Human–Robot Interaction

An Introduction

Christoph Bartneck, Tony Belpaeme, Friederike Eyssel,
Takayuki Kanda, Merel Keijsers, Selma Šabanović

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Contents

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Illustrations

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Tables

1

Introduction

1.1 About this book

There's a lot of talk about robots these days. Robots are in the news, on the movie screen, and even in our daily lives. Have you ever interacted with a robot? A vacuum-cleaning robot? A robotic toy, pet, or companion? Chances are that if you haven't, you will soon. Technology companies are eyeing the potential of personal robots, with start-ups as well as large multinationals readying themselves to revolutionize our world with robots.

But where is the field of robotics headed? What will, and should, our future with robots look like? How will robots find a place in our lives? These are still very open questions. A range of unknown, but exciting, futures awaits, in all of which robots support us, collaborate with us, transport us, or entertain us. If you've opened this book, you must be interested in seeing how this future might unfold. Perhaps you even want to get involved in shaping the robot revolution.

To get you started on this path, first of all, it is all about you: What kind of educational background do you have? Did you become fascinated by robots through your interest in engineering, psychology, art, or design? Or did you pick up this book because it rekindled a childhood fascination with robots? Human–robot interaction (HRI) is the endeavor that brings together ideas from a wide range of disciplines. Engineering, computer science, robotics, psychology, sociology, and design all have something to contribute to how we interact with robots. HRI lies at the confluence of these disciplines. As a computer scientist, it pays to know about social psychology; as a designer, there's value in dipping your toes in sociology.

If you have an engineering background, do you think you can build a robot that interacts with people, working only with other engineers? We, unfortunately, predict that you will not be able to do so. To design robots that people want to interact with, you need a good understanding of human social interaction. To reach such understanding, you need insight from people trained in the social sciences and humanities.

Are you a designer? Do you think you can design a socially interactive robot without working with engineers and psychologists? People's expectations about robots and their roles in everyday life are not just high, but they also vary a lot from person to person. Some people may tell you they want robots that will cook for them; others wish for a robot to do their homework, then have an intellectual conversation about the latest *Star Wars* movie. The prowess of robots as assistants, however, is still rather limited. Moravec's paradox, decades after being first expressed, still holds: anything that seems hard to people is relatively easy for machines, and anything a young child can do is almost impossible for a machine. As a designer, you would therefore need a good understanding of technological capabilities and of human psychology and sociology to create a design that is viable and realistic.

And last but not least, those of you who have training in psychology and sociology, do you want to just wait around for such robots to appear in our society? Wouldn't it be too late to start studying these technologies after they appear in our environment? Don't you want to have an impact on what they look like and how they interact? One thing you can do is start talking to friendly engineers and computer scientists, or have lunch with a designer. They will give your social science ideas some grounding in what is technically possible and help you find the areas in which your knowledge can have the most impact.

Just like the six of us writing this book, you will all need to work together. To do so in an effective way, you will need to understand the perspectives of HRI practitioners from different disciplines and be aware of the different kinds of expertise needed for developing successful HRI projects. In this book, we want to provide you with a broad overview of HRI topics central to the field and get you started on thinking about how you can contribute to them. We would like you to join us in expanding the boundaries of what is known and possible. Technology has progressed to a degree to which it is possible to build and program your own robot at little cost. Robots will be part of our future, so seize your chance to shape it. Go, read (this book!), create, test, and learn!

We assembled a team of leading experts from the wide spectrum of disciplines that contribute to HRI. All of our hearts beat for improving how humans and robots interact.

1.2 Christoph Bartneck

Christoph Bartneck is an associate professor and director of postgraduate studies at the Human Interface Technology Lab New Zealand (HIT Lab NZ) of the University of Canterbury. He has a background in industrial design and human–computer interaction, and his projects and studies have been published in leading journals, newspapers, and conferences. His interests lie in the fields of human-computer interaction, science and technology studies, and visual design. More specifically, he focuses on the effect of anthropomorphism on HAI. As a secondary re-

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

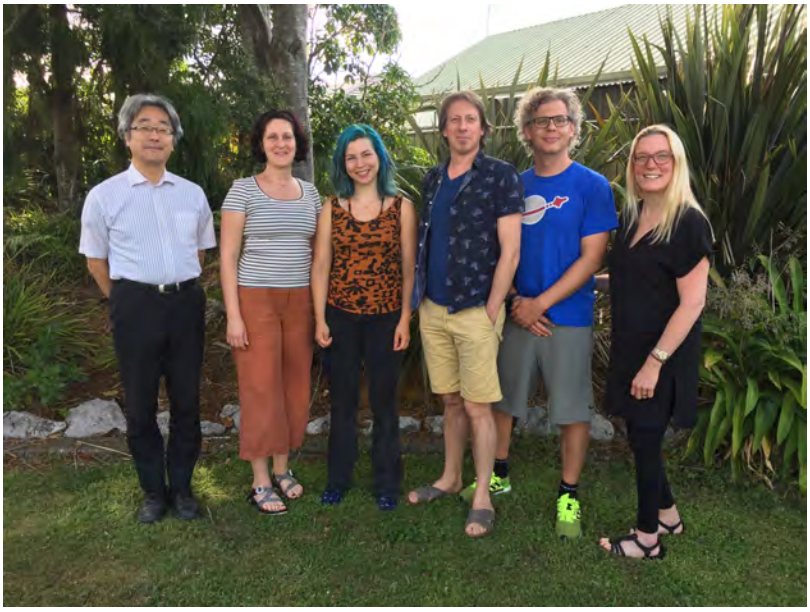

Figure 1.1 The authors of this book got together in Westport, New Zealand, in January 2018 to start the manuscript during a weeklong "Book Sprint". Writing and editing continued throughout the following year and a half through remote collaboration – many long Skype calls and emails.

search interest, he works on bibliometric analyses, agent-based social simulations, and the critical review of scientific processes and policies. In the field of design, Christoph investigates the history of product design, tessellations, and photography. The press regularly reports on his work, including *New Scientist*, *Scientific American*, *Popular Science*, *Wired*, the *New York Times*, *The Times*, the British Broadcasting Corporation (BBC), HuffPost, the *Washington Post*, *The Guardian*, and *The Economist*.

1.3 Tony Belpaeme

Tony Belpaeme is a professor at Ghent University, Belgium, and a professor of robotics and cognitive systems at Plymouth University, United Kingdom. He received his PhD in computer science from the Vrije Universiteit Brussel (VUB). Starting from the premise that intelligence is rooted in social interaction, Tony and his research team try to further the artificial intelligence of social robots. This approach leads to a spectrum of results, from theoretical insights to practical applications. He is involved in large-scale projects studying how robots can be used to support children in education, and he studies how brief interactions with robots can become long-term interactions and how robots can be used in therapy.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

1.4 Friederike Eyssel

Friederike Eyssel is a professor of applied social psychology and gender research at the Center of Excellence Cognitive Interaction Technology at Bielefeld University, Germany. Friederike is interested in various research topics ranging from social robotics, social agents, and ambient intelligence to attitude change, prejudice reduction, and the sexual objectification of women. Crossing disciplines, Friederike has published vastly in the field of social psychology, HAI, and social robotics and serves as a reviewer for more than 20 journals. Current third-party funded research projects (DFG, BMBF, FP7) address user experience and smart-home technologies and the ethical aspects associated with assistive technology and social robots in general.

1.5 Takayuki Kanda

Takayuki Kanda is a professor in informatics at Kyoto University, Japan. He is also the visiting group leader at Advanced Telecommunications Research (ATR) Intelligent Robotics and Communication Laboratories, Kyoto, Japan. He received his bachelor's degree in engineering, his master's degree in engineering, and his PhD in computer science from Kyoto University, Kyoto, Japan, in 1998, 2000, and 2003, respectively. He is one of the starting members of the Communication Robots project at the Advanced Telecommunications Research (ATR) Institute in Kyoto. He has developed a communication robot, Robovie, and applied it in daily situations, such as peer tutoring at an elementary school and as a museum exhibit guide. His research interests include HAI, interactive humanoid robots, and field trials.

1.6 Merel Keijsers

Merel Keijsers is a PhD student at the HIT Lab NZ, University of Canterbury. She has a research master's degree in statistics and in social and health psychology from the University of Utrecht. In her PhD program, she studies what conscious and subconscious psychological processes drive people to abuse and bully robots. Having a background in social psychology, she is mainly interested in the similarities and differences in how people deal with robots versus other humans.

1.7 Selma Šabanović

Selma Šabanović is an associate professor of informatics and cognitive science at Indiana University, Bloomington, where she founded and directs the R-House Human-Robot Interaction Lab. Her research combines studies of the design, use, and consequences of socially interactive

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

and assistive robots in different social and cultural contexts, including healthcare institutions, user homes, and various countries. She also engages in the critical study of the societal meaning and potential effects of developing and implementing robots in everyday contexts. She received her PhD in science and technology studies from Rensselaer Polytechnic Institute in 2007, with a dissertation on the cross-cultural study of social robotics in Japan and the United States. She currently serves as the Editor in Chief of the journal *ACM Transactions on Human-Robot Interaction*.

2

What Is Human–Robot Interaction?

What is covered in this chapter:

- the academic disciplines that come together in the field of human–robot interaction (HRI);
- the barriers created by the disciplines' different paradigms, and how to work around this;
- the history and evolution of HRI as a science;
- landmark robots in HRI history.

Human–robot interaction, or HRI, is commonly referred to as a new and emerging field, but the notion of human interaction with robots has been around as long as the notion of robots themselves. Isaac Asimov, who coined the term *robotics* in the 1940s, wrote his stories around questions that take the relationship between humans and robots as the main unit of analysis: "How much will people trust robots?"; "What kind of relationship can a person have with a robot?"; "How do our ideas of what is human change when we have machines doing humanlike things in our midst?" (see page 193 for more on Asimov). Decades ago, these ideas were science fiction, but nowadays, many of these issues have become a reality in contemporary societies and have become core research questions in the field of HRI.

Distinguishing physical and social interaction: One way to understand some key differences between the fields of HRI and robotics is that whereas robotics is concerned with the creation of physical robots and the ways in which these robots manipulate the physical world, HRI is concerned with the ways in which robots interact with people in the social world. For example, when the humanoid ASIMO (see Figure 2.1) goes up the stairs in a house or pushes a cart in an office, it is sensing and acting in the physical world alone and dealing with the physics of its own body and its environment. When ASIMO delivers coffee to a group of office workers or chases children around in a courtyard, it is dealing with the physical motions needed for those actions, but it must also address the

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

social aspects of the environment: where the children or the office workers are, how to approach in a way that is safe and that they consider appropriate, and the social rules of the interaction. Such social rules might be obvious to humans, such as acknowledging the other actors, knowing who is "it" in a game of tag, and saying "you're welcome" when someone says "thank you." But for a robot, all these social rules and norms are unknown and require the attention of the robot designer. These concerns make HRI questions different from those pursued in robotics alone.

As a discipline, HRI is related to human–computer interaction (HCI), robotics, artificial intelligence, the philosophy of technology, and design. Scholars trained in these disciplines have worked together to develop HRI, bringing in methods and frameworks from their home disciplines but also developing new concepts, research questions, and HRI-specific ways of studying and building the world.

What makes HRI unique? Clearly, the interaction of humans with social robots is at the core of this research field. These interactions usually include physically embodied robots, and their embodiment makes them inherently different from other computing technologies. Moreover, social robots are perceived as social actors bearing cultural meaning and having a strong impact on contemporary and future societies. Saying that a robot is embodied does not mean that it is simply a computer on legs or wheels. Instead, we have to understand how to design that embodiment, both in terms of software and hardware, as is commonplace in robotics, and in terms of its effects on people and the kinds of interactions they can have with such a robot.

A robot's embodiment sets physical constraints on the ways in which it can sense and act in the world, but it also represents an affordance for interaction with people. The robot's physical makeup elicits people to respond in a way similar to that in which they interact with other people. The robots' human-likeness enables humans to use their existing experience of human–human interaction in human–robot interaction. These experiences can be very useful to frame an interaction, but they can also lead to frustration if the robot cannot live up to the users' expectations.

HRI focuses on developing robots that can interact with people in various everyday environments. This opens up technical challenges resulting from the dynamics and complexities of humans and the social environment. This also opens up design challenges—related to robotic appearance, behavior, and sensing capabilities—to inspire and guide interaction. From a psychological perspective, HRI offers the unique opportunity to study human affect, cognition, and behavior when confronted with social agents other than humans. Social robots, in this

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 2.1 Honda developed the Asimo robot from 2000 through 2018. (Source: Honda)

context, can serve as research tools to study psychological mechanisms and theories.

When robots are not just a tool but, rather, collaborators, companions, guides, tutors, and all kinds of social interaction partners, HRI research considers many different relationships with the development of society, both in the present and in the future. HRI research includes issues related to the social and physical design of technologies, as well as societal and organizational implementation and cultural sense-making, in ways that are distinct from related disciplines.

2.1 The focus of this book

HRI is a large, multidisciplinary field, and this book provides an introduction to the problems, processes, and solutions involved. This book enables the reader to gain an overview of the field without becoming overwhelmed with the complexities of all the challenges that we are facing, although we do provide references to the most relevant literature, which interested readers might want to investigate at their leisure. This book provides a much-needed introduction to the field so that students, academics, practitioners, and policy makers can become familiar with the future of how humans will interact with technology.

This book is an introduction, and as such, it does not require extensive knowledge in any of the related fields. It only requires the reader's curiosity about how robots and humans can and should interact with each other.

After introducing the field of HRI and how a robot works in principle, we focus on the robots' designs. Next, we address the different interaction modalities through which humans can interact with robots, such as through speech or gestures. The processing and communication of emotions is the next challenge we introduce before reflecting on the role that robots play in the media. The research methods chapter introduces the unique issues that researchers face when conducting empirical studies of humans interacting with robots. Next, we cover the application areas of social robots and their specific challenges before discussing ethical issues around the use of social robots. The book closes with a look into the future of HRI.

2.2 HRI as an interdisciplinary endeavor

HRI is multidisciplinary and problem-based field by nature and by necessity. HRI brings together scholars and practitioners from various domains: engineers, psychologists, designers, anthropologists, sociologists, and philosophers, along with scholars from other application and research domains. Creating a successful human–robot interaction requires collaboration from a variety of fields to develop the robotics hardware and software, analyze the behavior of humans when interacting with robots in different social contexts, and create the aesthetics of the embodiment and behavior of the robot, as well as the required domain knowledge for particular applications. This collaboration can be difficult due to the different disciplinary jargon and practices. The common interest in HRI among this wide variety of participants, how-

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

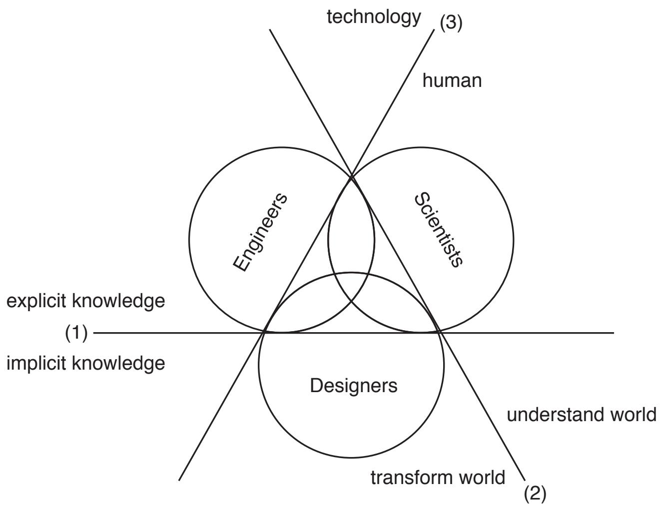

Figure 2.2 Barriers between the disciplines.

ever, is a strong motivation for familiarizing oneself with and respecting the diverse ways of acquiring knowledge.

HRI is, in this multidisciplinary sense, similar to the field of human–computer interaction (HCI), although dealing with embodied interactions with social agents differentiates HRI from HCI.

The various disciplines differ from each other in terms of their shared beliefs, values, models, and exemplars (Bartneck and Rauterberg, 2007). These aspects form a "paradigm" that guides their community of theorists and practitioners (Kuhn, 1970). Researchers within a paradigm share beliefs, values, and exemplars. The difficulties of working together on a shared project find their base in three barriers (see Figure 2.2) between designers [D], engineers [E] and scientists (in particular social scientists) [S]:

1. knowledge representation (explicit [S, E] versus implicit [D]);
2. view on reality (understanding [S] versus transforming reality [D, E]); and
3. main focus (technology [E] versus human [D, S]).

Barrier 1: Engineers [E] and scientists [S] make their results explicit by publishing in journals, books, and conference proceedings or by acquiring patents. Their body of knowledge is externalized and described to other engineers or scientists. These two communities revise their published results through discussion and control tests among peers. On the other hand, designers' [D] results are mainly represented by their concrete designs. The design knowledge necessary to create these designs lies within the individual designer, mainly as implicit knowledge, often referred to as intuition.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Barrier 2: Engineers [E] and designers [D] transform the world into preferred states (Simon, 1996; Vincenti, 1990). They first identify a preferred state, such as the connection between two sides of a river, and then implement the transformation, which in our example would be a bridge. Scientists [S] mainly attempt to understand the world through the pursuit of knowledge covering general truths or the operation of general laws.

Barrier 3: Scientists [S] and designers [D] are predominantly interested in humans in their role as possible users. Designers are interested in human values, which they transform into requirements and, eventually, solutions. Scientists in the HCI community are typically associated with the social or cognitive sciences. They are interested in the users' abilities and behaviors such as perception, cognition, and action, as well as the way these factors are affected by the different contexts in which they occur. Engineers [E] are mainly interested in technology, which includes software for interactive systems. They investigate the structure and operational principles of these technical systems to solve certain problems.

Not every HRI project can afford to have dedicated specialists from all these disciplines. HRI researchers often need to wear several hats, trying to gain expertise in a variety of topics and domains. Although this approach may reduce the problems of finding common ground, it is quite limiting. We often do not know what we do not know. It is therefore important to either engage with all or many of the involved disciplines directly or at least communicate with experts in the respective fields. As the field of HRI grows and matures, it has also been expanding to include more and more different disciplines, frameworks, and methods (e.g., historians, philosophers), which can require an even more expansive set of knowledge requirements. In this case, we suggest also getting used to reading broadly, not just in your own discipline or subdomain of HRI but also in related fields, to understand how your own work fits into the bigger picture. When developing specific HRI applications, it is also crucial to collaborate with domain experts, including potential users and stakeholders, in the design–from the beginning of the project–to make sure to ask relevant questions, use appropriate methods, and be aware of the potential broader consequences of the research to the application domain.

2.3 The evolution of HRI

The concept of "robot" has a long and rich history in the cultural imagination of many different societies, going back thousands of years to tales of humanlike machines, the later development of automata that reproduce certain human capabilities, and more recent science-fiction narratives about robots in society. Although these cultural notions of

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

robots may not always be technically realistic, they color people's expectations of and reactions to robots.

The first mention of "social robot" in print was in 1935, when it was used as a derogatory term for a person having a cold and distant personality.

> Toadying and bootlicking his autocratic superiors, he is advanced to preferment. He is a business success. But he has sacrificed all that was individual. He has become a social robot, a business cog. (Sargent, 2013)

In 1978, the first mention of "social robot" was made in the context of robotics. An article in *Interface Age* magazine described how a service robot, in addition to skills such as obstacle avoidance, balancing, and walking, would also need social skills to operate in a domestic setting. The article calls this robot a "social robot."

