* to understand the robot (and in this case agency and intentionality as well).

PERSONAL ATTACHMENT TO BEHAVIOURS — The excitement and pride that people showed regarding their creations was a pleasant surprise, and raised questions regarding the impact of enabling people to customize their robot. For example, can this affect pride, attachment, usage, or perceived robot success, or alleviate issues of fear, worry, or unease with the robot. There is also the question of the participant taking responsibility for the resulting actions, having implications on, for example, how forgiving they are with mistakes.

ROBOT-HUMAN COLLISIONS WERE A PROBLEM — We were surprised at the level of concern shown regarding the robots physically touching (colliding with) people despite it being clear that the robot posed no real threat. This concern emerged even for the *burglar* case where the robot was supposed to attack the person, and even though most participants explicitly trained the robot to collide with the person. We are curious as to how much this is related to the fact that the participants themselves trained the behaviour, perhaps creating a sense of personal responsibility for the robot's actions.

THE *PUPPET MASTER* ALGORITHM — That no effect of behaviour type was found on satisfaction, training time, or ability to identify a behaviour suggests that the *puppet master* algorithm is successful across a wide range of behaviour types, although our small sample size and arbitrary selection of behaviours makes it difficult to draw strong conclusions on this point. This lack of effect also has implication beyond the algorithm, i. e., that for the more-general SBD case it seems participants did not spend any more or less time training (or observing generated results) based on the behaviour type. This was unexpected as we anticipated that times would be different depending, for example, on how engaging or detailed a particular behaviour is perceived to be. Our results only speak of the general case and it may be that this relationship exists on a per-participant basis but is just not consistent across participants.

The problem of robot jitter also emerged as a major concern, despite attempts to improve the problem in the *puppet master* robotic algorithm over the animation one. We feel that the reason for this is in the low speed and response time of the robot, resulting in a much slower control loop: the robot could not recover from mistakes quickly enough thus exaggerating noise in the system.

7.6.2.8 *Designer Study Discussion — Interfaces*

REFLECTIONS ON *SURFACE PUPPET MASTER* AND THE *BROOMSTICK* INTERFACE — Both interfaces were successful in their goals, with minor complaints about the crowdedness of the table and of the difficulty of moving the *broomstick* interface. It was surprising that no effect was found of interface type (*Surface puppet master* versus *broomstick*) on any measure, for example, as we suspected the tabletop may be faster in terms of training time as less effort is required to move over the space, or the *broomstick* to be longer as perhaps the direct interface is more engaging. We further expected this hypothetical increased engagement to be reflected in the satisfaction scores although it was not. The only difference found between interfaces was that people had more complaints about the *Surface puppet master* than the *broomstick* (related to crowdedness) although this did not have any measured impact.

SOUNDS WERE NOT USED — Although this study showed how the *puppet master* algorithm was able to incorporate the addition of sounds, there remains the problem of understanding why many participants did not use the robotic sounds during training. There appears to be

interest, as many expressed regret that they forgot about sounds, and so we have the question of whether this is due to interface design problems (e. g., perhaps the sound usage was not clear, easy to use, or obvious) or through core issues with the basic idea, that demonstrating sounds to a robot may not make sense to people. Perhaps it is related to mixing modalities, where it may be easier to focus on one method at once. For the people who did use the sounds, they claimed that they were necessary, some requested a wider range, and said that they brought emotive and a human dimension to the characters. It remains future work to further explore the role that sounds can take in anthropomorphism and interaction.

7.6.2.9 Designer Study Discussion — Methodology

THE NEED FOR EMOTION AND PERSONALITY THEORY — The less-than-perfect matching rate, combined with thoughts echoed in participant written comments, suggested difficulty in differentiating between different behaviours, although no detailed or specific relationships were revealed through statistical tests. However, a closer look at Table 7.6, page 229 suggests a rough clustering of matches, such as a large overlap between *polite* and *stalker*, and the *burglar* and *happy*. This is similar to the results found in the animation study (Section 7.5), and supports future-work exploration into how interactive, stylistic behaviours overlap, commonalities between them, and how, for example, they fit into existing theories on the relationships between various personality types and emotions.

MECHANICAL VERSUS NATURAL, ORGANIC MAY NOT MAKE SENSE — There was no emergent consistent opinion on whether the behaviours were *mechanical* or *natural*, *organic*, although feedback suggests that this may be due to unclear framing of the questions, and generally difficult-to-define concepts. The primary finding of *polite* being seen as more mechanical than the *burglar* and *happy* points to the question of how much the behaviour type versus *puppet master* algorithm results impacts these measures, as *polite* can easily be construed as mechanical, for example, politeness often involves suppressing emotion and being careful (robotic?) about one's actions. A related finding is how the question of mechanical went beyond the movement path, for example, to include the jitter problem, as well as motor noises and the design and shape of the robot. This points to the wide range of factors within the *holistic interaction context*, for example, related to how one participant complained of the cleaning imagery portrayed by the robot and *broomstick*.

In this section we introduced our *stylistic locomotion* and *puppet master* designer study, outlined our analysis, and presented a detailed discussion regarding the findings of this study. The discussion in particular summarizes the many findings of this study into selection of modularized lessons and questions for future work. In the next section we detail the related observer study, and end the section with reflections on the overall effort.

7.6.3 Observer Study

For the observer study the participants did not demonstrate behaviours, but simply observed the robot interacting with an experimenter and reported on their observation. The entire experiment protocol, summarized below, is given in full detail in Appendix D, Section D.5.

7.6.3.1 Observer Study Design

We had one independent variable for this study: the behaviour which the participant was observing, manipulated within subjects. The same variable was used for both tasks in the study, but manipulated differently as explained below.

In the first task, we used 4 behaviours (one each of *happy*, *polite*, *burglar*, and *stalker*) selected subjectively from the designer studies as the overall *best* by the experimenters. In the second task, 12 additional behaviours were used: one of each of the four behaviours (*happy*, *polite*, *burglar*, *stalker*) from each of the three cases (programmer study, designer study with the *broomstick*, designer study with the *Surface puppet master*). We used 16 unique behaviours in all, and in both cases the order of the behaviour presentation was shuffled (fixed across participants).

Twelve participants were recruited for this study from the general university population, aged 19–36 ($M=26.3$), 7 male and 5 female, and were paid \$20 CAD for their time (the experiment took roughly one hour).

The experimenters kept notes on the participants' comments and feedback during the think-aloud exercises (explained below). In addition, video data was analyzed for the four cases where it was available: half (six) participants requested not to be video taped, and two video-taped sessions did not have audio due to technical problems.

7.6.3.2 *Observer Study Tasks*

There were three tasks in this study. First, the participant observed 4 different behaviours for 4 m each while performing a think-aloud exercise, and were encouraged to speak freely about their impressions of what they were seeing — we called this the *open-ended* phase.

For the second task, we also asked the participant to observe behaviours, but in this case we informed the participant of the categories of the behaviours and asked them to attempt to classify what they were seeing. In this case, 12 additional behaviours were presented for 45 s each. We called this the *matching* phase. For all the above cases, the experimenter walk pattern was the same, given in Appendix D, Figure D.3, page 368.

For the third task (the *in-situ* task), the participant could optionally wear the Vicon-tracked shoes (Figure 7.17c, page 194) and interact directly with the robot, using the original four behaviours from the open-ended phase. This stage was videotaped and participants were encouraged to *think aloud*.

7.6.3.3 *Observer Study Procedure*

We conducted this study using a structured protocol including the use of informed consent forms and questionnaires. Participants first completed a pre-test questionnaire which enquired about demographics and predisposition toward robots, artistic experience, and general technical ability.

The participants observed all cases for the first task, continuing straight to the second task where they completed a questionnaire after each case. The third task was optional, after which we administered the post-test questionnaire.

7.6.3.4 *Observer Study Open-Ended Phase Results*

DESCRIPTIVES — For all behaviours, participants used basic descriptives regarding how quickly or slowly the robot was moving, how closely it stayed to the person, what it does when the person moves or when it collides with the person, or even how the robot moved better / worse in a certain floor region or in a certain direction (e. g., horizontal versus vertical). There was a particular emerging theme of participants relating to how well the robot is “tracing the walking pattern,” particularly for the *stalker* and *polite* cases, as well as a theme of relating to how well the robot “does a good job of staying in the boundaries.” Some participants

used their model of how they felt they understood the visual tracking to explain the robot's movements, for example, that the robot had to visually see the shoe markers to follow the person, or in one case the robot was "probably hitting the back [of the experimenter's shoes] because the silver balls [Vicon markers] in the front of the shoes."

ANTHROPOMORPHISM — There was a clear trend of anthropomorphism mixed into the descriptives, for example, many people called the robot "he" (no participant was observed calling the robot "she"). Several participants related this to the beeping noises the robot made, for example, "it seems for me that when he's thinking about what to do, he beeps." Only one person was observed expressing animosity to the idea of robots having emotions or human-like personalities.

JITTER AS A PERSONALITY TRAIT — The most common comment from participants was the robot's jitter in its movement. Some simply noted characteristics such as the robot "shakes a lot" and explained this as flaws in the algorithm such as the robot "lost track of something" and "seems looking for something, seems it didn't discover [the experimenter]." However, most participants explained this using agency, for example, that the robot "seems to be very indecisive on the movements," or that "it seems frustrated when it jitters". Confusion was a particularly predominant comment, for example, that the robot is "trying to follow but is very confused." One participant said (in relation to jitter) that the robot is "somewhat a dog, he smells something some times," and another said that the robot is "very A.D.D. [Attention-Deficit Disorder], it gets distracted."

BEHAVIOUR DESCRIPTIONS — For the *stalker* behaviour (first to be observed) there were very few comments. Several participants voiced confusion such as "not really sure what it's [robot is] doing." There were some, however, who described using language such as the robot is "trying to hide, trying to follow," or "it is trying to avoid?"

For the *burglar* behaviour (second to be observed), many observations were descriptive such as the robot is "trying to hit [the experimenter]," "trying to get ahold of him because it keeps jumping on him" or "its clearly not trying to avoid." The keyword *aggressive* emerged as a theme, for example, "it's aggressive, as if its fighting for territory," or "definitely more aggressive." One participant said that the robot was "disturbing. This time it seems a bit

disturbing, disturbing,” where follow-up conversation revealed that “disturbing” was used in the context of bothering the person. One participant mused that the robot is trying to say “nothing here for you! What are you doing here?” maybe he wants to say ‘please pay attention to me,’” and the same participant said that “sometimes ... he wants to be friendly, sometimes he is disturbing him.”

For the *polite* behaviour (third to be observed), the descriptive feedback included, for example, that the robot was “not hitting him,” was “trying to follow as closely as it can” or “does nothing but just follow him,” and several people attempted to make sense of the exact path the robot took and the area it covered. Overall this behaviour was seen as simpler and less busy, as one person put it, “seemed to be moving more smoothly but isn’t as confused as much as before.” Many participants attributed the movements to personality, such as the robot “seemed to be moving more slowly but isn’t as confused as much as before,” that it was “approaching the person more carefully,” and that it “seems more polite this time.” One participant said that it “somehow look like a police man / security guard, walking around campus.” Some participants drew a parallel and found similarities between this behaviour and the *stalker*, and one participant said that the robot is “trying to hide. Trying to follow without his knowledge.”

For the *happy* behaviour (fourth to be observed), the feedback was the most mixed of any case. On the one side, some people noted that the robot was “not quite as violent as the second one [*burglar*], thats for sure,” or “I think its in a good mood,” “very opposite to number two.” Several others found the robot “aggressive, but not as much as the second round.” Some found the robot to “seem scared,” nervous: “after hitting [the experimenter] seems to look frightened.” One participant said that the robot was “giving [them] a nervous feeling.” Many people commented on the sounds for the *happy* behaviour, particularly that the meaning was unclear. One person felt that it “sounds like when a battery is going low,” and another noted that “to me, sounds very neutral, doesn’t sound like a good sound or a bad sound. It’s very neutral.” One participant said that they “found it a little hard to associate the sounds to certain behaviour when [they] didn’t know what the behaviours were.”

7.6.3.5 Observer Study Matching Phase Results


For the second phase of the experiment, participants matched the behaviours created under different conditions to their behaviour type. The results are summarized in Figure 7.28.

BEHAVIOUR MATCHING RESULTS — The overall match success (across all participants and all behaviours) was 54% (SD=16%, min=25%, max=83%) and the per-interface rates were *Surface puppet master* at 54%, *broomstick* at 45%, and programmer at 62%. Comparing the per-participant match success rates between *Surface puppet master* / *broomstick* / programmer did not reveal a significant effect of interface type on matching success using a Friedman's ANOVA ($\chi^2(2)=2.06, p=.358$). Using the logic that we would expect 25% correct on average given purely random answers, given that for each behaviour the chances of getting it correct by random chance is $\frac{1}{4}$ or 25%, we applied a one-sample *t* test to compare our distribution to this expected mean. The result ($t(11)=6.13, p<.001$) suggests that it is unlikely that the 54% average was reached by chance.


Particular groupings also emerged from the matching data. Figure 7.28b shows how *stalker* and *polite* were often mistaken for each other (and much less often mistaken for other behaviours), while *burglar* and *happy* had much-less well defined distributions. These tables also highlight how certain behaviours performed especially well (such as the *Surface puppet master stalker* and the programmer *burglar*) and how some performed especially poorly, such as the *broomstick burglar*, which only one person guessed correctly. One participant mentioned (in the post-test questionnaire) that “some of the behaviours were very obvious, a few were very confusing (given the option that we had to pick from four behaviours).”

PER-BEHAVIOUR QUESTIONNAIRE — Participant responses to the per-behaviour (twelve in all) questions are given in Figure 7.29a. For each question, we performed a Friedman's ANOVA to explore if there is a significant effect of which behaviour was shown on how people responded to each question. This analysis does not consider groupings per behaviour type or creation-interface type due to the difficulty of performing statistical tests on the two-way non-parametric dependent factors (Field, 2009). The ANOVAs did not reveal a significant effect for the *human-controlled* question ($\chi^2(11)=8.8, p=.637$) or the *felt mechanical* question ($\chi^2=14.0, p=.231$). The remaining questions had significant results, as shown in Figure 7.29b.

We did not apply post hoc tests, for example, pairwise *t* tests, as a Bonferroni adjustment over the 66 pairwise combinations would demand significance at $p<.00075$, less than the sensitivity of our analysis software. Rather, we present the average rankings used in the Friedman's ANOVA tests in Figure 7.29b to provide some insight into what possible relationships may exist. The rankings for a participant are calculated by ordering their responses from



(a) grouped by creation type, diagonals are correct responses



(b) grouped by behaviour, correct responses are rows aligned with given behaviour type

Figure 7.28: how observers classified behaviours shown to them, twelve per participant

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
It was difficult to classify the behavior	19	43	21	3	30	17	11
I found the behavior to be engaging	10	7	10	17	40	40	20
It felt like a human was controlling the robot	41	41	16	10	16	13	7
The behaviour fell into the categories I was given	5	13	6	9	41	44	26
The behavior felt mechanical	27	34	21	8	23	19	12

(a) frequency table of participants' responses to per-behaviour questions

	tabletop				broomstick				programmer				Friedman's ANOVA
	stalker	burglar	polite	happy	stalker	burglar	polite	happy	stalker	burglar	polite	happy	
it was difficult to classify the behaviour	5.13	7.17	7.25	8.25	8.38	6.92	7.96	7.96	2.79	4.33	5.38	6.50	$\chi^2(11)=35.5$ $p<0.001$
I found the behavior to be engaging	7.67	5.88	5.04	5.79	3.88	7.42	4.92	5.33	8.29	9.71	6.83	7.25	$\chi^2(11)=32.8$ $p=0.001$
the behavior fell into the categories I was given	8.17	6.17	5.71	5.33	4.25	6.50	5.21	4.79	9.13	9.33	7.58	5.83	$\chi^2(11)=33.2$ $p<0.001$

(b) average ranks of how participants rated each behaviour on a given question, where a higher number was a higher score

Figure 7.29: participant responses given on the rapid questionnaire given per each of the twelve behaviours

lowest to highest, and numbering them ranked one to twelve; our presented averages are the means of these rankings (per behaviour) across participants.

7.6.3.6 *Observer Study In-Situ Phase Results*

For the third phase of the experiment participants were offered the option to wear the tracked slippers and interact directly with the robot themselves: eleven of the twelve participants opted to do so. Most participants expressed excitement when this option was presented, with one person saying “I want to feel it,” and most were very animated and involved in interacting with the robot’s behaviour. Unlike the first open-ended phase, at this point the participants were aware of the intent of the behaviour they were seeing, and readily played along. For example, one commonality is that several participants, when interacting with the *burglar* behaviour, would quickly move away from the robot when it was chasing them, and some even exaggerated their movements to play along with the robot. Comments include “the way it moves, the sound, all makes it creepy like a stalker,” “the robot seems happy!”, “‘happy to see you’ is just a pet of child who really feels happy to see me,” and the robot “looks and feels like a polite machine.” Further, several participants said during this phase that overall the robot “reminds [them] of a dog.”

One particularly interesting case was an experienced electrical engineer who had, throughout the earlier phases of the study, clearly voiced animosity toward the idea of robots having personalities: once this person wore the shoes, however, they were laughing, talking to the robot as they may an animal, and used anthropomorphic language to describe what was happening such as “he’s doing a good job” (for the *burglar*).

The post-test questionnaire explicitly asked participants’ opinions on directly interacting with the robot. Many participants simply re-iterated their interest, for example, “definitely wanted to” and “awesome! I think they are a wonderful creation of man.” One person felt that “the behaviours and displayed intelligence of the robot was very impressive.” Some explained that they were able to test the robot’s capabilities and direct interaction gave them better insight, for example, “I found it very entertaining trying to predict the behaviour of the robot and seeing how it reacted,” and “it was cool to interact with the different personalities. I could get a better idea of some of the personalities when I interacted with it, compared to simply watching.” In relation to the quality of the behaviours, the following comment reflects the

tone of the bulk of the feedback: “for the most part the behaviours seemed very natural and I was able to believe the robot had a personality of its own.”

There were several issues or complaints which emerged with this direct-interaction phase. One prominent theme was the limitations of the robot, as the participants would move much more quickly than the experimenter did during the observation phases, and the robot could not modify its behaviour or catch up in time to interact properly with the person. Some participants felt that the behaviours we chose “were not as natural” as some of the previous ones they have seen. One participant said “I was waiting for the robot to interact with me,” while another found all behaviours except the *happy* boring to interact with.

7.6.3.7 *Observer Study Other Results — Sounds and General Feedback*

When explicitly asked about the sounds, several participants commented that they found that “the sounds (beeps) were helpful” and “important” for identifying behaviours, and were “a nice touch” that “is very important to give natural feeling to users.” One participant commented “when I was unable to determine the robot’s behaviour by its actions, I relied on the sound to determine if the robot was ‘happy’ or ‘angry;” and another said that “some robots appeared more happy because the sound sounded more upbeat.” One in particular said that “it’s good, but could be further improved,” and several participants noted that the “happy tone was not clearly happy,” although no one explicitly commented on the *burglar* attack sound.

Some participants commented that they “did not actually pay attention to sounds :(.“ Quite unexpectedly, one participant answered that “the sound of the robot should be less mechanical,” referring to the motor noises made as the robot moved rather than the beeps. One participant also (mistakenly) felt that the sounds reflected the robot’s “identification process [of the human] (proximity, direction, changes, ...).”

Finally, when asked for overall general feedback, several participants expressed gratitude as this “was a novel experience for [them]” and that they “had fun,” and it was “interesting to see how [the] robot expresses different personalities.” Several people complained that “frequent stutters undermined the experience,” asked us “to make the robot more smooth,” and one person said they “like the personalities that were more ‘fluid’ and didn’t have as many jerky movements.” One person said that they “felt on some of them it spent a lot of time ‘thinking’ what to do.” Related to this, “some of the more mechanical movements just

seemed too artificial, and associating an emotion with them was difficult,” and one person interpreted the jitter as “occasionally the robot could not find the human.”

7.6.3.8 *Observer Study Discussion*

Here we outline key themes that emerged from the data and the experience of conducting this study.

NOT AWARE OF DEMONSTRATION — The observer condition was an important part of this set of *stylistic locomotion* and *puppet master* studies as this is the only case where the participants were not aware of (let alone involved in) the SBD context. The act of training can perhaps be linked to learning, agency, and intentionality, a link which would be missing in this study. This further means participants did not have any direct personal connection to the quality or results of the behaviours and no experience considering how the robot should perform a given behaviour. Despite this, participants commonly attributed agency and intentionality to the robot and its actions, described interaction using anthropomorphic language such as “he [the robot] feels like,” and expressed that they could believe the “robot has a personality of its own;” several people even related to the robot as being like a pet dog.

JITTER PERCEIVED AS CONFUSION — As in the designer study case, people used the idea of confusion and uncertainty to explain and understand the robot jitter problem. That this emerged without the context of learning speaks to the core communication quality and meaning behind the robot’s physical act of jittering during actions.

WELL-BEHAVED ROBOTS — We did not expect the issue of robots staying inside the boundaries to emerge. This could perhaps be related to how well-behaved the robot is perceived as being in relation to established rules; the boundaries were explained to participants at the beginning of the study. The question remains as to why this was not raised as an issue in the designer case, although perhaps this is simply a result of the more-open think-aloud exercise used here. Related to this, we only found one mention of a participant being concerned about the robot-human collisions.

MECHANICAL VERSUS NATURAL, ORGANIC MAY NOT MAKE SENSE — In this study participants directly noted the lack of clarity regarding our questions on if the robot was *mechanical* or *human-controlled*, and several participants in particular explicitly linked the questions to the robot's construction (not only its behaviour). It also became evident that these questions may reflect on the type of behaviour rather than the quality of the results, for example, a very successful generation of *polite* may still seem mechanical.

SOME BEHAVIOURS MORE DIFFICULT TO CLASSIFY — Results from the Friedman's ANOVAs given in Figure 7.29b show how some behaviours were seen as being more difficult to classify, more engaging, and more suited to the categories than others. If we could deduce more information about these relationships (perhaps by conducting follow-up studies), then it may shed some light on the differences between our behaviours and how they relate to *puppet master's* capabilities.

INTERACTING VERSUS OBSERVING THE ROBOT — We find the fact that eleven of the twelve participants volunteered to wear the shoes and interact with the robot to support the idea that our behaviours were engaging, particularly for those who expressed animosity toward robots having emotions or personalities. Note that the participants could have left the evaluation early and had no obligation or incentive to stay beyond interacting with the robot, although we must account for the novelty factor of robots. Feedback from this phase further points to differences between a participant observing the robot interacting with another person (most of the study) and experiencing the interaction first hand, such as how they feel about and interpret the behaviours. This suggests future work on, for example, considering how the meaning of behaviours changes for direct interaction versus indirect observation.

AMBIGUITY OF SOUNDS — As the *happy* and *unhappy* robot sounds were not introduced in this case as they were in the designer case participants did not have predisposition toward their meaning in use. The *unhappy* sounds used during the *burglar* case were rarely commented on, so we make the assumption that they worked reasonably well or at least did not raise confusion. That *happy* sounds were often not clearly understood raises questions regarding the generic nature of sounds such as the flexibility of their interpretation, although in our case it is likely simply a poor choice of sound. Despite this, that many people claimed

the sounds were helpful for understanding a behaviour and added a “natural feeling” to the robot supports the idea that sound plays an important role in the meaning of a behaviour.

ACCURATE DESCRIPTIONS — Reflecting on the quality of behaviour generation, during the open-ended observation phase there were themes emerging of people recognizing the behaviour and some people constructed fairly-correct stories of what was happening. However this was the exception rather than the rule, and most participants gave descriptives of what the robot was doing (while still using anthropomorphic language such as “he is trying to move to...” and “he is confused now”). The fact that people accurately described the behaviours supports the capability of the *puppet master* algorithm to capture basic properties of interaction, for example, colliding with a person, following at a particular distance at a particular speed, repeating particular motions at appropriate times, etc. Further, while 54% success rate at matching was lower than hoped for, statistical tests support this as being different from random selection and so this suggests at a level of success of the *puppet master* algorithm.

In this section we have presented our formal evaluation on how people react to and perceive robots that communicate using *stylistic locomotion*, including a detailed analysis and discussion on our results. Overall this study has helped highlight important aspects of SBD as well as the *puppet master* algorithm, and many of the findings provide insight into core social HRI questions of how people treat robots as social entities and apply *the social stock of knowledge* to understand them. In the next section we conclude our series of studies, relating findings from all the studies presented in this chapter.

7.6.4 *Stylistic Motion and Puppet Master Studies: Reflections*

In this section we summarize all of our *stylistic motion* and *puppet master* studies — the animation-based designer and observer, the robot-based programmer, two designers, and observer — and synthesize the results into an overview set of reflections.

VIABILITY OF USING STOCK OF SOCIAL KNOWLEDGE — These studies support the viability of our approach of tackling complex interaction problems (such as communicating or programming interactive, stylistic robot behaviours) by leveraging peoples’ existing skill sets.

Our study results show how people understand that robots can communicate through *stylistic locomotion*, and readily accept the idea of programming such behaviour by demonstration. Participants were able to understand and use both systems with virtually no training. The general accessibility of this technique is further supported by how we (informally) did not find any relation between success or use and technical or artistic ability, or experience or disposition toward robots. Problems which did arise were for the most part centred around interface issues or quality problems with the *puppet master* algorithm.

PEOPLE UNDERSTAND THE INTRICACIES OF TEACHING — Our results highlighted how participants already have an intimate understanding of the challenges and difficulties of teaching style, for example, identifying how the approach relies on the robot's interpretation of the intention behind the teaching rather than the intention itself, and that it helps if the person is a *good* teacher that identifies the needs of the learner. From this starting point, the participants demonstrated that they understood that they must exaggerate and emphasize intention to expect the learner to understand. For our technical participants, this included suggestions of how to mesh our SBD approach with more structured methods. For example, by including the person in an iterative loop and offering a mechanism to allow them to specify which parts of the demonstration are important. This supports our claim that people are experts at (at least informal) teaching, a skill that robotic design can leverage.

INCLUSION OF PROGRAMMERS — The inclusion of experienced programmers was important for several reasons. It concretely demonstrated how SBD makes sense even for people who do have a technical understanding of what is happening and the capability to use traditional means (i. e., programming) to create the behaviour. Second it provided a programmed-behaviour comparison point to use in our studies against SBD-created behaviours. The simple fact that the SBD behaviours created in 50 s (average) by untrained members of the general university population can compete with behaviours programmed in roughly 30 m each by experienced programmers speaks to the success of the *puppet master* algorithm to use *the social stock of knowledge*, even considering potential confounds such as biases in our API and programmers' backgrounds.

INTERPRETATIVE FLEXIBILITY — The range and depth of participant comments, such as how the task was related to cleaning given the *broomstick* interface, or how the physical design of the robot impacted how the motions are interpreted, speaks to the importance of considering the *holistic interaction context*. This is the case for both the use of our design interfaces as well as interaction with the robot itself, and sometimes has made it difficult to determine what has shaped a particular participant response. One such point is the serendipitous finding of how people nearly-universally applied agency and their understanding of living things to interpret the robot jitter as robot confusion. Another example is how several participants related the robot to a pet dog, supporting ideas put forth in Section 3.2 about how people relate robots to living things of similar intelligence, for example, animals.

PUPPET MASTER ALGORITHM STILL NEEDS WORK — The lower-than-hoped-for behaviour identification and matching results highlight that our SBD system still needs improvement. For example, despite algorithm improvements over the animation case the remaining jitter problem (although effective in communicating uncertainty) speaks to required improvements in *puppet master* algorithm. Other improvements include replacing the slow robot with a faster and more capable one, and expanding our limited range of sounds.

TEACHING CHANGES INTERACTION — Clear differences between the designer and observer phases, such as the issue of robot-human collisions only being voiced in the designer study, raises questions relating to how the act of teaching a robot influences interaction. Differences suggest, for example, a sense of responsibility for the resulting quality and robot's actions may emerge, or perhaps increased tolerance (and rationalization) for mistakes. One such example is how jitter seemed to be more of a complaint for the observers than the demonstrators: perhaps this is related to how designers explicitly took responsibility for weak generation results.

ANIMATED JITTER WAS NOT CONFUSION — Having done a very similar study with animated entities provides a useful comparison point for considering how people interact with robots. For example, while jitter was a problem in both the animation and robot studies, only in the robot study did we find that people attributed this to the entity (robot) as confusion;

in the animation study it was just described as distracting and annoying “jitter.” This perhaps relates to the power of robots (in comparison to animated entities) to elicit agency.

RELATIONSHIP BETWEEN BEHAVIOURS — Both the animated and robot studies highlighted clustering in people’s identification of behaviours (e. g., as with *afraid* and *stalker*), and point to the need to further explore the theory behind emotion and personality characterizations. We did not do this previously in our work as our primary focus was to have reasons and tasks to encourage people to engage our interfaces; the need for this exploration has emerged from our studies.

Overall, the results from these studies address our overarching goals as outlined in Section 7.6, page 217. We believe that in particular, our exploratory fifth goal (to observe people interacting with robots and reflect on the social HRI experience) has provided us with many insights into how people perceive and interact with robots, particularly in relation to the social task of teaching. Below, we conclude the chapter.

7.7 STYLISTIC LOCOMOTION AND PUPPET MASTER: CONCLUSIONS

In this chapter we presented the ideas of robots that communicate by adding an element of style to their actions, and robots that learn style directly from demonstrations by people. Focusing on the style of a robot’s interactive locomotion path, we titled these *stylistic locomotion* and *puppet master*, respectively.

We further presented various interface designs and implementations for both the *stylistic locomotion* and *puppet master*, including: the *mouse GUI*, the *animation table*, *robot locomotion*, the *Surface puppet master*, and the *broomstick*. These interfaces serve as examples of how Style-by-Demonstration (SBD) can be integrated into interaction design, and of how such designs can be implemented. Further, we presented our original *puppet master* algorithm for both animated entities and robots, illustrating a method for realizing SBD.