Ever since the concept of "robot" emerged, first in fiction and later as real machines, we have pondered the relationship between robots and people and how they could interact with each other. Every new technological or conceptual development in robotics has forced us to reconsider our relationship with and perception of robots.

When the first industrial robot, the Unimate, was installed at General Motors' Inland Fisher Guide Plant in Ewing Township, New Jersey, in 1961, people did consider how they would interact with the robot, but they were more concerned about the place robots would take among human workers. People who saw behavior-based robots for the first time could not help but marvel at the lifelike nature of the robot. Simple reactive behaviors (Braitenberg, 1986) implemented on small mobile robots produced machines that seemed injected with the very essence of life. Scurrying and fidgeting around the research labs of the 1990s, these robots evoked humanlike character traits and fundamentally changed our idea of how intelligence, or at least the appearance of intelligence, could be created (Brooks, 1991; Steels, 1993). This led to the creation of robots that used fast, reactive behavior to create a sense of social presence.

An early example of a social robot is Kismet (see Figure 2.3). Developed at the Massachusetts Institute of Technology in 1997, Kismet was a robot head-and-neck combination mounted on a tabletop box. Kismet could animate its eyes, eyebrows, lips, and neck, allowing it to pan, tilt, and crane its head. Based on visual and auditory input, it reacted to objects and people appearing in its visual field. It extracted information on visual motion, visual looming, sound amplitude, and emotion from speech prosody, and it responded by animating its facial expressions, ears, and neck and by babbling in a nonhuman language

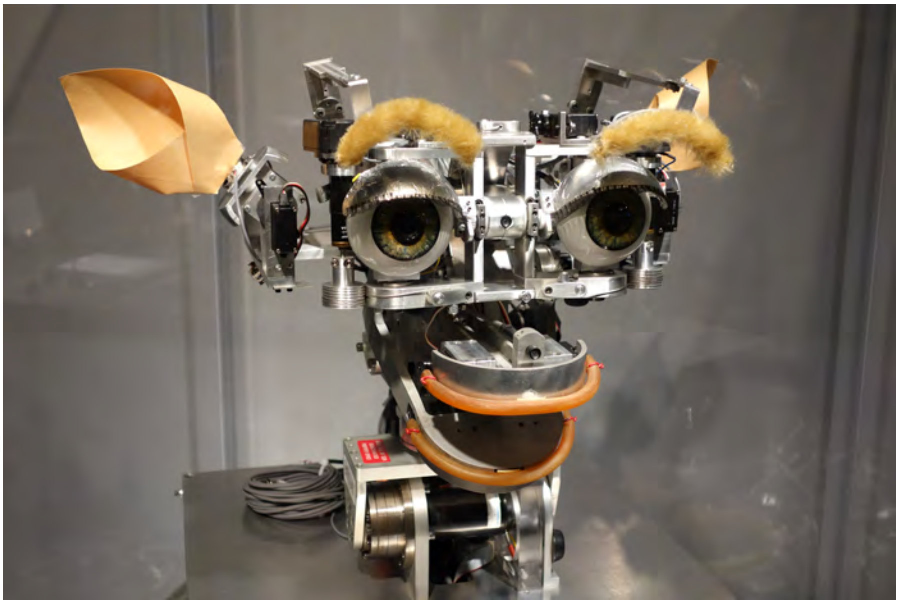

Figure 2.3 Kismet (1997–2004), an early example of social human–robot interaction research from the Massachusetts Institute of Technology. (Source: Daderot)

(Breazeal, 2003). Kismet was surprisingly effective at presenting a social presence, even though the control software only contained a small selection of social drives. It did so not only with its hardware and software architectures but also by taking advantage of human psychology, including what is known as the "baby schema," a predisposition to treat things with big eyes and exaggerated features in social ways despite their lack of fully functional social skills.

Like many robots in the early days of social robotics and HRI, Kismet was a bespoke robot, available to researchers in only one laboratory and requiring constant effort by students, postdocs, and other researchers to keep up and build up the robot's capabilities. These limitations understandably constrained the number of people and the range of disciplines that could participate in HRI in the field's early days. More recently, HRI research has been bolstered by the availability of reasonably priced commercial platforms that can be readily purchased by laboratories. These have expanded both the replicability and comparability of HRI research across labs, as well as the range of people who can engage in the discipline.

A number of robots have had a significant influence on the field. The Nao robot, developed by Aldebaran Robotics (now Softbank Robotics Europe), is perhaps the most influential robot in the study of social robotics (see Figure 2.4). First sold in 2006, the small humanoid robot, due to its affordability, robustness, and ease of programming, became a widespread robot platform for studying HRI. The robot, because of its size, is also highly portable, allowing for studies to be run outside the lab.

Figure 2.4 Nao (2006–present), a 58-cm-tall humanoid robot, currently the most popular research platform in social robotics.

The Keepon robot, developed by Hideki Kozima, is a minimal robot consisting of two soft yellow spheres to which a nose and two eyes are added. The robot can swivel, bend, and bop, using motors worked into the base of the robot (Kozima et al., 2009); see Figure 2.5. Keepon was later commercialized as an affordable toy, and through some moderate hacking, it can be used as a research tool for HRI. Studies with the Keepon robot convincingly demonstrated that a social robot does not need to appear humanlike; the simple form of the robot is sufficient to achieve interaction outcomes where one might assume the need for more complex and humanlike robots.

Figure 2.5 Keepon (2003–present), a minimal social robot developed by Hideki Kozima. The robot was later commercialized as an affordable toy. (Source: Hideki Kozima, Tohoku University)

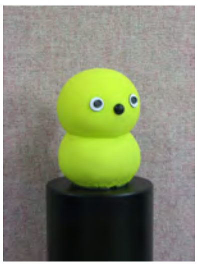

The Paro companion and therapy robot (see Figure 2.6), shaped like a baby seal, has been particularly popular in the study of socially assistive robots in eldercare, as well as other scenarios. Paro has been commercially available in Japan since 2006 and in the United States and Europe since 2009 and is a robust platform that requires almost no technical competence to operate. Paro has therefore been used by various psychologists, anthropologists, and health researchers, both to study the potential psychological and physiological effects on people and to explore ways in which robots might be adopted in healthcare organizations. The simplicity of the robot's operation and its robustness enable its use in many different contexts, including in long-term and naturalistic studies. At the same time, the fact that it is a closed platform—which does not allow robot logs or sensor data to be extracted from the robot or allow the robot's behaviors to be changed—poses some limitations for HRI research.

The Baxter robot, sold by Rethink Robotics until 2018, is both an

Figure 2.6 Paro (2003–present), a social robot made to resemble a baby harp seal. Paro is provided as a social companion robot. (Source: National Institute of Advanced Industrial Science and Technology)

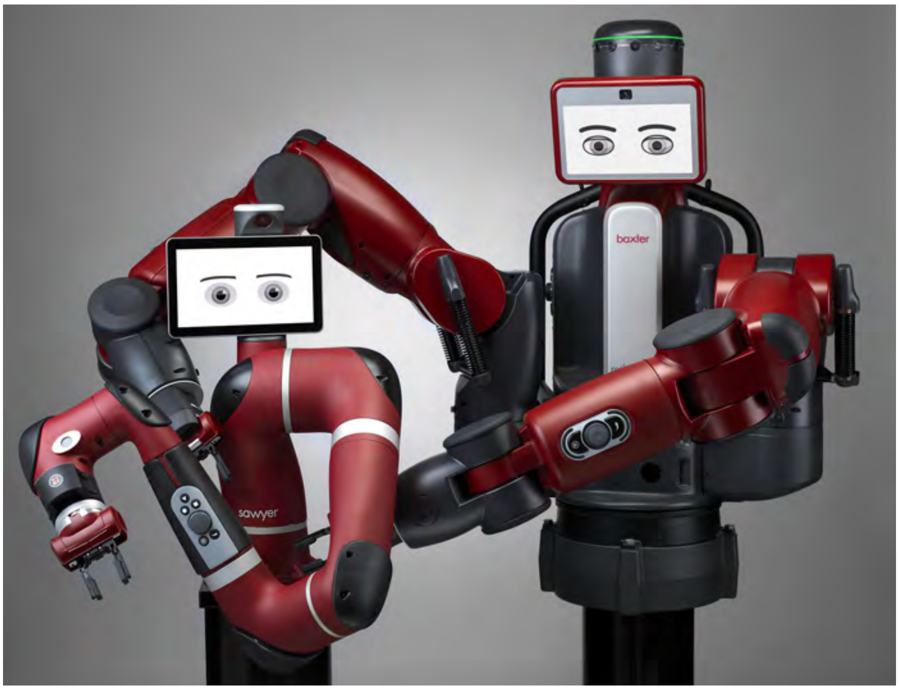

Figure 2.7 Baxter (2011–2018) and Sawyer (2015–2018), industrial robots with compliant arms by Rethink Robotics. Baxter was the first industrial robot to include social interaction features on an industrial manipulator. (Source: Rethink Robotics, Inc.)

industrial robot and a platform for HRI (Figure 2.7). The robot's two arms are actively compliant: in contrast to the stiff robot arms of typical industrial robots, Baxter's arms move in response to an externally applied force. In combination with other safety features, the Baxter robot is safe to work near, which makes it suitable for collaborative tasks. In addition, Baxter has a display screen mounted at head height on which the control software can display facial animations. Baxter's face can be used to communicate its internal state, and its eye fixations communicate a sense of attention to the human co-worker.

Although the availability of affordable commercial robots with open application interfaces caused a proliferation of HRI studies, a second development has allowed for in-house-built social robots. New develop-

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 2.8 InMoov (2012–present) can be built using rapid-prototyping technology and readily available components. The InMoov robot is an open-source social robot.

ments in mechatronic prototyping mean that robots can be modified, hacked, or built from scratch. Three-dimensional (3D) printing, laser cutting, and the availability of low-cost single-board computers have made it possible for researchers to build and modify robots in a short time and at minimal cost, for example, InMoov (see Figure 2.8) or Ono.

As you can see, the variety of robot hardware opens up endless research questions that can be addressed from a multidisciplinary perspective. Unlike other disciplines, HRI places particular emphasis on investigating the nature of social interactions between humans and robots, not only in dyads but also in groups, institutions, and sooner or later, in our societies. As will become clear in this book, technological advancements are a result of joint interdisciplinary efforts that have important societal and ethical implications. Keeping these in mind by doing human-centered research will hopefully lead to the development of robots that are widely accepted and that serve humans for the greater good.

Figure 2.9 Kaspar (2009–present) is a "minimally expressive" robot, built using brackets, servo motors, and a surgical silicon mask. Kaspar is used in autism therapy.

Questions for you to think about:

- The HRI field draws insights from many other fields, but what other fields could benefit from research in HRI?
- Are you a designer, engineer, or social scientist? Try to imagine a situation in which you are collaborating with others to construct a robot (e.g., if you are an engineer, you are now working with a designer and a social scientist on this endeavor). How is your way of working different from the approaches the other teammates might use?
- What is the main difference between the disciplines of HRI and HCI, and what makes HRI unique as a new field?

3

How a Robot Works

This chapter is written for readers who have limited technical background about intelligent interactive robotics. More specifically, what is covered in this chapter:

- the basic hardware and software components that a robot consists of;
- the techniques we can apply to make a robot ready for interacting with people.

As a way of thinking about how a robot works, let us role play by imagining being a robot. We might think we can do a lot of things, but we soon find out our capabilities are severely limited. If we are newly built robot, without appropriate software, our brains are completely empty. We cannot do anything—move, know where we are, understand what is around us, even ask for help. We find the experience of being a robot rather strange and difficult to imagine. The main source of strangeness is that the new robot's brain is nothing like a human brain, not even an infant's. The robot has no basic instincts, no goals, no memory, no needs, no learning capabilities, and no ability to sense or act. To make a robot system, we need to integrate, and at least partially develop, hardware and software together to enable the robot to sense and act in the world.

In this chapter, we look at the common components of a robot and how they are connected to enable participation in interaction. Section 3.1 explains basic ideas about the components needed to build a robot. Section 3.2 explains the types of hardware. Section 3.3 introduces sensors, such as cameras, range finders, and microphones, and Section 3.4 introduces actuators. Finally, Section 3.5 explains the software that accompanies the hardware elements, which addresses the perception (e.g., computer vision), planning, and action control of the robot.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

3.1 The making of a robot

To build a robot, one of the first steps is to establish connections between the robot's sensors, computer, and motors so that the robot is able to sense, interpret what it senses, plan actions, and then act them out. Once the robot is connected, to a camera, for example, its computer can read the data the camera provides. But the camera image is nothing more than a large table of numbers, similar to the following table:

9	15	10
89	76	81
25	34	29

From these numbers, can you guess what the robot is seeing? Perhaps a ball, an apple, or a fork? Assuming that each value in the table represents the lightness value of one sensor element in the camera, we can translate those numbers to a graphic that is more meaningful to humans (see Figure 3.1), but the graphic remains meaningless to the robot.

You might be able to see a line in the image shown in Figure 3.1, but a robot has no understanding of what a line is. This line might be the edge of a cliff from which the robot could fall and damage itself. But the robot does not have a concept of height or gravity. It would not comprehend that it could fall if it crossed this line. It does not know that if it fell, it would likely come to rest upside down. It would not even recognize that its arm would be broken. In other words, even concepts that are vitally important for interacting with and surviving in the world around us that are innate in humans have to be explicitly programmed in a robot.

Figure 3.1 The camera's data translated into a grid of grayscale pixels.

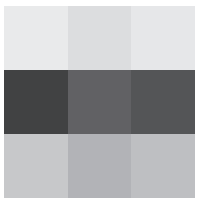

A robot, in essence, is a computer with a body. Any functionality needs to be programmed into the robot. A problem that all robots have to deal with is that although their sensors and motors are sufficient for operating in this world, their intelligence is not. Any concept of interest to roboticists needs to be internalized, that is, programmed into the robot. This requires a lot of time and effort and often involves many cycles of trial and error. The analogue world out there is converted into a digital world, and translating tables of numbers into meaningful information and meaningful responses is one of the core goals of artificial intelligence. Being able to identify a face from a large table of values, recognizing if a person has been seen before, and knowing that person's name are all skills that require programming or learning. Thus, the progress of human–robot interaction (HRI) is constrained by the progress that is made in the field of artificial intelligence. Robotics engineers integrate sensors, software, and actuators to enable the robot to make sense of and interact with its physical and social environment. An engineer might, for example, use accelerometer sensors, which can detect acceleration and the Earth's gravitational pull, to read the orientation of the robot and determine if it has fallen. A cliff sensor, consisting of a small infrared light source pointing down and a light sensor, can be used by the robot to avoid falling down a staircase.

Typical problems that robot engineers have to solve for the robot include the following:

- What kind of body does the robot have? Does it have wheels? Does it have arms?
- How will the robot know its location in space?
- How does the robot control and position its body parts—for example, arms, legs, wheels?
- What does the space around the robot look like? Are there obstacles, cliffs, doors? What does the robot need to be able to perceive about this environment to move safely?
- What are the robot's goals? How does it know when it has achieved them?
- Are there people around? If so, where are they, and who are they? How will the robot know?
- Is a person looking at the robot? Is someone talking to it? If so, what does the robot understand from these cues?
- What is the human trying to do? What does the person want the robot to do? How can we make sure the robot understands this?
- What should the robot do, and how should the robot react?

To address these questions, HRI researchers need to build or choose appropriate hardware and an appropriate morphology for the robot, and then develop relevant programs—the software—that can tell the robot what to do with its body.

3.2 Robot hardware

At the time of this writing, a number of robots have been produced for the consumer market. Although not all of them may have become domestic staples, these commercial robots are often suitable platforms for HRI research. Commercially available robots provide a variety of body types, including animal-like, humanoid, and more mechanical.

Aibo, an example of an animal-like robot, looks like a dog with a somewhat mechanical appearance (see Figure 3.2) and has the ability see, hear, feel touch, make sounds, wag its ears and tail, and move around on its four legs. The first Aibo models were sold in 1999, and sales were discontinued in 2006. Eleven years later, sales of new models started again.

Pepper, on the other hand, is an adolescent-size humanoid (see Figure

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 3.2 Aibo ERS-1000 robot (2018–present). (Source: Sony)

3.3). Some stores use Pepper to attract visitors and market wares and services. The company that produces Pepper also has the smaller Nao humanoid (see Figure 2.4) available for consumer purchase.

A more mechanical-looking robot, the K5 security guard robot is commercially available in the United States and is one of the few robots that are meant to be used outdoors.

Robots that were not explicitly designed to be used for HRI can nevertheless still be used or even modified for HRI studies. The most commercially successful home robot is still the iRobot Roomba vacuum-cleaning robot, millions of which have been sold around the world. Roombas not only are an interesting agent for use in studying the public's relationship with robots (Forlizzi and DiSalvo, 2006) but have also been modified and hacked for HRI research. iRobot also makes a programmable version of the Roomba, the Create, which lacks the vacuuming component and is used in research and educational applications of robots.

Telepresence robots can also be used as platforms for HRI research. Many different types exist on the market, including mobile versions such as the Beam and desktop versions like Kubi. Small mobile robots carrying a screen displaying a friendly face are being developed, soon to be ready for release in the consumer market.

Although commercially available robot hardware provides a wide variety of morphologies and sensing and programming capabilities, every robot is limited in what it can do; its appearance and capabilities constrain the interactions it can engage in. Researchers, therefore, also conceive and build their own robots, which range from simple desktop and mobile platforms with or without a manipulator to very humanlike android robots. The choice of a particular morphology for a robot to be used in HRI research often depends on the capabilities needed for the expected task (e.g., whether it needs to be able to pick up objects), the type of interaction (e.g., petlike interactions can benefit from an animal-like robot), and people's expectations and perceptions of differ-

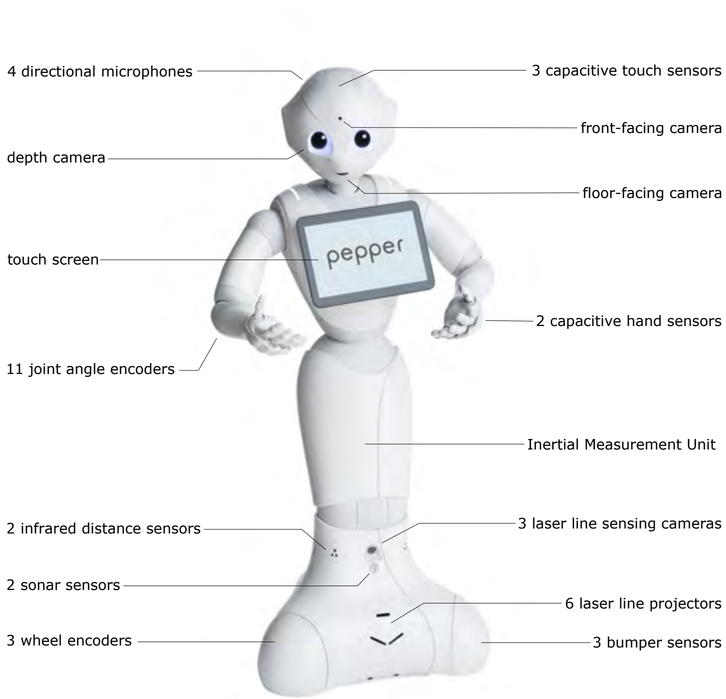

Figure 3.3 Pepper robot (2014–present) and its sensors (Source: Softbank Robotics and Philippe Dureuiltoma)

ent morphologies (e.g., humanoids may be expected to behave and be intelligent in ways similar to humans).

3.3 Sensors

Most social robots are equipped with sensors that allow them to gauge what is happening in their environment. Many commonly used sensors are related to the three most commonly used modalities in human interaction—vision, audition, and touch—but robots are not at all limited to human modes of sensing. It is often helpful, therefore, to consider what types of information the robot needs to perceive and what the most accurate and expedient ways are for it to do so, rather than focusing on reproducing human capabilities.

3.3.1 Vision

Camera

A camera consists of lenses that focus an image onto a sensor surface. The sensor surface is implemented using either a charge-coupled device (CCD) or, more often, a complementary metal-oxide-semiconductor (CMOS) technology. The basic element of a camera is a light sensor consisting mainly of silicon that converts light into electrical energy.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

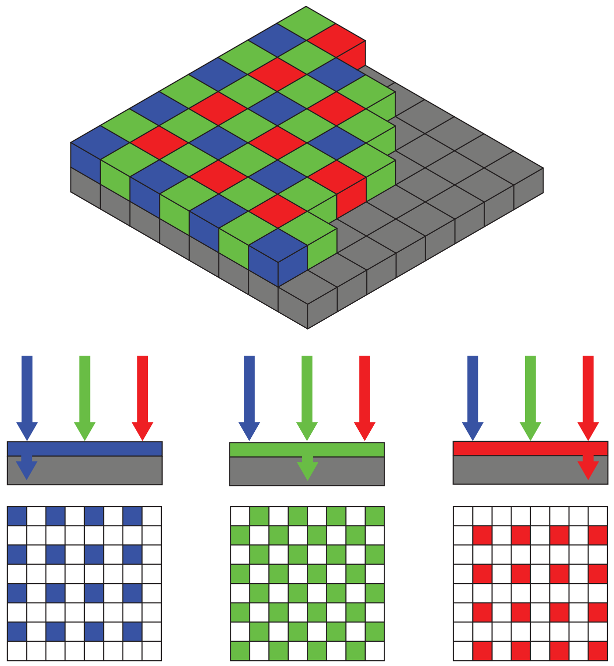

Figure 3.4 Array of CCDs in RGB camera.

A camera consists of an array of millions of these light sensors. Typically, color in a camera image is represented using three values, red (R), green (G), and blue (B). Hence, a camera is commonly referred to as an RGB camera. The sensors on the sensor surface are not sensitive to the color of the light hitting them; they are only sensitive to light intensity. To make an RGB camera, small color filters are placed on top of the sensor surface, with each filter letting through only red, green, or blue light (see Figure 3.4). Cameras are the richest and most complex sensors available to robots, and through its wide adoption in digital cameras and smartphones, the RGB camera has become miniaturized and very cheap.

In computer vision research, investigators often put cameras in the environment to facilitate accurate vision. Although this is one of the realistic approaches to yielding stable performance from computer vision, in the HRI setting, it is sometimes discouraged because people can feel uncomfortable around cameras. For example, in a project in which elderly people were being assisted in their home by a robot, the engineers would have loved to have cameras on the robot and in the home because it would have allowed the robot to accurately track and interact with people. However, the elderly participants were quite firm in their refusal of the installation and use of cameras, forcing the team to use localization beacons and laser range finders instead (Cavallo et al., 2014).

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Most cameras have a more restricted field of view than that of humans. Whereas people can see more than 180 degrees, a typical camera might only see 90 degrees, thus missing a lot of what is going on in the periphery. A robot with a single camera will have a limited field of view and might have to rely on other sensors, such as laser range finders or microphones, to give it a sense of what is going on around it.

Most importantly, the camera image needs to be processed using computer-vision algorithms in order for the robot to be able to respond to its visual environment (see Section 3.5.4).

Depth sensors

Just as human vision uses stereo vision, knowledge about objects, and self-motion to figure out the distance to objects, so can computer-vision algorithms be used to extract a three-dimensional (3D) image from two-dimensional (2D) information. Stereo cameras have been the technology of choice for a long time, but in recent years, technologies have emerged that allow us to see depth directly, without the need for computer vision. These "depth sensors" output a "depth image" or RGBD image (with D standing for *depth*), a map of distances to objects in view of the camera.

Typically, a depth sensor can measure the distance to objects a few meters away. Depending on the strength of the emitted infrared light, most depth sensors only work reliably indoors. There are several ways of making such depth sensors. One of the typical mechanisms is time of flight (TOF), in which a device transmits invisible infrared light pulses and measures the time taken between the moment when it transmitted the light and the moment when it received the light's reflection. Because the speed of light is so high, the camera would need to record the timing of the returning light with a precision that is out of reach of current electronics hardware. Instead, the camera emits pulses of infrared light and measures the phase difference between the light leaving the camera and the light returning to the camera. The Microsoft Kinect One, the second iteration of Microsoft's game controller, is based on this principle (see Figure 3.5). Despite being developed as a game controller, it was quickly adopted by robot builders and is now widely used to give robots a sense of depth. Combined with appropriate software, the Kinect sensor can also perform skeleton tracking, which is helpful for figuring out where people are, what they are doing, and even how they are feeling. Smaller devices are now available that return RGBD images based on a range of different technologies, including TOF, structured light, and stereo vision.

Figure 3.5 The Microsoft Kinect Azure DK for Windows sensor. (Source: Used with permission from Microsoft)

Laser range finders

In order to measure distances at longer ranges, researchers frequently use a laser range finder, also known as light detection and ranging (LIDAR). A typical laser range finder can measure distances to objects up to 30 meters away, and it samples the environment between 10 and 50 times per second. The accuracy of laser range finders is within a few centimeters. The basic mechanism of this type of sensor is also TOF. A laser range finder transmits a single beam of infrared laser light and measures the distance by measuring the time between the moment it transmits the laser beam and the time it receives its reflection. Typically, the transmitter and receiver are on a rotating platform, sweeping the laser beam around the environment. Thus, the device only measures distance in a single 2D plane (i.e., the plane of rotation of the rotating platform).

Robots can have range finders mounted at different heights to scan for objects on a horizontal plane. Range finders close to the ground can sense objects on the floor and people's legs, whereas range finders that are set higher up can be used to sense objects on a table or counter (see Figure 3.6)

3.3.2 Audio

Microphones are commonly used devices for auditory sensing. A microphone converts sound into electrical signals. Microphones have different sensitivity profiles; some are omnidirectional, picking up all sounds in the environment, whereas others are directional, only picking up sounds in a cone-shaped area in front of the microphone. Combining multiple microphones into an array allows us to use "beam-forming" techniques, which can separate sound signals coming from a specific direction from

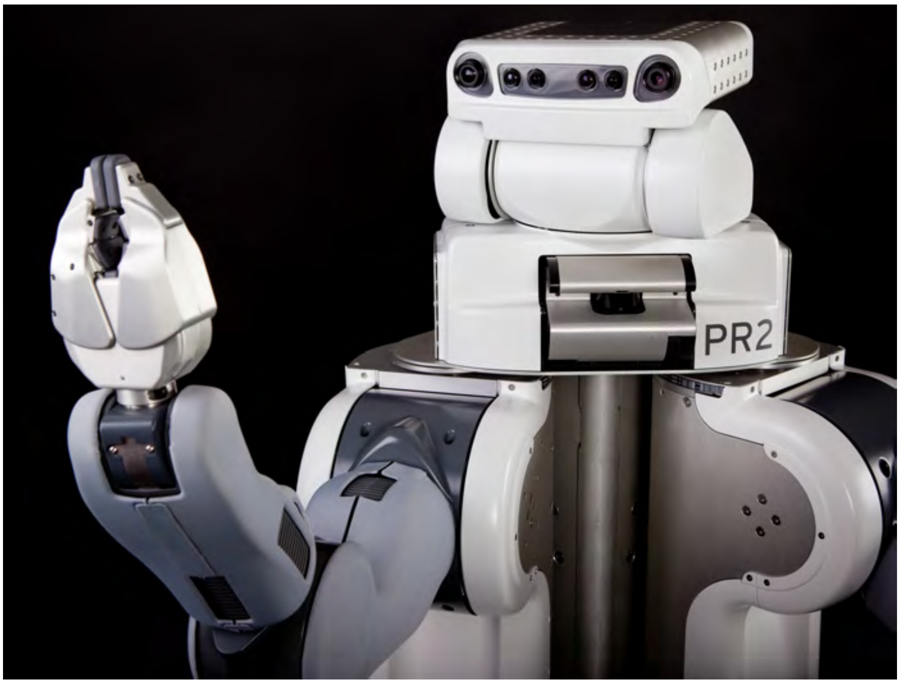

Figure 3.6 The PR2 robots (2010–2014): Can you tell where the range finder is? (Source: Willow Garage)

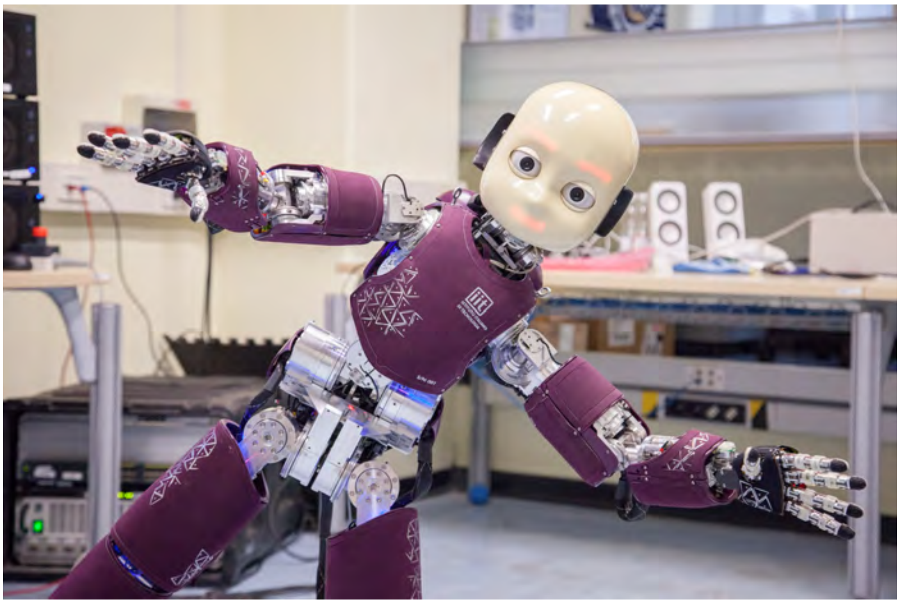

Figure 3.7 The iCub (2004–present) humanoid has capacitive tactile sensors worked into its fingers, palms, and torso. (Source: IIT Central Research Lab Genova)

ambient noises. Microphone arrays are used for sound source localization, that is, getting an accurate reading on the angle of a given sound source with respect to its position in relation to the microphone array.

3.3.3 Tactile sensors

Tactile sensors can be important in HRI, for example, when the robot is physically guided by the user. Many different implementations exist, from physical buttons or switches to capacitive sensors such as those found on touch screens (see Figure 3.7).

The most commonly used tactile sensor is a mechanical push switch. It is often used together with a bumper. When a robot collides with an object, the switch is closed, allowing the robot to detect the collision. Pressure sensors and capacity sensors, like the ones reading your finger's position on a touch screen, can also be used to detect physical contact with the environment. Pressure sensors can be implemented using a range of technologies but usually contain a material that changes its electrical properties (resistance or capacitance) when force is applied (see Figure 3.7). Pressure sensors can help robots recognize whether and how hard they are touching a person or object. They are also very useful for enabling robots to pick up and handle objects appropriately. Tactile sensors can also be used to allow the robot to know whether someone is touching it, and the robot can be programmed to respond accordingly. For example, the seal-like Paro robot has a tactile sensor net all over its body that allows it to sense the location and pressure with which a person is touching it and react by cooing for soft strokes and crying out after a harder hit.

3.3.4 Other sensors

Various other sensors exist, many of which can be relevant to HRI. Light sensors read the amount of light falling on the sensor and can be used to sense a sudden change in light, signaling that something has changed in the environment. When combined with a light source, they can be used to detect objects. A simple and very effective obstacle sensor combines an infrared light-emitting diode (LED) light with an infrared light sensor; when light bounces back from objects in front of the sensor, it can determine the distance to objects. This not only is used to detect obstacles in front of the robot but can also be used to sense when people are approaching the robot.

In recent years, the inertial measurement unit (IMU) has become a popular sensor. It combines three sensors—an accelerometer, a gyroscope, and a magnetometer—and is used the read the rotation and motion of the sensor or, more accurately, the rotational and translational acceleration. Recent advances in micro-electrical manufacturing have allowed these sensors to be miniaturized down to a few millimeters. They have become ubiquitous in mobile phones and miniature drones, and when used in a robot, they allow the robot to sense if it falls or to keep track of where it has moved over time.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Far infrared sensors (FIRs) are cameras that are sensitive to long-wavelength infrared light, which is emitted by warm bodies. They can be used to detect the presence of people, as used in burglar alarms, or when integrated into an FIR camera, they can be used to record an image of the temperature of the room. FIR sensors are still expensive and are mainly used for thermal imaging, but eventually, they may allow the robot to see people at night or in cluttered environments.

It is important to realize that, unlike our own senses, sensors do not necessarily need to be mounted on the robot. A robot might rely on a ceiling-mounted camera to interpret the social environment, or it could use a wall-mounted microphone array to localize who is speaking. The whole environment could, in a sense, be considered part of a robot system.

3.4 Actuators

An actuator converts electrical signals into physical movements. A system with one actuator typically realizes motion either on one straight line or on one rotational axis. This means that the system has one degree of freedom. By combining multiple motors, we can develop a robot that has motion with multiple degrees of freedom, allowing for navigation of a 2D plane or gesturing with human-like arms.

3.4.1 Motors

The standard actuator for robots is a direct-current (DC) servo motor (see Figure 3.8). It typically consists of a DC motor and a microcontroller, with a sensor such as a potentiometer or an encoder, which outputs the absolute or relative position of the motor's output axis. To control the speed, the controller typically sends pulse-width modulation (PWM) signals to the DC motor. PWM is an on/off pulse, literally switching the motor on for a few milliseconds and then back off. This is done several times per second (up to 100 times per second), and the duration of the on phase against the off phase (known as the duty cycle) determines the speed at which the motor rotates. The PWM signal controls the speed of the motor, and the controller sets the position of the motor. This is done through feedback control, where the controller continuously reads the position of the motor and adjusts the motor's PWM and direction to reach or maintain a desired position. For motors used in a robot's arms and head, the controller typically performs position control to rotate the motor toward a given commanded angle. For motors used in wheels on a mobile base, the controller typically performs velocity control to rotate the motor at the commanded velocity.

Figure 3.8 Connecting servo motors to each other allows robots to move around in various ways, such as in this robot arm. (Source: Trossen Robotics)

Robots can have different configurations and numbers of motors depending on the body shape and the functions they are meant to perform. Commercially available cleaning robots, such as Roomba, typically have two motors driving the wheels and one tactile sensor for moving around the room. Thus, Roomba has two degrees of freedom (DOFs). A simple nodding robot may have one motor to control its head direction, meaning that it has one DOF. A better-equipped humanoid may have three DOFs for its head, controlling pan, tilt, and yaw; two arms with four to seven DOFs; a mobile base with at least two motors; and sensors for visual, auditory, and tactile sensing. A robot arm, such as the KUKA (see Figure 3.9), must have at least six DOFs to manipulate an object. Three DOFs are necessary to locate its end effector (e.g., hand) to be in a position within a reachable range of the object, and another three DOFs are needed to grasp the object from any direction. A human arm can be approximated as an arm having seven DOFs, with an additional redundant one DOF beyond the necessary six DOFs for manipulation. To grasp objects, a robot arm must have some type of end effector attached at the end. A 1-DOF gripper can be used to grasp an object, but more complex robot hands can have as many as 16 DOFs. Android robots, designed to closely resemble humans, typically have many more DOFs (e.g., 50 DOFs) and are able to control their facial expressions and other bodily movements in relatively nuanced ways compared to simpler robots.

Motors come in many different sizes, speeds, and strengths and thus have differing power needs. It is therefore important to consider from

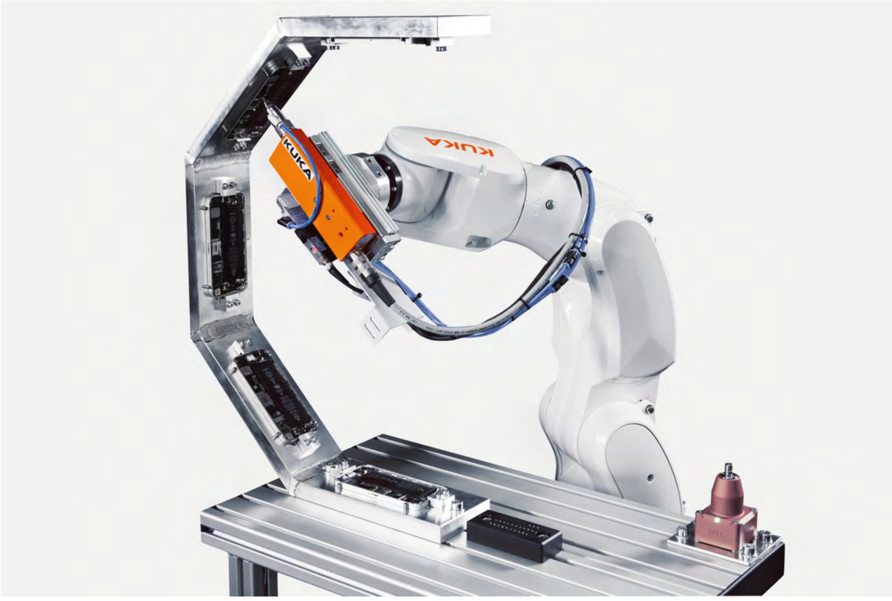

Figure 3.9 Kuka robot arm. (Source: Kuka)

early on in the design process how the motor specifications relate to the robot's design and what kinds of actions a robot will need to make, such as whether it will need to pick up a 1-kilogram bag or just needs to wave its arms; how big the robot can be while still fitting in well with its environment; how quickly it needs to respond to stimuli; and whether it needs to have a portable power bank or can be plugged into the wall.

Figure 3.10 RoboThespian (2005–present) uses pneumatic actuators to achieve the acceleration required to deliver a convincing theatrical performance. The robot can run for around a day on a scuba tank's worth of compressed gas, although it can also be attached to a compressor. (Source: Photo copyright Engineered Arts)

3.4.2 Pneumatic actuators

A pneumatic actuator uses a piston and compressed air. Air is delivered from a compressor or from a vessel containing high-pressure air, which needs to be attached to the robot in some way. Pistons typically can extend and contract, depending on which valves are opened to let in the compressed air. As opposed to electric motors, pneumatic actuators produce linear motion, which is somewhat similar to human muscle motion, and are able to produce accelerations and speeds that are difficult to achieve using electric motors. Hence, they are often preferred for humanoid robots and android robots that need to gesticulate at humanlike acceleration and velocity (see Figure 3.10). The compressors that they need to operate can be quite loud, so it is important to consider how to give the robot access to compressed air without marring the interaction experience.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

3.4.3 Speakers

To generate sounds and speech, standard loudspeakers are used. Speakers are perhaps the cheapest actuator on the robot, but in terms of HRI, they are indispensable. Where to place a speaker or speakers in the robot's body is an important factor to consider when designing a robot that will interact with people. Takayama (2008) showed that the relative height from which the voices of a user and an agent interacting with each other are projected can influence who is seen to be dominant in the interaction.