We presented a set of formal studies and design critiques that target people’s experiences and interactions with robots for both the *stylistic locomotion* and *puppet master* interaction scenarios. In addition, these studies reflect on our particular interface designs and implementations themselves.

This work has directly addressed our overarching research questions as presented in Section 1.3.5. The core idea of SBD, our interface designs, as well as our interface implementations all inform social HRI researchers on ways to design their robots to take advantage of people's social skills (question 2). Many results from our evaluations also address this question, for example, that jitter may be used to convey confusion, or that the act of teaching may change overall perceptions of the robot.

Our *puppet master* algorithm serves as a tool for social HRI researchers to achieve SBD (question 3), an algorithmic solution which we believe can be extended well beyond our targeted application of *stylistic locomotion*. In addition, we believe that our particular evaluation methodology and methods employed to target social HRI, and our successes and failures, serve as experience which social HRI researchers can build on to create their own evaluations, for example, our particular mixed use of qualitative and quantitative, and our use of non-parametric statistics for an exploratory study. We have structured our evaluation experiences formally as one of our heuristics (Chapter 8). Finally, our evaluation results (including our comparisons between robotic and animation implementations) help to illustrate some of the dynamics of how agency is both manifested with robots and how it changes interaction, for example, pushing people to perceive jitter as confusion (question 1).

Overall, our exploration has highlighted how people can apply their understanding of the style of actions, and of how to demonstrate style to others, to interaction with robots. Thus we have shown how robot design can use *the social stock of knowledge* to make complex HRI problems easy and accessible by using people's existing skills.

Part III

CONCLUSIONS

HEURISTICS FOR SOCIAL HRI

Social Human-Robot Interaction (HRI) is an emerging field and there is currently no standard social HRI-targeted body of knowledge or literature, no set of standard practises, and no set of accepted tools which social HRI designers can turn to for direction. In this chapter we draw from our overall experiences with social HRI, including our theoretical exploration into sociology, robot intentionality and agency, and experiences designing, implementing, and evaluating social HRI interfaces, and summarize them as guidelines — grounded in our explorations — for the consideration of the social aspects between people and robots. We present a series of heuristics, “general formulations that serve to guide investigation” (WordWebOnline, 2010, “heuristics”), that serve as practical tools to help re-conceptualize the robot design problem and actively guide researchers to focus on social HRI. They are:

- H1 employ agency and anthropomorphism
- H2 improve social accessibility
- H3 design for specific interpretations
- H4 use real robots
- H5 get people involved

We present each heuristic in detail below. H5, *get people involved*, also includes a detailed method for exploring social HRI experience possibilities. We follow with a demonstration of these heuristics by applying them directly to our interfaces that we presented throughout this dissertation. Altogether this chapter forms one of the first social HRI-specific research tool-sets.

8.1 HEURISTICS

In this section we present general social HRI guidelines in the form of explicit design recommendations. These are summarized findings which have emerged from our research,

expressed here as practical, toolbox-style design heuristics, to allow future use by social HRI researchers. Each guideline is a subsection below.

8.1.1 (H1) *Employ Agency and Anthropomorphism*

People tend to treat robots as social entities, anthropomorphise them, and attribute intentionality and agency to them. Social HRI designers should expect these trends to emerge and should take advantage of them in their interaction designs. As this tendency exists regardless of designer intention, explicitly accounting for them is a powerful way to gain control over how people interpret and use an interface, and provides an existing mechanism and channel for interaction which people are ready to accept. Such interface integration does not have to be blatant or conspicuous, for example, animated facial expressions or active speech, but can also be mild and more passive such as minor inflections in movements, or primarily static design such as subtle eyes and life-like morphology.

In addition to considerable evidence of tendencies toward agency and anthropomorphism in related work, this heuristic emerged in part from our original theoretical framework which offers new explanations for this tendency and outlines the magnitude of impact that this has on the overall interaction experience. Further, this heuristic draws from our own work in designing, implementing, and in particular evaluating the interfaces presented throughout this dissertation, demonstrated the prominent nature of these tendencies and their ability to shape overall interaction. For example, we found this in how people anthropomorphised the simple white-disk Roomba in the *stylistic locomotion* and *puppet master* systems, and how *cartoon artwork* made the robot feel more fun and personal.

8.1.2 (H2) *Improve Social Accessibility*

Leverage the *social stock of knowledge* to make robotic interfaces accessible. While this general approach makes sense for Human-Computer Interaction (HCI) as well, peoples' tendency to already treat robots as social entities makes this approach particularly relevant for robots. Robots can tap into the wide range of abilities that people use to understand and interact with other people, animals, and their general everyday world, to reduce the learning and mental load required to interact with them. Leveraging peoples' existing social skills and

understanding is a powerful way to make complex robot interaction and control problems accessible. In addition, robots' physical nature dramatically expands on what is possible with more traditional technologies.

Leveraging *the social stock of knowledge* can mean adding an abstract or difficult-to-define layer to the robot's interface, such as human-like expression or vague emotional representations. However, we point out that abstraction — and even vagueness — can be beneficial when it is something people already understand and can access and accept (*robot expressionism*, Section 3.5). It is more important to focus on generating clear comprehension and understanding in an attempt to improve accessibility than accurate representation of robotic or sensor state.

This heuristic originally emerged from our theoretical exploration into robots' particular embodiment and how this relates to how people understand interaction — designing robots to leverage this understanding makes sense. It further emerged from our own work where in all of our own interfaces we have made difficult robot communication and interaction problems easy and accessible through using the *social stock of knowledge*. The *dog-leash* robot leverages the familiar scenario of leading an animal on a leash, *cartoon artwork* leverages familiar communication from culture, *stylistic locomotion* leverages our understanding of style embedded in motion and *puppet master* leverages our social teaching abilities. All interfaces also use the tendency to treat the robot as a living thing, and our many extensive evaluations provide strong evidence that this approach indeed was successful in all cases, adding validity to this heuristic.

8.1.3 (H3) Design for Specific Interpretations

Robots are still novel technologies which are subject to a large degree of *interpretative flexibility*. As such, the term *robot* can have dramatically different meanings to different people, and is rapidly changing and evolving. This means that it is difficult to determine or predict where people will draw experience from when interacting with a new robotic interface or how they will perceive it. Assumptions should be made with care regarding how people will respond to robots, and it is important to consider what other interpretations of the robot people may have.

As such we argue that social HRI designers should target specific interpretations for the design of interaction (and the robot) to encourage people to interact in particular ways. Failure to do so, for example, by focusing on technical achievement only, can increase the likelihood that participants attribute the robot with false assumptions and expectations — designing for specific interpretations can reduce this possibility.

This heuristic emerged directly from our own work and user studies, for example, in how interpretations of our robots used for *stylistic locomotion* and *puppet master* varied greatly, and in how the interpretation had a sizable impact on overall interaction. The *dog leash* robot was an attempt at applying this heuristic (detailed below) as it was designed as a kind of dog or animal, building on people's tendencies to associate robots as dogs (as illustrated in our other projects) to further encourage our target interpretation. Further, this idea was informed by our theoretical standpoint of how *interpretative flexibility* is particularly important and broad for interaction with robots.

8.1.4 (H4) Use Real Robots

Due to the difficulty of programming robots it is common to use virtual robot simulations (i. e., no physical robot involved) for proof-of-concepts, and often even for evaluation of interaction design (Gockley et al., 2006; Wang and Lewis, 2007). However, whenever there is a person involved, the integration of real, physical robots fundamentally changes interaction; robots should be used whenever possible in the design of social HRI interfaces. While the exact nature and properties of this impact are not yet entirely clear, results from our work strongly suggest that the difference between using virtual simulation and physical robots is significant. Following, we believe that interaction studies done with virtual robots may yield flawed results which can be quite different from a final real-robot interface implementation.

Our approach, requiring the integration of physical robots in social HRI design and evaluation, may seem unreasonable as there are many practical reasons why researchers resort to simulated robots: robots are expensive, there are many difficult behavioural, vision and engineering problems in the way of realizing even simple real-world interaction, and in some cases, the behaviour a researcher wants to test (from a social HRI perspective) is not possible with current technology. When real robots cannot be used, we recommend acting and puppetry as a compromise to keep interaction in the real-world physical context. This domain

of low-fidelity social HRI prototyping is still largely unexplored, but we believe, based on a few experiences we had in our design group as well as in classroom exercises (not reported in this dissertation) that using hand-puppets as well as acting for the rapid prototyping of social HRI interfaces can provide a helpful low-fidelity prototyping tool for social HRI.

This heuristic emerged from both our theory and practice. Our explanation and discussion on why interaction with robots is fundamentally unique is anchored in the fact that they are physically-dynamic social actors in the real world. Further, we have strong evidence in our own experiences with implementing and evaluating our interfaces. For example, for both *stylistic locomotion* and *puppet master* we found prominent differences in the interaction experience between the animated and robotic implementations of very similar interfaces, and our evaluation of *touch and toys* revealed an overarching theme of the participant focusing on the real robot behind the interface.

8.1.5 (H5) *Get People Involved*

People, in general, are experts in social interaction, and so we especially recommend that social HRI designers leverage this expertise for evaluating and studying their designs and interfaces: conduct user studies with the general public. In our own experiences through our many evaluations we found participants' insight to be extremely informative, and often surprising, highlighting new research questions and directions. Thus this demonstrates the importance of getting people involved with new robotic interfaces, observing the interactions, and learning from the experience. Unfortunately, there is still a lack of social HRI-specific evaluation methodology and techniques, and so the question of how to best perform evaluations remains open. We contribute to answering this problem with a new method for aiding social HRI evaluation, detailed below.

8.1.5.1 *Exploring Social Interaction Possibilities*

Through our own experiences designing social HRI evaluation and conducting analysis, we found it particularly helpful to stop and explore alternate interaction experience possibilities for a particular scenario or finding. One goal in doing this was to try and keep our exploration and discussion anchored to the social aspects of interaction between people and robots. In

this section we show how our three social HRI perspectives (Section 3.6.2, P1: visceral, P2: social mechanics, P3: social structures) can provide such a social anchor for exploration.

For example, when an evaluator uncovers a particular finding, they can directly use the three perspectives as a probing tool, asking “how does this finding impact the interaction experience on the three perspectives?” As a detailed example, given a finding that people do not like to interact with a given domestic robot when guests are over, then the following hypothetical statements could be considered. Perhaps the robot’s P2-type communication, or people’s P1-type reactions to the robot, are intimate and inappropriate in group situations. Maybe the P3 integration into the home makes it uncomfortable to interact with the (lesser) robot in front of guests. Regardless, how does not interacting with the robot around guests impact the long-term P1-type reactions, direct P2 interactions, or P3 integration?

Below we formalize this probing approach into a concrete method for exploring social interaction possibilities. This surrounds the development and use of what we call an *interaction experience map*, a rigorous exploration of interaction possibilities and outcomes for a given social HRI scenario. The results from this exploration can be a useful tool for both evaluation design and analysis, as it could be consulted to explore alternative outcomes in an interaction scenario or to help explain unexpected observations or results. This approach has emerged from how we informally used our three perspectives in our own social HRI projects, refined into a method for presentation as given below.

Our method is based on using the three perspectives as probes into interaction experience possibilities to construct the *interaction experience map*. This has two components: a) the three perspectives serve as direct brainstorming and sensitizing tools to encourage breadth and a social focus, and b) both the human- and robot-centric views are explicitly and simultaneously considered. This latter idea is represented visually in Figure 3.7, page 83. The human-centric view considers how the person feels about and interprets the interaction experience, and the robot-centric view considers how the robot itself, including its design, behaviour and actions, influences the experience. Below we detail the process for constructing the *interaction experience map* from each perspective.

8.1.5.2 Human-Centred View

For the human-centred view, the researcher starts by brainstorming interaction scenario possibilities between a person and the particular robot or interface. The evaluator generates

a list of high-level scenarios that could conceivably take place, for example, a person trying to have an extended conversation with the robot even though the robot does not intelligently respond, or the person completely ignoring the robot, and so forth.

Then, for each scenario listed in the first step, P_1 , P_2 , P_3 can be used as probes to consider the interaction experience possibilities within the scenarios, and to sensitize the exploration to the particular social considerations. Following our example, a person conversing with a robot exhibits (social mechanics) P_2 elements of conversation and gestures, but they may also have visceral P_1 -type reactions when the robot does not respond as expected. This may include frustration and annoyance, which the person may externalize by means of body language (P_1 or P_2), communication which the robot may be able to detect. One emerging question is how does being frustrated with the robot, and being unable to have an in-depth conversation, influence how the robot is ultimately used and integrated into its target environment (P_3)?

For each idea and social reaction, we encourage the experimenter to consider alternate possibilities as a means to generate additional interaction possibilities, focusing on the three perspectives. For example, rather than being frustrated with a limited robot, the person may find the robot silly and the situation humorous (P_1), or the robot insistent (P_2) and perhaps intimidating. Following, each of these alternates can be then constructed into additional possible interaction scenarios, for example, perhaps the robot will be perceived as humorous and the person will use the robot for its entertainment value. Finally, the process can loop in an iterative fashion and these new interaction scenarios can be again analyzed using the three perspectives. This entire process is outlined on the left of Figure 8.1.

8.1.5.3 *Robot-Centred View*

Simultaneous to the human-oriented exploration, a similar process is followed for the robot-centred case. First, the experimenter brainstorms robot design characteristics that they expect may influence the interaction experience. For example, the fact that the robot has a face, makes loud noises when it moves, or even that it is the colour red.

Then, for each characteristic that was identified, the experimenter considers how people may react to it, and thus, how it may influence the interaction experience. Here the three perspectives can be used as probes for exploration, for example, people may find the robot's red colour to mean warning or danger (on P_2 or perhaps P_1), and the robot being noisy may severely hinder its deployment success as it may clash with existing P_3 social structures.

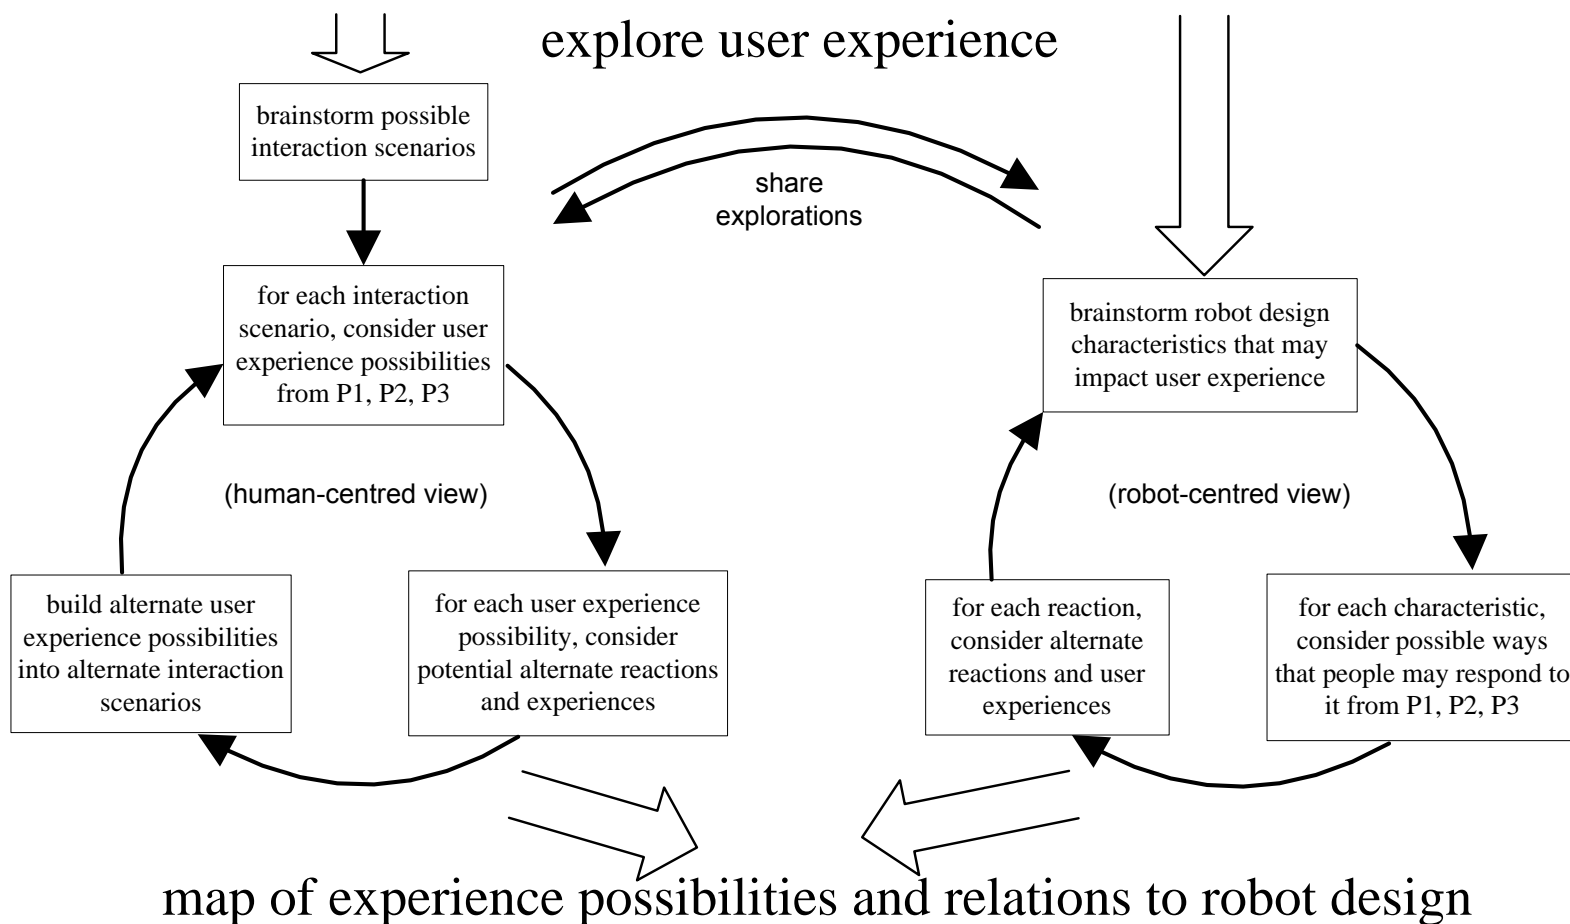


Figure 8.1: example process of using the three perspectives to fuel an exploration into experience possibilities

For each reaction possibility discovered above, consider alternate ways that the interaction experience may be affected. For example, the red colour may be seen as being festive or warm (P₁), or the noise may be perceived as a friendly quirk of the robot (P₁ or P₂), or perhaps that it represents the robot complaining while working (P₂). Finally, the next step is to use these alternate experience possibilities to re-think and re-brainstorm which characteristics may impact experience, and this leads to another iteration of the entire process (as outlined on the right side of Figure 8.1). For example, now that we have considered that the red, noisy robot may be seen as a festive robot with a quirky, fun sound, we can consider which other design aspects could support this identity, such as perhaps the particular face of the robot or the way that it moves. In this case, the three perspectives have helped keep the focus on the social, human-oriented aspects of evaluation while exploring interaction possibilities.

8.1.5.4 *Bringing it Together*

As highlighted in Figure 8.1, ideas and discoveries should be shared between the two simultaneous human- and robot-centred processes. This leads to the flexibility that we see as being inherent in this process, despite the structured and directional method presented in Figure 8.1: by no means do we suggest that the experimenter constrain their brainstorming to the process we present here.

This flexibility also matches our own experiences of map exploration, the foundations of where this process began, where the map-building process is primarily a guide and aid to brainstorming. In practise, we jumped between various methods of design and brainstorming, using components of our method when we felt they were particularly useful. While this process can be followed structurally, particularly as a way to start exploration, in practise we see it as something the experimenter can turn to for hints and ideas for pushing the brainstorming to new directions, particularly in relation to the three perspectives. This is highlighted by the fact that, as presented, this process has no explicit end and could conceivably yield a very large map. It is up to the sense and judgement of the experimenter to decide which possibility directions to pursue and which ones to cut.

The overall result of this process is a tool that provides the experimenter with a method for exploring interaction experience possibilities while staying grounded in social human-oriented concerns. We further envision that the three perspectives can be used directly during evaluation as sensitizing tools to consider observations from different angles, particularly for

more exploratory techniques such as field studies and unstructured interviews, and could also be integrated directly into data collection instrumentation, for example, as note paper which highlights P₁, P₂, P₃.

In this section we have presented five heuristics for social HRI: (H1) employ agency and anthropomorphism, (H2) improve social accessibility, (H3) design for specific interpretations, (H4) use real robots, and (H5) get people involved. We continue our discussion below by applying these heuristics to our own interface designs and implementations presented throughout this dissertation.

8.2 EVALUATING OUR INTERFACE DESIGNS USING OUR HEURISTICS

In this section we perform an evaluation of our interface designs using the heuristics presented above. One purpose of this application is as an illustrative example of how the heuristics can be used with concrete instances, and this also serves as an opportunity for us to reflect back on our interfaces in light of our overall findings.

8.2.1 *A Dog-Leash Interface for Leading a Robot*

This robot was created with the explicit decision to design for a specific interpretation (H3): a dog-like robot (Chapter 4) that knows how to follow you. Our approach directly takes advantage of tendencies to anthropomorphise and attribute agency (H1) and reinforces them by creating a robotic-dog conceptual target for the anthropomorphism. We should consider how to alter the currently-mechanical visual design to further encourage this tendency.

The dog-leash interface directly works with *the social stock of knowledge* as people are familiar from everyday life and culture with the concept of walking a dog (or other animal) on a leash. Our study illustrated how people can readily apply this knowledge to interact with our robot, creating a perception of accessibility and usability (H2), reinforced by our evaluation result that people were able to complete all tasks with reasonable ease without training (just a demonstration).

Finally, this project used real robots (H4), increasing the validity of our findings, and we involved the general public in our evaluations (H5), the benefits of which were reflected in our detailed evaluation discussion.

8.2.2 *Cartoon Artwork for Expression: Bubblegrams and Jeeves*

Both *bubblegrams* and *Jeeves* (Chapter 6) take advantage of tendencies to anthropomorphise and attribute agency (H1): *bubblegrams* use thought bubbles as if the robot itself is thinking and *Jeeves* uses facial expressions and human-like hats. By fitting to the tendencies toward attributing agency we encourage people to readily accept the characteristics as a natural part of the interface, a hypothesis supported by the successful results of our design critiques. Both interfaces further leverage *the social stock of knowledge* in that cartoon artwork (including thought bubbles) is deeply ingrained in our society and well understood by the general public: this helps make the interfaces accessible and easily understood (H2).

As both interfaces use real robots (H4) we can expect a reasonable level of genuineness in our observations of people's responses to them. Although we involved real people (H5), these were lab-member computer scientists, and so perhaps their perspectives are very similar to ours — we should perform a study with the general public.

We designed the robot for a specific interpretation (H3) of being like a cartoon character. Our studies support the success of our approach and how the cartoon-artwork focused perceptions toward fun and simple entities. Our particular design appears to help direct peoples' anthropomorphic and agency tendencies toward their perceptions of cartoon characters.

8.2.3 *Stylistic Locomotion and Puppet Master*

Our various interfaces for *stylistic locomotion* are designed to build agency and anthropomorphism (H1) in that they attempt to communicate with intrinsically social styles (e. g., *happy*). These styles further improve social accessibility (H2), as they are rooted in *the social stock of knowledge* and give the robot a means to communicate with people that does not require the person to learn.

We intentionally did not design for specific interpretations (H3), as the robot was intended to be diverse in the range of *stylistic locomotion* that it could convey. This can explain, for example, why some participants saw the robot as a cleaning device (due perhaps to the cleaning-robot model used). Perhaps designing more holistically for style (rather than movements only) would have narrowed the interpretations. We could have considered how to design the robot to be seen as something that communicates with stylistic movements (e. g.,

through morphological anthropomorphic design), although this would have contradicted our study design to test for emerging anthropomorphism.

The *puppet master* interfaces — and the overall Style-by-Demonstration (SBD) concept — take advantage of tendencies to anthropomorphise and attribute agency (H1) by encouraging people to understand the robot as something that can naturally learn as a person or animal might. This is further designed to leverage *the social stock of knowledge* as people are generally experts at casual informal teaching, an approach that we believe improves perceptions of the interface being accessible and usable (H2). We did not design for specific interpretations (H3), and as such interpretations of the robot ranged from being a small dog to being a cleaning device (in this case perhaps in relation to the *broomstick*). Designing explicitly for the interpretation of an entity that can learn would help alleviate these problems.

Both the *stylistic locomotion* and *puppet master* projects use robots (H4) in half of the cases. We found disparities between the animation and robot interface cases, such as differences in attribution of intentionality or responsibility of actions, highlighting the importance of *using real robots* (H4): some important findings emerged only when robots were involved.

We also got people from the general public involved with our interfaces (H5), which was very informative, and study results supported our claims of ease-of-use and using *the social stock of knowledge* as people required minimal (or no) training to use the interfaces.

In this section we applied our heuristics to our own work interfaces as a means of illustrating how the heuristics can be applied to social HRI designs and interfaces, as well as to validate that our own work fits the heuristics' recommendations.

8.3 SUMMARY: HEURISTICS FOR SOCIAL HUMAN-ROBOT INTERACTION

In this chapter we took a step back from our designs, implementations and evaluations, and summarized our lessons learnt in the form of heuristics: (H1) employ agency and anthropomorphism, (H2) improve social accessibility, (H3) design for specific interpretations, (H4) use real robots, and (H5) get people involved. This chapter serves as a significant contribution that directly addresses our research questions (Q3): our heuristics are a list of concrete recommendations for social HRI researchers that we expect can be used as simple tools to help other social HRI researchers quickly and easily apply our findings to their own work.

CONCLUSIONS AND FUTURE WORK

This dissertation detailed our in-depth exploration of the field of social Human-Robot Interaction (HRI). We conclude our discussion by re-iterating our overarching research questions, summarizing our efforts to address them, and our overall contributions. We follow with a discussion on the expected scope of our contributions, and finish the chapter with an outline of where this dissertation points in terms of future work for social HRI.

9.1 RESEARCH QUESTIONS

The core goal leading this dissertation was our desire to *define and substantiate the understanding of the field of social HRI, as a means of exploring, both theoretically and practically, how the inclusion of social aspects impacts the field of HRI*. Below we list our more-targeted research questions that consider this overall challenge from different perspectives, as presented in Chapter 1. We present our three research questions here to set the tone for this concluding chapter, and to remind the reader of the overarching goals underlying our contributions.

- Q1 *What does the tendency to treat robots as social entities mean for interaction between a person and a robot?*
- Q2 *How can robots be designed to leverage this tendency in their interaction and interface designs?*
- Q3 *Which methodologies, structured techniques, taxonomies, and heuristics can be developed and used for social HRI?*

9.2 RESEARCH CONTRIBUTIONS

In this section we detail how we have addressed our research questions directly through our research contributions. The primary contribution of this thesis is to serve as the first

formal, thorough treatment of social HRI which researchers can use as a starting point to understand what social HRI is, as well as how to *design*, *implement*, and *evaluate* social HRI instances. Below we enumerate the targeted contributions that make up this larger goal, and explicitly relate them to our research questions. We structure this discussion around our four contributions as listed in Section 1.5.1: 1. *social HRI fundamentals*, 2. *social HRI heuristics*, 3. *social HRI designs and implementations*, and 4. *social HRI evaluations*.

9.2.1 *Social HRI Fundamentals*

The objective of this contribution was to provide the first set of explanations, vocabularies, and tools for social HRI as a domain. We developed this through detailed theoretical analysis using ideas from sociology, social psychology, and philosophy, as well as our own social HRI experiences. We list the specific contributions below.

1. *a theoretical framework for social HRI* — Our framework serves as a thorough, packaged, and grounded resource for understanding social HRI. This framework identifies the unique properties of robots, the broad context within which social interaction with robots occurs, and provides new social HRI-targeted vocabulary. This is a significant contribution as it is the first thorough formalization of social HRI. (Section 3.6, page 80)
2. *definition of social HRI* — The theoretical framework above contains a formal and thorough definition of social HRI (as far as we know, the first such definition). This serves as a short, concise description which researchers can refer to. We include a definition of the word *robot* for social HRI purposes, an explanation of the practical utility of designing robots that leverage tendencies to treat them as social entities, and relate to *the social stock of knowledge* for understanding how robots can leverage people's existing skills and understanding. (Section 3.6.1, page 80)
3. *holistic interaction context* — We present this concept to point to the importance of considering the broad context for understanding a person's perceptions when interacting with a robot. This idea further serves as a sensitizing tool encouraging researchers to keep a broad perspective when they explore how people interact with their robotic interfaces. A key point of the *holistic interaction context* concept is that the robot, due

to *active agency*, plays a prominent role in shaping the interaction context, similar to how a living entity may. (Section 3.2.3, page 59)

4. *perspectives for describing social HRI interaction experience* — Our three social HRI perspectives are the first tools that target a broad spectrum of social HRI interaction experience. As illustrated by example in Section 3.6.2, our perspectives serve as a mechanism to describe, classify, and compare HRI projects, and as a means to decompose social HRI interaction experience into components for more-targeted consideration. (Section 3.6.2, page 82)
5. *factors shaping robot acceptance* — These factors provide insight into what impacts people’s perceptions of robots and their acceptance of them. These emerged from our thorough analysis of the special social HRI case of domestic robots using established social psychology research. In particular, this helped illustrate important components of the *holistic interaction context*. (Section 3.6.3, page 89)
6. *application of sociology and social psychology to social HRI* — We brought various concepts and methodologies from sociology and psychology to applied them both HRI and social HRI, including the ideas of social constructionism and Actor Network Theory (ANT), the *social stock of knowledge*, and several models of technology acceptance. These have proved to be particularly useful for describing and analyzing social HRI, and we expect will be relevant to researchers beyond our application in this dissertation. (Section 2.3.4, page 41, Section 3.3, page 61, Section 3.4, page 67)
7. *robot expressionism* — Our *robot expressionism* concept is an elegant representation of the power of abstraction for representing a robot’s state to people. We applied this idea directly to develop our cartoon artwork interfaces, and to help explain our *stylistic locomotion* interfaces. The link between expressionism and robots is further strengthened by the idea of *active agency*, and that people tend to treat robots as social actors, as expressionism generally deals with human experience. (Section 3.5, page 78)

In this section we outlined our contributions in relation to social HRI fundamentals. By laying out the theory of what social HRI is and why it is important we have helped to answer our question of what the tendency to treat robots as social entities means for interaction (research question 1). Further, our new social HRI-specific definition, perspectives, and

new vocabulary have helped to answer our question of which tools we can create to aid researchers in social HRI design and implementation (research question 2). In the section below we discuss our social HRI heuristics contribution.