3.5 Software

All the currently available robots are controlled by software running on one or several computers. The computers receive data from sensors and periodically send commands to the actuators. Some robots do all processing on-board, but many robots will offload processing to other computers. In more recent robot software, the speech recognition, computer vision, and storage of user data often happen in the cloud, transmitted by internet-connected software services, typically operating on a pay-per-use basis. The advantage of cloud-based computing is that the robot has access to much more computing power and storage than it could ever carry on-board. Smart speakers, such as Google Home and Amazon Alexa, rely on cloud-based computing. However, a disadvantage is that when a robot relies on cloud-based computing, it needs robust communication with the cloud server, which is not necessarily guaranteed, particularly when a robot is mobile. Thus, time-critical computing and computing used to guarantee safety (e.g., emergency stop) are usually done on-board.

3.5.1 Software architecture

A robot is much more than a computer with a body. A computer operates in a clean, digital environment, whereas a robot needs to interface with the messy, buzzing confusion of the real world. Not only does it need to make sense of the world, but it also needs to do so in real-time. This environment requires a radically different approach to robot software.

Architecture models

How should software for a robot be organized? A first rule of thumb, which is applicable to nonrobot software as well, is that messy program code should be avoided. Researchers and developers must aim to modularize software. One typical approach is to follow the "sense-plan-act"

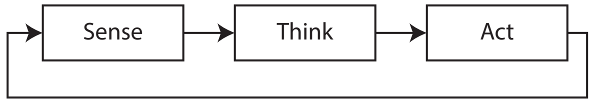

Figure 3.11 Sense-plan-act model.

model (see Figure 3.11), in which inputs from sensors are processed using software modules specific to perception, which then convert sensor streams into high-order presentations. For example, audio recordings of speech are converted into a text transcription, or camera images are analyzed to report on the location of faces. Next, there is a section that deals with "planning," which plans the robot's next actions using information gleaned from the sensing process, then outputs commands to modules for action.

For instance, a person-finding perception module reports on the location of people detected in a 2D camera image and also returns the size of the heads, indicative of how close people are to the robot. Next, the planning module computes the head orientation for the robot to face the nearest speaker and sends a command to move the head to the output modules. The output modules then calculate which angle is needed for the robot's neck motors and send these to the low-level motor controllers.

The sense-plan-act approach is also known as the *deliberative approach* because the robot deliberates its next action. Quite often, we want a robot to respond quickly to external events, without spending a lot of time pondering what to do next. In this case, we often program simple "behaviors" for the robot (Brooks, 1991). Behaviors are tightly coupled sensor–action processing loops, which immediately respond to an external event. These can be used to make an emergency stop when the robot is about to drive down the stairs, but they can serve equally well in social interaction. When a loud bang is heard, or when a face appears in view, we want the robot to respond as fast as possible. Act first; think later. Often, there are dozens of behaviors running on the robot, and mechanisms exist to mediate between which behaviors are active and which are not. One such mechanism is the subsumption architecture, which organizes behavior into hierarchies, allowing a behavior to activate or inhibit others (Brooks, 1986) (see Figure 3.12).

With this approach, even though the robot does not have an explicit "representation" of the world, it can still behave in an apparently intelligent way. For instance, if a cleaning robot uses two behaviors in parallel, one that avoids the wall and another that makes it have a slight pull to the right, the resulting, or emergent, behavior is that of wall following. Even though wall following wasn't programmed explicitly, it emerges from the interaction between two simpler behaviors. The vacuum robot Roomba has been developed with such an idea in mind.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

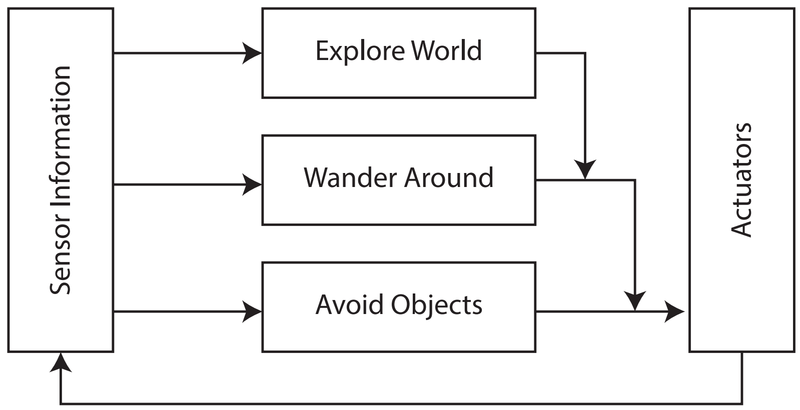

Figure 3.12 The subsumption behavior-based architecture.

In HRI studies, we typically find ourselves looking for a middle ground between the deliberative and reactive approaches. We want a reactive control layer, which responds quickly to subsecond social events, followed by a deliberative layer, which formulates a coherent response to slower elements of the interaction, such as conversation.

In light of this, it is important to develop software that can be decomposed into a number of smaller modules. Even if the complete wealth of a sense-plan-act model is not needed, it is still common practice to separate modules into perception, planning, and action.

Planning is diverse in terms of components and complexity and depends heavily on the robot and the application. A cleaning robot may need to compute the next location to clean, whereas a companion robot may need to make a decision on how it should initiate a conversation with a user. The software on a Roomba vacuum will therefore be radically different from that on a Pepper humanoid. For interactive robots, various forms of HRI knowledge will be embedded into the various software modules.

Action modules take care of the actuation and social output of the robot, such as nonverbal utterances, speech, hand gestures, and locomotion. For instance, the speech-synthesis module may receive text and convert this into spoken words together with timing information that allows the robot to accentuate its speech with appropriate gestures.

3.5.2 Software-implementation platform

Software typically runs on an operating system (e.g., Windows or Linux) and typically on some implementation platform. Robot Operating System (ROS) is a platform commonly used in the robotics and HRI communities. It deals with communications between sensors and modules and offers libraries and tools to support frequently used robot abilities, such as localization and navigation. ROS has a large community of users, who often share modules on public software repositories.

3.5.3 Machine learning

Some tasks can be learned rather than being programmed explicitly. The practice of letting a robot learn a skill is called *machine learning.* There are various machine-learning techniques, as described next.

Training data

Machine learning requires data from which the robot can learn. This training data set should contain a large number of examples of the thing to be learned, which may be data from sensors or text and generally has been manually annotated by people. For instance, there can be a data set with camera images of human faces, and for each image, the emotion of the person is labeled, such as "neutral," "smile," or "angry." Typical data sets contain hundreds of thousands or even millions of examples.

Feature extraction

To aid machine learning, sensor data are often preprocessed by converting the sensor data into a more suitable representation and by extracting salient features from the data. This process is called feature extraction. There are many algorithms to extract features from raw sensor input. For instance, edge detection highlights the pixels in an image where the intensity abruptly changes, and a segmentation algorithm identifies regions in an image where the colors are all similar, which can indicate a face, hair, or an eye (see Figure 3.13).

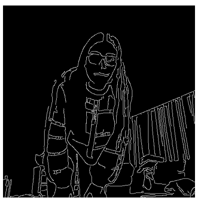

Figure 3.13 Canny edge detection of a user operating the buttons on a robot.

Features are, in essence, numbers. Often these features are placed into a feature vector, a row of numbers ready for processing. For instance, one could count up the number of pixels detected as an edge and use it as one of the variables of the feature vector. Researchers often manually analyze their data sets and identify salient features. For instance, with careful observation, one might find that a child fidgets more than an adult does; once such a feature is found, one can add variation of motion to the feature vector.

Classification based on training

There are a number of machine-learning approaches. One often-used approach is *classification.* In classification, an algorithm decides, based on training data, what class an unknown data point belongs to. For example, given a camera image of a person, the classifier decides what emotion the person's face shows.

Suppose we can compute a one-dimensional (1D) feature vector representing people's height and have a data set with two classes, "child"

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

or "adult" (i.e., each data point in the training data will have a label saying whether the data point is a "child" or "adult"). The classifier learns a threshold value from the training data set (e.g., 150 cm) to distinguish the two classes.

In this case, the feature vector contains only a single feature, the height of the user. We call this a 1D feature vector. Machine-learning algorithms typically work with thousands of features and try to recognize up to thousands of classes. Classification errors often happen. For instance, a tall child or short adult would probably be classified incorrectly.

Machine-learning algorithms perform better when having access to more data. Ideally, we want machine-learning algorithms to "generalize," meaning they correctly handle data that they have never been exposed to. However, sometimes machine learning produces an algorithm that "overfits." When this happens, the algorithm does really well on the data it has been trained on, but it performs poorly when confronted with new problems.

Deep learning

Deep learning, also known as deep neural networks (DNNs), is a machine-learning technique made possible through the increased availability of computational power. Deep learning relies on artificial neural networks with a large number of layers of interconnected artificial neurons—hence the name "deep." It takes a large amount of computational power to train DNNs, but recent progress in using parallel computing and graphical-processing units (GPUs) has allowed us to train these networks within a matter of days.

DNNs do not require careful feature extraction by hand. Instead, DNNs discover the relevant features from the data by themselves. A drawback is that DNNs require huge amounts of data, typically millions of data points. For instance, Google collected an enormous data set, containing more than 230 billion data points, to train its speech-recognition algorithm.

The need for large data sets is a significant challenge for HRI because it is difficult to collect large amounts of data in which humans and robots are interacting. The complexity of deep learning also makes it difficult to know exactly what the network bases its decisions on (e.g., we may not know what features it has identified or how it decided to use these features to come to a classification), which can be particularly problematic for HRI outside of the laboratory when we need to trust that the system will be robust, safe, and predictable. If the robot does something wrong, we need to be able to figure out how to debug and correct the system, as in the case of an autonomous Uber vehicle that had trouble classifying a person crossing the road and ran over the person as a result (Marshall and Davies, 2018).

3.5.4 Computer vision

Computer vision is an important area for HRI. In essence, computer vision interprets a 2D array of numbers when working with single images, or a series of 2D images recorded over a period of time when working with video data. Computer vision can be rather straightforward and still very effective in the context of HRI. Motion detection, for example, can be achieved by subtracting two camera images taken just a fraction of a second apart. Any pixels that captured motion will have a nonzero value, which in turn can be used to calculate the region with the most motion. When used on a robot, a motion detector lets the robot orient itself toward the areas with the most motion, providing the illusion that the robot is aware of things moving, which, in the context of HRI, often involves people gesturing or talking.

Another computer-vision technique relevant to HRI is processing faces. The ability to detect faces in an image has advanced and can be used, for example, to let the robot look people in the eye. Face recognition (i.e., identifying a specific person in an image) is still a challenge, however. Impressive progress has been made in recent years, mainly fueled by the evolution of deep learning, and it is now possible to reliably recognize and distinguish between hundreds of people when they are facing the camera. But face recognition typically fails when the user is seen from the side.

Skeleton tracking is another technique relevant to HRI. In skeleton tracking, the software attempts to track where the user's body and limbs are. This technique was first used in gaming on the Microsoft Xbox console, with software specific to the Kinect RGBD sensor, but is now a staple in many HRI applications. Several software solutions exist, but recently, deep learning has enabled the reading of skeletons of dozens of users in complex scenarios from a single simple camera image, without the need for an RGBD sensor. The software for this, called OpenPose, is now freely available and often used in HRI studies (Cao et al., 2017).

There are many commercial and free software solutions that offer a range of out-of-the-box computer-vision functionality. OpenCV is perhaps the best-known offering; it is a free software library, developed over 20 years, that can be used for facial recognition, gesture recognition, motion understanding, object identification, depth perception, and motion tracking, among others.

Because computer vision often requires a considerable amount of computational power, which is not realistic on small or cheaper robots, sometimes the computer-vision process is addressed on the cloud. In this case, the video stream of the robot is sent over an internet connection to servers on the cloud. There are commercial-based cloud solutions for face recognition, person identification, and image classification being sold on a per-use basis.

3.6 Limitations of robotics for HRI

There are several limitations of robotics, some of which are specific to HRI and some of which apply to robotics in general. One general challenge is that a robot is a complex system that needs to translate between the analogue world and the digital internal computation of the robot. The real world is analogue, noisy, and often very changeable, and the robot first needs a suitable digital representation of the world, which the software then uses to make decisions. Once a decision is made, this is translated back into analogue actuation, such as speaking a sentence or moving a leg.

Another major challenge applicable to all of robotics is that of learning. Currently, machine learning needs to iterate through millions of examples to slowly nudge itself toward performing a task with a reasonable level of skill. Despite speedups due to advances in DNNs and GPUs, at the time of writing, computers need days or often weeks to learn, and this is only when all the learning can happen internally, for example, in simulation or using prerecorded data. Learning from real-time data that a robot samples from the world is still virtually impossible. Related to this is the challenge of "transfer," or the performance of one skill transferring to another. For example, people can learn to play one game of cards and will then be able to transfer that knowledge to quickly pick up another game of cards with different rules. Machine learning typically struggles with this task and needs to start the learning of a new challenge from scratch.

The seamless integration of the various systems on a robot is also a major challenge. Speech recognition, natural-language understanding, social-signal processing, action selection, navigation, and many other systems all need to work together in order to create convincing social behavior in a robot. On simple robots, this is manageable, but on more complex robots, the integration and synchronization of these various skills are still beyond our grasp. Face detection, emotion classification, and sound-source localization might each work well in isolation, but bringing the three together to make the robot respond in a humanlike manner to people approaching the robot is still a challenge. Greeting people who smile at the robot, looking up when the door slams, or ignoring people who show no interest in the robot sound easy, but it is difficult to build such behavior that consistently works well. The challenge becomes formidable once further skills are added. Conversational robots, which aim to interact with people using natural language in addition to using their full suite of sensors to react in an appropriate manner, are only now being attempted in research labs across the world. It is unlikely that a robot will be built in the next decade that can handle a conversation as well as people can.

Robots and artificial intelligence (AI) systems in general struggle with semantics: they often do not truly understand what happens around them. A robot might seem to respond well to a person approaching it and asking for directions, but this does not mean that the robot understands what is happening—that the person is new to the space, or where the directions it gives actually lead to. Often, the robot has been programmed to face people when they come near and to respond to key words it hears. Real understanding is, at the moment, still exclusive to humans. Although there are research projects on imbuing AI systems with a sense of understanding (Lenat, 1995; Navigli and Ponzetto, 2012), there are not yet robots that can use their multimodal interaction with the world to understand the social and physical environment.

The reasons why AI has not yet achieved a humanlike general intelligence level are manifold, but conceptual problems were identified right from the outset. Searle (1980) pointed out that digital computers alone can never truly understand reality because they only manipulate syntactical symbols that do not contain semantics. In his *Chinese Room thought experiment*, a slip of paper with Chinese symbols is slid under the door of a room. A man inside the room reads the symbols and comes up with a response by applying a set of rules he finds in a book full of instructions containing more Chinese characters. He then writes the response in the form of other Chinese characters and slides it back under the door. The audience behind the door might be under the impression that the man in the room understands Chinese, whereas in reality, he just looks up rules and has no understanding of what those symbols really mean. In the same manner, a computer also only manipulates symbols to come up with a response to input. If the computer's response is of humanlike quality, does that mean the computer is intelligent?

According to Searle's line of argument, IBM's chess-playing computer Deep Blue does not actually understand chess, and DeepMind's AlphaGo does not understand the game of Go. Both programs may have beaten human masters of the game, but they did so only by manipulating symbols that were meaningless to them. The creator of Deep Blue, Drew McDermott, replied to this criticism: "Saying Deep Blue doesn't really think about chess is like saying an aeroplane doesn't really fly because it doesn't flap its wings" (1997). That is, he debated that as far as it functions as it is supposed to, a new machine or AI does not need to replicate all the details of humans, animals, or birds. This debate reflects different philosophical viewpoints about what it means to think and understand and is still under way today. Similarly, the possibility of developing general AI remains an open question. All the same, progress has been made. In the past, a chess- or Go-playing machine would have been regarded as intelligent. But now it is regarded as the feat of a calculating machine—our criteria for what constitutes an intelligent machine have shifted along with the capabilities of machines.

In any case, no sufficiently intelligent machine has yet been built that would provide a foundation for many of the advanced application scenarios that have been imagined for robots. Researchers often fake the intelligence of the robot by applying the Wizard-of-Oz method (see p. 156). The requirements of HRI often imply unrealistic assumptions about what can be achieved with current technology, and novice research and the public should be aware of the limitations of robotics and AI.

3.7 Conclusion

Robots are made from multiple software modules connected with sensors and actuators. Software design requires HRI knowledge, and conversely, HRI researchers need to have a basic understanding of software in order to provide useful knowledge for future HRI developers. For a robot to be successful, the different components need to be chosen and integrated with an eye toward the specific HRI application and its needs. Despite limitations, however, robots can be designed to interact successfully with humans in various types of short-term, and sometimes longer, interactions.

Questions for you to think about:

- Chapters 2 and 3 introduce various robot types that are available on the market. What sensors do these robots have? What actuators do they have? What hardware components do you think are crucial?
- Imagine a scenario where you want to use a smart social robot. Which sensors and actuators should it have? What skills should the robot have, and is software available to deliver these skills?
- What kind of data set would be needed to train a machine-learning algorithm for a new interaction capability of a robot, such as distinguishing your face from others?

Future reading:

- For basic AI:
  Stuart Russell and Peter Norvig. *Artificial intelligence: A modern approach.* Pearson, Essex, UK, 3rd edition, 2009. ISBN 978-0136042594. URL `http://www.worldcat.org/oclc/496976145`
- For basic robotics:
  Maja J. Matarić. *The robotics primer.* MIT Press, Cambridge, MA, 2007. ISBN 9780262633543. URL `http://www.worldcat.org/oclc/604083625`
- For diverse topics in robotics:
  Bruno Siciliano and Oussama Khatib. *Springer handbook of robotics.* Springer, Berlin, 2016. ISBN 9783319325507. URL `http://www.worldcat.org/oclc/945745190`

4

Design

What is covered in this chapter:

- How a well-designed robot can lift interactions to the next level (physical design);
- How people do not treat robots as an assembly of plastic, electronics, and code but, rather, as humanlike entities (anthropomorphism);
- How HRI research draws on psychological theories of anthropomorphism to design and study people's interactions with robots;
- Design methods and prototyping tools used in human–robot interaction.

How does a pile of wires, motors, sensors, and microcontrollers turn into a robot that people will want to interact with? Although it sounds like magic, the trick of turning metal and plastic into a social interaction partner is in the iterative and interdisciplinary process of robot design.

Robot design is a fast-growing field of research and practice in human–robot interaction (HRI), and the need to develop robots that are able to interact with humans challenges existing ways of designing robots. To date, most robots are developed by engineers, and their ability to interact with humans is then tested later on by social scientists. This process of design starts from the inside and builds up to the outside—solving technical issues first and designing the robot's appearance and behavior to fit. For example, a mobile platform such as a TurtleBot (see Figure 4.1) might be used as a starting point, with the desired sensors and actuators added to the body later on. If time allows, a casing could be designed to cover up all the technology. The robot's appearance and the specific social interaction capabilities then have to be built on top of this technical infrastructure. This common approach to robot building is also known as the "Frankenstein approach": we take whatever technology is available and put it together to get certain robotic functions. A lack of consideration of the social context of use within the design process can lead to surprising effects in robot interaction, however.

Alternative, more holistic approaches to robot design start by con-

Figure 4.1 A TurtleBot2 (2012–present) platform. (Source: Yujin Robot)

sidering who will use it, where, and how. Based on the characteristics of the users and context of use, one can then decide on specific robot design features, such as appearance, interaction modalities, and level of autonomy. This might be termed a more "outside-in" mode of developing robots, in which the design process starts from the interaction that we expect the robot to be engaged in, which will determine its outside shape and behaviors. Once the design has been settled upon, we work all the technology into it.

Designers are trained to approach the design of artifacts in this way (see Figure 4.2 for an example) and are able to make valuable contributions (Schonenberg and Bartneck, 2010). The unique contributions include the aesthetics of the robots, but designers also have the skill to create thought-provoking robots that challenge our understandings of the roles of humans and robots.

This form of robot design often requires incorporating expertise from several disciplines—for example, designers might work on developing specific concepts for the design, social scientists may perform exploratory studies to learn about the potential users and context of use, and engineers and computer scientists need to communicate with the designers to identify how specific design ideas can be realistically instantiated in working technology (Šabanović et al., 2014). HRI design can take advantage of existing robots, designing specific behaviors or use tasks for them that fit particular applications, or it can involve the development of new robot prototypes to support the desired interactions. In either case, HRI design both takes advantage of existing design methods and develops new concepts and methods specifically suited to the development of embodied interactive artifacts (i.e., robots).

Figure 4.2 Mythical robots designed from the outside to the inside. First, the shape of the robots was sculptured before fitting the technology into it.

4.1 Design in HRI

4.1.1 Robot morphology and form

A common starting point for designing HRI is to think of what the robot is going to be doing. There is a debate about whether form follows function, in which the shape of an object is largely determined by its intended function or purpose, or if the reverse holds true. However, in HRI, form and function are inherently interconnected and thus cannot be considered separately.

Contemporary HRI designers have several different forms of robots to choose from. Androids and humanoids most closely resemble humans in appearance, but they have a lot to live up to in terms of capabilities. Zoomorphic robots are shaped like animals with which we are familiar (e.g., cats or dogs) or like animals that are familiar but that we do not typically interact with (e.g., dinosaurs or seals). HRI designers, eager to make robot appearances commensurate with their limited capabilities, also often design minimalist robots, which explore the minimal requirements necessary for inspiring social HRI, such as Muu (see Figure 4.4, left), or Keepon (see Figure 4.4, middle). The arguably most minimalistic robot is the busker robot, which consisted of a pair of animated sandals on top of a box with a signpost in front of it proclaiming "Naked Invisible Guy" (Partridge and Bartneck, 2013) (see Figure 4.4, right).

Figure 4.3 Robovie-MR2 (2010) is a humanoid robot controlled through a cell phone.

Recently, along with these organism-based robots, the HRI field has started considering "robjects," interactive robotic artifacts whose design is based on objects rather than living creatures (e.g., Robot Ot-

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

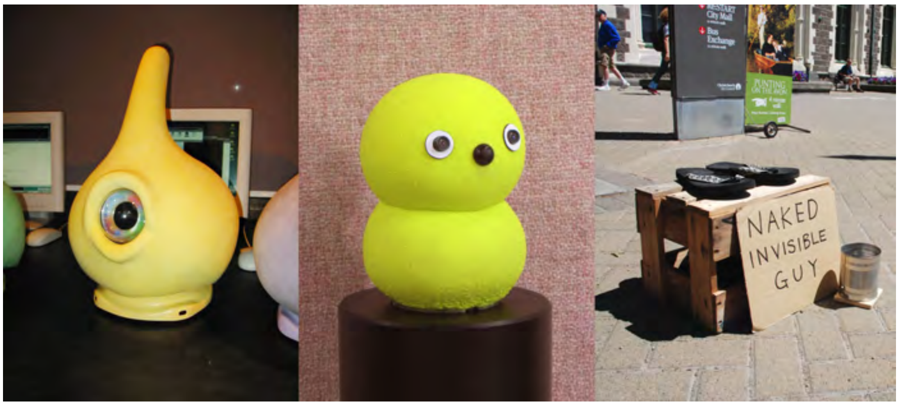

Figure 4.4 Zoomorphic and minimalistic robots: Muu (2001–2006), Keepon (2003–present) and Naked Invisible Guy.

Figure 4.5 Sociable Trash Box robots are an example of robjects—robotic objects with interaction capabilities. (Source: Michi Okada)

toman, social trashcans (see Figure 4.5), robotic piggy banks (Fink et al., 2014). Because the design space of robots is relatively large and considers questions regarding form, function, level of autonomy, interaction modalities, and how all those fit with particular users and contexts, an important aspect of design is figuring out how to make appropriate decisions about these various design aspects.

4.1.2 Affordances

Another important concept in HRI design is the notion of affordances. This notion was initially developed as a concept in ecological psychology (Gibson, 2014), where it referred to the inherent relationship between an organism and its environment. For example, a person might want to throw a rock when he or she sees it, but a mouse would want to hide behind it. This concept was amended by Don Norman (Norman, 2008) to describe the perceivable relationships between an organism and its environment that enable certain actions (e.g., a chair is something to sit on, but so is a stair).

A designer needs to design a product while making its affordances explicit. Furthermore, he or she needs to incorporate user expectations and cultural perceptions. For Norman, these "design affordances" are also an important way to develop common ground between robots and humans so that people can understand robot capabilities and limitations appropriately and adapt their interactions accordingly. A robot's appearance is an important affordance because people tend to assume that the robot's capabilities will be commensurate with its appearance. If a robot looks like a human, it is expected to act like a human; if it has eyes, it should see; if it has arms, it should be able to pick up things and might be able to shake hands. Another affordance can be the robot's interaction modalities. If a robot speaks, for example, saying "Hello," people will also expect it to be able to understand natural language and carry on a conversation. If it expresses emotions through facial expressions, people might expect it to be able to read their emotions. Other robotic affordances can be based on technical capabilities; for example, if it has a touch screen on its body, people might expect to interact with the robot through the touch screen. Because robots are novel interaction partners, the affordances used by designers are particularly important for signaling appropriate ways of engaging with them.

4.1.3 Design patterns

Because the focus of HRI is the relationship between humans and robots, the task of HRI design is not only to create a robotic platform but also to design and enable certain interactions between humans and robots in various social contexts. This suggests that the main units of design that need to be considered are not only the characteristics of individual robots (e.g., appearance, sensing abilities, or actuation) but also what Peter Kahn calls "design patterns" in HRI, inspired by Christopher Alexander's idea of design patterns in architecture (Kahn et al., 2008) Such patterns describe "a problem which occurs over and over again in our environment, and then describes the core of the solution to that problem, in such a way that you can use this solution a million times over, without ever doing it the same way twice" (Alexander, 1977, p. x).

Within HRI, Kahn et al. (2008) suggest that patterns should be abstract enough that you can have several different instantiations, that they can be combined, that less complex patterns can be integrated into more complex patterns, and that they serve to describe interactions with the social and physical world. For example, the didactic communication pattern (where the robot assumes the role of a teacher) could be combined with a motion pattern (where the robot initiates a movement and aligns it with the human counterpart of the interaction) to create a robotic tour guide. Kahn et al. suggest that HRI design patterns can be developed based on observation of human interactions, prior empirical knowledge about humans and robots, and designers' experiences with HRI, through an iterative design process. Some patterns they developed and have used in their designs are things like the "initial introduction" of the robot, or "in motion together," where the robot moves along with the person. Although Kahn et al.'s design patterns are not meant to be exhaustive, they emphasize the idea that the design should focus on the relationship between humans and robots.

4.1.4 Design principles in HRI

When combining the two ideas of design affordances and patterns in the process of HRI design, the usual design types that robots may be divided into, such as androids and humanoids, zoomorphic robots, minimally designed robots, or robjects, are no longer the main design focus or question. Instead, designers consider how different robot forms and capabilities fit into or express particular HRI design patterns and how they can be designed as affordances that appropriately signal the robot's interaction capabilities and purpose. With this in mind, HRI researchers have suggested some of the following principles to consider when developing the appropriate robot forms, patterns, and affordances in HRI design.

Matching the form and function of the design: If your robot is humanoid, people will expect it to do humanlike things—talk, think, and act like a human. If this is not necessary for its purpose, such as cleaning, it might be better to stick to less anthropomorphic designs. Similarly, if it has eyes, people will expect it to see; if it talks, they will expect it to be able to listen. People can also be prompted to associate specific social norms and cultural stereotypes with robots through design; for example, researchers have shown that people might expect a female robot to be more knowledgeable about dating or that a robot made in China would know more about tourist destinations in that country (Powers et al., 2005; Lee et al., 2005)

Underpromise and overdeliver: When people's expectations are raised by a robot's appearance or by introducing the robot as intelligent or companion-like, and those expectations are not met by its functionality, people are obviously disappointed and will negatively evaluate the robot. Sometimes these negative evaluations can be so serious that they affect the interaction. To avoid such problems, it is better to decrease people's expectations about robots (Paepcke and Takayama, 2010), which might have been increased by how robots are portrayed in society, as described in the "Robots in Society" chapter (see Chapter 11). This might even include not calling your design a robot because the word itself often connotes quite advanced capabilities to members of the public.

Interaction expands function: When confronted with a robot, people will, in effect, fill in the blanks left open by the design depending on

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

their values, beliefs, needs, and so on. It can thus be useful, particularly for robots with limited capabilities, to design them in a somewhat open-ended way. This allows people to interpret the design in different ways. Such an open-ended design approach has worked particularly well with, for instance, the seal-like robot Paro (see Figure 2.6). This baby seal robot invokes associations with pets that people have had, but it also does not get compared to animals they know, such as cats and dogs, which would inevitably lead to disappointment. As a consequence, Paro becomes a natural part of the interactions with humans and passes as a petlike character even though its capabilities are significantly below those of a typical domestic animal or that of an actual seal baby (Šabanović and Chang, 2016).

Do not mix metaphors: Design should be approached holistically—the robot's capabilities, behaviors, affordances for interaction, and so forth should all be coordinated. If you design a humanlike robot, people may find it disturbing if it has skin covering only some parts of its body. Similarly, if the robot is an animal, it may be strange for it to talk like an adult human or try to teach you mathematics. This is related to the Uncanny Valley (see p. 52) because inappropriately matched abilities, behaviors, and appearance often lead to people having a negative impression of the robot.

> Take a look at the two pictures in Figure 4.6. How do they make you feel? Although both of these android representations of the science-fiction writer Philip K. Dick are perhaps a bit strange and uncanny, the one that seems unfinished and shows the robot's insides also mixes design metaphors—the robot is both humanlike and machinelike, making it even more disturbing.

Like Kahn's design patterns, these design principles are not exhaustive but are meant to inspire thinking about how to approach designing HRI in a way that acknowledges and incorporates the interdependence between human and robot capabilities, the need for interaction partners to be intelligible to and support each other, and the effects of the context of interaction on its success.

4.2 Anthropomorphization in HRI Design

Have you ever found yourself yelling at your computer because it suddenly crashes while you are working on an essay that is due in just a few hours? You urge the computer to please bring it back again after restarting, gently touching the mouse after realizing that, indeed, the file reopens and you can continue. You sigh in relief because "Genius"—that's what you call your computer when no one is around to hear you—has not let you down. In fact, what you have pictured now is an

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 4.6 Philip K. Dick Robot (2005; rebuilt in 2010).

ordinary scenario of a person humanizing an object, anthropomorphizing it. What a tongue twister. But what's it about, in fact?

Anthropomorphization is the attribution of human traits, emotions, or intentions to nonhuman entities. It derives from ánthrōpos (meaning "human") and morphē (meaning "form") and refers to the perception of human form in nonhuman objects. We all experience anthropomorphism in our daily lives. "My computer hates me!"; "Chuck (the car) is not feeling well lately"; "That grater looks like it has eyes"—you've either heard or uttered the sentiment before. The latter is a special example of anthropomorphization called *pareidolia*, the effect of seeing humanlike features in random patterns or mundane objects. The *Viking 1* spacecraft took a photo of the Cydonia area on Mars on July 25, 1976 (see Figure 4.7). Many people saw a face on Mars's surface, which sparked many speculations about the existence of life on Mars. The National Aeronautics and Space Administration (NASA) sent its Mars Global Surveyor to the exact same location in 2001 to take higher-resolution photos under different lighting conditions, which revealed that the structure photographed in 1976 is certainly not a human face.

Anthropomorphization is a natural outgrowth of the significance of social interaction and social cognition in human life. It is also a main theme of design and research in HRI. We will discuss anthropomorphism here in some detail as a case study of a specific design theme in HRI that incorporates technical development, psychological study, and design to enable social HRI. A robot's level of anthropomorphism is one of the main design decisions that robot designers need to take

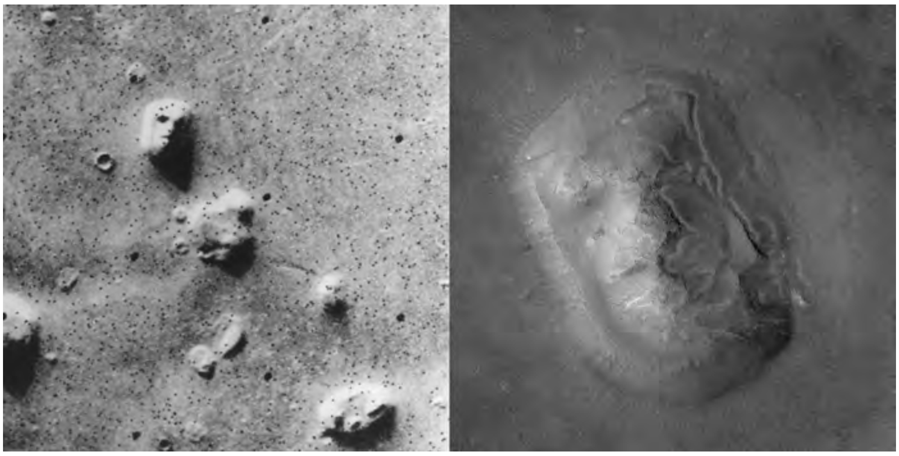

Figure 4.7 The face on Mars is an example of pareidolia. On the left is the photo from 1976, and on the right is the same structure photographed in 2001. (Source: NASA/JPL, NASA/JPL/MSSS)

into account because it influences not only the robot's appearance but also the functionality it needs to offer.

4.2.1 Anthropomorphization and robots

People's innate predisposition to anthropomorphize the things around them has become a common design affordance for HRI. In anthropomorphic design, robots are constructed to have certain humanlike characteristics (see Figure 4.9), such as appearance, behavior, or certain social cues, which inspire people to see them as social agents. At one extreme, android robots are designed to be as humanlike as possible; some have been fashioned as exact replicas of living humans, like a moving Madame Tussaud's wax figure (like Geminoid in Figure 4.8), or as representations of aggregated human features (e.g., Kokoro). Humanoid robots use a more abstract notion of human-likeness in their anthropomorphic designs. ASIMO, for example, has a human body shape (two arms and legs, a torso, and a head) and proportions, but it does not have eyes. Rather, its head resembles an astronaut's helmet. Nao similarly has a humanlike body, as well as two light-emitting diode (LED) eyes that can change in color to connote different expressions, but no mouth. Some other humanoids, such as Robovie, Wakamaru, and Pepper, are not bipedal but have arms and have heads with two eyes.

Nonhumanoid robots, however, can also have anthropomorphic features. The minimalist Keepon has two eyes and a symmetrical body, as well as displays of behavioral cues for attention and affect that inspire anthropomorphization. Google's autonomous car prototype has an almost cartoon-like appearance, with wide-set headlights and a button nose that suggest an anthropomorphic appearance. Finally, giving robots an animal-like appearance and/or behavior, for example, Pleo

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 4.8 The Geminoid HI 4 robot (2013), a replica of Hiroshi Ishiguro. (Source: Hiroshi Ishiguro)

Figure 4.9 People readily anthropomorphize all kinds of robots, with appearances ranging from minimalist to indistinguishable from the human form. From left to right: Keepon, Wakamaru (2005–2008), Nao (2008–present), ASIMO (2000–2018), and Kokoro's Actroid (2003–present) android.

(see Figure 10.5) and Roomba with a tail by Singh and Young (2012), can also be seen as a form of anthropomorphic design because it takes its inspiration from people's common anthropomorphization and social perception of animals.

Anthropomorphism has been key to animation designers for some time, only relatively recently sparking the interest of social psychologists. Disney's Illusion of Life (Thomas et al., 1995) has inspired several social robotic projects, such as Wistort et al.'s Tofu, which displays the animation principles of "squash" and "stretch" (Wistort and Breazeal, 2009), and Takayama et al.'s work with the PR-2 using animation to give the robot apparent goals, intentions, and appropriate reactions to events (Takayama et al., 2011). Animation principles such as anticipation and exaggerated interaction have also been applied to robot design, for example, in Guy Hoffman's Marimba player (Hoffman and Weinberg, 2010) and music companion robots (Hoffman and Vanunu, 2013). These anthropomorphic designs take advantage not only of appearance and form but also of behavior in relation to the environment and other actors to evoke ascriptions of human-likeness.

Anthropomorphism in robot design includes factors related to form and appearance as well as factors relating to behavior, but all rely on people's ability to imaginatively imbue robots with traits and abilities that go a bit beyond what they might in fact have.

4.2.2 Theorizing anthropomorphism

A psychological perspective

In the classic engineering-oriented literature on anthropomorphism, researchers have mainly focused on assessing the perceived appearance of the robot. Going beyond this notion, recent theorizing in psychology has provided a complementary perspective on the nature of anthropomorphism. The theoretical framework proposed by Nicholas Epley and colleagues (Epley et al., 2007) has been influential both in psychology and in robotics and serves to broaden our understanding of the notion of anthropomorphism, its causes, and its consequences. Epley and colleagues have suggested three core factors that determine anthropomorphic inferences about nonhuman entities: effectance motivation, sociality motivation, and elicited agent knowledge. Let us introduce these concepts briefly.

Firstly, effectance motivation concerns our desire to explain and understand the behavior of others as social actors. This might be activated when people are unsure about how to deal with an unfamiliar interaction partner. Most people are still relatively unfamiliar with robots as social interaction partners, so it is easy to imagine that being asked to socially engage with a robot could elicit effectance motivation in them, thus increasing their tendency to anthropomorphize robots. People might therefore attribute humanlike characteristics to robots to psychologically regain control over the novel situation they find themselves in. In this case, anthropomorphization can reduce the stress and anxiety associated with human–robot interaction.

Second, anthropomorphization of robots could also be caused by sociality motivation, particularly by people who lack social connections. In this case, people may turn to nonhuman entities as social interaction partners to address their feelings of situational or chronic loneliness. Supporting this idea, previous research has shown that people who have been made to feel lonely in an experimental situation, or who are chronically lonely, anthropomorphize robots to a greater extent than people who are sufficiently socially connected (Eyssel and Reich, 2013).

Lastly, elicited agent knowledge refers to the way in which people use their commonsense understanding of social interactions and actors to understand robots. For example, Powers et al. (2005) showed that people who considered women to be more knowledgeable about dating norms behaved toward male and female robots as if they also had dif-

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

fering competencies regarding dating; for instance, they used more time and words to explain dating norms to a male robot. This factor in particular can be used to guide the design and technical implementation of social robots for various tasks.