9.2.2 *Social HRI Heuristics*

The objective of our heuristics (Chapter 8) is to provide *hard and fast* tools for social HRI design and evaluation. These are derived from our experiences as outlined in this dissertation.

1. *general design recommendations* — We present concrete-and-simple design recommendations for social HRI: (H1) employ agency and anthropomorphism, (H2) improve social accessibility, (H3) design for specific interpretations, (H4) use real robots, and (H5) get people involved (Chapter 8, page 257). This contribution is significant as it is the first social HRI-targeting set of heuristics. We envision that researcher can use our design recommendations as simple tools, allowing them to apply our overall lessons-learned without having to investigate the depth of our work. To this end we apply our heuristics to our own interface instances to serve as an illustrative example of how our heuristics can be used directly with real-world social HRI designs and implementations.
2. *evaluation heuristic* — We present the first set of detailed guidelines for specifically aiding in social HRI evaluation, directly leveraging our three perspectives on social interaction with robots as driving tools, providing a road map for the design, execution, and analysis of social HRI evaluation and studies. (Section 8.1.5.1, page 261)

Here we summarized our social HRI heuristics contribution, a concrete answer to our question of which tools can be developed specifically for social HRI (research question 1). We follow below with a thorough outline of our social HRI designs and implementations.

9.2.3 *Social HRI Designs and Implementations*

In this dissertation we presented various interfaces which we designed and implemented to utilize people's existing skill sets and familiar interaction scenarios. Our goal was to simplify complex robot communication and control problems, making robots easier to understand

and easier to work with for specific task contexts. While our interfaces vary greatly in scope, in task content, and in technical complexity, they share the common goal of demonstrating how social HRI design can leverage *the stock of social knowledge*. Below we list our social HRI designs and implementations.

1. *dog-leash robot* — This is an interface design and implementation for using a leash interface to enable a person to easily lead a robot to accompany them as they walk about, designed to use people's existing understanding of how to lead an animal on a leash. (Chapter 4, page 97)
2. *bubblegrams* — We developed and presented the Mixed-Reality Integrated Environment (MRIE), a cartoon-artwork concept and taxonomy for social HRI (Section 6.1.3), and the *bubblegrams* interaction technique that uses people's familiarity with cartoon thought bubbles for interaction. We presented a technical solution to *bubblegrams*, including an original computer vision algorithm for live marker-less tracking of a Sony AIBO robotic dog. (Section 6.2, page 151)
3. *Jeeves* — *Jeeves* is a test-bed for interfaces that use people's existing understanding of cartoon artwork for robot communication. This includes a methodology that recommends *which* cartoon elements a robot can use, and also several ways that it can use them (Section 6.1.4.4). We presented a functional *Jeeves* Mixed Reality (MR) implementation with several non-interactive proof-of-concepts and one interactive, functional interface. (Section 6.3, page 157)
4. *stylistic locomotion* — We designed and implemented three robotic interfaces that communicate to people through socially-rooted (and thus naturally understood) interactive, stylistic locomotion: the *mouse GUI* (an on-screen mouse-based proof of concept), the *animation table* (a tabletop animation and Tangible User Interface (TUI) proof-of-concept), and the *robot locomotion* interface (a robotic interface which enables a person to physically interact with a *stylistic locomotion* robot). (Section 7.1, page 166)
5. *puppet master interfaces* — We designed and implemented several interfaces that enable people to directly and easily show robots how to move in particular styles, similar to how they may show other people; some of these interfaces were extensions of the *stylistic locomotion* project. This included the *mouse GUI* with sequential (one entity at a

time) demonstration and the *animation table* for simultaneous (both entities together) demonstration. Also, we presented a separate tabletop system with a motion-restricted TUI for robot-style demonstration, the *Surface puppet master*, and a *broomstick* interface for direct robot demonstration. (Section 7.2, page 177)

6. *puppet master Style-by-Demonstration (SBD) algorithm* — This algorithm was an extensive, original programming style by demonstration technique for realizing both *puppet master* and *stylistic locomotion*. We developed a core animation technique, and an extension to work with robots. (Section 7.3, page 195, Section 7.4, page 206)

Our original interface designs and implementations are among the first to be explicitly designed for social HRI, and add breadth to the existing landscape of social HRI possibilities. We hope that other researchers can both use our techniques and draw from them for their own designs. Further, our extensive implementations and complete systems highlight the technical feasibility of our particular social HRI designs, and the general approach of interfaces that support social interaction between people and robots. Thus these contributions substantiate an answer to our question of how robots can be designed to leverage social interaction in their designs (research question 2).

These systems further included the presentation of new implementation solutions and original algorithms for targeted social HRI scenarios, the *bubblegrams* vision algorithm and the *puppet master* SBD algorithm. These serve as tools to realize interaction techniques previously not immediately possible, and thus provide a concrete set of tools in answer to our question of which techniques can be developed for social HRI designers to use in their work (research question 3).

In this section we outlined our social HRI design and implementation contributions. Below, we present our evaluation-related contributions.

9.2.4 *Social HRI Evaluations*

We performed several extensive evaluations of our interfaces, some of the very first evaluations with an explicit social HRI focus, and some of the first occasions where people had the opportunity to interact with social HRI interface instances. Further, through conducting these we explored various evaluation techniques and methodologies in relation to applicability

to social HRI. As such, our reporting of these evaluations and their results is a fundamental contribution to the emerging social HRI domain.

1. *cartoon-artwork informal design critiques* — We performed in-lab design critiques for both the *bubblegrams* and *trash Jeeves* interaction scenarios. These highlighted early in our explorations the impact that social (and anthropomorphic) designs can have on perception and interaction. Further, this highlighted the power of light-weight evaluations for learning about interfaces, despite being informal evaluations with few participants. (Section 6.4, page 159)
2. *formal evaluation of the dog-leash interface (12 participants)* — This evaluation compared different relative locations of the robot (*behind, behind angle, front*) from a social HRI perspective and described the dynamics between the cases using both quantitative and qualitative techniques. Further, we shared our experiences with using standardized questionnaires for social HRI. (Section 4.4, page 106)
3. *formal evaluation of touch and toys (23 participants)* — Initially designed as a cut-and-dry task efficiency experiment, this evaluation highlights how social elements emerged to be prominent despite the non-social evaluation design focus. While the statistical (task efficiency) results were inconclusive, the qualitative discussion was an important contribution. (Chapter 5, page 117)
4. *formal evaluations of stylistic locomotion* — These (two) studies exposed many of the dynamics surrounding how people interact with robots that communicating via stylistic, interactive locomotion. This consisted of a formal evaluation of the *tabletop animation* interface (10 participants, Section 7.5, page 209), and a formal evaluation of the *robot locomotion* implementation (12 participants, Section 7.6.3, page 239).
5. *formal evaluations of puppet master* — These three studies detailed many of the interaction components, trade-offs, interface successes and failures surrounding the *puppet master* SBD approach and implementations. This also showed that the underlying *puppet master* algorithm can be used to realize interactive SBD. We conducted a formal evaluation of the sequential-training *tabletop animation* interface (10 participants, Section 7.5, page 209), and a formal design critique comparing the *broomstick* and robotic

puppet master interfaces to direct programming of social stylistic behaviour (4 participants, Section 7.6.1, page 219). We also conducted a formal evaluation of the robotic *puppet master* systems (24 participants total, Section 7.6.2, page 223) for the *Surface puppet master* (motion-constrained tabletop TUI, 12 participants) and the *broomstick* interface (12 participants).

The results throughout our studies show how our attempts to build robots that leverage peoples' existing knowledge from social interactions with the world (other people and animals) have served successful: people find our interfaces easy to use, the core interaction ideas make sense, and there is minimal or no training required. In addition to this, our results (particularly with *puppet master* studies) support our stance that people are very quick to apply intentionality and agency to robots, and we found it important to consider the unique nature of robots when designing, conducting, and analyzing the results of an evaluation. Overall, this validation constitutes an important component of our answer to the question of how robots can be designed to leverage social interaction tendencies (research question 2).

Overall, we found that we received more information and insight regarding the interaction experience through qualitative and description-oriented observation and analysis techniques, over more-quantitative direct measures. In particular, we found participant self-reflection (on their experience, emotional state, etc.) and the integration of open-ended questions and interviews to be particularly revealing. That said, we did find task and efficiency-oriented measurements useful for helping to stipulate context and socially-oriented qualities such as, for example, engagement and interest, boredom, distractions, ease-of-use, and general understanding. These techniques alone, however, were primarily useful as a part of the holistic description-oriented approach. This body of work and experience provides an answer to our question of which evaluation methodologies and techniques can be developed and used for social HRI, in particular evaluation (research question 3).

Finally, through providing detailed, thick social-oriented description of people's experiences with robots, we provide answers to the question of what the tendency toward social interaction with robots means for the overall interaction experience (research question 1).

Here we provided a thorough enumeration of our research contributions, presenting the highlights of each. In the next section, we discuss the scope of these contributions and how they can be applied and related to new challenges, beyond our own targeted goals.

9.3 SCOPE OF CONTRIBUTION

In this section we discuss the scope of our contributions in terms of their generalizability to other contexts. Particularly, we discuss how our contributions scale to social interaction settings and cultures beyond the ones we investigated, to robots and robotic interfaces beyond our particular implementations, as well as to other technologies.

9.3.1 *Application to Other Cultures*

We believe that much of our work in this dissertation can potentially transcend cultures and countries. There are many aspects of our work which arguably have a culture-specific angle to them, for example, our choice of cartoon artwork, locomotion styles, or even the dog-leash metaphor. However, even where culture-specific elements exist, in each of these respective chapters we have shown how, fundamentally, the idea speaks to basic human social perception and are fundamentally global. As such, we expect that our approaches will still largely be applicable to other cultures, although they may require fine tuning to other culture-specific social aspects.

Our evaluations were conducted primarily in the context of Canadian culture, although the dog-leash study was performed in Japan, and participants in our in-Canada studies had a large portion of international students: we had participants from various countries including Afghanistan, Brazil, China, India, Iran, Japan, Korea and Pakistan, just to name a few. As such, we feel confident that our work and lessons learnt from these evaluations will scale to other cultures and contexts, although further investigation (e. g., with actual evaluations) is required to make this a stronger and more valid statement.

Overall, we believe that our work speaks to many core human social principles, and in this section we have outlined how we believe our contributions scale to other cultures. Below we discuss how they scale to other interfaces.

9.3.2 *Application to Other Interfaces*

We believe that our contributions scale to other social HRI designs and robot platforms. Our specific interface designs, as outlined through the chapters, are presented as instances of

broader approaches, for example, *Jeeves* is merely a realization of the use of cartoon artwork for interaction. In this example, the core idea of using cartoon artwork can be applied to other realizations, be integrated with other methodologies (e. g., an android), and so forth. The same relationship exists with stylistic locomotion (robots communicate using style), *puppet master* (people are good at demonstrating style), and the dog leash (people can work with robots as they may with an animal).

Our specific implementation methods can be used for other robot interfaces, for example, our MR solutions, use of motion tracking, and overall system designs are not specific to our interfaces presented in this dissertation. Further, our specific novel algorithms (AIBO detection and the *puppet master* algorithm) can be adapted and applied to various scenarios. Our marker-less vision tracking algorithm could be re-targeted at other robotic platforms and morphologies, and we believe that the *puppet master* algorithm can be directly applied to other applications such as a robot handshake, gesture, or body postures.

Our theoretical framework and heuristics, while grounded in our experiences with our own interface designs, are inherently interface and platform independent. They were explicitly composed to be applicable to any robot or interface design. This has been illustrated, for example, both throughout the theoretical exploration, and our explicit application of our three perspectives to existing robotic interfaces (Section 3.6.2, page 82).

We believe that much of our work can scale to non-robotic technologies and the more-general Human-Computer Interaction (HCI) case. This is for our theory as well as our interface designs, implementations, and evaluations. The core of our theory surrounds our own particular definition of robot: machines that a) have a dynamic physical presence in the real world, and b) elicit a sense of agency and intentionality. Thus, our methods can be directly applied to any other technology which meets this criteria (i. e., anything which we would consider a robot with this definition, even if it is not usually considered as such). Further, we believe that our theory would likely apply in many cases where only one of these criteria is met, although we stop short of speculating what such a dynamic would be.

The sociology and social psychology work that we introduced in our theoretical framework, as well as our own additions, can be also used to explore social aspects of interaction with other technologies. We grounded our analysis in robots' unique *embodiment* in the human world, but ideas such as *the social stock of knowledge* and the *holistic interaction context* can be used to analyze social interaction with technologies in general.

Finally, aspects of our interface designs can be used for non-robotic interfaces. We demonstrated how the *puppet master* and *interactive locomotion* approaches also make sense for animation, although people reacted differently than they did for real robots. Another example is how cartoon artwork is used extensively in virtual environments, and can also be used by people for self-expression (Ng and Sharlin, 2010).

In this section we have detailed how we believe that our contributions as presented in this dissertation (the theory, interface designs, implementations, and evaluations, as well as our heuristics) can be generalized across cultures, to other robotic platforms and interfaces, and even to non-robotic interactions. Below, we outline the future next steps leading from our dissertation work.

9.4 FUTURE WORK

The contributions presented in this thesis constitute some of the very first social HRI theoretical frameworks, interface designs, implementations, evaluations, and heuristics. We see this as just the start of social HRI, and as such, our work raises and articulates many new questions to be addressed, as briefly outlined in the sections below.

9.4.1 *Social HRI Framework and Heuristics*

Both our social HRI theoretical framework and heuristics are new, and while emerging from our theoretical explorations and experiences, they have little validation outside the background work, argumentation, and the experience presented in this dissertation. Future analysis should continue to build on this theory, testing and relating the various components to future findings. It also remains to be seen how this work can be applied in practice for future social HRI research efforts; this application as well should continue to reflect on and improve the theories.

Much of this future work revolves around building a deeper understanding of the dynamics behind particular components of the theory such as designers using *the social stock of knowledge* (Section 3.1.2, page 53). This is also the case for our heuristics, for example, while we recommend social HRI researchers to use real robots, we cannot yet describe in detail why and how exactly interaction changes with this variable, we just posit that it does change.

While some of our theory is formulated as tools which can be applied (e. g., the three perspectives on breadth, and the heuristics), future work should continue to build concrete methodologies relating to the application of our theoretical framework. For example, while our factors of technology acceptance (Section 3.6.3, page 89) detail common influences on the perceptions of robots such as *media*, future efforts should construct these into a set of explicit tools for researchers to use. (In this example case, a method for exploring how a given robot relates to the *media* factor.)

Here we discussed the next steps for our social HRI theoretical framework and heuristics. Below we continue our discussion with future work related to social HRI evaluation.

9.4.2 *Social HRI Evaluation*

Future research should continue to explore existing evaluation techniques in terms of how they apply to the field of social HRI. While we presented our own exploration and heuristics in this dissertation, there are many additional methodologies which we did not explore for practicality of scope reasons. We recommend that researchers should carefully apply evaluation techniques in consideration of the context of social HRI.

In terms of our evaluations themselves, they were all conducted in fabricated environments (in-lab or open space) due to practicality constraints with the existing state of technology (e. g., we used expensive, localized tracking systems), and our limited resources. According to our own theory, this has important implications on the *holistic interaction context*. Future research should strive to, whenever possible, evaluate interaction with robots within a more valid context (e. g., using ethnographic methods). Not only will this improve the validity of evaluations, but will also provide an opportunity to reflect on higher-level social structure interaction which is difficult to do in lab settings. For example, can robots better integrate into existing social structures because of their socially-motivated design?

Finally, future research should make a point of performing inter-cultural and international studies to broaden social HRI perspectives. While we acknowledge the importance of more-targeted social HRI understanding, crossing cultural and national borders can help define the fundamental-to-humans (rather than, e. g., fundamental to Canadians, or Japanese) components of social HRI.

Here we have discussed important future work for evaluation in social HRI. In the next section we discuss next directions for our particular interface designs and implementations.

9.4.3 *Social HRI Interface Designs and Implementations*

As explained in our introduction, our particular selection of interface designs represents only a few of the many possible social HRI interface designs. Future work should continue to explore ways that robot designers can leverage *the social stock of knowledge*. Particularly, (using our perspectives from Section 3.6.2, page 82) in our interfaces our focus was primarily on visceral- (P₁) and social mechanics-level (P₂) interaction, although the discussions did touch on social structure (P₃) aspects of the designs. Future work should explicitly target social-structure (P₃) interaction, for example how the presence of a robot can impact the social structures of how others interact, to help better understand how robots can interact as social actors. In the remainder of this section we discuss interface-specific future work.

9.4.3.1 *Dog-Leash Interface*

The dog-leash project was an initial exploration into using a leash to lead a robot, and left many additional research questions unexplored. One question which future work should consider is the impact of using the leash versus no leash, perhaps through studies that compare similar interaction scenarios in both cases. While we argue that the leash provides an important physical link between the person and the robot, and also serves as an awareness mechanism for third-party observers (i. e., informs them that a person is leading a robot), these claims have not been formally investigated. A related question is the general one of social integration, and what other people think of someone leading a robot in their space.

Other remaining questions include: at what distance should the robot follow and at which location, should this be static or should it change based on context (and, which context?), and how do these factors relate to a person being right or left handed? Although we asked some of these questions in our evaluation, research should look to further understand the dynamics of interaction. Here we outlined remaining questions for our dog-leash interface; below we discuss our evaluation of the *touch and toys* project.

9.4.3.2 *Touch and Toys*

We reflect here briefly on some of the social HRI-specific questions which emerged from our study of the Touch and Toys interface. For example, future research should further explore the finding that people showed great concern for the robot-robot collisions, despite reassurance that no damage could be done. Participants further expressed that it was important to use real robots. These questions could be studied, for example, by comparing real robots versus virtual simulation studies. We continue our future work discussion below with reflections on our cartoon-artwork interfaces.

9.4.3.3 *Cartoon Artwork*

Future research on how robots can use cartoon artwork should consider how to improve our implementations beyond bulky hand-held or goggles-based MR, moving to mechanisms (such as on-robot displays or smart phones) that enable a more light-weight interaction. Also, the question of how a person will interact with cartoon elements needs to be explored, as our only solution provided was a hand-held tablet with a stylus.

Several important questions remain about how cartoon artwork relates to social HRI such as a robot's social integration. One challenge is understanding how to create and adjust cartoon content to fit the task and scenario at hand, for example: when should cartoon be used on when not? When is a technique subtle enough? When should a robot be distracting, using bright and distracting cartoons? When should a robot use more gentle cartoon elements? When should robots resort back to physical interaction?

Future research should also continue the exploration of how this technique relates to art for social HRI, as we started with (but by no means intend to limit to) the idea of *robot expressionism*. Related to this, our cartoon-artwork theory needs to be developed further: the MRIE needs to be expanded and integrated with the *Jeeves* cartoon model, and we need to more thoroughly examine which kinds of cartoon elements can be used by robots, and in which additional ways.

We did not perform formal evaluations of the cartoon artwork projects, although we did consider approaches that we present here. For *bubblegrams*, we envisioned (and created early sketched and low-fidelity prototypes, not presented in this manuscript) a task which encouraged the person to engage the interface, and to use the interactive bubbles to accomplish a goal. For example, a search and rescue scenario where the person had to interact with the

robot to complete a task, or a hospital-guide robot, where the person used the interactive bubble to locate a patient, and the robot led them to the correct room.

For *Jeeves*, our evaluation ideas centred around a comparison between a robot with cartoon elements and without. One example is a task where multiple cartoon-art using robots (iRobot Roombas) are set to clean an area, and slowly break down, requiring the person to go and *fix* the robot in a prescribed manner (the experiment was not planned to be realistic and can be implemented using a wizard-of-oz technique). In one case, the robot would indicate that it is breaking down with noises and erratic movements. In the second case, cartoon art elements would be added. This evaluation would encourage the person to interact with the robot, while setting the cartoon artwork as the independent variable, helping to isolate the impact cartoon artwork has on interaction.

9.4.3.4 *Stylistic Locomotion and Puppet Master*

Future research related to our stylistic locomotion should explore how to expand the core ideas, communicating through style of actions and SBD, beyond our current narrow focus of locomotion path to, for example, full-body gestures. Our initial exploration into enabling robot sounds in the robotic *puppet master* showed promising results. Further, we recommend that researchers bring emotion theory to this work, such as Plutchik's (1980) wheel of emotions, to better inform the selection of stylistic behaviours, and to have a better vantage point from which to analyze how people interact with them. This can further serve as a means to compare how people interpret other people's emotions to robotic synthetic ones.

Another open research question is, how does the meaning of a stylistic behaviour change between when the robot is directly interacting with the person (as in our robotic *puppet master* observation study), and when the the person is observing the robot interacting with someone else (as in our robotic *puppet master* designer study).

For *puppet master*, suggestions from our study participants pointed to the idea of combining the existing raw SBD approach with more-explicit components, such as enabling people to use some kind of emotional sliders or feedback correction mechanism. Also, research should consider how the light-weight nature of our SBD systems can serve as an aid in behavioural rapid prototyping. To aid in these considerations, researchers should examine how people teach other people similar kinds of behaviours, and build the robot systems scaffolded around these existing approaches in *the social stock of knowledge*.

The *puppet master* algorithm itself will need to be extended and adapted to accommodate many of the above suggestions. In relation to the current work, the algorithm needs to be improved to reduce the problem of jitter, and perhaps extended to include multiple reacting entities (not only one-on-one, as in the current algorithm).

In this section we have outlined the various next research steps to build on the foundations set in this dissertation. We conclude below with a few final words.

9.5 FINAL WORDS

In this dissertation we presented our foundational discussion and theoretical framework for social HRI, the design, implementation, and evaluation of several novel social HRI instances, and a set of social HRI-targeted heuristics. Through these original contributions we have defined and outlined the field of social HRI and connected it to existing related work in the domain. Our contributions highlight the importance — and indeed the necessity — of considering the broad social aspects of interaction between people and robots, and present answers and reasons behind why social interaction is naturally occurring and what it means for interaction. We have demonstrated through designing and implementing interfaces how peoples' social tendencies toward robots can be leveraged to create interfaces that abstract away complex and alien computer concepts, creating interaction that is easy to use and easy to understand. Finally, based on our thorough exploration we have presented several sets of heuristics and tools for social HRI design, analysis, and evaluation.

One overarching lesson that we learnt throughout this dissertation is the usefulness of the theory, design and implementation, and evaluation design cycle. Although we did not chronologically present the work as such, we performed several design cycles throughout our work, where the evaluation from one project led directly into the design of another (or a re-design of the same interface). Further, this has played a crucial role in the evolution of our overarching social HRI theory.

Our social HRI perspective as presented in this dissertation is but one part of a bigger picture of how people interact with robots. It is our hope that in addition to contributing to the emerging domain of social HRI, our work will fuel further developments relating to how people can interact with robots, and new interface designs that enable them to do so.

Part IV

APPENDIX



TOUCH AND TOYS STUDY MATERIALS

This appendix contains the study materials used for our study on the touch and toys project (Chapter 5, page 117). We provide the study materials as outlined below.

A.1 PROTOCOL AND QUESTIONNAIRES

A.1.1 participant consent form.

A.1.2 experiment protocol

A.1.3 pre-test questionnaire

A.1.4 one-robot questionnaire *

A.1.5 post-one-robot questionnaire

A.1.6 two-robot questionnaire *

A.1.7 post-two-robot questionnaire

A.1.8 three-robot questionnaire *

A.1.9 post-three-robot questionnaire

A.1.10 post-test questionnaire

* – these questionnaires are used twice. Once as shown, and once with the “toy interface” title text changed to “touch interface.”



Dr. Ehud Sharlin and Cheng Guo
 Department of Computer Science
 University of Calgary
 2500 University Drive
 Calgary, AB, CANADA T2N 1N4

Consent Form for Participants

Research Project: RICON User Study

Investigators: Dr. Ehud Sharlin, Cheng Guo and Jim Young

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail 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.

Note: The University of Calgary Conjoint Faculties Research Ethics Board has approved this research study.

Description of Research Project:

The purpose of this study is to explore the possibility of using Tangible User Interface and Touch Interface for remotely controlling multiple robots. The entire study is divided into three parts. You will be asked to control one, two and three robots to follow target points in each part of the study. Before each study, we will teach you how to use each interface and let you practice first. After you have grasped the concept, we will start the real experiment. The entire experiment will take 60 minutes. The experiment will be video taped and your comments about the experiment will be audio recorded.

Participation in this study will not put you at any risk or harm and is strictly voluntary. You choose to participate by playing the AIBO with two different controllers. You may choose to withdraw from the study at any time by simply not using the system any more. Any data collected to your withdrawal will still be available to the investigators for analysis. Personally identifiable information will only be used in papers or presentations with your explicit permission. If we wish to use any personally identifiable information, we will contact you with the particulars of the information we wish to use, and you may decide whether or

not you give us permission to use it. In this study, the personal information we will collect are your name, age and handedness which will be used only for identification purposes and grouping results. There are several options for you to consider if you decide to take part in this research. You can choose all, some or none of them. Please note that in any case we will not expose your name or identity. However, if you grant us permission, we may use your picture during interaction in academic publications/presentations about this research. Please put a check mark on the corresponding line(s) to grant me your permission to:

I grant permission to be audio taped:

Yes: ___ No: ___

I grant permission to be videotaped:

Yes: ___ No: ___

I grant permission to have quotations from my comments answers that are recorded during the study to be used in publications and/or presentations (note that your identity will never be associated with the quotations):

Yes: ___ No: ___

I grant permission to have video or still images of me used in publications and/or presentations:

Yes: ___ No: ___

If researchers wish to include information that may identify me, such as my picture or video, in reports of the data, I prefer the researchers to re-contact me for permission:

Yes: ___ No: ___ If Yes, Please leave your contact information:

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a 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 the study at any time. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

At the conclusion of the study and its analysis, we will post any resulting papers that we have written about them. You can view these by asking the investigator or by accessing the website: <http://groupplab.cpsc.ucalgary.ca/papers/index.html>

Electronic data will be stored in a secure manner, such as in a computer secured with a password. Hardcopies of data will be stored in a locked cabinet/room located at the University of Calgary Interactions Laboratory (Math Science building, room 680) with restricted access. Data will be kept for a minimum of three years and a maximum of 7 years. On disposal, electronic data will be erased and hardcopies will be shredded.

If you have further questions concerning matters related to this research, please contact:

Dr. Ehud Sharlin (403) 210-9499
 ehud@cpsc.ucalgary.ca

Ricon Experiment Protocol

<Remarks in brackets are directed for the administrator only>

1. Introduction

"Hello, my name is Cheng. Today, we will perform an experiment involving remote robot control. I'll briefly describe the concept of our project, talk about the prototype application that we have developed, and the procedure of the experiment."

"The goal of the experiment is to compare and contrast two different interaction techniques for controlling a group of robots. The two interaction techniques are Toy interface <show the toys to the participant> and Touch interface <briefly explain what it is>." hmmm. Make sure to introduce properly the table and how it works.

"The experiment consists of three parts, for the first part, you will be asked to navigate a single robot by following target points on the table. For the second and third part, you will control two and three robots respectively to complete the same task. During the experiment, you will be observed and data will be collected for further analysis. The entire experiment will be video taped. Also, we will audio record some of the questionnaires that will be given to you during and after the experiment. You may quit the experiment at anytime if you don't want to continue." Add a note here about privacy, anonymous data.

"Would you like to participate in our study?"

2. Signing the consent form (sign, not sing)

3. Participants are asked to complete the pre-study questionnaire

4. Training for experiment part 1

"Now, I am going to show you how to use the Toy/Touch interface for controlling a Roomba vacuum cleaner."

<Demonstrating one of the interfaces depending on the order>

<1) Tell the user that the robot's movement is imprecise>

<2) Tell the user that the green circle indicates that the robot has reached the target>

"Try it yourself. Once you feel comfortable with this interface, please tell me and we will start the real experiment. Please ask if you have any questions."

5. Experiment Part 1 Start




Figure B.1 – Task 1

<The arrows indicate the orientation of the robot. The number inside the circle indicates the order of steps. >

<Test administrator resets the robot (**Roomba**) at target 1>

"Now, I am going to use these images <Images of the actual robot printed on a piece of paper> to indicate the next target location and orientation of the robot. All you need to do is to make the robot to move to this location and align itself correctly. Once the robot stops (indicated by the green circle) then I will show you the next target location. We will repeat this process until I tell you the experiment is done."

"You have probably noticed that due to the imprecise movement of the robot, it does not move onto the point where you want it to be. In stead, it will just somewhere close to it. When we run the experiment, you don't have to worry about this problem. All you need to do is to move the toy/icon (depending on the interface) onto the location that I indicated on the white board. Any questions"

<Start the experiment> <hand-time how long they take>

“Great, now I want you to try another robot <**AIBO**> with the same task. I will let you try out this robot first and then we will repeat the previous experiment.”

<Repeat previous experiment>

Do a questionnaire before changing interfaces. We have one for touch and one for toy.

<Change the Interface and repeat the task again>

6. In-between study questionnaire.

“Please fill out the questionnaire.”

7. Training for experiment part 2

“Now, we are going to start the second part of our experiment. For this part, I am going to ask you to control two robots on the table.” – do the same thing, waypoints

<Use **Roomba & AIBO** or **AIBO & AIBO** depending on the order>

“Please try out the interface and we will start the experiment.”