These three determinants shed light on the psychological mechanisms underlying why we tend to humanize nonhuman entities. This includes the attribution of emotions, intentions, typically human traits, or other essentially human characteristics to any type of nonhuman entity, real or imagined (Epley et al., 2007). The basic assumption is that people use self-related or anthropocentric knowledge structures to make sense of the nonhuman things—or in our case, robots—around them. Human resemblance in appearance and behavior triggers anthropomorphic judgments, and people may thus attribute traits and emotions to a technical system despite the fact that the system, indeed, is merely a piece of technology. This, in turn, not only affects the social perception of robots but also the actual behavior displayed toward them during an interaction. Research by Reeves and Nass (Soash, 1999) has already demonstrated in the context of human–computer interaction (HCI) that anthropomorphization of computers and other media occurs automatically. Whether this holds true for robots, though, is currently still under empirical debate (Zlotowski et al., 2015). The three-factor model of anthropomorphism, however, has been thoroughly empirically tested and validated with social robots (Eyssel, 2017).

The Uncanny Valley

Mori (1970) made a prediction about the relationship between the anthropomorphism of robots and their likeability (see Figure 4.10). The idea is that the more humanlike robots become, the more likable they will be, until a point where they are almost indistinguishable from humans, at which point their likability decreases dramatically. This effect is then amplified by the ability of the robot to move.

Mori et al. (2012) translated the original paper to English in collaboration with Mori himself. It is important to note that Mori only proposed this idea and never did any empirical work to test his ideas. Moreover, Mori used the term 親和感 (shinwa-kan) to describe one of his key concepts. The translation of this concept to English remains challenging—it has been translated as likeability, familiarity, and affinity. Other researchers have approached the problem by asking participants about the eeriness of the robot instead. Unfortunately, Mori's theory has been used and abused to explain a huge number of phenomena without proper justification or empirical backup. It is often used to explain why certain robots are being perceived unfavorably, without studying the exact relationship between the features of the robot at hand and its likability. Anthropomorphism is a multidimensional concept, and reducing it to just one dimension does not model reality

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

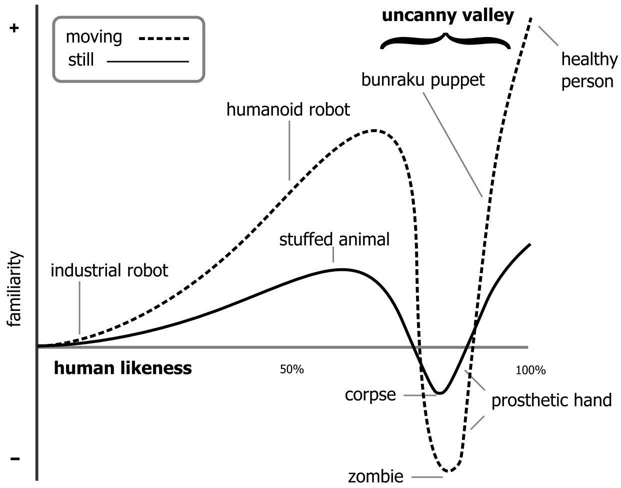

Figure 4.10 Mori's Uncanny Valley theory.

adequately. Moreover, the more humanlike robots become, the greater is the risk of getting a certain aspect of their appearance or behavior wrong and thereby decreasing the level of likability (Moore, 2012). A simple possible explanation of why humanlike robots are liked less than, for example, toy robots, is that the difficulty of designing a robot to perform to user expectations increases with its complexity.

4.2.3 Designing anthropomorphism

Robot designers treat anthropomorphism as a characteristic of the robot itself, whereas social scientists see anthropomorphism as something that a human attributes to the robot. Considering both of these together suggests that anthropomorphism is about the relationship between robot design and people's perceptions of robots.

Design approaches

To trigger anthropomorphic inferences, robot designers can take into account the dimensions of robot appearance and behavior, among many other aspects. By exploiting these aspects, they can achieve an immediate perception of the robot as more or less humanlike.

Robot appearance

Graphical illustration shows us that often only a few lines on a sheet of paper are needed to evoke the human form. In the same manner, anthropomorphism in robots can be very simple: just having two dots suggesting eyes and a simple nose or mouth is sufficient to suggest the robot is humanlike. This can be further enhanced by adding more human features, such as arms or legs, but these do not necessarily do very much to further increase the anthropomorphization. Although there

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

are many reasons why robots look increasingly humanlike, anthropomorphization can be achieved with only a minimal set of humanlike features. Whereas androids mimic human appearance in most ways, simple robots such as Keepon and R2D2 are already very effective at triggering people to anthropomorphize. Thus, a large body of research has documented how minimal design cues might be sufficient to elicit a humanlike perception.

Robot behavior

A second approach to increasing anthropomorphization is to design the behavior of an artifact such that people perceive humanlike characteristics in its behavior. Heider and Simmel (1944) showed how simple geometric shapes—triangles and circles—moving against a white background evoked people to describe their interactions in terms involving social relationships (e.g., these two are friends; this one is the attacker) and humanlike feelings and motivations (e.g., anger, fear, jealousy). Animators understand how motion, rather than form, can be extremely powerful for expressing emotions and intents. A surprisingly wide range of humanlike expressive behavior can be communicated through movement alone, without the need for humanlike form.

The Dot and the Line: A Romance in Lower Mathematics is a 10-minute animation film by Chuck Jones, based on a short book by Norton Juster. It tells the story of the amorous adventures of a dot, a line, and a squiggle. Even though the visuals are minimal, the viewer has no problem following the story. It is a prime example of how motion rather than form can be used to communicate character and intent.

Many robots are not humanoid in form or do not have humanlike features but are still anthropomorphized. A robot vacuum cleaner trying to wriggle its way out from under a table will be described as "being lost" or "not knowing what it wants," humanlike descriptions that have little to do with the actual perception and processing of the robot but that help us communicate to others what the robot is doing.

Robot builders can actively encourage anthropomorphization. One effective method is to increase the reaction speed of the robot to external events: a robot that immediately responds to touch or sound will be perceived as more anthropomorphic. Such *reactive behavior*, in which the robot responds quickly to external events, is an easy approach to increase anthropomorphization. The robot jolting when the door slams shut or looking up when touched on the head immediately conveys that it is both alive and responsive. *Contingency*, responding with behavior that is appropriate for the context of the interaction, can also be used to enhance anthropomorphization. When a robot detects motion, for example, it should briefly look toward the origin of the movement. If the event—such as a tree swaying in the wind—is irrelevant to the robot, it should look away again, but if it is relevant—such as a human waving hello to engage the robot in interaction—the robot should sustain its gaze.

Although robot builders will often prefer a combination of both form and behavior to inspire users to anthropomorphize their robots, certain types of robots may be limited in how humanlike they can be. Android robots, which appear virtually identical to people, are technically limited in their behavioral repertoire. On the other hand, the developers of many toy robots are under pressure to make the hardware as cheap as possible and thus opt for an effective combination of simple visual features and reactive behaviors. It is important to also take people's expectations into account; the more apparently humanlike the robot, the more people will expect in terms of humanlike contingency, dialogue, and other features.

Impact of context, culture, and personality

People's perceptions of anthropomorphic robot design are often affected by contextual factors. Some people are more likely than others to anthropomorphize things around them, and this can affect how they perceive robots, as previous research has shown (Waytz et al., 2010). A person's age and cultural background can also affect their likelihood of anthropomorphizing or their interpretation of the robot's social and interactive capabilities Wang et al. (2010).

The context in which the robot is used, furthermore, can support anthropomorphization. In particular, just putting a robot in a social situation with humans seems to increase the likelihood that people will anthropomorphize it. The collaborative industrial Baxter robot, when used in factories alongside human workers, was regularly anthropomorphized by them (Sauppé and Mutlu, 2015). Furthermore, it seems that people who work alongside robots prefer them to be designed in more anthropomorphic ways: people preferred that Roomba have the ability to display its emotions and intentions with a dog-like tail (Singh and Young, 2012). Workers using Baxter put hats and other accessories on it and wanted it to be more polite and chitchat with them (Sauppé and Mutlu, 2015). Workers in a car plant using a co-bot, which was named Walt (see Figure 10.12) and had been designed to have a blend of social features and features reminiscent of a vintage car, considered the robot to be a team member (El Makrini et al., 2018). Office workers who were given a break management robot gave it names and requested that it be more socially interactive (Šabanović et al., 2014).

Seeing other people anthropomorphize robots can also suggest that anthropomorphization is a social norm to be followed. Researchers found that older adults in a nursing home were more likely to engage socially with Paro, the seal-like companion robot, when they saw others interacting with it like a pet or social companion (Chang and Šabanović, 2015). Clearly, anthropomorphic inferences may emerge instantly upon a first encounter and likewise become reshaped as a function of long-term interaction and acquaintance with a technical system.

4.2.4 Measuring anthropomorphization

Along with identifying anthropomorphization of robots as a common occurrence in HRI, researchers also need to know how to measure its presence in an interaction. According to the influential three-factor model of anthropomorphism, anthropomorphism extends to nonhuman entities the attribution of mental and emotional states that are essentially human. HRI researchers seeking to assess the human-likeness of a robot's form or behavior draw from the large body of literature on measuring humanity attribution among humans. These days, the HRI community measures a variety of related constructs, including asking research participants about the extent to which they would attribute mind (i.e., agency and experience (Gray et al., 2007)) or human nature and human uniqueness, which are typical human traits (Haslam, 2006). Similarly, other research has assessed psychological anthropomorphism and asked whether people perceived a robot to be capable of experiencing uniquely human emotions (Leyens et al., 2001), intentions, or free will (Epley et al., 2008).

A measure for anthropomorphism specifically developed for HRI is the Godspeed questionnaire. It has been widely used in the field and has been translated into several languages (Bartneck et al., 2009). More recently, researchers have started developing additional related scales, such as the ROSAS scale (Carpinella et al., 2017) and the revised Godspeed questionnaire (Ho and MacDorman, 2010).

Although many of these measures rest on self-reports and questionnaires, other, more subtle behavioral indicators (e.g., language use, application of social norms that are used in human–human interaction, such as in proxemics) may also be used to investigate the consequences of implementing humanlike form and function in social robots. Enriching the repertoire of measurements from direct to more indirect approaches will be beneficial, not only for the current research in the field of social robotics but likewise as a form of external validation of theorizing in psychology.

4.3 Design methods

Design in HRI spans a variety of methods inspired by practice from various disciplines, from engineering to HCI and industrial design. Depending on the method, the starting point and focus of design may weigh more heavily on technical exploration and development or on exploring human needs and preferences, but the ultimate goal of design in HRI is to bring these two domains together to construct a successful HRI system.

The design process is often cyclical in nature, following this pattern:

1. Define the problem or question.
2. Build the interaction.
3. Test.
4. Analyze.
5. Repeat from step 2 until satisfied (or money and time run out).

4.3.1 Engineering design process

The engineering design method is, as the name suggests, commonly used in engineering. Starting from a problem definition and a set of requirements, numerous possible solutions are considered, and a rational decision is made on which solution best satisfies the requirements. Often, the function of an engineered solution can be modeled and then simulated. These simulations allow engineers to systematically manipulate all the design parameters and calculate the resulting properties of the machine. For well-understood machines, it is even possible to calculate the specific design parameters necessary to meet the performance requirements. If a new aircraft takes off for its maiden flight, engineers can be almost certain that it will fly. It is important to note, however, that they cannot be absolutely certain because the new aircraft will interact with an environment that is not completely predictable in all its detail. Enough is understood, though, to be very sure of the macroscopic properties of the environment, allowing the engineers to design an aircraft that crosses the boundary from simulation to actual prototype without any hiccups. However, validating a solution in simulation is not always possible. The simulation might not be able to capture the real world in sufficient detail. Or the number of design parameters can be so high that a complete simulation of all possible designs becomes computationally impossible because it would take a computer years to calculate how each solution performs.

Engineers working in HRI tried to design a robot to teach eight- and nine-year-olds what prime numbers are. They believed that the children's learning would benefit from having a very personal and friendly robot, so they programmed the robot to make eye contact, use the child's first name, and politely support the child during the

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

quite taxing exercises. They compared the friendly robot against a robot in which the software to maintain an engaging relation was switched off, expecting that robot to be the worse teacher. They were dumbfounded when the aloof robot turned out to be the better teacher by a large margin, showing how their preconceptions regarding robot design were firmly out of touch with the reality of using a robot in the classroom (Kennedy et al., 2015) (see Figure 4.11).

To make things even more difficult, some design problems can be ill-defined, or insufficient information is available about the requirements or the environment. In this case, designers may say that they are dealing with a "wicked design problem" (Buchanan, 1992), which has changing, incomplete, interdependent, or indeterminate requirements that make it difficult to follow a linear model of design thinking in which problem definition can be cleanly followed by a process of problem solution. HRI design often is such a wicked design problem because there is a lack of information about the appropriate behaviors and consequences of robots in social contexts. Another approach to take in this case is to focus not on producing the absolute best solution, but on producing satisficing solutions Simon (1996). Satisficing is a portmanteau of *satisfy* and *suffice*, meaning that the resulting solution will be just good enough for the purpose it is meant to serve. This is a common problem-solving approach in all human endeavors, and it is almost unavoidable in HRI, where technical capabilities may never reach the ultimate design requirement of the robot performing just as well or better than people.

4.3.2 User-centered design process

As mentioned previously, relying only on the engineering design method can guide HRI development only so far, particularly when the intended uses of HRI are in open-ended interactions and spaces, outside labs or tightly controlled factory environments. In the process of satisficing, we may all too often choose not to measure the things that matter but instead only care about what is easy to measure. One way to address these issues is to focus more specifically on the people who will use the robot and the contexts of use they inhabit throughout the design process. This can be done through user-centered design (UCD). UCD is not specific to HRI and is used in many other design domains, such as HCI, and is a broad term used to describe "design processes in which end-users influence how a design takes shape" (Abras et al., 2004). The users can be involved in many different ways, including through

Figure 4.11 Boy learning math with a robot.

initial analyses of their needs and desires that can help to define the design problem, by asking them to comment on potential robot design variations to see which ones are preferable, and through evaluations of various design iterations of the robot and of the final product to evaluate its success among different users and in different use contexts.

Developers are typically confronted with having to make design decisions for which there are no obvious answers. Do people prefer the robot to have a red torso or a blue torso? Will a chirpy voice on a retail robot invite more people into the store? To answer these questions, they often build prototypes of the different design options and test them with their target audience. By carefully eliciting responses from the users, often using methods similar to those used in experimental research (see Chapter 9), the developers can ensure that the preferences or differences that they observe are not just coincidences but are really caused by the design feature under consideration. The results then inform the developers in building the best design option, and the cycle continues with new problems or design decisions. It is important to run these cycles as early as possible because the cost of making changes to the system increases dramatically later in the process. The credo is "test early; test often."

Designers often focus mainly on the people they think will use their products directly (i.e., the primary users), such as the nurses and patients who interact with a drug-delivery robot. It is, however, also important for designers to consider people who might only intermittently come into contact with the artifact or use it through an intermediary (i.e., the secondary users), such as other medical staff who see the robot in the hallway, and those who will be affected by the use of the artefact (i.e., the tertiary users), such as people whose job might be replaced or changed due to the introduction of new robotic technology. These various people involved in and affected by the robot's uses are called stakeholders, and an initial step in the design process can involve doing some research to identify who the relevant stakeholders are. Once the stakeholders are identified, the designers can then involve them in the design process through a variety of user-centered methods, which can include needs and requirements analyses, field studies and observations, focus groups, interviews and surveys, and user testing and evaluations of prototypes or final products (Vredenburg et al., 2002). You can learn more about several of these methods in Chapter 9 of this book.

Figure 4.12 Snackbot (2010), a system developed at Carnegie Mellon University to study robots in real-world settings.

Carnegie Mellon University's Snackbot was designed through a user-centered process that involved taking into consideration the robot, people, and the context. It was iteratively performed over 24 months and involved research on where people could already get snacks in the building to establish need, initial technology feasibility and interaction studies, multiple prototypes, and further studies of how the robot was used and the effects of different forms of dialogue and robot behaviors on user satisfaction. (Lee et al., 2009) (see Figure 4.12)

4.3.3 Participatory design

Recently, HRI researchers have started applying more collaborative and participatory design approaches to HRI. Both collaborative and participatory methods seek to include the potential users and other stakeholders, or people who might be affected by robots, in the process of making decisions about appropriate robot design from early on in the design process. This is clearly distinct from the notion of bringing users in at the evaluation stage, where the design is partially or fully formed and users' input is largely used to test particular factors and assumptions already expressed in the design. In this way, participatory design recognizes the expertise people have about their everyday experiences and circumstances.

Participatory design has been present in the design of other computing technologies, particularly information systems, since the 1970s, when it was used to enable workers in organizations to participate in the design of software and other technologies that they would use in their work later on. Participatory design in HRI has been working on developing ways for users to become engaged in the process of making design decisions about robots—for instance, by testing and developing particular behaviors for robots, designing robot applications for their local environments, and conceptualizing how existing robotic capabilities can potentially address their needs and fit into their everyday contexts. DiSalvo et al. (2008) performed one of the early participatory design projects in HRI in their "neighborhood networks" project.

Here, community members used a robotic prototype provided by the researchers to develop environmental sensors for their neighborhood. In another participatory project, roboticists and visually impaired community members and designers worked together in a series of workshops to develop appropriate guidance behaviors for a mobile PR-2 robot (Feng et al., 2015). Participatory design has also been used in various healthcare and educational applications for HRI (see, e.g., Šabanović et al., 2015).

Participatory design is always challenging, but working on participatory design with robots has its particular difficulties. One is the fact that people have many different preconceptions about robots but little knowledge about the technology involved in making them, which leads to unrealistic design ideas. At the same time, designers have little knowledge of the day-to-day lives and experiences of people in many of the applications in which HRI is most needed (e.g., eldercare). While working with older adults and nursing home staff to develop assistive robots for older adults with depression, Lee et al. (2017) and Winkle et al. (2018) focused on supporting a process of mutual learning between HRI researchers and participants, which allowed both sides to explore and teach each other about their different areas of expertise. This also helped support participants' learning to start thinking about design beyond just designing for themselves. Participatory design is still new in HRI, but with more and more applications being envisioned for diverse populations and everyday contexts, it is becoming an increasingly important component of the HRI design methods toolkit.

4.4 Prototyping tools

Although it is possible to develop simple robot prototypes from generally available materials such as cardboard or found objects, several prototyping kits and tools for creative interactive technologies have recently become available on the market. These make it possible for a wide variety of people with different levels of technical expertise and economic resources to try their hand at robot design. They also enable more rapid and iterative development of robot designs by making the representation of interaction a simpler thing to create.

Perhaps the earliest type of kit that could be used for developing different robot designs was the first-generation LEGO Mindstorms system, which provided bricks for building and specialized bricks for programming and actuating simple robot prototypes. Bartneck and Hu (2004) used LEGO robots to illustrate the utility of rapid prototyping for HRI, and the first case studies had already appeared in 2002 (Klassner, 2002).

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 4.13 LEGO Mindstorms (1998–present) was the brainchild of Seymour Papert, an MIT professor who was an avid proponent of using computers to support child learning.

The Vex Robotics Design System[1] is also widely known and used, and its advanced version is the kit of choice for the popular FIRST Robotics Competitions.[2] More recent additions to the array of kits available are Little Bits, which provides easy-to-use plug-and-play electronic bricks, including sensors and actuators, among others, that can be used to quickly and easily create interactive prototypes.

The Arduino microcontroller[3] is very affordable and has a large hobbyist community providing open-source designs and code, as well as a wide array of peripherals (sensors, motors, LEDs, wireless units, etc.) that allow for more flexibility in design but require more technical know-how.

Other equipment, such as the Raspberry Pi[4] single-board computer and affordable and even portable three-dimensional (3D) printers not only make HRI prototyping easier but also may even be said to be making it accessible to the masses (or at least to college students).