8. Experiment Part 2 Start

“The experiment procedure is going to be the same as the previous experiment. I will show you the target point of both robots on the white board. You follow the waypoints until the experiment ends. When both robots stop on the target location, I will reveal the next location. Any questions? If no, then let's start.”




Figure B.2 – Task 2

<Change the robot set and allow the participant to practice until he/she is comfortable>

<Repeat the same task again>

<Change the Interface and repeat the task again>

9. In-between study questionnaire.

“Now we are done with the second part of the experiment. Please fill out the questionnaires.”

10. Training for experiment part 3

“Let's start the last part of the experiment. This time, I am going to give you three robots for you to control. Please try out the interface and see if you have any questions.”

<Let the user practice with 3 robots>

“Just like the previous two experiments, I will show you the target location of each robot, you make them to move to their locations. Any questions?”



Figure B.3 – Task 3 Formation 1

<First formation>

<The green rectangle marks the start position. The red rectangle marks the target position>



Figure B.4 – Task 3 Formation 2

<Second formation>

<The participant has to change from the first formation to the second formation>



Figure B.5 – Task 3 Formation 3

<Third formation>

<The participant has to change from the second formation to the third formation>



Figure B.6 – Task 3 Formation 4

<Fourth formation>

<The participant has to change from the third formation to the fourth formation>



Figure B.7 – Task 3 Formation 5

<Fifth formation>

<The participant has to change from the fourth formation to the fifth formation>

<Change the Interface and repeat the task again>

11. In-between study questionnaire.

“Please fill out the questionnaires.”

12. Post-study questionnaire & debriefing

“Thank you very much for your participation today. Now, you have done all of the experiment. We’d like to know your overall feeling about this experiment.”

<Ask the participant to fill out the post-study questionnaires>

Make sure to include interview 1-on-1 time in this

13. Pay the participant

Pre-Study Questionnaire

Age: _____ Handedness: Left-Handed or Right-Handed or Ambidextrous
Gender: Male Female

1. How familiar are you with touch-screen interfaces?

1	2	3	4	5
Never Seen Before	Never Used Before	Somewhat Familiar	Very Familiar	Expert

2. Do you have experience with remote robot control?

Yes No

If "Yes", then how familiar are you with robot remote controlling interface?

1	2	3	4	5
Only Used Once	Used a few Times before	Somewhat Familiar	Very Familiar	Expert

3. Do you consider yourself to be good at working with your hands? For example, good at woodworking or sewing?

Yes No

Do you consider yourself to have good hand-eye coordination?

Yes No

4. How often do you play video games?

1	2	3	4	5
Never	Yearly	Monthly	Weekly	Daily

One Robot Remote Control – TOY INTERFACE

To what extent do you agree / disagree with the following statements?
 (if you feel there is no difference between the Roomba and AIBO, just fill out one set)

With the toy interface...	Roomba							AIBO						
	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
...it was easy to move the robot to the target location (this does not include the orientation of the robot). Comments:	1	2	3	4	5	6	7	1	2	3	4	5	6	7
...it was difficult to rotate the robot in the way that I wanted. Comments:	1	2	3	4	5	6	7	1	2	3	4	5	6	7
...I had precise control over the robot movement.	1	2	3	4	5	6	7	1	2	3	4	5	6	7

Comments:							
...the robots generally did not react as expected. Comments:	1	2	3	4	5	6	7

1	2	3	4	5	6	7	

Any comments on the difference between controlling the two different robots?

Any additional thoughts or comments?

(1 robot)

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
Overall, I preferred the touch interface Comments:	1	2	3	4	5	6	7
Overall, I preferred the toy interface. Comments:	1	2	3	4	5	6	7

Two Robots Remote Control – TOY INTERFACE

To what extent do you agree / disagree with the following statements?
 (if you feel there is no difference between the robot configurations, just fill out one set)

AIBO & AIBO

Roobma & AIBO

With the toy interface...

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
...it was easy to move the robot to the target location (this does not include the orientation of the robot). Comments:	1	2	3	4	5	6	7
...it was difficult to rotate the robot in the way that I wanted. Comments:	1	2	3	4	5	6	7
...I had precise control over the robot movement.	1	2	3	4	5	6	7

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
	1	2	3	4	5	6	7
	1	2	3	4	5	6	7
	1	2	3	4	5	6	7

Comments:							
...the robots generally did not react as expected. Comments:	1	2	3	4	5	6	7
...it was confusing to monitor both robots at the same time. Comments:	1	2	3	4	5	6	7
...it was easy to control the two robots at the same time Comments:	1	2	3	4	5	6	7
...I worked with both robots at the same time, operating them simultaneously Comments:	1	2	3	4	5	6	7

	1	2	3	4	5	6	7
	1	2	3	4	5	6	7
	1	2	3	4	5	6	7
	1	2	3	4	5	6	7

...I worked with one robot at a time, operating them sequentially. Comments:	1	2	3	4	5	6	7
...I often used both of my hands at the same time. Comments:	1	2	3	4	5	6	7

1	2	3	4	5	6	7	
1	2	3	4	5	6	7	

Any comments on the difference between the two robot configurations?
Any additional thoughts or comments?

(2 robot)

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
Overall, for two robots I preferred the touch interface Comments:	1	2	3	4	5	6	7
Overall, for two robots I preferred the toy interface Comments:	1	2	3	4	5	6	7

Were there any particular changes or differences that you encountered with two robots that you did not find with the one robot?

Three Robots Remote Control – TOY INTERFACE

To what extent do you agree / disagree with the following statements?

With the toy interface...

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
...it was easy to move the robot to the target location (this does not include the orientation of the robot). Comments:	1	2	3	4	5	6	7
...it was difficult to rotate the robot in the way that I wanted. Comments:	1	2	3	4	5	6	7
...I had precise control over the robot movement. Comments:	1	2	3	4	5	6	7
...the robots generally did not react as expected. Comments:	1	2	3	4	5	6	7
...it was easy to form the group formations. Comments:	1	2	3	4	5	6	7

...it was confusing to monitor all three robots at the same time. Comments:	1	2	3	4	5	6	7
...it was easy to control the three robots at the same time. Comments:	1	2	3	4	5	6	7
...I worked with all three robots at the same time, operating them simultaneously. Comments:	1	2	3	4	5	6	7
...I worked with one robot at a time, operating them sequentially.. Comments:	1	2	3	4	5	6	7
...I often used both of my hands at the same time. Comments:	1	2	3	4	5	6	7

Any additional comments or thoughts?

Any additional comments or thoughts?

(3 robot)

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
...overall, for three robots I preferred the touch interface. Comments:	1	2	3	4	5	6	7
...overall, for three robots I preferred the toy interface. Comments:	1	2	3	4	5	6	7

Were there any changes or differences with the three robot case that you did not notice or find with the one and two robot cases?

Post-Study Questionnaire

To what extent do you agree / disagree with the following statements?

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
I found the graphical feedback on the table easy to understand. Comments:	1	2	3	4	5	6	7
The graphical feedback on the table was unnecessary. Comments:	1	2	3	4	5	6	7

Please describe the benefits that you noticed, if any, of the toy interface.

Please describe the benefits that you noticed, if any, of the touch interface.

Please describe the problems that you noticed, if any, of the toy interface.

Please describe the problems that you noticed, if any, of the touch interface.

Would you have rather done this experiment on a standard desktop PC? Why or why not?

You were controlling real robots. Instead, we could have done this with a simulation. Would this have been better? Why or why not?

Where else – besides robot control – could you imagine using the interfaces you used today?

DOG-LEASH ROBOT STUDY MATERIALS

B.1 MATERIALS AND QUESTIONNAIRES

This section contains the following documents. The measurement of Negative Attitudes towards Robots Scale (NARS) questionnaires are included from (Nomura et al., 2006) with minor format modifications. The the Self-Assessment Manikin (SAM) questionnaires and instructions are included from (Morris, 1995), with our original translations.

B.1.1 consent form

B.1.2 pre-test questionnaire

B.1.3 measurement of Negative Attitudes towards Robots Scale (NARS)*

B.1.4 the Self-Assessment Manikin (SAM) instruction sheet

B.1.5 per-condition SAM

B.1.6 post-test questionnaire

* – this same questionnaire was used pre and post test.

Signed Consent Form for Participants

Name of Researchers:

Juliane Reichenbach, (Psychology, TU Berlin) juliane.reichenbach@gmail.com
James Young, (Computer Science, University of Calgary) jim.young@ucalgary.ca
Dr. Takeo Igarashi (Computer Science, University of Tokyo)

Title of Project: *Dog leash robot*
Sponsored by JSPS

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. Thank you very much for your involvement in this project!

Purpose of the Study

The purpose of this study is to evaluate a new system of robotic behaviors. At Igarashi Lab in 東京大学 / ERATO we developed a new interface called the dog-leash robot, a technique where people can lead a robot using a leash as they may lead a dog. In this study we want to explore how people react to this new interface and perceive various differing robot leading and following styles.

What Will I Be Asked To Do?

We will present you with a robot which will display different behaviors. You will do various tasks with the robot and answer questionnaires.

The study will take approximately one hour from start to finish.

You will receive 4400Yen for your time. This includes money for transfer to ハウスケア横浜.

Participation in this study is voluntary. You are free to withdraw from the study at any time.

At the conclusion of the study and its analysis, we will post any resulting papers that we have written about them. You can view these by asking the investigator.

What Type of Personal Information Will Be Collected?

All data that is collected in this study will only be used anonymously for statistical purpose. In this study we will record personal data like your sex, age, and computer experience. We may also, with your explicit permission, videotape all, or parts of your session, and, as such will limit your anonymity since videotaped portions of your participation may be shown in public presentations (at conferences). Researchers will be unable to control any further use of images after they are presented in a public forum, which may result in these images being reposted in some unknown context, including possibly on the internet.

Your name will be recorded for the cash payout and for this informed consent form only. This information will not be correlated with or matched to the study results in any way.

What Happens to the Information I Provide?

Collected information, including audio or video recordings, will only be used in academic papers and presentations, and will be presented as anonymous data using generic identifiers.

Signatures (written consent)

I grant permission to be videotaped: yes no

I grant permission to show parts of the video recordings on conferences: yes no

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: _____

A copy of this consent form has been given to you to keep for your records and reference.

私は実験の記録ビデオを撮られることに同意します。 はい いいえ

私は実験の記録ビデオが学術目的での研究会や国際会議で使用されることに同意します。 はい いいえ

署名 _____

1Q Participant ID

Date 記入日

Please answer the following questions.
以下の質問にお答えください。

Gender 性別 male 男 female 女

Age 年齢

What is your nationality? 国籍

What is your major in university? 大学での専攻

Please rate your technical computer ability. コンピュータに関する知識について

1 None 全くない 2 3 4 I can install new software ソフトウェアをインストールできる程度 5 6 7 I have a technical degree コンピュータに関する知識は十分持っている

Do you play computer games?

どのくらいの頻度でゲーム（テレビゲームやニンテンドーDSなどの携帯ゲーム機、携帯電話でのゲームを含む）をしますか？

1 Yes, every day 毎日します 2 Yes, several times per week 週に何度かします 3 Yes, once per week 週に一度します 4 Yes, several times per month 月に何度かします 5 Yes, once per month 月に一度します 6 Yes, less than once per month 月に一度するかもしれない程度です 7 No, never 全くしません

If yes, which games do you play and which platforms (Wii, Playstation ...) do you use?

ゲームをすると答えた方（上の質問の6以下の方が対象）に質問です。どのゲーム機でゲームをしますか？(Wii, Playstation, ニンテンドーDS, PSP, 携帯電話など)

Do you have any experience with robots? Yesはい Noいいえ

ロボットを使ったことがありますか？

If yes, how much (months, years)

どのくらいの期間ですか？

and what kind? (private, work)

どのようなロボットですか？

それは家庭で使用しましたか？

それとも仕事ですか？

1Q Participant ID

Date 記入日

Do you know any of the following robots? 下のリストにあるロボットを知っていますか？

	No 知らない	Heard about it / seen it in the media 聞いたことがある/雑誌等で見ただことがある	Seen the real robot 実物を見たことがある	Played / Worked with it 触ったことがある/使ったことがある	I own one 持っている
Aibo					
Kismet					
Mars Explorer					
Asimo					
Lego Mindstorms					
Roomba					
R2D2					

Do you have pets? ペットを飼っていますか？

Yesはい Noいいえ

If yes, what kind of pet? (はいの場合) 種類をお答えください。

.....

Have you ever walked a dog on a dog leash before? 犬の散歩をしたことがありますか？

Never 一度もない Sometimes たまにする Often 良くする

What is your preferred hand? どちらの腕が開き手ですか？

Left 左手 Right 右手

Do you have a drivers licence for a car? 車の運転免許を持っていますか？

Yesはい Noいいえ

Do you have experience with other advanced machinery? (e.g., forklift, boat, etc.)

他種の運転免許を持っていますか？(フォークリフトやボードなど)

Yesはい Noいいえ

<p>ロボットに対する気持ちに関して、14 個の文章が以下に書かれています。</p> <p>それぞれの文書に対して、「全くそう思わない」、「それほどでもない」、「どちらともいえない」、「まあそうだ」、「全くそう思う」のいずれかにチェックを付けて下さい。</p> <p>正しい答、又は間違った答はありません。</p> <p>幾つかの文章は他の文章と似通っています。そのことに関して気にしないでください。</p> <p>時間をかけずに、第一印象で答えてください。</p>		全くそう思わない	それほどでもない	どちらともいえない	まあそうだ	全くそう思う
1	もしロボットが本当に感情を持ったら不安だ。					
2	ロボットが生き物に近づくと、人間にとってよくないことがありそうな気がする。					
3	ロボットと会話すると、とてもリラックスできるだろう。					
4	就職してロボットを利用するような職場にまわされるかもしれないと考えると、不安になる。					
5	ロボットが感情を持ったら、親しくなれるだろう。					
6	感情的な動きをするロボットを見ると、気分がいやされる。					
7	ロボットと聞いただけで、もうお手上げの気持ちだ。					
8	人が見ている前でロボットを利用すると、恥をかきそうだ。					
9	人工知能とか、ロボットによる判断といった言葉を聞くと不愉快になる。					
10	私は、ロボットの前に立っただけで、とても緊張してしまうだろう。					
11	ロボットに頼りすぎると、将来、何か良くないことが起こりそうな気がする。					
12	ロボットと会話をすると、とても神経過敏になるだろう。					
13	ロボットが子供の心に悪い影響を与えないか心配だ。					
14	これからの社会は、ロボットによって支配されてしまいそうな気がする。					

御協力、ありがとうございました。

Here a 14 statement about robots. You will probably find that you agree with some of the statement and disagree with others, to varying extents. Please indicate your reaction to each of the statements. There is no right or wrong answer, just give your first impression.		Strongly disagree	disagree	neutral	agree	strongly agree
1	I would feel uneasy if robots really had emotions.					
2	Something bad might happen if robots developed into living beings.					
3	I would feel relaxed talking with robots.					
4	I would feel uneasy if I was given a job where I had to use robots.					
5	If robots had emotions I would be able to make friends with them					
6	I feel comforted being with robots that have emotions.					
7	The word "robot" means nothing to me.					
8	I would feel nervous operating a robot in front of other people.					
9	I would hate the idea that robots or artificial intelligences were making judgements about things.					
10	I would feel very nervous just standing in front of a robot.					
11	I feel that if I depend on robots too much, something bad might happen.					
12	I would feel paranoid talking to a robot.					
13	I am concerned that robots would be a bad influence on children.					
14	I feel that in the future society will be dominated by robots.					


3Q Participant ID

Date 記入日

SAM

During this study you will meet SAM. This is SAM:

今回の実験を通してあなたは、何度か SAM と我々が呼んでいるキャラクターを見る機会があります。これが SAM のイメージです。



SAM stands for SELF ASSESSMENT MANIKIN. SAM tries to depict you and your emotions. You will use SAM today to help describe your emotional state.

SAM とは Self Assessment Manikin (自己評価のためのマネキン) という意味であり、あなたと、あなたの心の状態を表すものです。今回の実験ではこの図を使って、あなたの心の状態を示してもらいます。

An emotion can be seen from two perspectives:

心の状態とは以下の二つの側面があると考えられます。

- Pleasure - Displeasure: positive or negative feeling
快 - 不快 (快不快度) : 肯定的か否定的な気持ち
- Degree of Arousal: excited - calm
覚醒度 : 興奮している - 落ち着いている

Examples: 例えば、

fear = Pleasure: negative, Arousal: high
怖い = 快不快度 : 否定的 (不快)、覚醒度 : 高い

relaxation = Pleasure: positive, Arousal: low
リラックス = 快不快度 : 肯定的 (快)、覚醒度 : 低い

enthusiasm = Pleasure: positive, Arousal: high
熱狂 = 快不快度 : 肯定的 (快)、覚醒度 : 高い

So you will be asked to give your rating of your current emotional state in 2 rows:

この例のように、あなたには今回の実験で、以下の2つの指標で自分の心の状態を示してもらいます。


3Q Participant ID

Date 記入日

1. Pleasure: 1. 快不快度

The first row shows SAM going from a smiling happy figure to a frowning unhappy figure. This row depicts feelings going from positive feelings, like "contented" or "proud" or "secure", to negative feelings like "discontented" or "humiliated" or "insecure".

一つ目の指標は SAM が「笑っていて幸せそうな状態」から、「辛そうにしていて不幸そうな状態」までを表しています。これは「満足して」「自信があり」「安心している」というポジティブな状態から、「不満があり」「屈辱を感じ」「不安を感じる」というネガティブな状態を表しています。



"contented"
"proud"
"secure"
満足している
自信がある
安心している


"discontented"
"humiliated"
"insecure"
不満がある
屈辱を感じる
不安を感じる

2. Arousal: 2. 覚醒度 :

The second row shows SAM going from excited with eyes open to relaxed, sleepy with eyes closed.

This row depicts your arousal from "anxious" or "excited", to "quiet" or "relaxed".

二つ目の指標は「目が大きく開いて、興奮している状態」から「リラックスするように目を閉じて、眠そうな状態」までを表しています。この覚醒度は「興奮している」「熱望している」ことを示している状態から、「静かに落ち着いている」「リラックスしている」状態までを示しています。



"anxious"
"excited"
興奮している
熱望している

"quiet"
"relaxed"
静かで落ち着いている
リラックスしている

3Q Participant ID

Date 記入日

Here is an example of how it will look overall:

以下は SAM がどのように使われるのかを示した例です。



Please mark how you feel at the moment by putting a **mark in both lines** (pleasure and arousal)

例)

いまの心の状態を以下の二つの指標（快不快度、覚醒度）の中で適切なものにチェックをしてください。

Please describe your current emotional state using SAM:

以下の SAM の図を使ってあなたの現在の心の状態を説明してください。



4Q Participant ID



Robot behavior

Date 記入日

How do you feel now?

Please tell us about you current emotional state using SAM.

ロボットの動作を見て、どのように感じられたでしょうか？
いまのあなたの気持ちの状態を SAM を使って教えてください。



Please rate your emotional state on these scales:

以下のスケールに基づいてあなたの心の状態を評価してください：

Anxious 不安な	1	2	3	4	5	6	7	Relaxed 落ち着いた
Agitated 動揺している	1	2	3	4	5	6	7	Calm 冷静な
Quiescent 平穏な	1	2	3	4	5	6	7	Surprised 驚いた
Unpleasant 不愉快な	1	2	3	4	5	6	7	Comfortable 快適な

Overall I enjoyed the interaction. 全体を通して、私は実験を楽しめたと思う：

Disagree strongly 全く違うと思う	1	2	3	4	5	6	7	Agree strongly 強くそう思う
------------------------------	---	---	---	---	---	---	---	--------------------------

Comments: 理由を教えてください：

4Q Participant ID

Robot behavior

Date 記入日

Please rate your impression of the robot on these scales:

以下のスケールに基づいてこのロボットの印象を評価してください。

Dislike 嫌い	1	2	3	4	5	6	7	Like 好き
Unfriendly 親しみにくい	1	2	3	4	5	6	7	Friendly 親しみやすい
Unkind 不親切な	1	2	3	4	5	6	7	Kind 親切な
Unpleasant 不愉快な	1	2	3	4	5	6	7	Pleasant 愉快的な
Awful ひどい	1	2	3	4	5	6	7	Nice 良い
Aggressive 攻撃的な	1	2	3	4	5	6	7	Non-aggressive 友好的な
Not Controllable 制御不可能	1	2	3	4	5	6	7	Controllable 制御可能な
Not Predictable 予測不可能	1	2	3	4	5	6	7	Predictable 予測可能な
Not Autonomous 自動的ではない	1	2	3	4	5	6	7	Autonomous 自動的な

4Q Participant ID

Robot behavior

Date 記入日

You just experienced one possible robot behavior. How was it?
今回みたロボットの動作について、どのように感じましたか？

What did you like?
どの点が良いと思いましたか？

What did you not like?
どの点が悪いと思いましたか？

If you had to think of a name for this following behaviour, how would you call it?
今回みたロボットの動作に名前を付けるとしたら、どのような名前を付けますか？

4Q Participant ID

Robot behavior

Date 記入日

Please circle one answer.
I liked the robot's behavior. 以下の質問に関して、適したものを選択してください。
Disagree strongly 1 2 3 4 5 6 7 Agree strongly
全く違うと思う 強くそう思う

Comments: 理由を教えてください：

I felt in control of the robot. ロボットをしっかりコントロールできたと感じた。
Disagree strongly 1 2 3 4 5 6 7 Agree strongly
全く違うと思う 強くそう思う

Comments: 理由を教えてください：

I felt safe using the robot. ロボットを安全に使うことができたと思った。
Disagree strongly 1 2 3 4 5 6 7 Agree strongly
全く違うと思う 強くそう思う

Comments: 理由を教えてください：

The robot did what I wanted it to do. ロボットが自分の思い通りに動作したと思った。
Disagree strongly 1 2 3 4 5 6 7 Agree strongly
全く違うと思う 強くそう思う

Comments: 理由を教えてください：

5Q Participant ID

Date 記入日

You have experienced 3 different possible following behaviors of the robot.
Did you prefer any following position?
Please try to rank order them.

実験へのご協力ありがとうございます。
今回の実験では3種類のロボット動作を経験して頂きましたが、3つの動作のうち、どの動作が一番好ましいと感じたでしょうか？
以下の例のように順位を付けてください。

Example 例：

	Strawberry いちご	Apple りんご	Banana バナナ
I liked ... best 好きな順番	3	1	2

Please rank order the robot's behaviors. ロボットの行動について順位を付けてください。

	follow directly behind 真後ろにロボットが付いてくる動作	follow behind at the side 斜め後ろにロボットが付いてくる動作	lead in front あなたの前をロボットが行く動作
I liked ... best. 好ましい順番			
I felt in control the most with ... 操作できたと感じる順番			
I felt the robot did what you wanted it to do the most for ... ロボットが自分の思い通りに動いたと感じる順番			

Comments: 理由を教えてください：

Imagine your friend gets a new robot like the one you just tested. Imagine the robot can be set to do only one following behavior, and your friend has to decide for one following position. Which following position would you recommend your friend? Why?
もし、あなたの友人が、あなたが今日使ったロボットを購入するとして、今日のロボットが3種類の動作のうち1種類しか購入できないとしたら、あなたはどれをお勧めしますか？また、それはなぜですか？

5Q Participant ID

Date 記入日

If you ever consider buying a household robot, do you think you would care about the style of the robot's behaviors, such as you saw in today's study?

家庭用ロボットの購入を考える際に、ロボットの動作や振る舞いについて考慮しますか？たとえば、今日使ってもらったロボットについてはどのような点が気になりましたか？

If you could change the robots behaviour, what would you change?

もし、ロボットの動作を変更できるとしたら、どのように変更したいですか？

If possible, which additional features would you add?

もし、ロボットに機能を追加することができたとしたら、どのような機能を追加したいですか？

Do you have any additional positive comments on the robot in today's study?

実験で使用したロボットについて、良いと思った点がありましたら、記入してください。

5Q Participant ID

Date 記入日

Do you have any additional negative comments (problems, room for improvement, etc) on the robot in today's study?

実験で使ったロボットについて、悪いと思った点（問題点や改善点など）がありましたら、記入してください。

Finally, any last comments, ideas, suggestions, etc?

最後に、全体を通したコメントやご意見などがあれば記入してください。

What role do you think will robots play in our lives in the future?

将来、ロボットは私たちの生活の中でどのような役割を担うと思いますか？

Where do you think will robots be used in the future?

将来、ロボットはどのような場所で使われると思いますか？

<input type="checkbox"/> Factories	<input type="checkbox"/> 工場
<input type="checkbox"/> Hazardous locations (contaminated areas, battlefield...)	<input type="checkbox"/> 危険地帯（核汚染された地域や、戦場など）
<input type="checkbox"/> Remote locations (deep sea, space...)	<input type="checkbox"/> 人が立ち入れない場所（深海や宇宙空間など）
<input type="checkbox"/> Homes	<input type="checkbox"/> 家庭
<input type="checkbox"/> Offices	<input type="checkbox"/> 職場
<input type="checkbox"/> Schools	<input type="checkbox"/> 学校
<input type="checkbox"/> Hospitals	<input type="checkbox"/> 病院

B.2 TASK INSTRUCTIONS

This section includes the instructions as presented to the participant

Your task (1.1)

A) go to A and **pick up pen box**

B) walk around R

go through O and Q,

then walk around P and

pick up 3 pens at B

C) go to C and **pick up stamp**

D) go to D and **deliver 1 pen**

E) go through O and Q,

then walk around R,

and **deliver pen box and 1 pen** at E

Your task (1.2)

A) go to A and **deliver 1 pen**

B) walk around R

go through O and Q,

then walk around P and

deliver stamp at B

C) go to C and **pick up ruler**

D) go to D **pick up postcard**

E) go through O and Q,

then walk around R,

and **deliver post card and ruler** at E

Your task (2.1)

A) go to A and **pick up chopstick**

B) walk around R

go through O and Q,

then walk around P and

pick up chopstick box at B

C) go to C and **pick up Lilo & Stich blanket**

D) go to D and **deliver chopstick box and chopstick**

E) go through O and Q,

then walk around R,

and **deliver Lilo & Stich blanket** at E

and **pick up postcard, ruler and pen box** at E

Your task (2.2)

A) go to A and **deliver pen box**

B) walk around R

go through O and Q,

then walk around P and

pick up stamp at B

C) go to C and **deliver postcard and ruler**

D) go to D and **pick up chopstick box and chopstick**

E) go through O and Q,

then walk around R,

and **deliver chopstick box and chopstick** at E

Your task (3.1)

A) go to A and **pick up paper roll**

B) walk around R

go through O and Q,

then walk around P and

deliver stamp at B

C) go to C and **pick up ruler and postcard**

D) go to D and **deliver paper roll**

E) go through O and Q,

then walk around R,

and **deliver post card** at E

Your task (3.2)

A) go to A and **pick up pen and pen box**

B) walk around R

go through O and Q,

then walk around P and

pick up stamp at B

C) go to C and **pick up pen**

D) go to D and **pick up pen and paper roll**

E) go through O and Q,

then walk around R,

and **deliver everything you have on your box** at E

STYLISTIC LOCOMOTION AND PUPPET MASTER STUDIES, ANIMATION

This appendix contains the study materials used for our animation-based studies on our *stylistic locomotion* and *puppet master* interfaces (Section 7.5, page 209). We provide the study materials for the designer and then the observer study.

C.1 DESIGNER STUDY

These documents pertain only the designer study. The observer study is explained in Section C.2. The outline of this section is provided below:

C.1.1 participant consent form

C.1.2 experiment protocol

C.1.3 pre-test questionnaire

C.1.4 post-test questionnaire

C.1.5 behaviour-matching sheet

C.1.6 behaviour description sheet*

* – we have only included one behaviour description sheet, for the *lovers* character condition. In the study we used five identical sheets with this label changed: one for the each of the *lovers*, *bully*, *playful friend*, *stalker*, and *afraid* conditions.



UNIVERSITY OF
CALGARY

Dr. Ehud Sharlin and James E. Young
Department of Computer Science
University of Calgary
2500 University Drive
Calgary, AB, CANADA T2N 1N4

Consent Form for Participants

Research Project: Exploratory study of Reactions to Social Computerised Entities

Investigators: Dr. Ehud Sharlin and James E. Young

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail 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.

Description of Research Project:

The purpose of this study is to evaluate a new system for creating animated computer characters by demonstration. We will ask you design animated characters using our system. To start, we will explain how our system works and show you how to design the animations. The entire experiment will take around one hour and you will be paid \$15 for your participation. This experiment has been reviewed and approved by the Conjoint Faculties Research Ethics Board and is being conducted to partially fulfill Ph.D. research requirements for James Young.

Participation in this study will not put you at any risk or harm and is strictly voluntary. You choose to participate by using our system. You may choose to withdraw from the study at any time by simply not using the system any more and you will still receive your \$15 payment. Any data collected to your withdrawal will still be available to the investigators for analysis. In this study, the only personal information we will collect are your sex and age which will be used only for identification purposes and grouping results. This information will only be used in papers and presentations, and will be presented as anonymous data using generic identifiers. The experiment will not be videotaped.

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a 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 the study at any time. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

At the conclusion of the study and its analysis, we will post any resulting papers that we have written about them. You can view these by asking the investigator or by accessing the website: <http://groupiab.cpsc.ucalgary.ca/papers/index.html>

Electronic data will be stored in a secure manner, such as in a computer secured with a password. Hardcopies of data will be stored in a locked cabinet/room with restricted access.



UNIVERSITY OF
CALGARY

Data will be kept for a minimum of three years and a maximum of 7 years. On disposal, electronic data will be erased and hardcopies will be shredded.

If you have further questions concerning matters related to this research, please contact:
Dr. Ehud Sharlin (403) 210-9499 ehud@cpsc.ucalgary.ca

If you have any concerns about the way you've been treated as a participant, please contact Bonnie Scherrer in the Research Services Office, University of Calgary at (403) 220-3782; email bonnie.scherrer@ucalgary.ca.