Designers also incorporate other existing technologies into robot design, including smartphones. Even an average smartphone these days has sufficient computing power to control a robot. Furthermore, it has many built-in sensors (microphone, camera, gyro sensor, accelerometer) and actuators (screen, speaker, vibration motor). Robovie-MR2 is an early example of integrating a smartphone into a robot to control all

[1] `https://www.vexrobotics.com`
[2] `https://www.firstinspires.org/`
[3] `https://www.arduino.cc`
[4] `https://www.raspberrypi.org`

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

of its functions (see Figure 4.3). Hoffman calls this the "dumb robot, smartphone" approach to social robot design (Hoffman, 2012).

Available technologies for prototyping continue to develop, fueled at least in part by ongoing efforts to engage more students, hobbyists, and even potential users in technology design.

4.5 Culture in HRI design

As not only an interdisciplinary but also an international field of research, HRI design has been particularly interested in the question of cultural effects on perceptions of and interactions with robots. Culture, the different beliefs, values, practices, language, and traditions of a group of people, plays into robot design both in the form of factors introduced by designers and the context in which users interpret different HRI designs.

Researchers commonly make connections between cultural traditions and the design and use of robots, particularly contrasting the norms, values, and beliefs in the East and West: animist beliefs have been used to explain the perceived comfort of Japanese and Korean populations with robots (Geraci, 2006; Kaplan, 2004; Kitano, 2006), whereas human exceptionalism has been suggested as a source of Westerners' discomfort with social and humanoid robots (Geraci, 2006; Brooks, 2003). Holistic and dualistic notions of mind and body (Kaplan, 2004; Shaw-Garlock, 2009) and individualist and communitarian social practices (Šabanović, 2010) have been identified as design patterns represented in the design of robots and potential human interactions with them.

In addition to these generalized connections between culture and robotics, HRI researchers have been studying cultural differences in and effects on people's perceptions of and face-to-face encounters with robots. In a comparison using Dutch, Chinese, German, U.S., Japanese, and Mexican participants, it was found that U.S. participants were the least negative toward robots, whereas the Mexican participants were the most negative. Against expectations, the Japanese participants did not have a particularly positive attitude toward robots (Bartneck et al., 2005). MacDorman et al. (2009) showed that U.S. and Japanese participants have similar attitudes toward robots, suggesting that such factors as history and religion may affect their willingness to adopt robotic technologies. Survey evaluations of the seal-like robot Paro (see Figure 2.6) by participants from Japan, the United Kingdom, Sweden, Italy, South Korea, Brunei, and the United States found that participants generally evaluated the robot positively but identified different traits as most likable according to their country of origin (Shibata et al., 2009).

In the context of human–robot teamwork, Evers et al. (2008) found that users from China and the United States responded differently to robots and that human team members found robots more persuasive when they used culturally appropriate forms of communication (Lindblom and Ziemke, 2003). Findings from two generative design studies with participants in the United States and South Korea, which asked users to think about robotic technology in their own homes, showed that user expectations of and needs for robotic technologies are related to culturally variable conceptions of the home as relation oriented in Korea and more functionally defined in the United States (Lee et al., 2012). The growing body of work on cross-cultural differences in HRI and their potential design implications identifies that cultural considerations should be taken into account when designing robots both for international and local uses.

4.6 From machines to people, and the in between

As the previous discussion shows, designing human–robot interactions involves making many decisions about the form, function, and desired effects of robots. HRI designers, however, also bring deeper philosophical, ethical, and even political commitments into their work. Although these can be unconsciously brought into HRI research, we think it is useful for HRI scholars to consciously engage with these concerns in the course of their robotics research and development.

One of the most basic decisions that robotics researchers make is the type of robot they want to work on—is it meant to resemble a human or be more like a machine? Another decision can involve the main goals of the work—is it focused on producing technical developments, understanding humans, or perhaps developing HRI systems that can be used for specific applications and contexts of use? These decisions have significance beyond just the design and use of the robot, however. One could argue that the creation of robots by their designers, in particular those in which robotic copies of actual people are created, is an immortality project. Such projects are “symbolic belief systems that promise that the individual will not be obliterated by the demise of his or her physical body” (Kaptelinin, 2018). Hiroshi Ishiguro’s work on android copies of living human persons is a case in point, in which the robotic copy can aim to stand in the place of that specific person, both in current and ostensibly future interactions. Ishiguro himself describes how he feels his own identity is interconnected with the robot, which persists as a replica of his past and younger self that he now feels the pressure to emulate (Mar, 2017). But the relationship between machine-like robots and designers can be just as deep. Describing his work with industrial robots, Japanese roboticist Masahiro Mori defined the relationship between humans and machines as being “fused together in an interlocking entity” (Mori, 1982). This close relationship has direct consequences for the form and function of the robot on the one side and the designer on

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

the other side, as well as on the future consequences and uses of the robot in society.

Robot design can also be guided by a personal commitment to specific social and philosophical values, such as improving access to resources for broader populations, increasing participation in the design of and decision-making about robots, or contributing to the solution of pressing social issues. Roboticist Illah Nourbakhsh described how his personal values affect his robotic projects as follows:

> One way out is to say my work is purely theoretical, who cares how somebody applies it? I didn't want to do that. I wanted to say my work involves theoretical components, but I'm taking it all the way to seeing a real result in the physical world. And furthermore, I want it to be socially positive in some measure... I want to work on something so socially positive that not only do I hope everyone uses it, but I want to see at least one used case to fruition. Then you have this feedback loop from real-world application back to engineering design. (Šabanović, 2007, p. 79)

In this way, the choice of what type of HRI project to pursue and the goals to focus on in design can reflect personal or collective values (e.g., of the research group or of project collaborators). After all, time is limited and valuable, so it makes sense to consciously choose what we hope to make of it.

One of the authors finds inspiration for his design in the work of Robert M. Pirsig (see Figure 4.14), who put it this way:

> The real [aesthetics] lies in the relationship between the people who produce the technology and the things they produce, which results in a similar relationship between the people who use the technology and the things they use. (Pirsig, 1974)

Pirsig emphasizes the crucial role of obtaining peace of mind in order to arrive at good design as the barrier between the designer and the object to be designed dissolves:

> So the thing to do when working on a motorcycle, as in any other task, is to cultivate the peace of mind which does not separate one's self from one's surroundings. When that is done successfully then everything else follows naturally. Peace of mind produces right values, right values produce right thoughts. Right thoughts produce right actions and right actions produce work which will be a material reflection for others to see of the serenity at the centre of it all. (p. 305)

Once peace of mind is achieved and the barrier between the object and the designer is broken down, the design work can start. This work is similar to that of artists. It takes patience, care, and attentiveness to

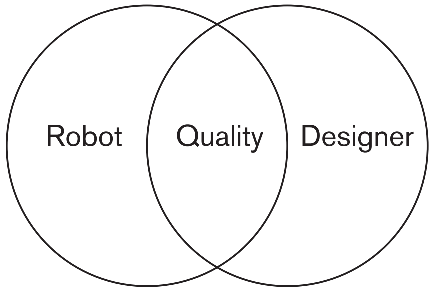

Figure 4.15 Quality in the design of robots.

what you are doing. A good indicator of whether the design is progressing in a good direction is the inner peace of the designer. If you are in harmony with what you design, then the robot and your thoughts change together in a state that is often described as "flow" (Csikszentmihalyi and Csikszentmihalyi, 1988). The materials and the inner state of mind will come to rest at the point where the design is complete and good. Peace of mind, according to Pirsig, is not only the prerequisite for good design work, but it also accompanies good design work and is also its final goal:

> Peace of mind isn't at all superficial to technical work. It's the whole thing. That which produces it is good work and that which destroys it is bad work. The specs, the measuring instruments, the quality control, the final checkout, these are all means toward the end of satisfying the peace of mind of those responsible for the work. What really counts in the end is their peace. (p. 302)

The connection between the robot and its designer is far deeper than you may assume. Robert M. Pirsig spent his whole life working out *The Metaphysics of Quality*, in which he argues that there is no fundamental difference between the designer and the object he or she designs. What connects them is "quality" (see Figure 4.15).

Figure 4.14 Robert M. Pirsig (September 6, 1928–April 24, 2017) is the author of *The Metaphysics of Quality*, which has inspired many designers.

Considering the peace of mind of the designer might sound strange at first, but Pirsig argued that in the moment of the perception of quality, there is no division of objects and subjects. In the moment of such pure quality, the subject and the object are one (Pirsig, 1974, p. 299). Artists might be familiar with the experience of unity with their work, and the work of designers and engineers might be enhanced if they, too, would be more sensitive to this connection.

4.7 Conclusion

Designing robots requires multidisciplinary expertise, often by means of a team, and a process that takes the users and the interaction context into consideration. Various prototyping tools are available to quickly build and test robots. Once the users and their interactions with the

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

robot are understood, the robot needs to be designed from the outside in—starting with the potential users and use context to develop design concepts and the technical specifications for the robot. HRI designs also express, whether consciously or unconsciously, the social and ethical values of the designers.

The robots' anthropomorphism is one of the most important design considerations in contemporary HRI. We provided a detailed description of the construct of psychological anthropomorphism as a prime opportunity for a fruitful exchange between disciplines, leading to a broader overall understanding of the concept in the social sciences and robotics. Beyond the theoretical and methodological gains from investigating anthropomorphism, HRI studies have also shown the importance of considering humanlike form and function in robot design for perceived interaction quality, HRI acceptance, and enjoyment of the interaction with humanlike robots.

Questions for you to think about:

- Find examples of pareidolia in your environment.
- Think about the features of a humanlike robot in terms of "design affordances." Which affordances should be considered in humanlike robots?
- Try to think about "design patterns" for social robots that greet people daily. Find and describe repeatedly reused patterns in behavior.
- Imagine you have to design a robot. Consider the necessary steps, taking a participatory design approach.
- Discuss the role of user expectations in robot design. What are important points to consider if you want to market your robot?
- What is your opinion: Should a social robot have very few humanlike cues, or should it be highly anthropomorphic in design (e.g., like an android)? Which robot would be accepted more by people in general? Why?
- Think about a robot that you might want to have in the near future. Picturing this robot, try to think about a way to encourage more anthropomorphization based on its behavior. Which behaviors should the robot show to be perceived as humanlike?

Future reading:

- Brian R. Duffy. Anthropomorphism and the social robot. *Robotics and Autonomous Systems*, 42(3):177–190, 2003. ISSN 0921-8890. doi: 10.1016/S0921-8890(02)00374-3. URL `https://doi.org/10.1016/S0921-8890(02)00374-3`
- Nicholas Epley, Adam Waytz, and John T. Cacioppo. On seeing

human: A three-factor theory of anthropomorphism. *Psychological Review*, 114(4):864–886, 2007. doi: 10.1037/0033-295X.114.4.864. URL `https://doi.org/10.1037/0033-295X.114.4.864`

- Julia Fink. Anthropomorphism and human likeness in the design of robots and human-robot interaction. In Shuzhi Sam Ge, Oussama Khatib, John-John Cabibihan, Reid Simmons, and Mary-Anne Williams, editors, *Social robotics*, pages 199–208, Berlin, Heidelberg, 2012. Springer. ISBN 978-3-642-34103-8. doi: 10.1007/978-3-642-34103-8_20. URL `https://doi.org/10.1007/978-3-642-34103-8_20`
- Peter H. Kahn, Nathan G. Freier, Takayuki Kanda, Hiroshi Ishiguro, Jolina H. Ruckert, Rachel L. Severson, and Shaun K. Kane. Design patterns for sociality in human-robot interaction. In *The 3rd ACM/IEEE International Conference on Human-Robot Interaction*, pages 97–104. ACM, 2008. ISBN 978-1-60558-017-3. doi: 10.1145/1349822.1349836. URL `https://doi.org/10.1145/1349822.1349836`
- Travis Lowdermilk. *User-centered design: A developer's guide to building user-friendly applications.* O'Reilly, Sebastopol, CA, 2013. ISBN 978-1449359805. URL `http://www.worldcat.org/oclc/940703603`
- Don Norman. *The design of everyday things: Revised and expanded edition.* Basic Books, New York, NY, 2013. ISBN 9780465072996. URL `http://www.worldcat.org/oclc/862103168`
- Robert M. Pirsig. *Zen and the art of motorcycle maintenance: An inquiry into values.* Morrow, New York, NY, 1974. ISBN 0688002307. URL `http://www.worldcat.org/oclc/41356566`
- Herbert Alexander Simon. *The sciences of the artificial.* MIT Press, Cambridge, MA, 3rd edition, 1996. ISBN 0262691914. URL `http://www.worldcat.org/oclc/552080160`

5

Spatial Interaction

What is covered in this chapter:

- The importance of the spatial placement of agents in social interaction;
- Basic understanding about human proxemics (i.e., how people manage space around them in a social context);
- How a robot manages the space around it, including interactions such as approaching, initiating interaction, maintaining distance, and navigating around people;
- How the properties of spatial interaction can be used as cues for robots.

In 2012, Exertion Games Labs released a drone exercise companion called Joggobot (see Figure 5.1). Runners who feel like they can use a little extra motivation or companionship during their run but don't have a personal trainer or a friend to join them can now have a drone accompany them during their exercise laps. One of the critical features of Joggobot is its placement in space during the run: right in front of the runner, like a carrot tempting a running horse. This position wasn't chosen on a whim. The developers studied where the drone should ideally be in relation to the runner (i.e., above, following, leading, on the side) and how much of a distance it should keep in order to maximize motivation (Graether and Mueller, 2012). They found that having the drone flying behind the jogger made people feel like they were being chased, which decreased their enjoyment in exercising. Users much preferred to take on the chasing role themselves.

Figure 5.1 The Joggobot Drone (2012). (Source: Photo provided by Eberhard Gräther and Florian "Floyd" Mueller)

This example shows that the placement of the robot with respect to the user is an important aspect of human–robot interaction (HRI). When only taking the need for collision avoidance into account while deciding on the optimal location or path of a robot, one might inadvertently create robot behavior that is considered uncomfortable, rude, or inappropriate. When a Roomba vacuum cleaner treats people as "obstacles" and keeps bumping into them as it tries to avoid them, it can comically seem to be "humping" their feet. Thus, when planning a robot's placement in space, it is important to take into account people's

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

preferences and the social norms that exist regarding such placement in relation to others.

5.1 Use of space in human interaction

When space is available, individuals are strongly expected to adhere to social distance norms. Most people feel that it is inappropriate for a stranger to sit down right next to them in an otherwise empty bus. However, when taking the bus during rush hour, we are forced to step into others' personal space, and it is acceptable to sit or stand close to people. Even though it is not considered impolite to stand next to someone on a busy commute, people still often feel uncomfortable, avoiding eye contact and quickly repositioning themselves at a greater distance when more space becomes available.

5.1.1 Proxemics

Cultural anthropologists coined the term *proxemics* to describe how people take up space in relation to others and how spatial positioning influences attitudes, behaviors, and interpersonal interaction.

Hall et al. (1968) describe four distance zones in their original work: intimate distance, personal distance, social distance, and public distance (see Figure 5.3). When the available space is (relatively) unlimited, these distances indicate the psychological closeness between people (see Figure 5.2).

Figure 5.2 Commuters during rush hour on the Tokyo underground having their personal space violated. We often deal with this by avoiding the gaze of others.

As the name suggests, intimate distance is reserved for close personal relationships or the sharing of private information. Intimate distance ranges roughly from a few centimeters to about half a meter, depending on one's age and culture. Together with personal distance (which ranges from about half a meter to 1.2 meters), these zones make up the personal space of a person: the amount of space that people generally consider theirs to take up. Under normal circumstances, only friends, relatives, and partners are expected to come this close. For less personal relationships, such as acquaintances or colleagues, one is expected to maintain social distance, which ranges between 1.2 and about 4 meters between persons. Finally, public distance starts at around 4 meters, which is the distance people are expected to keep between them in relatively impersonal settings, such as public speaking at a conference.

Hall considered people's use of space as an often-overlooked dimension of cultural experience and noted that people from different cultures have varying personal proxemic preferences and expectations. For example, in "high-contact cultures" such as those of South America, people will frequently enter each other's personal space and touch, whereas in "low-contact cultures," such as the United States, touching a stranger may be construed as assault. Hall wittily observes that

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

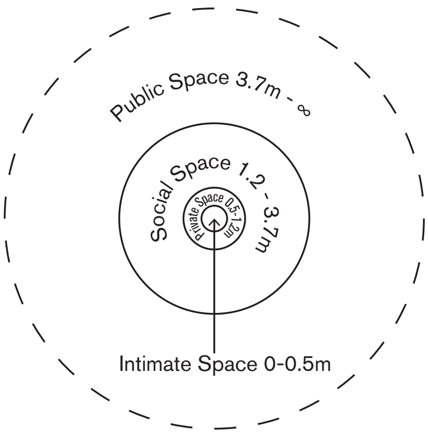

Figure 5.3 Intimate, personal, social, and public distance according to Hall et al. (1968).

North Americans visiting South America will find themselves "barricaded behind desks, using chairs and typewriters to keep the Latino at what is to us a comfortable distance."

Sometimes slight breaches of proxemic norms are made on purpose by individuals, for instance, to create more psychological closeness or perhaps to intimidate. For example, a man who nonchalantly places his arm first on the backrest of the sofa where his date is sitting and then cautiously inches closer and closer is making a transition from personal distance to the intimate zone. The friend who touches your arm when you are telling a personal story does the same, although with a different underlying motive. However, these moves have to be made very cautiously and under continuous assessment and reassessment of the reaction of the other person. Few people would be charmed if the hopeful suitor had abruptly placed himself right on his date's lap. Likewise, when we attempt to comfort a colleague by giving a hug at the wrong moment, the interaction can turn awkward rather quickly. This is because the meaning of spatial interaction cues is highly contextual. Unlike the friendly moves just mentioned, an investigator questioning a suspect may "get in the suspect's face" by moving as close to him or her as possible to seem more threatening.

Not only the distance at which we interact with each other but also our placement in relation to interaction partners are bound by social norms. For example, researchers found that people who sat next to each other were more cooperative, whereas people sitting opposite each other behaved more competitively. During conversations, people usually position themselves at an angle to each other (Cook, 1970). The way in which people place themselves with respect to each other is therefore an important aspect of the dynamics of interaction (Williams and Bargh, 2008).

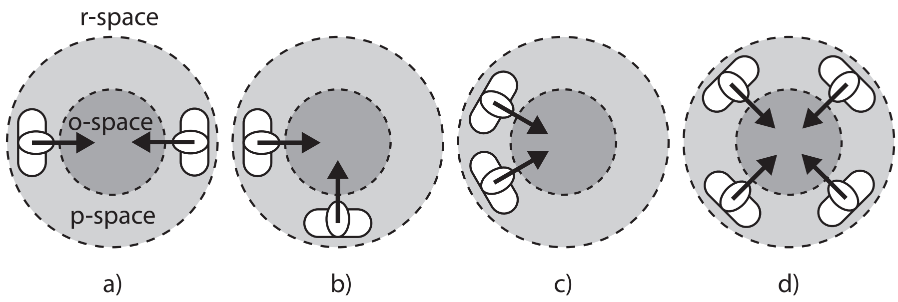

Figure 5.4 Kendon's F-formations come in several variants, all of which include the components of o-, p-, and r-space, namely (a) the face-to-face, (b) the L, (c) the side-by-side, and (d) the circular formation.