	Please one	circle	Please initial
I agree to participate in the activities explained above.	YES NO		
I agree to let my conversation during the study be directly quoted, anonymously, in presentation of the research results.	YES NO		

Participant's Signature _____ Date _____

Investigator and/or Delegate's Signature _____ Date _____

A copy of this consent form has been given to you to keep for your records and reference.

Experiment Procedures

Introduction

Hello. My name is Jim. Thank you very much for your participation today. With your help, I'll be conducting a user evaluation involving animated characters. For example, you may be familiar with such characters from your favorite cartoon show or from a video game. These are the kinds of characters we will be dealing with today. Before starting the evaluation, I'll briefly describe the concept of our project, talk about the prototype application that we have developed, and lay out how the experiment will be run. Throughout the study, you can quit at any time for any reason, and can still keep the money.

Conceptual Description

Our system deals with the behaviors of animated characters. Behaviors are the personalities, emotions, and ways that characters act, in particular in reaction to other characters or input data.

Our project is intended to explore how these character behaviors can be designed through demonstration. That is, rather than using computer languages, we will get you to design a behavior by demonstrating it to our system. The question here is to see how and if our program can capture your demonstrated behavior. With this kind of approach, then, a person does not need to be a computer expert to create, or program, computerized character behaviors.

Application Description

The application we have developed here involves two characters, where one character reacts to the actions of another character. One character is human controlled (the main character), and while it moves around, another character (the reactor character) interacts with and responds to the main character. This reactor character can exhibit a particular emotion or personality towards the main character. Your involvement here will be to, using our system, give demonstrations of particular personalities that we give to you. Do you have any questions before we move forward?

Now that you have an understanding of the study, please review this consent form and sign if you are still interested.

Before moving forward, I would like to conduct a pre-study interview if it is okay with you.

<PRESTUDY INTERVIEW >

Training and Demonstration

Before you start, I will do one myself to show you how the system works, and to show you how to use our system.

<System description and explanation>

<TRAINING AND DEMO, <3 min >

- Point out the vicon, pucks, main puck, reactor puck.
- Clear vicon action
- Explain the steps of the system
 - > Training step
 - > Testing step
 - > Clear
- Do a non computerized example
- Do computer example (leap frog)
- Let the user play

<End Training>

Now I would like you to demonstrate behaviors. After you are done a particular behavior, let me know so I can save it, and I will give you a quick question sheet to answer regarding the behavior you created. You can also quit if you get frustrated or you feel the system does not allow you to make the desired behavior.

Core Study

<CREATE BEHAVIORS AND DO SHEETS, ATTACHED>

Thank you for creating the behaviors. For the next stage, we will present you with your behaviors one at a time in a random order. For each behavior, after testing it out, we would like you to please try to match each behavior with the ones you trained.

<REPLAY BEHAVIORS IN RANDOM ORDER>

Thank you for doing our study, now we will do the post-study interview.

<POST STUDY INTERVIEW>

Pre-Study interview

1. Sex: ___ Age ___

2. Please rate your technical computer ability.

Absolutely none	I can install new software	I have a technical Degree				
1.....	2.....	3.....	4.....	5.....	6.....	7

Any comments about this?

3. Do you have any computer programming experience? If Yes, please describe.

4. How comfortable are you when using a new piece of software?

5. Do you have any experience with creating animation? If yes, please elaborate.

6. Please describe any formal artistic training you have:

7. Do you have any general artistic experience? For example, do you paint or draw as a hobby?

8. Would you consider yourself an artist?

Post-study interview

How much do you agree with the following statements?

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
I enjoyed using the system Comments:	1	2	3	4	5	6	7
The system was often Frustrating to use Comments:	1	2	3	4	5	6	7
I was often disappointed by the resulting characters Comments:	1	2	3	4	5	6	7
The resulting characters were fun to play with Comments:	1	2	3	4	5	6	7

1. How did you feel about the character generated by the system? e.g., did they feel natural, mechanical, real, fake, etc?

2. What would you say are the positive, good parts of this system?

3. How about the shortfalls, negative parts?

4. Do you have any ideas for improvement?

5. What is your general impression of the system?

Scenarios:

- _____ Lovers
- _____ Bully
- _____ Playful Friend
- _____ Stalker
- _____ Afraid

. **Lover**
Behavior Description Sheet

1. You were satisfied with how well the system captured the behavior you were trying to demonstrate.

Strongly disagree	Disagree	Neither agree nor disagree	Agree	Strongly Agree
1	2	3	4	5

(additional comment space, if needed)

2. The generated behavior felt mechanical

Strongly disagree	Disagree	Neither agree nor disagree	Agree	Strongly Agree
1	2	3	4	5

(additional comment space, if needed)

3. The generated behavior felt human-controlled

Strongly disagree	Disagree	Neither agree nor disagree	Agree	Strongly Agree
1	2	3	4	5

(additional comment space, if needed)

4. What good things did you notice about the system in respect to this behavior?

5. What about the system shortfalls or problems in respect to this behavior

6. Any other comments / points?

C.2 OBSERVER STUDY

These documents pertain only the observer study. The designer study is explained in Section C.1. The outline of this section is provided below:

C.2.1 participant consent form

C.2.2 experiment protocol

C.2.3 pre-test questionnaire

C.2.4 post-test questionnaire

C.2.5 behaviour-matching sheet

C.2.6 behaviour description sheet*

* – we have only included one behaviour description sheet, although five identical (save the number at the top) sheets were provided to the participant. These were unlabelled as this was an exploratory exercise.



UNIVERSITY OF
CALGARY

Dr. Ehud Sharlin and James E. Young
Department of Computer Science
University of Calgary
2500 University Drive
Calgary, AB, CANADA T2N 1N4

Consent Form for Participants

Research Project: Exploratory study of Reactions to Social Computerised Entities

Investigators: Dr. Ehud Sharlin and James E. Young

This consent form, a copy of which has been given to you, is only part of the process of informed consent. It should give you the basic idea of what the research is about and what your participation will involve. If you would like more detail 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.

Description of Research Project:

The purpose of this study is to evaluate a new system for animated computer characters. We will ask you interact with animated characters using our system. To start, we will explain how our system works and show you how to interact with the animations. The entire experiment will take around one hour. The experiment will not be videotaped.

Participation in this study will not put you at any risk or harm and is strictly voluntary. You choose to participate by using our system. You may choose to withdraw from the study at any time by simply not using the system any more. Any data collected to your withdrawal will still be available to the investigators for analysis. Personally identifiable information will only be used in papers or presentations with your explicit permission. If we wish to use any personally identifiable information, we will contact you with the particulars of the information we wish to use, and you may decide whether or not you give us permission to use it. In this study, the only personal information we will collect are your sex and age which will be used only for identification purposes and grouping results.

Your signature on this form indicates that you have understood to your satisfaction the information regarding participation in the research project and agree to participate as a 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 the study at any time. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

At the conclusion of the study and its analysis, we will post any resulting papers that we have written about them. You can view these by asking the investigator or by accessing the website: <http://groupslab.cpsc.ucalgary.ca/papers/index.html>

Electronic data will be stored in a secure manner, such as in a computer secured with a password. Hardcopies of data will be stored in a locked cabinet/room with restricted access.



UNIVERSITY OF
CALGARY

Data will be kept for a minimum of three years and a maximum of 7 years. On disposal, electronic data will be erased and hardcopies will be shredded.

If you have further questions concerning matters related to this research, please contact:
Dr. Ehud Sharlin (403) 210-9499 ehud@cpsc.ucalgary.ca

If you have any concerns about the way you've been treated as a participant, please contact Bonnie Scherrer in the Research Services Office, University of Calgary at (403) 220-3782; email bonnie.scherrer@ucalgary.ca.

	Please one	circle	Please initial
I agree to participate in the activities explained above.	YES NO		
I agree to let my conversation during the study be directly quoted, anonymously, in presentation of the research results.	YES NO		

Participant's Signature _____ Date _____

Investigator and/or Delegate's Signature _____ Date _____

A copy of this consent form has been given to you to keep for your records and reference.

Experiment Procedures

Introduction

Hello. My name is Jim. Thank you very much for your participation today. With your help, I'll be conducting a user evaluation involving animated characters. For example, you may be familiar with such characters from your favorite cartoon show or from a video game. These are the kinds of characters we will be dealing with today. Before starting the evaluation, I'll briefly describe the concept of our project, talk about the prototype application that we have developed, and lay out how the experiment will be run. Throughout the study, you can quit at any time for any reason, and can still keep the money.

Application Description

The application we have developed here involves two characters, where one character reacts to the actions of another character. One character is human controlled (the main character), and while it moves around, another character (the reactor character) interacts with and responds to the main character. Your involvement here will be to simply control the main character and observe the reacting character. After you observe a character, we will ask you a few questions regarding the character. Now that you have an understanding of the study, please review this consent form and sign if you are still interested.

Before moving forward, please complete this pre-study questionnaire.

<PRESTUDY INTERVIEW >

First, let me explain the technical system.

<VICON and puck introduction>

<Explain shape / dim>

Core Study

Now, I will load a character into the system, and I want you to interact with it using this puck. Please spend a little bit of time to explore the character, and try to talk out loud about your observations (I will take notes). When you are done exploring a character, let me know and I will give you a short questionnaire where you will be asked to describe the character. Note that these graphics will be the same throughout the characters I show you.

<Set behavior>

<Could you please describe this character - verbally?>

<Written questionnaire>

<take puck off>

Core Study part 2

Now, I will bring up five other characters that were explicitly designed for particular personalities. For each character, please take your time to interact with it and match it as best you can to one of the five on the sheet here. As you proceed you can change your answer but you cannot go back to a previous character.

Thank you for doing the study with us.

<Post study questionnaire>

Pre-Study interview

1. Sex: ___ Age ___

2. Please rate your technical computer ability.

Absolutely none	I can install new software	I have a technical Degree				
1.....	2.....	3.....	4.....	5.....	6.....	7

Any comments about this?

3. Do you have any computer programming experience? If Yes, please describe.

4. How comfortable are you when using a new piece of software?

5. Do you have any experience with creating animation? If yes, please elaborate.

6. Please describe any formal artistic training you have:

7. Do you have any general artistic experience? For example, do you paint or draw as a hobby? Play an instrument?

8. Would you consider yourself an artist?

Post-study interview

How much do you agree with the following statements?

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
I enjoyed the experiment Comments:	1	2	3	4	5	6	7
The puck and table system was often Frustrating to use Comments:	1	2	3	4	5	6	7
The characters' actions were often confusing Comments:	1	2	3	4	5	6	7
The characters were fun to interact with Comments:	1	2	3	4	5	6	7

1. How did you feel about the characters presented to you? e.g., did they feel natural, mechanical, real, fake, etc?

2. Overall, what, if anything, did you like about the characters?

3. Overall, what, if anything, did you dislike about the characters?

4. Do you have any comments on the puck / table system?

5. Any other points or comments? Was there anything in particular about the experiment you would like to mention? Any suggestions for improvement.

Scenarios:

- _____ Lovers
- _____ Bully
- _____ Playful Friend
- _____ Stalker
- _____ Afraid

. 1

Character Description Sheet

1. Please give two or three descriptive words (key words) to describe this character

2. The reacting character felt ...

Extremely mechanical			Somewhat mechanical			Not mechanical at all
1	2	3	4	5	6	

(additional comment space, if needed)

3. The reacting character felt like...

A person was controlling it			Somewhat life-like			Not life-like at all
1	2	3	4	5	6	

(additional comment space, if needed)

4. What, if anything, did you like about the character?

5. What, if anything, did you not like about the character?

6. Any other comments / points?

STYLISTIC LOCOMOTION AND PUPPET MASTER STUDIES, ROBOTS

This appendix contains the study materials used for our studies on the robotic *stylistic locomotion* and *puppet master* interfaces (Section 7.6, page 217). We first provide the programmer-study Application Programming Interface (API) description, diagrams of the walk paths, the study materials for both the broomstick and tabletop conditions of the *designer* study, and finally the materials for the observer study.

D.1 PROGRAMMER STUDY API

The API we provided to the programmer revolved around a function, `behaviorStep`, which the main system would call at regular intervals (15 Hz). At each call the system would provide the programmer with the current location of the both the person and the robot, and the programmer would return the next robot command to the system.

We provided the programmers with a base class (Listing D.1) that gave them information about the system, robot, and the space, and asked them to override the `behaviorStep` function in their own class. To make this easier we provided them with a skeleton (Listing D.2). Programmers were expected to work directly with the `EntityState` (Listing D.3) and `RobotCommand` (Listing D.4) classes, as well as the `RoombaSong` (Listing D.5) enumeration which defining robot output sound sequences.

Further, we provided the programmer with a basic utilities class (Listing D.6), and also an example where a robot will turn to always look at the person, and make periodic sounds (Listing D.7). These listings are all given below.

Listing D.1: base class given to programmers, they were asked to override the behaviorStep function

```

public abstract class RobotBehavior {

    public final static int CYCLES_PER_SECOND = Globals.TARGET_FPS;
    public final static long NANoseconds_PER_CYCLE = Globals.SLICE_NS;
    public final static int SPACE_WIDTH_MM = Globals.SPACE_WIDTH_IN_MM;
    public final static int SPACE_HEIGHT_MM = Globals.SPACE_HEIGHT_IN_MM;
    public final static int ROBOT_MAX_SPEED_MM_PER_S = EntityController.MAX_VEL;

    //...

    // called each system cycle
    abstract public RobotCommand behaviorStep(EntityState personState, EntityState robotState);
}

```

Listing D.2: skeleton class which overrides behaviorStep, as given to the programmers

```

public class SkeletonBehavior extends RobotBehavior {

    @Override
    public RobotCommand behaviorStep(EntityState personState,
        EntityState robotState) {

```

```
        // put your code here  
    }  
}
```

Listing D.3: the EntityState class which provided data to the programmer

```
public class EntityState {  
  
    // if two entity states are value-wise equal  
    @Override  
    public boolean equals(Object obj);  
  
    public double x;  
    public double y;  
    public double angle;  
    public RoombaSong action;  
}
```

Listing D.4: the RobotCommand which the programmer used to give a command

```
public class RobotCommand {
```

```

public double turnSpeed;
public double moveSpeed;
public RoombaSong action;

public RobotCommand(double TurnSpeed, double MoveSpeed, RoombaSong Action);
}

```

Listing D.5: the RoombaSong enumeration for selecting robot sound output

```

public enum RoombaSong {
    None,
    Sad,
    Happy;
}

```

Listing D.6: utility functions provided with the programmer API

```

public class Utilities {

    private static double TWOPI = Math.PI*2;

    // normalize an angle to the range 0 <= a <= 2*PI
    public static double normalizeAngle(double angle);
}

```



```

// polar to cartesian: X
public static double getXFromPolar(double angle, double radius);

// polar to cartesian: Y
public static double getYFromPolar(double angle, double radius);

// return acute bisector of two angles
public static double differenceAngles(double a, double b)
}

```

Listing D.7: example behaviour provided to programmers

```

public class LookToPerson extends RobotBehavior {

    @Override
    public RobotCommand behaviorStep(EntityState personState, EntityState robotState) {

        double turnSpeed = calculateTurn(personState, robotState);
        RoombaSong action = calculateSong(turnSpeed);

        return new RobotCommand(turnSpeed, 0, action);
    }
}

```

```
}

private static double CLOSE_ENOUGH_ANGLE = 0.2;
protected double calculateTurn(EntityState personState, EntityState robotState) {

    double angleToPerson = Math.atan2(personState.y - robotState.y, personState.x - robotState.x);
    double mustTurnToPerson = Utilities.differenceAngles(angleToPerson, robotState.angle);

    if (Math.abs(mustTurnToPerson) < CLOSE_ENOUGH_ANGLE)
        return 0;

    if (Math.abs(mustTurnToPerson) > Math.PI/2) { // if more than quarter circle
        if (mustTurnToPerson > 0) return 1;
        else return -1;
    }
    return (mustTurnToPerson / (Math.PI/2) ); // closer to right look dir, less speed
}

private int framesSinceLastSound = 0;
private RoombaSong lastSong = RoombaSong.None;
private RoombaSong calculateSong(double turnSpeed) {
    RoombaSong song;
```

```
framesSinceLastSound++;  
if (framesSinceLastSound > RobotBehavior.CYCLES_PER_SECOND) { // once per second  
    framesSinceLastSound = 0;  
    if (turnSpeed == 0)  
        song = RoombaSong.Happy;  
    else  
        song = RoombaSong.None;  
    lastSong = song;  
}  
return lastSong;  
}  
}
```


D.2 WALKING PATHS

Here we show the direction cards used by the experimenter to outline the path they should walk. The red Xs denote landmarks as were drawn on the actual floor of the working area. The black lines and arrows denote the walking direction and path, and the walker always started in the bottom-right corner.

Figure D.1a is the training path used for robots that follow the person, the *polite* and *stalker* behaviours, with the generation path shown in Figure D.1b. For robots that share a space with people (not necessarily follow, *happy* and *attack*), Figure D.2a shows the walk path used during training with the one for generation given in Figure D.2b. Figure D.3 shows the path used when the participant was shown their created behaviours in a random order.

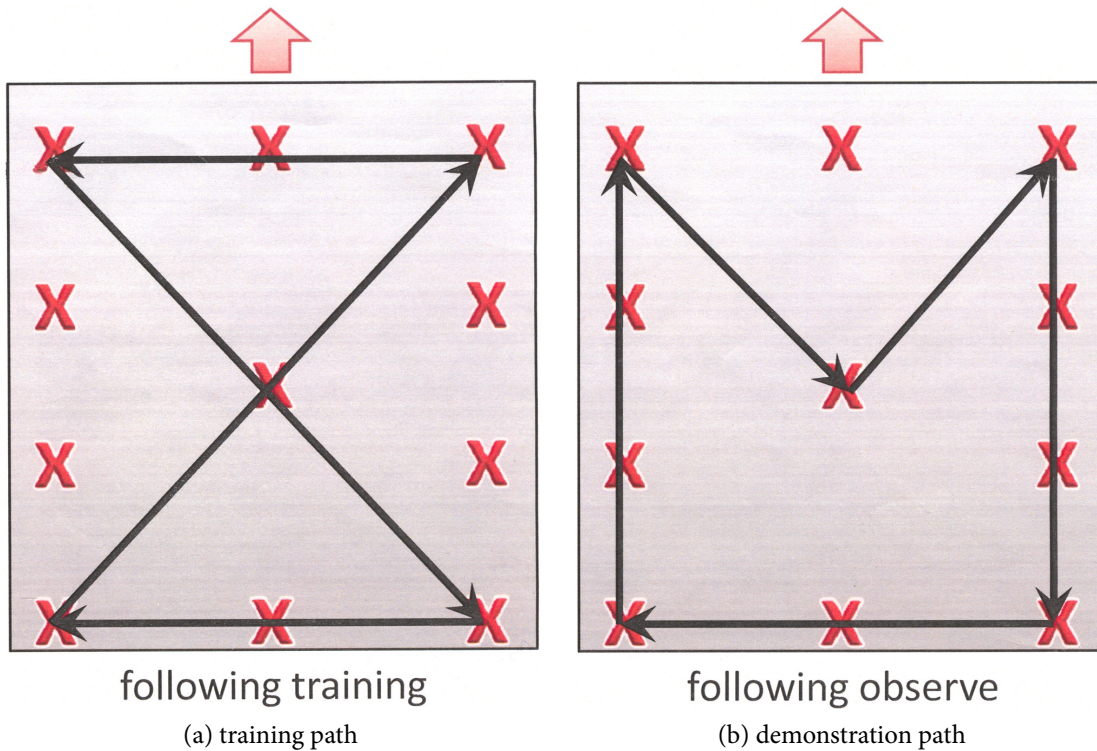


Figure D.1: walk paths used for robots that follow




Figure D.2: walk paths used for robots that share their space

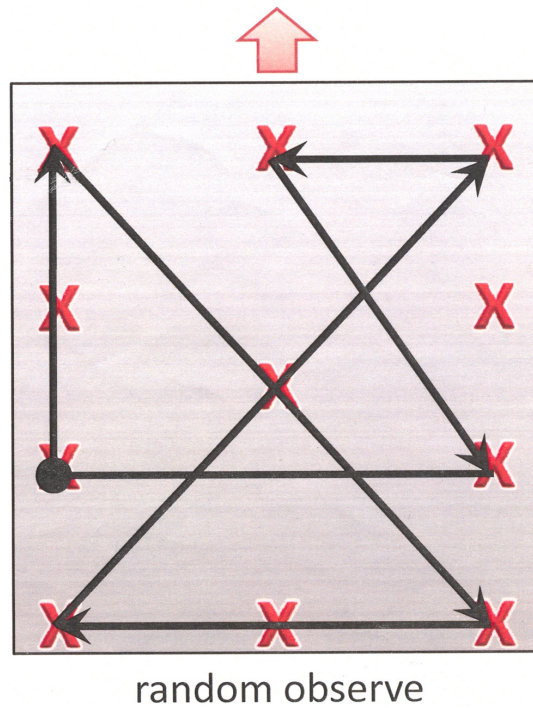


Figure D.3: walk paths used for observing randomized behaviours

D.3 DESIGNER STUDY: BROOMSTICK CONDITION

These documents pertain only the broomstick designer condition of the *puppet master* study.

The outline of this section is provided below:

D.3.1 participant consent form

D.3.2 experiment protocol

D.3.3 pre-test questionnaire

D.3.4 post-test questionnaire

D.3.5 behaviour-matching sheet

D.3.6 per-behaviour questionnaire*

* – we have only included one example questionnaire. In the study we used four identical sheets for the each of the *polite*, *stalker*, *burglar*, and *happy* scenarios.



Signed Consent Form for Participants

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

Dr. Ehud Sharlin, James Young, Daniel Van Dale (Computer Science, University of Calgary)
Dr. Takeo Igarashi (Computer Science, University of Tokyo)
(403)210-9502
ehud@cpsc.ucalgary.ca, jim.young@ucalgary.ca

Title of Project:

Designing social robots that convey emotive movement

Sponsored by:

Natural Sciences and Engineering Research Council of Canada (NSERC)

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. Thank you very much for your involvement in this project!

This research study has been approved by the University of Calgary Conjoint Faculties Research Ethics Board.

Purpose of the Study:

The purpose of this study is to evaluate a new system that enables people to demonstrate movement styles to a robot, as well as to investigate how people react to and perceive the movement styles and patterns of robots. This study will be used in the context of informing a PhD project.

What Will I Be Asked To Do?

We will introduce you to a new system for demonstrating movement styles to robots, and ask you to demonstrate particular movement styles to a particular robot. Further, you will be asked to observe the resulting behaviour of the robots and to reflect on if the robot effectively "learned" the style you demonstrated. We will provide a demonstration at the beginning of the study. This process will take approximately one hour from start to finish, including the study, verbal interviews, and the completion of several questionnaires. You will receive \$20 for your time.

Are there Risks or Benefits if I Participate?

Participation in this study will not put you at risk or harm and is strictly voluntary. You may refuse to participate, refuse to participate in parts of the study, or may withdraw from the study at any time for any reason. In any case, you will still receive your \$20 payment.

You are free to withdraw from the study at any time. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

At the conclusion of the study and its analysis, we will post any resulting papers that we have written about them. You can view these by asking the investigator or by accessing the website:
<http://grouplab.cpsc.ucalgary.ca/papers/index.html>

What Type of Personal Information Will Be Collected?

In this study, should you agree to participate, we will only record your sex and age, which will be only be used anonymously for statistical and analysis purposes. We may also, with your explicit permission, videotape all, or parts of your session, and, as such will limit your anonymity since videotaped portions of your participation may be shown in public presentations. Researchers will be unable to control any further use of images after they are presented in a public forum, which may result in these images being reposted in some unknown context, including possibly on the internet. Should you chose to withdraw from the study after a portion is completed, then any information collected before you withdraw may be used unless you expressly indicate that you wish otherwise.

Your name will be recorded for the cash payout and for this informed consent form only. This information will not be correlated with or matched to the study results in any way.

What Happens to the Information I Provide?

Collected information, including audio or video recordings, will only be used in academic papers and presentations, and will be presented as anonymous data using generic identifiers. Electronic data will be stored in a secure manner, such as in a computer secured with a password. Hardcopies of data will be stored in a locked cabinet/room with restricted access. Data will be kept for a minimum of three years and a maximum of 7 years. On disposal, electronic data will be erased and hardcopies will be shredded.

Signatures (written consent)

I grant permission to be videotaped: Yes: ___ No: ___

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: _____

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

A copy of this consent form has been given to you to keep for your records and reference.

Experiment Procedures

Hello, my name is Jim. Thank you very much for helping with our study. To summarize, this study involves robots and how robots behave around people. Behaviors are the personalities, emotions, and ways that the robots act, in this case, in reaction to a person.

Our project is intended to explore how these robot behaviors can be designed through demonstration. That is, rather than using computer languages, we will get you to design a robotic behavior by demonstrating it directly to the robot. With this kind of approach, then, you do not need to be a computer expert to create, or program, robot behaviors.

The application we have developed here involves a person and a robot, where the robot is supposed to react to the person. The person moves naturally, while the robot reacts appropriately to the person's movements. Your involvement here will be to, using our system, give demonstrations of particular behaviors that we give to you. Do you have any questions before we move forward?

Before we begin, I have a document here that I need you to read over. The purpose of this document is to give a brief overview of what will you will be asked to do during the study, and to give you a chance to decide if this is something you are comfortable doing. We will require a signature from you stating that you understand what we expect you to do, as well as to confirm what sorts of data collection you are comfortable with. These signatures do not mean that you are obliged to participate, however, and you are still completely free to stop at any time if you so wish. If you have any questions at any time then please feel free to ask anything at all. Thank you very much.

<GIVE FORM>

Before getting into the study, I would like you to fill out a pre-study questionnaire if that is okay with you.

<PRE-STUDY QUESTIONNAIRE>

To start, I will do a brief training and demonstration session to show how the system works, and let you give it a rough try.

-- SYSTEM DESCR and EXPL

Introduce Daniel, the person the robot will follow.

Point out the robots, the real one and broomstick.

Broomstick button sounds

Show the space boundaries, robot cannot (tries very hard not to) leave – tries to run away, gets stuck.

Point out vicon markers

Verbally explain the steps of the system (training with broomstick, watch real robot try to learn – starts right away)

The person moves in a path and loops the path as long as necessary

During training, you need to be clear on when your demonstration starts and ends. If you sit and wait and say, okay, im done – that waiting will be part of the training. I will demonstrate this. Training can be short, 10 seconds, or long, minutes, up to you.

When the robot starts up, the person will loop through a different – but similar – movement.

Do a fake example simply of following,

Re-iterate start and end, with example

Let the participant play with the roomba in the space

Please be patient with the roomba – he cannot move so

fast, as fast as you move the stick

--END TRAINING

Now I would like you to demonstrate to the robot. I will give you descriptions of the behavior to demonstrate. You will demonstrate to the robot, have an opportunity to observe the robot's behavior – if you are not happy, you can re-demonstrate to the robot or just move on – up to you. When you are finally done, I will ask you to fill out a quick questionnaire regarding the behavior you created before moving on. Do you have any questions?

CORE STUDY

-- Camera On, flash participant ID

For each behavior, I will describe the scenario and show you a card of how the person will move. The person will repeat the movements until you are finished the demonstration.

Scenarios:

The first batch are following, where the robot follows the person

(1) Polite, careful following as a robot may follow a Doctor in a hospital

Save, Questionnaire

(2) Stalker who doesn't want to be noticed

Save, Questionnaire

The next are a person entering the robots space.

This is how the person will move:

(3) A burglar enters to steal something and the robot acts

aggressively toward it.

Save, Questionnaire

(4) An owner comes home and the robot is happy to see them

Save, Questionnaire

Thank you for helping in creating the behaviors. For the next stage, we will present you with your behaviors one at a time in a random order, with the person moving in an entirely new movement path. For each behavior, after observing, we would like you to please try to match it with the trained behavior.

(give matching sheet)

More observ instructions

<REPLAY BEH IN RANDOM ORDER>(5,6,3,4,2,1)

Thank you for doing our study, we will now to the post-study questionnaire.

<POST-STUDY QUESTIONNAIRE>

1. Sex: ___ Age___

2. Please rate your technical computer ability.

Absolutely none	I can install new software	I have a technical Degree
1.....	2.....	3.....
4.....	5.....	6.....
7.....		

Any comments about this?

3. Do you have any computer programming experience? If Yes, please describe.

4. How comfortable are you when using a new piece of software?

5. Do you have any experience with robots? If Yes, please describe.

6. Please describe any formal artistic training you have.

7. Do you have any acting or theatre experience, even amateur? If yes, please elaborate.

8. Do you have any (brief) reflections on the near future of robots? That is, where you think technology will develop and how it may affect your life.

9. Do you have any opinions or thoughts on the idea of “teaching” to a robot? If so, could you briefly reflect on them?

Post-study questionnaire

How much do you agree with the following statements?

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
I enjoyed the overall demonstration and observation of behaviors Comments:	1	2	3	4	5	6	7
I was often disappointed by the resulting robot behaviors Comments:	1	2	3	4	5	6	7
The process of demonstrating to a robot was often frustrating Comments:	1	2	3	4	5	6	7
The resulting robot behaviors conveyed the style and personality I intended Comments:	1	2	3	4	5	6	7

1. Overall, can you reflect on the idea of teaching a robot by demonstrating to it?

2. Overall, can you comment on how successful you thought the robot was at learning from your demonstration?

3. Can you reflect on the use of a broomstick to demonstrate to a robot?

4. What did you think of the robot sounds, e.g., were they an important part of the robot behaviors?

5. How did you feel about the quality of the characters generated by the system? e.g., did they feel natural, mechanical, real, fake, etc?

6. How important of a role do you think that robots will play in our lives in the future?

7. If you ever consider buying a household robot, do you think you would care about the personality and style of the behaviors, such as we worked with in today's study?

8. Do you have any additional positive reflections on the robots in today's study?

9. Do you have any additional negative reflections (problems, room for improvement, etc) on the robots in today's study?

Designer-Broom—PostTestQuestionnaire

10. Any additional ideas for improvement?

11. Finally, any last comments, ideas, suggestions, etc

Scenarios:

- _____ polite follow
- _____ stalker
- _____ burgler
- _____ happy to see you

Designer-table—perBehaviorQuestionnaire

Behavior Description Sheet

1. You were satisfied with how well the system captured the behavior you were trying to demonstrate.

Strongly disagree	Disagree	Neither agree nor disagree	Agree	Strongly Agree
1	2	3	4	5

(additional comment space, if needed)

2. The resulting robot behavior felt overly mechanical

Strongly disagree	Disagree	Neither agree nor disagree	Agree	Strongly Agree
1	2	3	4	5

(additional comment space, if needed)

3. The resulting robot behavior felt natural, organic, possibly human-controlled

Strongly disagree	Disagree	Neither agree nor disagree	Agree	Strongly Agree
1	2	3	4	5

(additional comment space, if needed)

Designer-table—perBehaviorQuestionnaire

4. I think it makes sense to teach robots this behavior by demonstration

Strongly disagree	Disagree	Neither agree nor disagree	Agree	Strongly Agree
1	2	3	4	5

(additional comment space, if needed)

5. What good things did you notice about the overall approach in respect to this behavior?

6. What about the shortfalls or problems in respect to this behavior

7. Any other comments / points?

D.4 DESIGNER STUDY: SURFACE PUPPET MASTER CONDITION

These documents pertain only the *Surface puppet master* designer study. The outline of this section is provided below:

D.4.1 experiment protocol

D.4.2 post-test questionnaire

Various forms are identical to those used in the broomstick condition (Section D.3, page 369), and so are not listed here: the participant consent form, pre-test questionnaire, behaviour-matching sheet and per-behaviour questionnaires.

Experiment Procedures

Hello, my name is Jim. Thank you very much for helping with our study. To summarize, this study involves robots and how robots behave around people. Behaviors are the personalities, emotions, and ways that the robots act, in this case, in reaction to a person.

Our project is intended to explore how these robot behaviors can be designed through demonstration. That is, rather than using computer languages, we will get you to design a robotic behavior by demonstrating it directly to the robot. With this kind of approach, then, you do not need to be a computer expert to create, or program, robot behaviors.

The application we have developed here involves a person and a robot, where the robot is supposed to react to the person. The person moves naturally, while the robot reacts appropriately to the person's movements. Your involvement here will be to, using our system, give demonstrations of particular behaviors that we give to you. Do you have any questions before we move forward?

Before we begin, I have a document here that I need you to read over. The purpose of this document is to give a brief overview of what will you will be asked to do during the study, and to give you a chance to decide if this is something you are comfortable doing. We will require a signature from you stating that you understand what we expect you to do, as well as to confirm what sorts of data collection you are comfortable with. These signatures do not mean that you are obliged to participate, however, and you are still completely free to stop at any time if you so wish. If you have any questions at any time then please feel free to ask anything at all. Thank you very much.

<GIVE FORM>

Before getting into the study, I would like you to fill out a pre-study questionnaire if that is okay with you.

<Give money and do collection sheet>

<PRE-STUDY QUESTIONNAIRE>

To start, I will do a brief training and demonstration session to show how the system works, and let you give it a rough try.

-- SYSTEM DESCR and EXPL

Explain the table and simulation vs real robots and real people

Introduce Daniel, the person the robot will follow.

Point out the real robot

Show the space boundaries, robot cannot (tries very hard not to) leave – tries to run away, gets stuck.

Point out vicon markers

Move to surface,

Show how you demonstrate to the surface / robot button sounds

Verbally explain the steps of the system (training with puck, watch real robot try to learn)

The person moves in a path and loops the path as long as necessary

During training, you need to be clear on when your demonstration starts and ends. If you sit and wait and say, okay, im done – that waiting will be part of the training. I will demonstrate this. Training can be short, 10 seconds, or long, minutes, up to you.

When the robot starts up, the person will loop through a different – but similar – movement.

Do a fake example simply of following,

Re-iterate start and end, with example

Let the participant play with the puck on the table

Please be patient with the roomba – he cannot move so fast, as fast as you move the puck

---END TRAINING

Now I would like you to demonstrate to the robot. I will give you descriptions of the behavior to demonstrate. You will demonstrate to the robot, have an opportunity to observe the robot's behavior – if you are not happy, you can re-demonstrate to the robot or just move on – up to you. When you are finally done, I will ask you to fill out a quick questionnaire regarding the behavior you created before moving on. Do you have any questions?

CORE STUDY

-- Camera On, flash participant ID

For each behavior, I will describe the scenario and show you a card of how the person will move. The person will repeat the movements until you are finished the demonstration.

Scenarios:

The first batch are following, where the robot follows the person

(1) Polite, careful following as a robot may follow a Doctor in a hospital

Save, Questionnaire

(2) Stalker who doesn't want to be noticed

Save, Questionnaire

The next are a person entering the robots space.
This is how the person will move:

(3) A burglar enters to steal something and the robot acts aggressively toward it.

Save, Questionnaire

(4) An owner comes home and the robot is happy to see them

Save, Questionnaire

Thank you for helping in creating the behaviors. For the next stage, we will present you with your behaviors one at a time in a random order, with the person moving in an entirely new movement path. For each behavior, after observing, we would like you to please try to match it with the trained behavior.

(give matching sheet)

More observ instructions

<REPLAY BEH IN RANDOM ORDER>(5,6,3,4,2,1)

Thank you for doing our study, we will now to the post-study questionnaire.

<POST-STUDY QUESTIONNAIRE>

Post-study questionnaire

How much do you agree with the following statements?

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
I enjoyed the overall demonstration and observation of behaviors Comments:	1	2	3	4	5	6	7
I was often disappointed by the resulting robot behaviors Comments:	1	2	3	4	5	6	7
The process of demonstrating to a robot was often frustrating Comments:	1	2	3	4	5	6	7
The resulting robot behaviors conveyed the style and personality I intended Comments:	1	2	3	4	5	6	7

1. Overall, can you reflect on the idea of teaching a robot by demonstrating to it?

2. Overall, can you comment on how successful you thought the robot was at learning from your demonstration?

3. Can you reflect on the use of a puck on a table to demonstrate to a robot?

Designer-Table—PostTestQuestionnaire

4. What did you think of the robot sounds, e.g., were they an important part of the robot behaviors?

5. How did you feel about the quality of the characters generated by the system? e.g., did they feel natural, mechanical, real, fake, etc?

6. How important of a role do you think that robots will play in our lives in the future?

Designer-Table—PostTestQuestionnaire

7. If you ever consider buying a household robot, do you think you would care about the personality and style of the behaviors, such as we worked with in today's study?

8. Do you have any additional positive reflections on the robots in today's study?

9. Do you have any additional negative reflections (problems, room for improvement, etc) on the robots in today's study?

Designer-Table—PostTestQuestionnaire

10. Any additional ideas for improvement?

11. Finally, any last comments, ideas, suggestions, etc

D.5 OBSERVER STUDY

These documents pertain only the observer study. The outline of this section is provided below:

D.5.1 participant consent form

D.5.2 experiment protocol

D.5.3 pre-test questionnaire

D.5.4 post-test questionnaire

D.5.5 per-behaviour questionnaire*

* – we have only included one example questionnaire. In the study we used twelve identical sheets for the each of the *polite*, *stalker*, *burglar*, and *happy* scenarios across the tabletop, broomstick, and programmer studies.



Signed Consent Form for Participants

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

Dr. Ehud Sharlin, James Young, Daniel Van Dale (Computer Science, University of Calgary)
Dr. Takeo Igarashi (Computer Science, University of Tokyo)
(403)210-9502
ehud@cpsc.ucalgary.ca, jim.young@ucalgary.ca

Title of Project:

Designing social robots that convey emotive movement

Sponsored by:

Natural Sciences and Engineering Research Council of Canada (NSERC)

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. Thank you very much for your involvement in this project!

This research study has been approved by the University of Calgary Conjoint Faculties Research Ethics Board.

Purpose of the Study:

The purpose of this study is to evaluate a new system of robotic behaviours, and how people react to and perceive the movement styles and patterns of robots. This study will be used in the context of informing a PhD project.

What Will I Be Asked To Do?

We will present you with a robot which will display a given behaviour, and ask you to reflect upon and describe particular movement styles of the robot. There will be questionnaires that you can use to fill out the descriptions. We will provide a demonstration at the beginning of the study. This process will take approximately one hour from start to finish, including the study, verbal interviews, and the completion of several questionnaires. You will receive \$20 for your time.

Are there Risks or Benefits if I Participate?

Participation in this study will not put you at risk or harm and is strictly voluntary. You may refuse to participate, refuse to participate in parts of the study, or may withdraw from the study at any time for any reason. In any case, you will still receive your \$20 payment.

You are free to withdraw from the study at any time. Your continued participation should be as informed as your initial consent, so you should feel free to ask for clarification or new information throughout your participation.

At the conclusion of the study and its analysis, we will post any resulting papers that we have written about them. You can view these by asking the investigator or by accessing the website:
<http://grouplab.cpsc.ucalgary.ca/papers/index.html>

What Type of Personal Information Will Be Collected?

In this study, should you agree to participate, we will only record your sex and age, which will be only be used anonymously for statistical and analysis purposes. We may also, with your explicit permission, videotape all, or parts of your session, and, as such will limit your anonymity since videotaped portions of your participation may be shown in public presentations. Researchers will be unable to control any further use of images after they are presented in a public forum, which may result in these images being reposted in some unknown context, including possibly on the internet. Should you chose to withdraw from the study after a portion is completed, then any information collected before you withdraw may be used unless you expressly indicate that you wish otherwise.

Your name will be recorded for the cash payout and for this informed consent form only. This information will not be correlated with or matched to the study results in any way.

What Happens to the Information I Provide?

Collected information, including audio or video recordings, will only be used in academic papers and presentations, and will be presented as anonymous data using generic identifiers. Electronic data will be stored in a secure manner, such as in a computer secured with a password. Hardcopies of data will be stored in a locked cabinet/room with restricted access. Data will be kept for a minimum of three years and a maximum of 7 years. On disposal, electronic data will be erased and hardcopies will be shredded.

Signatures (written consent)

I grant permission to be videotaped: Yes: ___ No: ___

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: _____

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

A copy of this consent form has been given to you to keep for your records and reference.

Experiment Procedures

Hello, my name is Jim. Thank you very much for helping with our study. To summarize, this study involves robots and how robots behave around people. Behaviors are the personalities, emotions, and ways that the robots act, in this case, in reaction to a person.

Our project is intended to explore how people react to robots that act in various ways. Our robots are quite simple, as they show their behavior using only their movements and sounds (for example, no facial expressions).

The application we have developed here involves a person and a robot, where the robot is supposed to react to the person. The person moves naturally, while the robot reacts to the person's movements. Your involvement here will be to watch our robots interacting with this person and give quick feedback on what you think. Do you have any questions before we move forward?

Before we begin, I have a document here that I need you to read over. The purpose of this document is to give a brief overview of what will you will be asked to do during the study, and to give you a chance to decide if this is something you are comfortable doing. We will require a signature from you stating that you understand what we expect you to do, as well as to confirm what sorts of data collection you are comfortable with. These signatures do not mean that you are obliged to participate, however, and you are still completely free to stop at any time if you so wish. If you have any questions at any time then please feel free to ask anything at all. Thank you very much.

<Give money and do collection sheet>

<GIVE FORM>

Before getting into the study, I would like you to fill out a pre-study questionnaire if that is okay with you.

<PRE-STUDY QUESTIONNAIRE>

To start, I will give a quick introduction to the system that we are using

-- SYSTEM DESCR and EXPL

Introduce Daniel, the person the robot will follow.

Point out the real robot

Show the space boundaries, robot cannot (tries very hard not to) leave – tries to run away, gets stuck.

Point out vicon markers

Explain that I can load movement patterns and behaviors into the robot, and it will act differently.

---END TRAINING

CORE STUDY

-- Camera On, flash participant ID

Now, for the first part of our study we will show you four different robot behaviors that interact with Daniel. For each one, we will play it for exactly four minutes. During this time, we would like you to talk out loud about what you think about the behavior. Then, we will load the next behavior.

-- load 4 main behaviors.

Now, we will change a little. So the following behaviors fall into four categories – A polite follow, a Stalker (spy), a robot attacking a burglar, and a robot happy to see the person. We will give you 75 seconds per behavior and you will fill out a quick form about the behavior

GIVE 12 forms, give a chance to read over

■ Load and produce 12 behaviors.

Now, if you would like, we will bring the original four behaviors back and give you a chance to interact with the robot yourself. You will need to wear these slippers. If you are uncomfortable with shared footwear, we use disinfectant between users. Also, we

offer you several options – NEW socks to put above your other socks, and, if you want, you can try pushing your shoes into, but it may be uncomfortable to walk.

Is this something you would like to try?

(2 min per each)

Thank you for doing our study, we will now to the post-study questionnaire.

<POST-STUDY QUESTIONNAIRE>

Observer--PreTestQuestionnaire

1. Sex: ___ Age___

2. Please rate your technical computer ability.

Absolutely none I can install new software I have a technical Degree

1.....2.....3.....4.....5.....6.....7

Any comments about this?

3. Do you have any computer programming experience? If Yes, please describe.

4. Do you have any experience with robots? If Yes, please describe.

5. How comfortable are you when using a new piece of software?

Observer--PreTestQuestionnaire

6. Please describe any formal artistic training you have. For example, any actual classes such as painting, music (piano, etc), dancing, pottery, etc.

7. Do you have any acting or theatre experience, even amateur? If yes, please elaborate.

8. Do you have any (brief) reflections on the near future of robots? That is, where you think technology will develop and how it may affect your life.

Observer-PostTestQuestionnaire

Post-study questionnaire

1. How did you feel about the quality of the robot behaviors you saw today? e.g., did they feel natural, mechanical, real, fake, etc?
2. If you participated, how did you feel about interacting directly with the robot?
3. What did you think of the robot sounds, e.g., were they an important part of the robot behaviors?

Observer-PostTestQuestionnaire

4. How important of a role do you think that robots will play in our lives in the future?
5. If you ever consider buying a household robot, do you think you would care about the personality and style of the behaviors, such as what you saw in today's study?
6. Do you have any additional positive reflections on the robots in today's study?

Observer-PostTestQuestionnaire

7. Do you have any additional negative reflections (problems, room for improvement, etc) on the robots in today's study?

8. Finally, any last comments, ideas, suggestions, etc?

Observer—Per-behavior-questionnaire

Per-behavior questionnaire. Behavior # ---

1. This robot behavior is: (circle one)

Polite Follow Stalker

Aggressive to Burglar Happy to see you

How much do you agree with the following statements?

	strongly disagree	disagree	somewhat disagree	no opinion	somewhat agree	agree	strongly agree
It was difficult to classify the behavior Comments:	1	2	3	4	5	6	7
I found the behavior to be engaging Comments:	1	2	3	4	5	6	7
It felt like a human was controlling the robot Comments:	1	2	3	4	5	6	7
The behavior fell into the categories I was given Comments:	1	2	3	4	5	6	7
The behavior felt mechanical Comments:	1	2	3	4	5	6	7

BIBLIOGRAPHY

- Icek Ajzen. The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2):179–211, December 1991. (Cited on page 64.)
- Icek Ajzen and Martin Fishbein. *Understanding Attitudes and Predicting Social Behavior*. Prentice Hall, Inc., Upper Saddle River, NJ, USA, 1980. (Cited on pages 62, 64, 69, 61, and 68.)
- Kenji Amaya, Armin Bruderlin, and Tom Calvert. Emotion from motion. In *Proceedings of Graphics Interface 1996. GI '96, Toronto, Canada, May 22–24, 1996*, pages 222–229, 1996. (Cited on pages 169 and 165.)
- Brenna Argall, Yang Gu, and Brett Browning. The first segway soccer experience: Towards peer-to-peer human-robot teams. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 321–322, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1121241.1121296. (Cited on page 33.)
- Isaac Asimov. *I, Robot*. Grafton Books, London, UK, 1968. (Cited on page 74.)
- Rowel Atienza and Alexander Zelinsky. Active gaze tracking for human-robot interaction. In *Proceedings of the IEEE International Conference on Multimodal Interfaces, 2002. ICMI '02, Pittsburgh, US, October 14–16, 2002*, pages 261–267, Los Alamitos, CA, USA, 2002. IEEE Computer Society, IEEE Computer Society Press. (Cited on page 30.)
- Christopher G. Atkeson, Joshua G. Hale, Frank E. Pollick, Marcia Riley, Shinya Kotosaka, Setefan Schaal, Tomohiro Shibata, Gaurav Tevatia, Ales Ude, Sethu Vijayakumar, and Mitsuo Kawato. Using humanoid robots to study human behavior. *IEEE Intelligent Systems*, 15(4):46–56, 2000. (Cited on page 38.)
- Michael Baker, Robert Casey, Brenden Keyes, and Holly A. Yanco. Improved interfaces for human-robot interaction in urban search and rescue. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, 2004. SMC '04, The Hague, The Netherlands, October 10–13, 2004*, pages 2960–2965, Los Alamitos, CA, USA, 2004. IEEE Computer Society, IEEE Computer Society Press. (Cited on page 22.)
- Christoph Bartneck and Jodi Forlizzi. A design-centred framework for social human-robot interaction. In *IEEE International Workshop on Robot and Human Interactive Communication, 2004. ROMAN '04, Kurashiki, Okayama, Japan, September 20–22, 2004*, pages 581–594, Los Alamitos, CA, USA, 2004. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ROMAN.2004.1374827. (Cited on page 35.)
- Christoph Bartneck, Michel van der Hoek, Omar Mubin, and Abdullah Al Mahmud. “Daisy, Daisy, give me your answer do!”: switching off a robot. In *Proceedings of the 2nd ACM/IEEE*

- Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 217–222, New York, NY, USA, 2007a. ACM Press. doi: 10.1145/1228716.1228746. (Cited on pages 26, 27, and 84.)
- Christoph Bartneck, Marcel Verbunt, Omar Mubin, and Abdullah Al Mahmud. To kill a mockingbird robot. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 81–87, New York, NY, USA, 2007b. ACM, ACM Press. doi: 10.1145/1228716.1228728. (Cited on pages 26, 27, 59, 84, and 58.)
- Christoph Bartneck, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. My robotic doppelgänger – a critical look at the uncanny valley. In *IEEE International Workshop on Robot and Human Interactive Communication, 2009. ROMAN '09, Toyama, Japan, September 27–Oct 2, 2009*, pages 269–276, Los Alamitos, CA, USA, 2009a. IEEE Computer Society, IEEE Computer Society Press. (Cited on pages 36 and 37.)
- Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Suzana Zoghbi. Measurement instruments for the anthropomorphism, animacy, likability, perceived intelligence, and perceived safety of robots. *International Journal on Social Robotics*, 1(1):71–81, 2009b. doi: 10.1007/s12369-008-0001-3. (Cited on pages 108, 111, 112, 113, 104, 107, and 109.)
- Joseph Bates. The role of emotion in believable agents. *Communications of the ACM*, 37(7): 122–125, July 1994. doi: 10.1145/176789.176803. (Cited on pages 45, 46, 47, 166, 169, 162, and 165.)
- Peter L. Berger and Thomas Luckmann. *The social construction of reality: a treatise in the sociology of knowledge*. Anchor Books, New York, NY, USA, 1966. (Cited on pages 41, 54, and 63.)
- Cindy L. Bethel, Christine Bringes, and Robin R. Murphy. Non-facial and non-verbal affective expression in appearance-constrained robots for use in victim management: robots to the rescue! In *Proceedings of the 4th ACM/IEEE Conference on Human-Robot Interaction, 2009. HRI '09, San Diego, California, USA, March 11–13, 2009 (Video Abstract)*, pages 191–192, New York, NY, USA, 2009. ACM, ACM Press. doi: 10.1145/1514095.1514130. (Cited on pages 32, 84, and 83.)
- Hugh Beyer and Kkaren Holtzblatt. *Contextual Design: Defining Customer-Centered Systems*. Morgan Kaufman Publishers, imprint of Elsevier, San Francisco, CA, USA, 1998. (Cited on page 47.)
- Raffaele Bianco, Massimiliano Caretti, and Stefano Nolfi. Developing a robot able to follow a human target in a domestic environment. In Amedo Cesta, editor, *Proceedings of the First RoboCare Workshop, 2003. RoboCare '03, Rome, Italy, October 30, 2003*, pages 11–14, Viale Marx, Italy, 2003. Consiglio Nazionale della Ricerche, Istituto di Scienze e Technologie della Cognizione. (Cited on pages 99 and 95.)
- Wiebe E. Bijker. Do not despair: there is life after constructivism. *Science, Technology, & Human Values*, 18(4):113–138, 1993. (Cited on pages 42 and 63.)

- Mark Billinghurst, Hirokazu Kato, and Ivan Poupyrev. The MagicBook: Moving seamlessly between Reality and Virtuality. *IEEE Computer Graphics and Applications*, 21(3):6–8, May/June 2001. doi: 10.1109/38.920621. (Cited on pages 138, 158, 134, and 154.)
- Mike Blow, Kerstin Dautenhahn, Andrew Appleby, Chrystopher L. Nehaniv, and David C. Lee. Perception of robot smiles and dimensions for human-robot interaction design. In *IEEE International Workshop on Robot and Human Interactive Communication, 2006. ROMAN '06, Hatfield, UK, September 6–8, 2006*, pages 469–474, Los Alamitos, CA, USA, 2006. IEEE Computer Society, IEEE Computer Society Press. (Cited on page 36.)
- William Bluethmann, Robert Ambrose, Myron Diftler, Scott Askew, Eric Huber, Michael Goza, Fredrik Rehnmark, Chris Lovchik, and Darby Magruder. Robonaut: A robot designed to work with humans in space. *Autonomous Robots*, 14(2–3):179–197, March 2004. doi: 10.1023/A:1022231703061. (Cited on page 33.)
- Bruce M. Blumberg and Tinsley A. Galyean. Multi-level direction of autonomous creatures for real-time virtual environments. In *Proceedings of the ACM SIGGRAPH International Conference and Exhibition on Computer Graphics and Interactive Techniques, 1995. SIGGRAPH '95, Los Angeles, CA, USA, August 6–11, 1995*, pages 47–54, New York, NY, USA, 1995. ACM, ACM Press. doi: 10.1145/218380.218405. (Cited on pages 178 and 174.)
- Kirsten Boehner, Rogério DePaula, Paul Dourish, and Phoebe Sengers. How emotion is made and measured. *International Journal of Human-Computer Studies (IJHCS)*, 65(4): 275–291, April 2007. doi: 10.1016/j.ijhcs.2006.11.016. (Cited on pages 16, 46, and 48.)
- Cynthia Breazeal and Brian Scassellati. A context-dependent attention system for a social robot. In *Proceedings of the International Joint Conference on Artificial Intelligence, 1999. IJCAI '99, Stockholm, Sweden, July 33–August*, pages 1146–1151, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers, imprint of Elsevier. (Cited on pages 178 and 174.)
- Cynthia Breazeal, Matt Berln, Andrew Brooks, Jesse Gray, and Andrea L. Thomaz. Using perspective taking to learn from ambiguous demonstrations. *Robotics and Autonomous Systems*, 54:385–393, 2006. (Cited on pages 32, 34, 139, and 135.)
- Cynthia L. Breazeal. *Designing Sociable Robots*. The MIT Press, Cambridge, MA, USA, 2002. Chapters 1, 2, 5.1, 10.1, 10.3, 13. (Cited on pages 139, 166, 178, 135, 162, and 174.)
- Cynthia L. Breazeal. Emotion and sociable humanoid robots. *International Journal of Human-Computer Studies (IJHCS)*, 59(1–2):119–155, July 2003a. doi: 10.1016/S1071-5819(03)00018-1. (Cited on pages 3, 4, 6, and 49.)
- Cynthia L. Breazeal. Toward sociable robots. *Robotics and Autonomous Systems*, 42(3–4): 167–175, March 2003b. doi: 10.1016/S0921-8890(02)00373-1. (Cited on page 35.)
- Cynthia L. Breazeal, Andrew G. Brooks, Jesse Gray, Guy Hoffman, Cory D. Kidd, Hans Lee, Jeff Lieberman, Andrea Lockerd, and David Mulanda. Tutelage and collaboration for humanoid robots. *International Journal of Humanoid Robotics (IJHR)*, 1(2):315–348, 2004. doi: 10.1142/S0219843604000150. (Cited on pages 34 and 86.)

- Cynthia L. Breazeal, Daphna Buchsbaum, Jesse Gray, David Gatenby, and Bruce Blumberg. Learning from and about others: Towards using imitation to bootstrap the social understanding of others by robots. *Artificial Life*, 11(1–2):31–62, January 2005. doi: 10.1162/1064546053278955. (Cited on page 43.)
- Harry Brenton, Marco Gilles, Daniel Ballin, and David Chatting. The uncanny valley: does it exist? In *19th British HCI Group Annual Conference: workshop on human-animated character interaction, Edinburgh, Sep 2005*, 2005. (Cited on page 37.)
- Andrew G. Brooks and Cynthia L. Breazeal. Working with robots and objects: revisiting deictic reference for achieving spatial common ground. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 297–304, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1121241.1121292. (Cited on pages 30 and 33.)
- Wolfram Burgard, Armin B. Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner, and Sebastian Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1–2):3–55, October 1999. doi: 10.1016/S0004-3702(99)00070-3. (Cited on page 28.)
- Robert E. Burnkrant and Alain Cousineau. Informational and normative social influence in buyer behavior. *Journal of Consumer Research*, pages 206–215, 1975. (Cited on page 65.)
- Daniel Byers and Odest C. Jenkins. HRI caught on film 2: Hands-free human-robot interaction. In Christoph Bartneck, editor, *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349873. (Cited on pages 169 and 165.)
- Michel Callon. The state and technical innovation: a case-study of the electric vehicle in france. *Research Policy*, 9:358–76, 1980. (Cited on page 42.)
- Michel Callon. Society in the making: the study of technology as a tool for sociological analysis. In Wiebe Bijker, Thomas P. Hughes, and Trevor Pinch, editors, *The Social Construction of Technical Systems*, pages 83–103. The MIT Press, 1987. (Cited on pages 42, 63, and 62.)
- Karel Capek. *Rossum's Universal Robots*. Pocket Books, New York, US, 1970. Originally appeared as a play, 1920. (Cited on page 52.)
- Daniel Chandler. Technological or media determinism. On-Line, <http://www.aber.ac.uk/media/Documents/tecdet/tecdet.html>, September 1995. Accessed March 9, 2010. (Cited on page 42.)
- Daniel Chandler. *Semiotics: The Basics*. Routledge, an imprint of the Taylor & Francis Group, London, UK, 2002. (Cited on pages 146 and 142.)
- Michael Chueh, Sanjay Joshi, Yi Lin Au Yeung, and Kim-Pang Lei. Towards a mobile-robot following controller using behavioral cues. In *Proceedings of the IEEE International Conference on Industrial Technology, 2006. ICIT '06, Mumbai, India, December 15–17, 2006*,

- pages 150–157, Los Alamitos, CA, USA, 2006. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ICIT.2006.372317. (Cited on pages 170 and 166.)
- Jon Clark, Ian McLoughlin, Howard Rose, and Robin King. *The Process of Technological Change: New Technology and Social Choice in the Workplace*. Cambridge University Press, Cambridge, UK, 1988. (Cited on page 63.)
- Wesley M. Cohen and Daniel A. Levinthal. Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1):128–152, March 1990. (Cited on pages 63 and 64.)
- Deborah Compeau and Christopher A. Higgins. Applications of social cognitive theory to training for computer skills. *Information Systems Research*, 6(2):118–143, 1995. (Cited on pages 66 and 65.)
- Auguste Comte. *A Dictionary of Sociology*. Oxford University Press, Oxford, UK, 3rd edition, 2005. (Eds. John Scott and Gordon Marshall). (Cited on page 61.)
- Andy Crabtree, Steve Benford, Chris Greenhalgh, Paul Tennent, Matthew Chalmers, and Barry Brown. Supporting ethnographic studies of ubiquitous computing in the wild. In *Proceedings of the 6th conference on Designing Interactive Systems, 2006. DIS '06, University Park, PA, US, June 26–28, 2006*, pages 60–69, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1142405.1142417. (Cited on page 47.)
- Mihalyi Csikszentmihalyi. *Flow: The Psychology of Optimal Experience*. HarperCollinsPublishers, New York, USA, 1990. (Cited on page 59.)
- Allen Cypher. Eager: programming repetitive tasks by example. In *ACM Conference on Human Factors in Computing Systems, 1991. CHI '91, New Orleans, Louisiana, USA, April 27–May 02, 1991*, pages 33–39, New York, NY, USA, 1991. ACM, ACM Press. doi: 10.1145/108844.108850. (Cited on pages 179 and 175.)
- Nils Dahlbäck, Arne Jönsson, and Lars Ahrenberg. Wizard of oz studies – why and how. *Knowledge-Based Systems*, 6(4):258–266, December 1993. doi: 10.1016/0950-7051(93)90017-N. (Cited on page 39.)
- Angela Davids. Urban search and rescue robots: From tragedy to technology. *IEEE Intelligent Systems*, 17(2):81–83, March/April 2002. doi: 10.1109/MIS.2002.10006. (Cited on page 20.)
- Fred D. Davis. *A Technology Acceptance Model for Empirically Testing New End-User Information Systems: Theory and Results*. PhD thesis, Sloan School of Management, MIT, 1986. (Cited on pages 62, 65, and 61.)
- Fred D. Davis, Richard P. Bagozzi, and Paul R. Warshaw. User acceptance of computer technology: A comparison of two theoretical models. *Management Science*, 35(8):982–1003, August 1989. (Cited on page 65.)

- Pieter M. A. Desmet. Measuring emotions: Development and application of an instrument to measure emotional responses to products. In Mark A. Blythe, Kees Overbeeke, Andrew F. Monk, and Peter C. Wright, editors, *Funology: from usability to enjoyment*, pages 111–123. Kluwer Academic Publishers, Norwell, MA, USA, 2005. (Cited on pages 16 and 45.)
- John Dewey. *Art as Experience*. Perigee Books, New York, USA, 1980. (Cited on pages 58 and 59.)
- Ruby Roy Dholakia. Gender and IT in the household: Evolving patterns of internet use in the united states. *Information Society*, 22(4):231–240, September 2006. (Cited on page 60.)
- Rüdiger Dillmann. Teaching and learning of robot tasks via observation of human performance. *Robotics and Autonomous Systems*, 47(2–3):109–116, June 2004. doi: 10.1016/j.robot.2004.03.005. (Cited on page 34.)
- Jonathan Dinerstein and Parris K. Egbert. Fast multi-level adaptation for interactive autonomous characters. *ACM Transactions on Graphics*, 24(2):262–288, April 2005. doi: 10.1145/1061347.1061352. (Cited on pages 178 and 174.)
- Jonathan Dinerstein, Parris K. Egbert, and Dan Ventura. Learning policies for embodied virtual agents through demonstration. In *Proceedings of 20th International Joint Conference on Artificial Intelligence, 2007. IJCAI '07, Hyderabad, India, January 6–12, 2007*, pages 1257–1262, Rochester Hills, MI, USA, 2007. IJCAI, IJCAI Press. (Cited on pages 178 and 174.)
- Alan Dix, Janet Finlay, Gregory D. Abowd, and Russell Beale. *Human-Computer Interaction*. Prentice Hall, Inc., Upper Saddle River, NJ, USA, second edition, 1998. (Cited on pages 15 and 45.)
- Mira Dontcheva, Gary Yngve, and Zoran Popović. Layered acting for character animation. *ACM Transactions on Graphics*, 22(3):409–416, July 2003. ISSN 0730-0301. doi: 10.1145/882262.882285. (Cited on pages 169, 179, 165, and 175.)
- Giovanni Dosi. Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change. *Research Policy*, 11(3): 147–162, June 1982. (Cited on page 61.)
- Paul Dourish. Seeking a foundation for context-aware computing. *Human Computer Interaction*, 16(2):229–241, December 2001a. (Cited on pages 40, 41, and 51.)
- Paul Dourish. *Where the Action Is: The Foundation of Embodied Interaction*. The MIT Press, Cambridge, MA, USA, 2001b. ISBN 0-262-04196-0. (Cited on pages 10, 23, 40, 58, 59, and 57.)
- Mauro Dragone, Brian R. Duffy, and Gregory M. P. O'Hare. Social interaction between robots, avatars & humans. In *IEEE International Workshop on Robot and Human Interactive Communication, 2005. ROMAN '05, Nashville, US, August 13–15, 2005*, pages 24–29, Los Alamitos, CA, USA, 2005a. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ROMAN.2005.1513751. (Cited on pages 139 and 135.)

- Mauro Dragone, Thomas Holtz, Brian R. Duffy, and Gregory M. P. O'Hare. Uniquitous realities through situated social agents. In *Proceedings of the 5th International Working conference on Computer animation and social agents, 2005. CASA '05, Hong Kong, China, October 17–19, 2005*. The Computer Graphics Society, 2005b. (Cited on pages 139 and 135.)
- Mauro Dragone, Thomas Holtz, and Gregory M. P. O'Hare. Mixing Robotic Realities. In *Proceedings of the 11th international conference on Intelligent User Interfaces, 2006. IUI '06, Sydney, Australia, January 29–February 1, 2006*, pages 261–263, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1111449.1111504. (Cited on page 23.)
- Jill L. Drury, Jean Scholtz, and Holly A. Yanco. Awareness in Human-Robot Interactions. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, 2003. SMC '03, Washington, D.C., USA, October 5–8, 2003*, volume 1, pages 912–918, Los Alamitos, CA, USA, 2003. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ICSMC.2003.1243931. (Cited on pages 20, 21, 22, 45, 49, 130, and 126.)
- Jill L. Drury, Laurel Riek, and Nathan Rackliffe. A decomposition of UAV-related situation awareness. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 88–94, New York, NY, USA, 2006a. ACM, ACM Press. (Cited on page 22.)
- Jill L. Drury, Holly A. Yanco, Whitney Howell, Brian Minten, and Jennifer Casper. Changing shape: improving situation awareness for a polymorphic robot. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 72–79, New York, NY, USA, 2006b. ACM, ACM Press. (Cited on page 22.)
- Ray E. Eberts. *User Interface Design*. Prentice Hall, Inc., Upper Saddle River, NJ, USA, 1994. (Cited on page 45.)
- Andy Field. *Discovering Statistics through SPSS: (and sex and drugs and rock 'n' roll)*. Sage Publications, Thousand Oaks, CA, USA, third edition edition, 2009. (Cited on pages 243 and 239.)
- James Fleck. The adoption of robots in industry. *j-PHTOL*, 15(1):4–11, 1984. (Cited on pages 52 and 61.)
- Terry Fong, Chuck Thorpe, and Charles Baur. Collaboration, dialogue, and human-robot interaction. In *Springer Tracts in Advanced Robotics. 10th International Symposium on Robotics Research. ISRR '01, Victoria, Australia, November 9–12, 2001*, volume 6, pages 255–270, Berlin, New York, Heidelberg, 2002. Springer-Verlag. (Cited on page 33.)
- Jodi Forlizzi. How robotic products become social products: an ethnographic study of cleaning in the home. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 129–136, New York, NY, USA, 2007. ACM, ACM Press. doi: 10.1145/1228716.1228734. (Cited on pages 6, 25, 27, 28, 88, and 87.)

- Jodi Forlizzi and Carl 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, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 258–256, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1121241.1121286. (Cited on pages 25, 27, 28, 47, 58, 59, 68, 69, 70, 72, 85, 87, 67, 84, and 86.)
- Phil Frei, Victor Su, Bakhtiar Mikhak, and Hiroshi Ishii. Curlybot: designing a new class of computational toys. In *ACM Conference on Human Factors in Computing Systems, 2000. CHI '00, The Hague, The Netherlands, April 1–6, 2000*, pages 129–136, New York, NY, USA, 2000. ACM, ACM Press. ISBN 1-58113-216-6. doi: <http://doi.acm.org/10.1145/332040.332416>. (Cited on pages 179 and 175.)
- Batya Friedman, Peter H. Kahn, Jr., and Jennifer Hagman. Hardware companions? – what online AIBO discussion forums reveal about the human-robotic relationship. In *ACM Conference on Human Factors in Computing Systems, 2003. CHI '03, Fort Lauderdale, USA, April 5–10, 2003*, pages 273–280, New York, NY, USA, 2003. ACM, ACM Press. doi: 10.1145/642611.642660. (Cited on pages 25, 59, 87, 88, and 86.)
- Susan R. Fussell, Sara Kiesler, Leslie D. Setlock, and Victoria Yew. How people anthropomorphize robots. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 145–152, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349842. (Cited on page 26.)
- Joel Garreau. Bots on the ground. Washington Post, WWW, http://www.washingtonpost.com/wp-dyn/content/article/2007/05/05/AR2007050501009_pf.html, Visited April 9th, 2008, May 6 2007. (Cited on pages 6, 26, 75, 87, and 86.)
- Bill Gaver, Tony Dunne, and Elena Pacenti. Design: Cultural probes. *Interactions*, 6(1):21–29, January / February 1999. doi: 10.1145/291224.291235. (Cited on page 47.)
- Tom Geller. Overcoming the uncanny valley. *IEEE Computer Graphics and Applications*, 28: 11–17, July / August 2008. doi: 10.1109/MCG.2008.79. (Cited on page 36.)
- Anthony Giddens. *Sociology*. Polity, Cambridge, UK, 2009. (Cited on page 41.)
- Björn Giesler, Tobias Salb, Peter Steinhaus, and Rüdiger Dillmann. Using augmented reality to interact with an autonomous mobile platform. In *Proceedings of the IEEE International Conference on Robotics and Automation, 2004. ICRA '04, New Orleans, USA, April 18–22, 2004*, pages 1009–1014, Los Alamitos, CA, USA, 2004. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ROBOT.2004.1307282. (Cited on pages 22, 23, 151, and 147.)
- Rachel Gockley, Jodi Forlizzi, and Reid Simmons. Interactions with a moody robot. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 186–193, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1121241.1121274. (Cited on pages 28, 29, 30, 85, 260, and 255.)

- Rachel Gockley, Jodi Forlizzi, and Reid Simmons. Natural person-following behavior for social robots. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 17–24, New York, NY, USA, 2007. ACM, ACM Press. doi: 10.1145/1228716.1228720. (Cited on pages 30, 85, 98, 170, 94, and 166.)
- Kevin Gold and Brian Scassellati. Using context and sensory data to learn first and second person pronouns. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 110–117, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1121241.1121262. (Cited on page 30.)
- Elena Gribovskaya and Aude Billard. Combining dynamical systems control and programming by demonstration for teaching discrete bimanual coordination tasks to a humanoid robot. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 33–40, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349828. (Cited on pages 33, 179, and 175.)
- W. Grimson, G. Ettinger, T. Kapur, M. Leventon, W. Wells, and R. Kikinis. Utilizing segmented MRI data in image-guided surgery. *International Journal of Pattern Recognition and Artificial Intelligence (IJPRAI)*, 11(8):1367–1397, February 1998. doi: 10.1142/S0218001497000639. (Cited on pages 138 and 134.)
- Victoria Groom, Jimmy Chen, Theresa Johnson, F. Arda Kara, and Clifford Nass. Critic, compatriot, or chump?: Responses to robot blame attribution. In *Proceedings of the 5th ACM/IEEE Conference on Human-Robot Interaction, 2010. HRI '10, Osaka, Japan, March 2–5, 2010*, pages 211–218, New York, NY, USA, 2010. ACM, ACM Press. doi: 10.1145/1734454.1734455. (Cited on page 85.)
- Cheng Guo and Ehud Sharlin. Exploring the use of tangible user interfaces for human-robot interaction: a comparative study. In *ACM Conference on Human Factors in Computing Systems, 2007. CHI '07, San Jose, California, USA, April 28 – May 3, 2007*, pages 121–130, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1357054.1357076. (Cited on pages 23 and 45.)
- Cheng Guo, James Everett Young, and Ehud Sharlin. Touch and toys: new techniques for interaction with a remote group of robots. In *ACM Conference on Human Factors in Computing Systems, 2009. CHI '09, Boston, USA, April 4–9, 2009*, pages 491–500, New York, NY, USA, 2009. ACM, ACM Press. doi: 10.1145/1518701.1518780. (Cited on pages 117 and 113.)
- Alfred Haar. Zur theorie der orthogonalen funktionensysteme (german). *Mathematische Annalen*, 69(3):331–371, 1910. (Cited on pages 154, 204, 150, and 200.)
- Dan Halbert. *Programming by Example*. PhD thesis, University of California Berkeley, 1984. (Cited on pages 179, 180, 175, and 176.)
- Lynne Hamill. Controlling smart devices in the home. *Information Society*, 22(4):241–249, September 2006. (Cited on pages 33, 68, 78, and 77.)

- Lynne Hamill and Richard Harper. Talking intelligence: a historical and conceptual exploration of speech-based human-machine interaction in smart homes. In *Proceedings of the 1st International Symposium on Intelligent Environments, 2006. ISIE '06, Cambridge, UK, April 5-7, 2006*, pages 121-128, Cambridge, UK, 2006. Microsoft Research Ltd., Microsoft Research Ltd. Press. (Cited on pages 33 and 77.)
- David Hanson, Andrew Olney, Ismar A. Pereira, and Marge Zielke. Upending the uncanny valley. In *Proceedings of the Twentieth national conference on artificial intelligence, 2005. AAAI '05, Pittsburgh, USA, July 9-13, 2005*, pages 1728-1729, Menlo Park, USA, 2005. Association for the Advancement of Artificial Intelligence, Association for the Advancement of Artificial Intelligence Press. (Cited on page 37.)
- Steve Harrison and Paul Dourish. Re-place-ing space: the roles of place and space in collaborative systems. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work, 1996. CSCW '96, Boston, US, November 16-20, 1996*, pages 67-76. ACM Press, New York, NY, USA, 1996. doi: 10.1145/240080.240193. (Cited on pages 58 and 57.)
- Kotaro Hayashi, Daisuke Sakamoto, Takayuki Kanda, Masahiro Shiomi, Satoshi Koizumi, Hiroshi Ishiguro, Tsukasa Ogasawara, and Norihiro Hagita. Humanoid robots as a passive-social medium: a field experiment at a train station. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10-12, 2007*, pages 137-144, New York, NY, USA, 2007. ACM, ACM Press. (Cited on pages 28 and 30.)
- Frank Hegel, Soeren Krach, Tilo Kircher, Britta Wrede, and Gerhard Sagerer. Theory of mind (ToM) on robots: a functional neuroimaging study. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12-15, 2008*, pages 335-342, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349866. (Cited on page 43.)
- Fritz Heider and Marianne Simmel. An experimental study of apparent behavior. *American Journal of Psychology*, 57:243-259, 1944. (Cited on pages 13, 168, 169, 164, and 165.)
- Aaron Hertzmann, Charles E. Jacobs, Nuria Oliver, Brian Curless, and David H. Salesin. Image analogies. In *Proceedings of the ACM SIGGRAPH International Conference and Exhibition on Computer Graphics and Interactive Techniques, 2001. SIGGRAPH '01, Los Angeles, CA, USA, August 12-17, 2001*, pages 327-340, New York, NY, USA, 2001. ACM, ACM Press. (Cited on pages 196, 201, 202, 203, 192, 197, 198, and 199.)
- Aaron Hertzmann, Nuria Oliver, Brian Curless, and Steven M. Seitz. Curve analogies. In *Proceedings of the ACM EUROGRAPHICS Workshop on Rendering, 2002. EGRW '02, Pisa, Italy, June 26-28, 2002*, pages 233-246, New York, NY, USA, 2002. ACM, ACM Press. ISBN 1-58113-534-3. (Cited on pages 179 and 175.)
- Uta Hinrichs, Sheelagh Carpendale, and Stacey Scott. Evaluating the effects of fluid interface components on tabletop collaboration. In *Proceedings of the International Working Conference on Advanced Visual Interfaces, 2006. AVI '06, Venezia, Italy, May 23-26, 2006*,

- pages 27–34, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1133265.1133272. (Cited on pages 120 and 116.)
- Chin-Chang Ho, Karl F. MacDorman, and Z. A. D. Dwi Pramono. Human emotion and the uncanny valley: a GLM, MDS, and isomap analysis of robot video ratings. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 169–176, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349845. (Cited on pages 84 and 83.)
- Guy Hoffman and Cynthia L. Breazeal. Collaboration in Human-Robot Teams. In *Proceedings of the American Institute of Aeronautics and Astronautics (AIAA) 1st Intelligent Systems Technical Conference, 2004. Chicago, IL, September 20–22, 2004*, Reston, VA, USA, 2004. American Institute of Aeronautics and Astronautics. (Cited on page 33.)
- Guy Hoffman and Cynthia L. Breazeal. Effects of anticipatory action on human-robot teamwork. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 1–8, New York, NY, USA, 2007. ACM, ACM Press. (Cited on page 33.)
- Kristina Höök. User-centered design and evaluation of affective interfaces. In *From Brows to Trust*, volume 7 of *Lecture Notes in Computer Science*, pages 127–160. Springer-Verlag, Berlin, New York, Heidelberg, 2005. (Cited on pages 16, 46, and 47.)
- Kristina Höök, Phoebe Sengers, and Gerd Andersson. Sense and sensibility: evaluation and interactive art. In *ACM Conference on Human Factors in Computing Systems, 2003. CHI '03, Fort Lauderdale, USA, April 5–10, 2003*, pages 241–248, New York, NY, USA, 2003. ACM, ACM Press. doi: 10.1145/642611.642654. (Cited on page 46.)
- Eva Hornecker and Jacob Buur. Getting a grip on tangible interaction: A framework on physical space and social interaction. In *ACM Conference on Human Factors in Computing Systems, 2006. CHI '06, Montréal, Québec, April 22–28, 2006*, pages 437–446, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1124772.1124838. (Cited on pages 58 and 57.)
- Helge Hüttenrauch and Kerstin Severinson Eklundh. Fetch-and-carry with CERO: Observations from a long-term user study with a service robot. In *IEEE International Workshop on Robot and Human Interactive Communication, 2002. ROMAN '02, Berlin, Germany, September 25–27, 2002*, pages 158–163, Los Alamitos, CA, USA, 2002. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ROMAN.2002.1045615. (Cited on page 87.)
- Takeo Igarashi, Tomer Moscovich, and John Hughes. Spatial keyframing for performance-driven animation. In *Proceedings of the ACM SIGGRAPH International Conference and Exhibition on Computer Graphics and Interactive Techniques, 2005. SIGGRAPH '05, Los Angeles, CA, USA, July 31–August 4, 2005*, pages 107–116, New York, NY, USA, July 2005a. ACM, ACM Press. doi: 10.1145/1073368.1073383. (Cited on pages 179 and 175.)
- Takeo Igarashi, Tomer Moscovich, and John F. Hughes. As-rigid-as-possible shape manipulation. *ACM Transactions on Graphics*, 24(3):1134–1141, July 2005b. ISSN 0730-0301. doi: 10.1145/1073204.1073323. (Cited on pages 179 and 175.)

- Katherine Isbister, Kia Höök, Michael Sharp, and Jarmo Laaksolahti. The sensual evaluation instrument: developing an affective evaluation tool. In *ACM Conference on Human Factors in Computing Systems, 2006. CHI '06, Montréal, Québec, April 22–28, 2006*, pages 1163–1172, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1124772.1124946. (Cited on pages 16, 45, 46, and 47.)
- Hiroshi Ishii and Brygg Ullmer. Tangible bits: Towards seamless interfaces between people, bits and atoms. In *ACM Conference on Human Factors in Computing Systems, 2007. CHI '07, Atlanta, GA, USA, March 22–27, 1997*, pages 234–241, New York, NY, USA, 1997. ACM, ACM Press. doi: 10.1145/258549.258715. (Cited on pages 41, 174, and 169.)
- Kentaro Ishii, Shengdong Zhao, Masahiko Inami, Takeo Igarashi, and Michita Imai. Designing laser gesture interface for robot control. In *Human-Computer Interaction – INTERACT 2009*, volume 5727/2009 of *Lecture Notes in Computer Science*, pages 479–492. Springer-Verlag, Berlin, New York, Heidelberg, 2009. doi: 10.1007/978-3-642-03658-3_52. (Cited on page 23.)
- Adam Jacoff, Elena Messina, Brian A. Weiss, Satoshi Tadokoro, and Yuki Nakagawa. Test arenas and performance metrics for urban search and rescue robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 2003. IROS '03, Las Vegas, USA, October 27–31, 2003*, pages 3396–3403, Los Alamitos, CA, USA, 2003. IEEE Computer Society, IEEE Computer Society Press. (Cited on page 22.)
- Waleed Kadous, Raymond Ka-Man Sheh, and Claude Sammut. Effective user interface design for rescue robotics. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 250–257, New York, NY, USA, 2006. ACM, ACM Press. (Cited on page 22.)
- Takayuki Kanda, Masayuki Kamashima, Michita Imai, Tetsuo Ono, Daisuke Sakamoto, Hiroshi Ishiguro, and Yuichiro Anzai. A humanoid robot that pretends to listen to route guidance from a human. *Autonomous Robots*, 22(1):87–100, January 2007a. doi: 10.1007/s10514-006-9007-6. (Cited on pages 178, 179, 174, and 175.)
- Takayuki Kanda, Rumi Sato, Naoki Saiwaki, and Hiroshi Ishiguro. A two-month field trial in an elementary school for long-term human-robot interaction. *IEEE Transactions on Robotics*, 23(5):962–971, October 2007b. doi: 10.1109/TRO.2007.904904. (Cited on page 87.)
- Saul M. Kassin. Heider and simmel (1944) revisited: causal attribution and the animated film technique. *Review of Personality and Social Psychology*, 3:125–150, 1982. (Cited on pages 169 and 165.)
- Charles C. Kemp, Cressel D. Anderson, Hai Nguyen, Alexander J. Trevor, and Zhe Xu. A point-and-click interface for the real world: laser designation of objects for mobile manipulation. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 241–248, New York, NY, USA, 2008. ACM, ACM Press. (Cited on page 22.)

- Sara Kiesler and Pamela Hinds. Introduction to This Special Issue on Human-Robot Interaction. *Human Computer Interaction*, 19(1/2):1–8, 2004. doi: 10.1109/TSMCA.2005.850577. (Cited on pages 3, 4, 6, and 20.)
- Marcus Kleinehagenbrock, Sebastian Lang, Jannik Fritsch, Frank Lömer, Gernot A. Fink, and Gerhard Sagerer. Person tracking with a mobile robot based on multi-modal anchoring. In *IEEE International Workshop on Robot and Human Interactive Communication, 2002. ROMAN '02, Berlin, Germany, September 25–27, 2002*, pages 423–429, Los Alamitos, CA, USA, 2002. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ROMAN.2002.1045659. (Cited on pages 99 and 95.)
- Boris Kluge, Christian Köhler, and Erwin Prassler. Fast and robust tracking of multiple moving objects with a laser range finder. In *Proceedings of the IEEE International Conference on Robotics and Automation, 2001. ICRA '01, Seoul, Korea, May 21–26, 2006*, volume 2, pages 1683–1688, Los Alamitos, CA, USA, 2001. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ROBOT.2001.932853. (Cited on pages 99 and 95.)
- Nathan Koenig, Sonia Chernova, Chris Jones, Matthew Loper, and Odest Chadwicke Jenkins. Hands-free interaction for human-robot teams. In *Proceedings of the ICRA 2008 Workshop: Social Interaction with Intelligent Indoor Robots (SI3R), in Proceedings of the IEEE International Conference on Robotics and Automation, 2008. ICRA '08, Pasadena, USA, May 19–23, 2008*, pages 35–41, Los Alamitos, CA, USA, 2008. IEEE Computer Society, IEEE Computer Society Press. (Cited on page 22.)
- Kurt Koffka. *Principles of Gestalt Psychology*. Routedledge, an imprint of the Taylor & Francis Group, London, UK, 1935. (Cited on page 5.)
- Johannes Kopf, Chi-Wing Fu, Daniel Cohen-Or, Oliver Deussen, Dani Lischinski, and Tien-Tsin Wong. Solid texture synthesis from 2d exemplars. *ACM Transactions on Graphics*, 26(3):2:1–2:9, 2007. doi: 10.1145/1276377.1276380. (Cited on pages 202 and 198.)
- Russel Kruger, Sheelagh Carpendale, Stacey Scott, and Anthony Tang. Fluid integration of rotation and translation. In *ACM Conference on Human Factors in Computing Systems, 2005. CHI '05, Portland, OR, USA, April 2–7, 2005*, pages 601–610, New York, NY, USA, 2005. ACM, ACM Press. doi: 10.1145/1054972.1055055. (Cited on pages 119, 125, 115, and 121.)
- David Kurlander, Tim Skelly, and Davil Salesin. Comic chat. In *Proceedings of the ACM SIGGRAPH International Conference and Exhibition on Computer Graphics and Interactive Techniques, 1996. SIGGRAPH '96, New Orleans, LA, USA, August 4–9, 1996*, pages 225–236, New York, NY, USA, 1996. ACM, ACM Press. (Cited on pages 139 and 135.)
- Henry A. Landsberger. *Hawthorne Revisited*. Ithaca Press, Ithaca, NY, USA, 1958. (Cited on page 48.)
- Paul Lapedes, Ehud Sharlin, and Mario Costa Sousa. Three dimensional tangible user interface for controlling a robotic team. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 343–350, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349867. (Cited on page 22.)

- Bruno Latour. *Science in Action: How to follow Scientists and Engineers Through Society*. MiltonKeys: Open University Press, 1987. (Cited on page 42.)
- Hyunjeong Lee, Hyun Jin Kim, and Changsu Kim. Autonomous behavior design for robotic appliances. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 201–208, New York, NY, USA, 2007. ACM, ACM Press. doi: 10.1145/1228716.1228744. (Cited on page 47.)
- Jehee Lee and Kang Hoon Lee. Precomputing avatar behavior from human motion data. In *ACM SIGGRAPH / Eurographics Symposium on Computer Animation, 2004. SCA '04, Grenoble, France, August 27–29, 2004*, pages 79–87, Germany, 2004. Eurographics Association Press, Eurographics Association. ISBN 3-905673-14-2. doi: 10.1145/1028523.1028535. (Cited on pages 179 and 175.)
- Min Kyung Lee, Jodi Forlizzi, Paul E. Rybski, Frederick Crabbe, Wayne Chung, Josh Finkle, Eric Glaser, and Sara Kiesler. The Snackbot: Documenting the design of a robot for long-term human-robot interaction. In *Proceedings of the 4th ACM/IEEE Conference on Human-Robot Interaction, 2009. HRI '09, San Diego, California, USA, March 11–13, 2009*, pages 7–14, New York, NY, USA, 2009. ACM, ACM Press. doi: 10.1145/1514095.1514100. (Cited on pages 84 and 83.)
- Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. Crowds by example. *Computer Graphics Forum*, 26(3):655–664, September 2007. doi: 10.1111/j.1467-8659.2007.01089.x. (Cited on pages 179 and 175.)
- Daniel T. Levin, Stephen S. Killingsworth, and Megan M. Saylor. Concepts about the capabilities of computers and robots: a test of the scope of adults' theory of mind. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 57–64, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349831. (Cited on page 43.)
- David Levy. *Love and Sex with Robots: The Evolution of Human-Robot Relationships*. Harper-CollinsPublishers, New York, USA, 2007. (Cited on pages 75 and 74.)
- Martijn Liem, Arnoud Visser, and Frans Groen. A hybrid algorithm for tracking and following people using a robotic dog. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 185–192, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349847. (Cited on pages 169 and 165.)
- Andrea Lockerd and Cynthia L. Breazeal. Tutelage and Socially Guided Robot Learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 2004. IROS '04, Sendai, Japan, September 28–October 2, 2004*, volume 4, pages 3475–3480, Los Alamitos, CA, USA, 2004. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/IROS.2004.1389954. (Cited on pages 34, 86, 179, 180, 175, and 176.)
- Karl F. MacDorman, Takashi Minato, Michihiro Shimada, Shoji Itakura, Stephen Cowley, and Hiroshi Ishiguro. Assessing human likeness by eye contact in an android testbed. In

- Proceedings of the 27th Annual Meeting of the Cognitive Science Society, 2005. CogSci '05, Stresa, Italy, July 21–23, 2005*, Mahwah, NJ, USA, 2005. Cognitive Science Society, Inc., Lawrence Erlbaum Associates, Inc. (Cited on pages 31 and 37.)
- Pattie Maes. Artificial life meets entertainment: lifelike autonomous agents. *Communications of the ACM*, 38(11):108–114, November 1995. doi: 10.1145/219717.219808. (Cited on pages 178 and 174.)
- Edwin Mansfield. The diffusion of industrial robots in japan and the united states. *Research Policy*, 18(4):183–192, August 1989. (Cited on pages 62 and 61.)
- Nicolai Marquardt, James E. Young, Ehud Sharlin, and Saul Greenberg. Situated messages for asynchronous human-robot interaction. In *adjunct proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (Late-Breaking Abstracts), 2009. HRI LBA '09, San Diego, California, USA, March 11–13, 2009*, pages 301–302, New York, NY, USA, 2009. ACM, ACM Press. (Cited on pages 23, 143, and 139.)
- Patrizia Marti, Alessandro Pollini, Alessia Rullo, and Takanori Shibata. Engaging with artificial pets. In *Proceedings of the annual conference of the European Association of Cognitive Ergonomics, 2005. EACE '05, Chania, Greece, September 29–October 1, 2005*, pages 99–106, New York, NY, USA, 2005. ACM, ACM Press. (Cited on pages 59, 84, and 83.)
- Kieran Mathieson. Predicting user intentions: Comparing the technology acceptance model with the theory of planned behavior. *Information Systems Research*, 2(3):173–191, September 1991. (Cited on pages 62, 65, 71, 61, and 64.)
- Daisuke Matsui, Takashi Minato, Karl F. MacDorman, and Hiroshi Ishiguro. Generating Natural Motion in an Android by Mapping Human Motion. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 2005. IROS '05, Edmonton, Canada, August 2–6, 2005*, pages 1089–1096, Los Alamitos, CA, USA, 2005. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/IROS.2005.1545125. (Cited on pages 179 and 175.)
- David Maulsby, Saul Greenberg, and Richard Mander. Prototyping an intelligent agent through wizard of oz. In *ACM Conference on Human Factors in Computing Systems, 1993. CHI '93, Amsterdam, The Netherlands, April 24–29, 1993*, pages 277–284, New York, NY, USA, 1993. ACM, ACM Press. doi: 10.1145/169059.169215. (Cited on page 39.)
- David L. Maulsby, Ian H. Witten, and Kenneth A. Kittlitz. Metamouse: specifying graphical procedures by example. In *Proceedings of the ACM SIGGRAPH International Conference and Exhibition on Computer Graphics and Interactive Techniques, 1989. SIGGRAPH '89, Boston, MA, USA, July 31 – August 4, 1989*, pages 127–136, New York, NY, USA, 1989. ACM, ACM Press. doi: 10.1145/74333.74346. (Cited on pages 179 and 175.)
- Scott McCloud. *Understanding Comics—The Invisible Art*. HarperPerennial – A Division of HarperCollinsPublishers. Originally published in 1993 by Kitchen Sink Press, reprinted by arrangement, New York, USA, First HarperPerennial edition, 1994. (Cited on pages 135, 146, 131, and 142.)

- A. McMeekin, K. Green, M. Tomlinson, and V. Walsh. *Innovation by Demand: an Interdisciplinary Approach to the Study of Demand and its Role in Innovation*. Edward Elgar, Cheltenham, 2002. (Cited on page 5.)
- MerriamWebsterOnline. Merriam-Webster OnLine dictionary. WWW, www.merriam-webster.com/, Accessed 2010. (Cited on pages 135, 166, 131, and 162.)
- Marek P. Michalowski, Selma Sabanovic, and Hideki Kozima. A dancing robot for rhythmic social interaction. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 89–96, New York, NY, USA, 2007. ACM, ACM Press. doi: 10.1145/1228716.1228729. (Cited on pages 32, 84, 139, and 135.)
- Paul Milgram and Fumio Kishino. A taxonomy of mixed reality visual displays. *IEICE Transactions on Information and Systems*, E77-D(12):1321–1329, December 1994. (Cited on pages 137, 142, 133, and 138.)
- Takashi Minato, Michihiro Shimada, Hiroshi Ishiguro, and Shoji Itakura. Development of an android for studying human-robot interaction. In *Proceedings of the Innovations in Applied Artificial Intelligence, Seventeenth International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems 2004. IEA/AIE '04, Ottawa, Canada, May 17–20, 2004*, pages 424–434, New York, NY, USA, 2004. ACM, ACM Press. doi: 10.1007/b97304. (Cited on page 37.)
- Takashi Minato, Yuichiro Yoshikawa, Tomoyuki Noda, Shuhei Ikemoto, Hiroshi Ishiguro, and Minoru Asada. Cb²: A child robot with biomimetic body for cognitive developmental robotics. In *Proceedings of the IEEE-RAS/RSJ International Conference on Humanoid Robots, 2007. Humanoids '07, Pittsburgh, USA, November 29–December 1, 2007*, Los Alamitos, CA, USA, 2007. IEEE Computer Society, IEEE Computer Society Press. (Cited on page 31.)
- Carlos Montalvo. *Environmental policy and technological innovation: Why do firms adopt or reject new technologies?* Edward Elgar Publishing, Cheltenham, UK, 2002. (Cited on page 62.)
- Michael Montemerlo, Sebastian Thrun, and William Whittaker. Conditional particle filters for simultaneous mobile robot localization and people-tracking. In *Proceedings of the IEEE International Conference on Robotics and Automation, 2002. ICRA '02, Washington, D.C., USA, May 11–15, 2002*, volume 1, pages 695–701, Los Alamitos, CA, USA, 2002. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ROBOT.2002.1013439. (Cited on pages 99 and 95.)
- Hans P. Moravec. *Robot: mere machine to transcendet mind*. Oxford University Press, 1998. (Cited on pages 7 and 20.)
- Masahiro Mori. Bukimi no tani: The uncanny valley (in Japanese.). *Energy*, 7:33–35, 1970. English translation provided at CogSci '05 workshop: Toward Social Mechanisms of Android Science, Views of the Uncanny Valley. WWW, <http://www.androidscience.com/>

- theuncannyvalley/proceedings2005/uncannyvalley.html. (Cited on pages 36, 37, 84, and 83.)
- Jon D. Morris. SAM: the Self-Assessment Manikin. an efficient cross-cultural measurement of emotional response. *Journal of Advertising Research*, 1:63–68, November 1995. (Cited on pages 108, 109, 110, 311, 104, 105, 106, and 307.)
- Claudia Muhl and Tukie Nagai. Does disturbance discourage people from communicating with a robot? In *IEEE International Workshop on Robot and Human Interactive Communication, 2009. ROMAN '09, Toyama, Japan, September 27–Oct 2, 2009*, pages 1137–1142, Los Alamitos, CA, USA, 2007. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ROMAN.2007.4415251. (Cited on pages 30, 86, and 85.)
- Bilge Mutlu, Toshiyuki Shiwa, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. Footing in human-robot conversations: How robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE Conference on Human-Robot Interaction, 2009. HRI '09, San Diego, California, USA, March 11–13, 2009*, pages 61–68, New York, NY, USA, 2009. ACM, ACM Press. doi: 10.1109/ROMAN.2007.4415251. (Cited on pages 30 and 85.)
- David Myers. *Social Psychology*. McGraw-Hill, New York, USA, 4th edition, 1993. (Cited on page 61.)
- Clifford Nass and Youngme Moon. Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1):81–103, 2000. doi: 10.1111/0022-4537.00153. (Cited on pages 25, 59, and 58.)
- Wai Shan Ng and Ehud Sharlin. Tweeting halo: Clothing that tweets. In *Adjunct proceedings of the 23rd annual ACM symposium on user interface and software technology, 2010. UIST '10, New York, NY, USA, October 3–6, 2010*, New York, NY, USA, 2010. ACM, ACM Press. (extended abstract, to appear). (Cited on pages 136, 279, 132, and 274.)
- Tatsuya Nomura, Tomohiro Suzuki, and Takayuki Kanda. Measurement of negative attitudes toward robots. *Interaction Studies*, 7(3):436–454, 2006. (Cited on pages 108, 311, 104, and 307.)
- Donald A. Norman. *The Design of Everyday Things*. Doubleday, New York, NY, USA, 1988. (Cited on pages 5, 45, 49, and 83.)
- Donald A. Norman. *Emotional Design: Why We Love (or Hate) Everyday Things*. Basic Books, New York, NY, USA, 2004. (Cited on pages 5, 7, 8, 58, 83, 130, 166, 177, 57, 126, 162, and 173.)
- Donald A. Norman. *The Design of Future Things*. Basic Books, New York, NY, USA, 2007. (Cited on page 6.)
- Illah R. Nourbakhsh, Clay Kunz, and Thomas Willeke. The mobot museum robot installations: A five year experiment. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 2003. IROS '03, Las Vegas, USA, October 27–31, 2003*, pages 3636–3641, Los Alamitos, CA, USA, 2003. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/IROS.2003.1249720. (Cited on pages 28 and 30.)

- Takayoshi Ootake, Ken-ichi Fukaya, Akira Syouzu, and Satoshi Nagai. Guided running of mobile robot through string that is connected to force sensor (japanese: 力覚センサを介したひも操作による移動ロボットの誘導走行). *Hokkaido University-Academy Department of Engineering Research Report (Japanese: 北海学園大学工学部研究報告)*, 35:85–90, 2008. ISSN 02865262. (Cited on pages 98 and 94.)
- Nuno Otero, Aris Alissandrakis, Kerstin Dautenhahn, Chrystopher Nehaniv, Dag S. Syrda, and Kheng Lee Koay. Human to robot demonstrations of routine home tasks: exploring the role of the robot's feedback. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 177–184, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349846. (Cited on pages 179 and 175.)
- Elena Pacchierotti, Henrik I. Christensen, and Patric Jensfelt. Design of an office guide robot for social interaction studies. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, 2006. IROS '06, Beijing, China October 9–15, 2006*, pages 4965–4970, Los Alamitos, CA, USA, 2006. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/IROS.2006.282519. (Cited on page 85.)
- Ivan P. Pavlov. *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*. Oxford University Press, Oxford, UK, 1927. Translated and Edited by G. V. Anrep. (Cited on page 69.)
- Ken Perlin and Athomas Goldberg. Improv: a system for scripting interactive actors in virtual worlds. In *Proceedings of the ACM SIGGRAPH International Conference and Exhibition on Computer Graphics and Interactive Techniques, 1996. SIGGRAPH '96, New Orleans, LA, USA, August 4–9, 1996*, pages 205–216, New York, NY, USA, 1996. ACM, ACM Press. ISBN 0-89791-746-4. doi: 10.1145/237170.237258. (Cited on pages 179 and 175.)
- Rosalind W. Picard. Affective computing for HCI. In *Proceedings of the HCI International Conference on Human-Computer Interaction, 1999. HCI '99, Munich, Germany, August 22–26, 1999*, pages 829–833, Mahwah, NJ, USA, 1999. Lawrence Erlbaum Associates, Inc. (Cited on pages 24, 41, and 45.)
- Trevor J. Pinch and Wiebe E. Bijker. The social construction of facts and artifacts: Or how the sociology of science and the sociology of technology might benefit each other. *Social Studies of Science*, 14:399–441, August 1984. (Cited on pages 42 and 52.)
- Robert Plutchik. A general psychoevolutionary theory of emotion. In Klaus R. Scherer and Paul Ekman, editors, *Emotion: Theory, research, and experience*, volume 1, pages 3–33. Academic Press, 1980. First presented at a conference on “The nature and function of emotion,” Bad Homburg, West Germany, July 6–10, 1981. (Cited on pages 171, 283, 167, and 278.)
- D. G. Premack and G. Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1:515–526, 1978. (Cited on page 43.)

- Hayes Solos Raffle, Amanda J. Parkes, and Hiroshi Ishii. Topobo: a constructive assembly system with kinetic memory. In *ACM Conference on Human Factors in Computing Systems, 2004. CHI '04, Vienna, Austria, April 24–29, 2004*, pages 647–654, New York, NY, USA, 2004. ACM, ACM Press. doi: <http://doi.acm.org/10.1145/985692.985774>. (Cited on pages 179 and 175.)
- Byron Reeves and Clifford Nass. *The 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. (Cited on pages 24, 58, 59, 166, 236, 162, and 232.)
- Craig W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the ACM SIGGRAPH International Conference and Exhibition on Computer Graphics and Interactive Techniques, 1987. SIGGRAPH '87, Anaheim, CA, USA, July 27–31, 1987*, pages 25–34, New York, NY, USA, 1987. ACM, ACM Press. doi: 10.1145/37401.37406. (Cited on pages 178 and 174.)
- Justin Richer and Jill L. Drury. A video game-based framework for analyzing human-robot interaction: characterizing interface design in real-time interactive multimedia applications. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 266–273, New York, NY, USA, 2003. ACM, ACM Press. doi: 10.1145/1121241.1121287. (Cited on pages 22 and 45.)
- Everett M. Rogers. *Diffusion of Innovations*. Free Press Publishers, New York, NY, USA, 1995. (Cited on pages 60, 63, 66, 71, 73, 65, and 72.)
- Daisuke Sakamoto, Takayuki Kanda, Tetsuo Ono, Hiroshi Ishiguro, and Norihiro Hagita. Android as a telecommunication medium with a human-like presence. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 193–200, New York, NY, USA, 2007. ACM, ACM Press. doi: 10.1145/1228716.1228743. (Cited on page 31.)
- Daisuke Sakamoto, Koichiro Honda, Masahiko Inami, and Takeo Igarashi. Sketch and run: a stroke-based interface for home robots. In *ACM Conference on Human Factors in Computing Systems, 2009. CHI '09, Boston, USA, April 4–9, 2009*, pages 197–200, New York, NY, USA, 2009. ACM, ACM Press. (Cited on pages 22 and 23.)
- Elizabeth B.-N. Sanders. Converging perspectives: Product development research for the 1990s. *Design Management Journal*, 3(4):49–54, 1992. (Cited on page 46.)
- Shin Sato and Shigeyuki Sakane. A Human-Robot Interface Using an Interactive Hand Pointer that Projects a Mark in the Real Work Space. In *Proceedings of the IEEE International Conference on Robotics and Automation, 2000. ICRA '00, San Francisco, California, USA, April 24–27, 2000*, Los Alamitos, CA, USA, 2000. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ROBOT.2000.844117. (Cited on pages 22 and 23.)
- Joe Saunders, Chrystopher L. Nehaniv, and Kerstin Dautenhahn. Teaching robots by moulding behavior and scaffolding the environment. In *Proceedings of the 1st ACM*

- SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 118–125, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1121241.1121263. (Cited on page 33.)
- Brian Scassellati. Theory of mind for a humanoid robot. *Autonomous Robots*, 12(1):13–24, January 2002. doi: 10.1023/A:1013298507114. (Cited on page 43.)
- Brian J. Scholl and Patrice D. Tremoulet. Perceptual causality and animacy. *Trends in Cognitive Sciences (TRENDS)*, 4(8):299–309, August 2000. doi: 10.1016/S1364-6613(00)01506-0. (Cited on pages 169 and 165.)
- Tibor Scitovsky. *The Joyless Economy: an Inquiry into Human Satisfaction and Consumer Dissatisfaction*. Oxford University Press, Oxford, UK, 1976. (Cited on pages 62 and 63.)
- Brennan P. Sellner, Laura M. Hiatt, Reid Simmons, and Sanjiv Singh. Attaining situational awareness for sliding autonomy. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 80–87, New York, NY, USA, 2006. ACM, ACM Press. (Cited on page 22.)
- Phoebe Sengers and Bill Gaver. Staying open to interpretation: engaging multiple meanings in design and evaluation. In *Proceedings of the 6th conference on Designing Interactive Systems, 2006. DIS '06, University Park, PA, US, June 26–28, 2006*, pages 99–108, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1142405.1142422. (Cited on page 47.)
- Ehud Sharlin, Benjamin Watson, Yoshifumi Kitamura, Fumio Kishino, and Yuichi Itoh. On tangible user interfaces, humans and spatiality. *Personal and Ubiquitous Computing*, 8(5): 338–346, 2004. ISSN 1617-4909. doi: 10.1007/s00779-004-0296-5. (Cited on pages 173, 174, 169, and 170.)
- Helen Sharp, Yvonne Rogers, and Jenny Preece. *Interaction Design: beyond human-computer interaction*. John Wiley & Sons, Inc., Hoboken, NJ, USA, 2nd edition, 2007. ISBN 0470018666. (Cited on pages 15 and 45.)
- Masahiro Shiomi, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. Interactive humanoid robots for a science museum. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 305–312, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1121241.1121293. (Cited on pages 28, 29, and 30.)
- Masahiro Shiomi, Daisuke Sakamoto, Kanda Takayuki, Carlos Toshinori Ishi, Hiroshi Ishiguro, and Norihiro Hagita. A semi-autonomous communication robot: a field trial at a train station. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 303–310, New York, NY, USA, 2007. ACM, ACM Press. doi: 10.1145/1349822.1349862. (Cited on page 28.)
- Toshiyuki Shiwa, Takayuki Kanda, Michita Imai, Hiroshi Ishiguro, and Norihiro Hagita. How quickly should communication robots respond? In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands*,

- March 12–15, 2008, pages 153–160, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349843. (Cited on pages 29 and 30.)
- Elaine Short, Justin Hart, Michelle Vu, and Bran Scassellati. No fair!!: an interaction with a cheating robot. In *Proceedings of the 5th ACM/IEEE Conference on Human-Robot Interaction, 2010. HRI '10, Osaka, Japan, March 2–5, 2010*, pages 219–226, New York, NY, USA, 2010. ACM, ACM Press. doi: 10.1145/1734454.1734546. (Cited on pages 86 and 85.)
- Candace L. Sidner, Christopher Lee, Louis-Philippe Morency, and Clifton Forlines. The effect of head-nod recognition in human-robot conversation. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 290–296, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1121241.1121291. (Cited on pages 30 and 85.)
- Roger Silverstone. Beneath the bottom line: households and information and communication technologies in an age of the consumer. *PICT Policy Research Papers*, 17, 1991. Economic and Social Research Council, Swindon. (Cited on page 60.)
- Roger Silverstone and David Morley. Families and their technologies: two ethnographic portraits. In *Household Choices*, pages 74–83. London, UK, Futures Publications, 1990. (Cited on page 60.)
- Emrah Akin Sisbot, Aurelie Clodic, Rachid Alami, and Maxime Ransan. Supervision and motion planning for a mobile manipulator interacting with humans. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 327–334, New York, NY, USA, 2008. ACM, ACM Press. (Cited on page 23.)
- Peter Squire, Greg Trafton, and Raja Parasuraman. Human control of multiple unmanned vehicles: effects of interface type on execution and task switching times. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, 2006. (Cited on page 22.)
- Maria Staudte and Matthew W. Crocker. Visual attention in spoken human-robot interaction. In *Proceedings of the 4th ACM/IEEE Conference on Human-Robot Interaction, 2009. HRI '09, San Diego, California, USA, March 11–13, 2009*, pages 77–84, New York, NY, USA, 2009. ACM, ACM Press. doi: 10.1145/1514095.1514111. (Cited on pages 30 and 85.)
- Carol J. Steiner. A philosophy for innovation: The role of unconventional individuals in innovation success. *Journal of Product Innovation Management*, 12(5):431–440, 1995. (Cited on pages 62 and 61.)
- P. Stoneman. *Technological Diffusion and the Computer Revolution: The UK Experience*. Cambridge University Press, Cambridge, UK, 1976. (Cited on page 63.)
- Anselm Strauss and Juliet Corbin. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Sage Publications, Thousand Oaks, London, New Delhi, 1998. (Cited on pages 46, 47, and 48.)

- Kristen Stubbs, Pamela J. Hinds, and David Wettergreen. Autonomy and common ground in human-robot interaction: A field study. *IEEE Intelligent Systems*, 22(2):42–50, March/April 2007. (Cited on page 30.)
- Ja-Young Sung, Lan Guo, Rebecca E. Grinter, and Henrik I. Christensen. “my Roomba is Rambo”: Intimate home appliances. In *Proceedings of the International Conference on Ubiquitous Computing, 2007. UBIComp '07, Innsbruck, Austria, September 16–17, 2007*, volume 4717/2007 of *Lecture Notes in Computer Science*, pages 145–162. Springer-Verlag, Berlin, New York, Heidelberg, 2007. doi: 10.1007/978-3-540-74853-3_9. (Cited on pages 25, 26, 28, 47, 58, 59, 84, 85, 87, 83, and 86.)
- Ja-Young Sung, Rebecca E. Grinter, Henrik I. Christensen, and Lan Guo. Housewives or technophiles?: understanding domestic robot owners. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 129–136, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349840. (Cited on pages 27, 47, 88, and 87.)
- JaYoung Sung, Henrik I. Christensen, and Rebecca E. Grinter. Robots in the wild: understanding long-term user. In *Proceedings of the 4th ACM/IEEE Conference on Human-Robot Interaction, 2009. HRI '09, San Diego, California, USA, March 11–13, 2009*, pages 45–52, New York, NY, USA, 2009a. ACM, ACM Press. (Cited on pages 25, 27, and 28.)
- JaYoung Sung, Rebecca E. Grinter, and Henrik I. Christensen. “pimp my Roomba”: designing for personalization. In *ACM Conference on Human Factors in Computing Systems, 2009. CHI '09, Boston, USA, April 4–9, 2009*, pages 193–196, New York, NY, USA, 2009b. ACM, ACM Press. doi: 10.1145/1518701.1518732. (Cited on pages 26, 88, and 87.)
- Susumu Tachi, Robert W. Mann, and Derek Rowell. Quantitative comparison of alternative soensory displays for mobility aids for the blind. *IEEE Transactions on Biomedical Engineering*, 30(9):571–577, September 1983. (Cited on pages 98 and 94.)
- Leila Takayama, Wendy Ju, and Clifford Nass. Beyond dirty, dangerous and dull: what everyday people think robots should do. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The Netherlands, March 12–15, 2008*, pages 25–32, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349827. (Cited on pages 88 and 87.)
- Leila Takayama, Victoria Groom, and Clifford Nass. I’m sorry, Dave: I’m afraid I won’t do that: social aspects of human-agent conflict. In *ACM Conference on Human Factors in Computing Systems, 2009. CHI '09, Boston, USA, April 4–9, 2009*, pages 2099–2108, New York, NY, USA, 2009. ACM, ACM Press. doi: 10.1145/1518701.1519021. (Cited on pages 86 and 85.)
- Fumihide Tanaka, Javier R. Movellan, Bret Fortenberry, and Kazuki Aisaka. Daily HRI evaluation at a classroom environment: reports from dance interaction experiments. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006. HRI '06, Salt Lake City, USA, March 2–4, 2006*, pages 3–9, New York, NY, USA, 2006. ACM, ACM Press. doi: 10.1145/1121241.1121245. (Cited on page 86.)

- Frank Thomas and Ollie Johnston. *The Illusion of Life: Disney Animation*. Disney Editions (Walt Disney Productions), first hyperion ed. edition, 1981. (Cited on pages 166, 167, 169, 162, 163, and 165.)
- Matthew Thorne, David Burke, and Michiel van de Panne. Motion doodles: an interface for sketching character motion. In *Proceedings of the ACM SIGGRAPH International Conference and Exhibition on Computer Graphics and Interactive Techniques, 2004. SIGGRAPH '04, Los Angeles, CA, August 8–12, 2004*, pages 424–431, New York, NY, USA, 2004. ACM, ACM Press. doi: 10.1145/1186562.1015740. (Cited on pages 179 and 175.)
- Patrice D. Tremoulet and Jacob Feldman. Perception of animacy from the motion of a single object. *Perception*, 29(8):943–951, May 2000. doi: 10.1068/p3101. (Cited on pages 169 and 165.)
- Munetoshi Unuma, Ken Anjyo, and Ryoza Takeuchi. Fourier principles for emotion-based human figure animation. In *Proceedings of the ACM SIGGRAPH International Conference and Exhibition on Computer Graphics and Interactive Techniques, 1995. SIGGRAPH '95, Los Angeles, CA, USA, August 6–11, 1995*, pages 91–96, New York, NY, USA, 1995. ACM, ACM Press. doi: 10.1145/218380.218419. (Cited on pages 203 and 199.)
- Liesbet van Zoonen. Gendering the internet: Claims, controversies and cultures. *European Journal of Communication*, 17(1):5–23, 2002. (Cited on page 60.)
- Alladi Venkatesh. Introduction to the special issue on “ICT in everyday life: Home and personal environments”. *Information Society*, 22(4):191–194, September 2006. (Cited on pages 62, 65, 69, and 61.)
- Viswanath Venkatesh and Susan Brown. A longitudinal investigation of personal computers in homes: Adoption determinants and emerging challenges. *MIS Quarterly*, 25(1):71–102, 2001. (Cited on pages 62, 65, 66, 70, 71, 73, 61, and 72.)
- Viswanath Venkatesh and Fred D. Davis. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Management Science*, 46(2):186–204, February 2000. (Cited on page 66.)
- Paul Viola and Michael Jones. Rapid Object Detection Using a Boosted Cascade of Simple Features. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2001. CVPR '01, Hawaii, US, December 9–14, 2001*, pages 511–518, Los Alamitos, CA, USA, 2001. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/CVPR.2001.990517. (Cited on pages 154, 157, 150, and 153.)
- Amy Volda, Rebecca E. Grinter, Nicolas Ducheneaut, W. Keith Edwards, and Mark W. Newman. Listening in: Practices surrounding iTunes music sharing. In *ACM Conference on Human Factors in Computing Systems, 2005. CHI '05, Portland, OR, USA, April 2–7, 2005*, pages 191–200, New York, NY, USA, 2005. ACM, ACM Press. doi: 10.1145/1054972.1054999. (Cited on page 47.)
- Eric Von Hippel. *Democratizing Innovation*. MIT Press, Cambridge, MA, 2005. (Cited on pages 5, 62, and 61.)

- Michael L. Walters, Dag S. Syrdal, Kerstin Dautenhahn, René te Boekhorst, and Kheng Lee Koay. Avoiding the uncanny valley: robot appearance, personality and consistency of behavior in an attention-seeking home scenario for a robot companion. *Autonomous Robots*, 24(2):159–178, February 2008. doi: 10.1007/s10514-007-9058-3. (Cited on page 37.)
- Jijun Wang and Michael Lewis. Human control for cooperating robot teams. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 9–16, New York, NY, USA, 2007. ACM, ACM Press. (Cited on pages 23, 260, and 255.)
- Joesph Weizenbaum. ELIZA – a computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1):36–45, January 1966. doi: 10.1145/365153.365168. (Cited on page 38.)
- Andrew Whiten. *Encyclopedia of Cognitive Science*, chapter Theory of Mind. John Wiley & Sons, Inc., Hoboken, NJ, USA, 2006. doi: 10.1002/0470018860.s00625. (Cited on pages 42 and 43.)
- Douglas J. Wiley and James K. Hahn. Interpolation synthesis of articulated figure motion. *j-IEEE-CGA*, 17(6):39–45, November 1997. doi: 10.1109/38.626968. (Cited on pages 179 and 175.)
- Robin Williams and David Edge. The social shaping of technology:. *Research Policy*, 25: 856–899, 1996. (Cited on pages 42, 61, and 63.)
- Langdon Winner. *The whale and the reactor: a search for limits in an age of high technology*, chapter Do Artifacts have politics?, pages 19–39. University of Chicago Press, Chicago, US, 1987. (Cited on pages 42, 63, and 62.)
- David Wolber. Pavlov: an interface builder for designing animated interfaces. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 4(4):347–386, 1997. doi: 10.1145/267135.267142. (Cited on pages 178 and 174.)
- WordWebOnline. WordWebOnline dictionary, WordWeb Software. WWW, www.wordwebonline.com, Accessed 2010. (Cited on pages 38, 54, 79, 257, 78, and 253.)
- Min Xin and Ehud Sharlin. Exploring human-robot interaction through telepresence board games. In *Advances in Artificial Reality and Tele-Existence*, volume 4282/2006 of *Lecture Notes in Computer Science*, pages 249–261. Springer-Verlag, Berlin, New York, Heidelberg, 2006. doi: 10.1007/11941354_26. (Cited on page 33.)
- Min Xin and Ehud Sharlin. Playing games with robots—a method for evaluating human-robot interaction. In N. Sarkar, editor, *Human-Robot Interaction*, pages 497–508. Advanced Robotic Systems International, Vienna, Austria, 2007. Invited submission. (Cited on page 33.)
- Fumitaka Yamaoka, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. How close?: model of proximity control for information-presenting robots. In *Proceedings of the 3rd ACM/IEEE Conference on Human-Robot Interaction, 2008. HRI '08, Amsterdam, The*

- Netherlands, March 12–15, 2008*, pages 137–144, New York, NY, USA, 2008. ACM, ACM Press. doi: 10.1145/1349822.1349841. (Cited on pages 30, 170, and 166.)
- Holly A. Yanco and Jill L. Drury. Classifying human-robot interaction: an updated taxonomy. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, 2004. SMC '04, The Hague, The Netherlands, October 10–13, 2004*, volume 3, pages 2841–2846, Los Alamitos, CA, USA, October 2004. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/ICSMC.2004.1400763. (Cited on pages 22 and 45.)
- Holly A. Yanco, Michael Baker, Robert Casey, Andrew Chanler, Munjal Desai, Dan Hestand, Brenden Keyes, and Philip Thoren. Improving human-robot interaction for remote robot operation. In *Proceedings of the Twentieth national conference on artificial intelligence, 2005. AAAI '05, Pittsburgh, USA, July 9–13, 2005*, volume 4, pages 1743–1744, Menlo Park, USA, 2005. Association for the Advancement of Artificial Intelligence, Association for the Advancement of Artificial Intelligence Press. (Cited on page 22.)
- James E. Young and Ehud Sharlin. Sharing spaces with robots: an integrated environment for human-robot interaction. In *Proceedings of the 1st International Symposium on Intelligent Environments, 2006. ISIE '06, Cambridge, UK, April 5–7, 2006*, pages 103–110, Cambridge, UK, April 2006. Microsoft Research Ltd., Microsoft Research Ltd. Press. (Cited on pages 139, 140, 142, 135, 136, and 138.)
- James E. Young, Ehud Sharlin, and Jeffrey E. Boyd. Implementing bubblegrams: The use of haar-like features for human-robot interaction. In *Proceedings of the second IEEE Conference on automation Science and Engineering, 2006. CASE '06, Shanghai, China, October 8–10, 2006*, pages 298–303, Los Alamitos, CA, USA, 2006. IEEE Computer Society, IEEE Computer Society Press. doi: 10.1109/COASE.2006.326897. (Cited on pages 139, 140, 151, 152, 157, 135, 136, 147, 148, and 153.)
- James E. Young, Min Xin, and Ehud Sharlin. Robot expressionism through cartooning. In *Proceedings of the 2nd ACM/IEEE Conference on Human-Robot Interaction, 2007. HRI '07, Washington, D.C., USA, March 10–12, 2007*, pages 309–316, New York, NY, USA, 2007. ACM, ACM Press. doi: 10.1145/1228716.1228758. (Cited on pages 12, 79, 139, 140, 135, and 136.)
- James E. Young, Ehud Sharlin, and Takeo Igarashi. The concept of a robot. In *Proceedings of the ACM/IEEE Human-Robot Interaction Pioneers Workshop, 2008. Pioneers '08, Amsterdam, The Netherlands, March 12, 2008*, New York, NY, USA, 2008. ACM, ACM Press. (Cited on page 6.)
- James E. Young, Richard Hawkins, Ehud Sharlin, and Takeo Igarashi. Toward acceptable domestic robots: Applying insights from social psychology. *Inagural Issue of the International Journal on Social Robotics.*, 1(1):95–108, January 2009. doi: 10.1007/s12369-008-0006-y. (Cited on pages 3, 4, 6, 36, 63, 166, 62, and 162.)
- James E. Young, Ehud Sharlin, and Takeo Igarashi. “what is Mixed Reality, anyway?” considering the boundaries of mixed reality in the context of robots. In Xiangyu Wang,

- editor, *Mixed Reality and Human-Robot Interaction*. Springer-Verlag, Berlin, New York, Heidelberg, 2010. (in press). (Cited on pages 137, 140, 142, 133, 136, and 138.)
- Shengdong Zhao, Koichi Nakamura, Kentaro Ishii, and Takeo Igarashi. Magic cards: a paper tag interface for implicit robot control. In *ACM Conference on Human Factors in Computing Systems, 2009. CHI '09, Boston, USA, April 4–9, 2009*, pages 173–182, New York, NY, USA, 2009. ACM, ACM Press. (Cited on page 23.)
- Tom Ziemke. Are robots embodied? In Christian Balkenius, Jordan Zlatev, Hideki Kozima, Kerstin Dautenhahn, and Cynthia Breazeal, editors, *Proceedings of the 1st International Workshop on Epigenetic Robotics - Modeling Cognitive Development in Robotic Systems, 2001. EPIROB '01, Lund, Sweden, September 17–18, 2001*, pages 75–83, Lund, Sweden, 2001. Lund University Cognitive Studies. (Cited on pages 38 and 39.)

COLOPHON

This thesis was typeset with L^AT_EX 2_ε

Final Version as of September 6, 2010 at 15:16.