5.1.2 Group spatial interaction dynamics

The importance of spatial dynamics goes beyond one-on-one interaction and is also salient in group interaction scenarios. The spatial orientation of people in a group in relation to others can make the group seem as if it were inviting more members or seeking to keep others out. For example, at a cocktail party, when people stand in a tight-knit circle, it can seem difficult to join in the conversation. However, if the group notices people wanting to join and opens up the circle so that there is space for new members to fill, it can be construed as an invitation to participate. This type of information can be useful for robots to gauge which groups of people they can approach in public spaces like museums or malls or if they want to affect the interaction dynamics of human groups.

Group spatial dynamics such as these were described by Adam Kendon as the "facing formation," or "F-formation" (Kendon, 1990) (see Figure 5.4). These formations are created through the positioning of two or more people in relation to each other, such that the areas of space that they are facing and on which they focus their attention are overlapping. The space between these people, which is "one to which they have equal, direct, and exclusive access," is termed the o-space. The group participants themselves are said to occupy the p-space, and they are surrounded by r-space. People can modify their positions to maintain this space or to include other participants in the group conversation, as in the previous example. Different configurations of the F-formation are possible, based on people's orientation to each other, and are termed the face-to-face, L-shape, and side-by-side formation for two people and the circular formation and other shapes for larger groups.

These group formations have been used to understand people's interactions with technology (Marshall et al., 2011) in general and with robots more specifically (e.g., Hüttenrauch et al., 2006; Yamaoka et al., 2010). In navigation around people, Pérez-Hurtado et al. (2016) found that a robot needs to be aware of people movements and cognizant of people engaged in conversation and not walk between them even if there is enough space.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

5.2 Spatial interaction for robots

Robots will often share physical space with humans. Some robots are mobile, moving over the ground or through the air. Some of them have arms and manipulators so that they can interact with objects and users. The placement and movement of such robots with respect to people must be considered when designing human–robot interactions. Robots that do not respect the personal space of the user will evoke negative reactions or even rejection and withdrawal by the user. Robot designers can attempt to increase acceptance of the robot by having it keep an appropriate distance (assuming that they can code the robot in such a way that it knows what the "appropriate distance" is at a given point in time and space) and adjusting its position to create a fitting interaction experience. For example, a security robot might initially keep a polite distance but enter a person's intimate space at some point in the interaction in an attempt to intimidate the person.

5.2.1 Localization and navigation

Before going into HRI, let us briefly explain basic techniques inherited from robotics that are required for a robot in order to engage in spatial interactions with humans. When a robot wants to interact with people, it needs to locate itself in space with regard to the people it aims to interact with. Thus, one of the basic techniques required for mobile robots is localization; a robot needs to know where it is. This is not a trivial problem. A typical robot is equipped with an odometer, a sensor that records the distance traveled by the robot's wheels. However, as the robot travels, these lose accuracy, and the robot therefore needs to correct the information that the odometry provides about its location. The typical solution to this is to let the robot build a map of its environment and then cross-reference information on its location and orientation from the odometry with information from other sensors, such as a laser range finder or camera, to locate itself on the map. This process is known as simultaneous localization and mapping, or SLAM (Davison et al., 2007; Thrun et al., 2005).

In addition to reporting the robot's location, localization can help the robot know what type of space it is in (e.g., whether it is in the living room or bathroom). However, it will not reveal anything about the whereabouts of any people in that space.

Identifying the location and orientation of people interacting with the robot thus is the next challenge. For detecting people at a short range, the robot will carry sensors such as two-dimensional (2D) cameras and depth cameras, that enable it to identify nearby people. The software processing the camera images can not only detect and track humans but also can report on the location of body parts such as arms, legs, and heads. For tracking people at longer distances, there are techniques that use laser range finders (also known as light detection and ranging [LIDAR]). A motion-capturing system is sometimes used. By placing reflective or fiducial markers on people and objects, motion capture can be used to identify and locate the markers (and by extension, the people or objects they were initially attached to). However, these marker-based approaches are difficult to use outside a lab setting. Finally, researchers can also mount sensors, such as cameras, in the environment to track people (Brscić et al., 2013).

Moving the robot through a crowded environment, also known as robot navigation, is a well-studied problem in mobile robotics. To avoid collisions between the robot and objects or people, techniques such as the dynamic window approach (DWA) are often used (Fox et al., 1997). The idea behind this technique is that a system computes its future location based on the current velocity of the robot while at the same time considering whether to keep or alter its velocity within the limitation of its actuation capability—and while calculating a future velocity that does not result in a collision. Over longer time scales, there are techniques based on path planning. In these techniques, if a given goal of a robot is not within immediate view of the robot, a path-planning algorithm computes a set of way-points or paths for the robot that will let it reach its goal. In HRI, most path-planning algorithms that work well for navigating around obstacles will result in socially inappropriate behavior when tried around people.

Localization and navigation can also take various elements of interaction with a user into account. For instance, Spexard et al. (2006) developed a robotic mapping technique that uses input from dialogue with users to learn about new places in an environment. To develop a human-friendly mapping technique, Morales Saiki et al. (2011) had a robot explore the environment while collecting visual landmarks to build a cognitive map from a humanlike perspective; this enabled the robot to generate route instructions that people could easily comprehend. Researchers have also worked toward developing techniques to understand human spatial descriptions, such as route directions. For instance, Kollar et al. (2010) developed a technique to associate a user's instructions and visual information about the environment to help the robot interpret the location mentioned by a user.

5.2.2 Socially appropriate positioning

Even though there are basic techniques for perception and navigation that allow robots to move around without colliding with obstacles, robots still often lack the capabilities to navigate in a socially appropriate way in the presence of other people. Suppose we want a robot to move through a corridor in an office building. What would happen if it considers people as obstacles? When a person walked toward the robot from the other end of the corridor, the robot would continue to move straight down the corridor until inches before colliding and then move out of the way. Although it would avoid the person eventually, this behavior is very different from what humans would do in a similar situation: we yield to each other well in time, nonverbally showing which side of the corridor we will walk on, and will avoid entering each other's personal space. Thus, a robot waiting until the last moment before moving out of the way may be seen as confrontational or aggressive, even though it still avoids collision with a person.

Most mapping techniques for robots only provide geometrical maps, where people are considered obstacles. They do not contain information on which direction people are facing, if they are having a conversation or just standing close to each other, or how people are moving. Hence, there are several techniques that allow a robot to acquire a more human-aware representation of its environment.

One of the focuses in investigating proxemics in HRI has been identifying appropriate interaction distances between users and robots. These include questions like the following: How close do people prefer to stand relative to a robot? And how close should a robot approach people before it is considered rude or inappropriate or makes people feel uncomfortable (see Figure 5.5). Walters et al. (2005) measured the distance at which people feel comfortable when they are approached by a robot. They reported that the majority of people prefer a personal or social distance when interacting with a robot, although some people prefer to stand even closer. Hüttenrauch et al. (2006) reported that people preferred that the robot stand at distances derived from human proxemics. Investigating interactions between a robot and a group of people, Kuzuoka et al. (2010) reported that a robot can change the conversational F-formations of the group by changing its body orientation, and they also found that movement of the robot's whole body was more effective than having the robot just move its head.

Figure 5.5 The drone calculates a probabilistic model of where the human will go and plans a safe route around (Fisac et al., 2018).

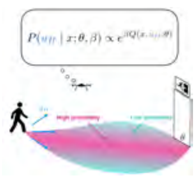

Relational position is also important when people and robots interact while they are moving. To enhance a robot's social acceptability, techniques have been developed for robot navigation based on human proxemics. For instance, when a robot follows a user from behind, the robot can either follow the same trajectory as the user, or it can move directly to the user's current location, which might be a shorter and faster pathway. Gockley et al. (2007) showed that users perceive the first behavior as more natural. Morales Saiki et al. (2012) developed a technique that allows a robot to navigate side by side with its user, for which they found it important for the robot to anticipate the user's future motion. Furthermore, people's perceived safety does not necessarily correspond to what a robot computes to be safe. For instance, in the corridor passing problem, it was found that a robot needs to

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

maintain enough distance to avoid entering a person's intimate zone (Pacchierotti et al., 2006). Alternatively, a robot can mimic how people avoid colliding into each other. Luber et al. (2012) and Shiomi et al. (2014), for example, developed a pedestrian model that implemented collision avoidance for dynamic environments. Considerations of comfort and perceived safety can also be integrated into path planning. Sisbot et al. (2007) developed a path planner for a mobile robot that plans how to reach a given goal while avoiding situations that might make people uncomfortable. The planner takes into account aspects such as whether people are sitting or standing and whether the robot might surprise them by suddenly appearing from behind an obstacle. Fisac et al. (2018) used a probabilistic model of a human walking to plan and execute a safe trajectory for an indoor drone (see Figure 5.5).

Planning a motion path that people will perceive as safe and comfortable is also necessary when only a part of the robot enters the user's personal space. For example, when a robot arm is used near a person, such as when a person and an industrial robot collaborate on a shared task, the robot must take the socially appropriate distance into account when computing a path for its end effector (e.g., hand) to reach its given goal (e.g., grasp an object or hand an object to a person) (Kulic and Croft, 2005). This may make the robot's movement inefficient from a purely functional standpoint, but it will lead to a more positive evaluation of the interaction by the user (Cakmak et al., 2011).

5.2.3 Spatial dynamics of initiating HRI

Every social interaction has to be initiated by someone, perhaps by hovering in the vicinity of the person you want to talk to at a cocktail party while orienting your body toward the person, for example, or by approaching a colleague to hand over the annual report. How you approach each other and how the approach is perceived have implications for the ensuing interaction.

Approaching behavior is generally expected to have positive effects on both parties in the interaction. The approacher makes an effort to attract and share attention, which signals interest in the person being approached. At the same time, initiating an interaction triggers neural activity in reward-related brain areas, resulting in positive affect in the initiator (Schilbach et al., 2010). Initiating interaction is, furthermore, a sign of being assertive and having faith in one's capability to conduct a successful social encounter. What may be more surprising is that this runs the other way too. People who approach others are seen by their peers as having more personal control (Kirmeyer and Lin, 1987).

Imagine the moment when a person meets a robot for the first time. Either of them could approach the other to initiate the interaction.

Whereas this can be rather trivial for a person, a robot needs to be carefully designed to appropriately initiate an interaction. Approaching behavior for robots has been studied from early on in the field of HRI. For instance, in a situation where a robot joins a queue, the robot needs to respect the personal space of other people who are also waiting (Nakauchi and Simmons, 2002). When a robot encounters people, it needs to switch its navigation mode from purely functional to considering social distance and spatial configuration (Althaus et al., 2004).

Initiating an interaction is also context and task dependent. Satake et al. (2009) show how a robot offering information about the stores in a mall will fail to initiate an interaction if the approach is poorly planned and executed. The planned trajectory needs to be both effective and acceptable to human visitors (Satake et al., 2009; Kato et al., 2015). Whereas approaching from the front was found to be desired when a robot was trying to initiate a conversation, approaching from the front when the robot was delivering an object to a person was less preferred and resulted in more failures (Dautenhahn et al., 2006b; Shi et al., 2013).

Some recent work incorporates machine learning to generate appropriate approaching behaviors that fit with a context. Liu et al. (2016) designed approaching and initiating behavior for a store clerk robot using a fully automated analysis of observed human behavior. The researchers first recorded how people moved and talked in a camera store scenario and then used machine learning to extract typical speech behavior and spatial formations. These behaviors were then transferred to the robot. A user study showed that the learned speech and motion behavior was considered to be socially appropriate by users.

Even in the case where a person approaches a robot, the robot should respond at just the right moment. If it fails to do so, the user could find the interaction unnatural and awkward and might even give up initiating interactions in the future (Kato et al., 2015). Human proxemics studies, particularly observational studies on the interactions of humans with either one another or with robots, can provide more contextually attuned and relevant models. For instance, Michalowski et al. (2006) developed a categorical model of human spatial interaction and engagement with a receptionist robot from observations of people's interactions with the robot. They defined the appropriate timing and types of behavior (e.g., turning toward a person, saying hello) that the robot could perform with people in different spatial zones, in order to both be perceived as more approachable and to successfully initiate an interaction when appropriate.

Social navigation has become particularly relevant in the context of self-driving cars. The story goes that the first self-driving cars at Google drove optimal trajectories following the highway code, but they

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

frequently startled other road users by driving too close or cutting them off. Only when politeness was explicitly added as an optimization criterion did the cars drive in a way that was socially acceptable.

5.2.4 Informing users of the robot's intent

Robot motion trajectories are often used to convey the intent and goal of the robot. Path-planning algorithms have been developed to explicitly convey information through the robot's trajectory. For instance, by slowly passing a few meters from a visitor, a mobile robot is able to express whether it is available for an interaction (Hayashi et al., 2012). Similarly, trajectories have been used as a means to allow a robot with few options to express itself, such as cleaning robots and drones, to communicate their intent to users (Szafir et al., 2015).

During handover in HRI, that is, when a robot hands an object to its user, users prefer a robot to behave with "legibility"—in a way that allows users to understand its goal and intention (Koay et al., 2007a). Hence, researchers have developed algorithms to control a robot arm to generate legible motions while reaching a given goal. A robot could hand over an object to a person in many different ways, but the most energy-efficient way may be incomprehensible to a person, so it is better to perform a motion that is easier to interpret (Dragan et al., 2013).

When a robot works closely with a person, it needs to have the capability to understand how the person is perceiving the space around him or her. An important related capability is spatial perspective-taking (Trafton et al., 2005). Imagine a situation where two people are working together. One might ask the other to pass an object by saying "give me that object." The referent of "object" will be obvious if there is only one object available. But what if there are several objects? For people, inferring the intended referent of "object" is often easy. We may use a complex set of cues, including gaze direction, body orientation, the prior context of the interaction, knowledge about the person and his or her preferences, task information, and other cues to disambiguate the request. For a robot, however, this can be rather complicated. Several approaches exist that allow the robot to take the perspective of the user. These often rely on geometric models that keep track of the location of people, robots, and objects and which of these are visible and reachable by whom (Lemaignan et al., 2017; Ros et al., 2010).

5.3 Conclusion

The study of spatial interaction in HRI is often inspired by our understanding of human proxemics, conversational relations, and relational positioning and approach behaviors, although we cannot expect the effects to always be the same. However, norms and understandings that

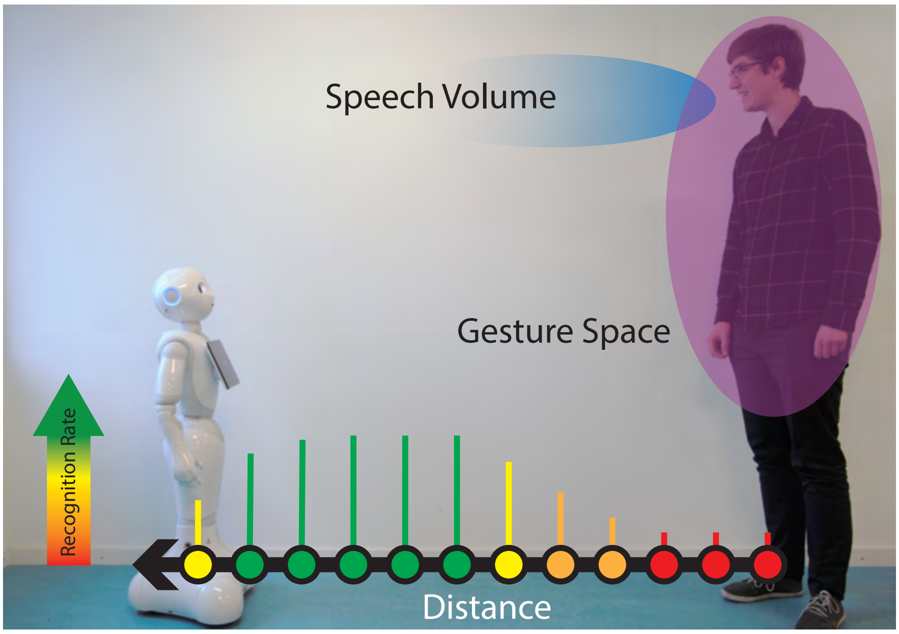

Figure 5.6 A lab setup for proxemic study of HRI.

are common knowledge for people—to the point where they may not even be aware of them anymore—often turn out to be not so trivial to incorporate into robot behavior. For instance, people will unconsciously and effortlessly adjust the distance to their conversation partner to an appropriate amount; however, a robot would need to conduct a careful computation to decide what distance it should keep during an interaction with its human counterpart. Even more difficulties are involved when the interaction is more complex, for example, when a robot has to approach a person, when it has to maintain spatial formation during a conversation, or when it has to navigate together with a person on the move. These considerations are important not only for achieving socially acceptable and comfortable HRI but also for ensuring that people understand the robot's intentions and can engage with robots safely in their physical space.

Questions for you to think about:

- Let's role play: To understand how much social information is involved in creating socially appropriate navigation, try to behave like a dumb robot that does not process any social information about space when interacting with a friend (maybe inform your friend beforehand, or "forget" to do so for a more natural response). What happened? How long could you keep this up?
- Think back on a situation when somebody violated your personal space. How did you notice? What was your reaction?

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

- Imagine you are an engineer building a robot. This robot will come to the market in Japan, Mexico, and the United States. Will the product be the same for every country? Will the robot's spatial navigation behaviors differ? If so, how?
- Think about the use of a robot in various daily situations (e.g., at home, at the office, and on a crowded train). Now, think about how you need to adapt the spatial-navigation behavior of the robot to fit each of these contexts. What would be important factors to consider in these different contexts?

Future reading:
Textbook to learn basic techniques for robot navigation:

- Howie M. Choset, Seth Hutchinson, Kevin M. Lynch, George Kantor, Wolfram Burgard, Lydia E. Kavraki, and Sebastian Thrun. *Principles of robot motion: Theory, algorithms, and implementation.* MIT Press, Cambridge, MA, 2005. ISBN 978-026203327. URL `http://www.worldcat.org/oclc/762070740`

More reading about space-related studies in HRI:

- Thibault Kruse, Amit Kumar Pandey, Rachid Alami, and Alexandra Kirsch. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems*, 61(12):1726–1743, 2013. doi: 10.1016/j.robot.2013.05.007. URL `https://doi.org/10.1016/j.robot.2013.05.007`
- Jonathan Mumm and Bilge Mutlu. Human-robot proxemics: Physical and psychological distancing in human-robot interaction. In *Proceedings of the 2011 ACM/IEEE International Conference on Human-Robot Interaction*, pages 331–338. ACM, 2011. ISBN 978-1-4503-0561-7. doi: 10.1145/1957656.1957786. URL `https://dl.acm.org/citation.cfm?doid=1957656.1957786`
- Satoru Satake, Takayuki Kanda, Dylan F. Glas, Michita Imai, Hiroshi Ishiguro, and Norihiro Hagita. How to approach humans? Strategies for social robots to initiate interaction. In *4th ACM/IEEE International Conference on Human-Robot Interaction*, pages 109–116. IEEE, 2009. ISBN 978-1-60558-404-1. doi: 10.1145/1514095.1514117. URL `https://doi.org/10.1145/1514095.1514117`
- Michael L. Walters, Kerstin Dautenhahn, René Te Boekhorst, Kheng Lee Koay, Dag Sverre Syrdal, and Chrystopher L. Nehaniv. An empirical framework for human-robot proxemics. *Proceedings of New Frontiers in Human-Robot Interaction*, 2009. URL `http://hdl.handle.net/2299/9670`

6

Nonverbal Interaction

What is covered in this chapter:

- The role of nonverbal communication in interactions between people—how communication is enhanced by facial expressions, hand gestures, body posture, and sounds;
- The importance of interpreting, using, and responding to nonverbal cues in the appropriate way, both to successful human–robot interactions and to generate a positive perception of robots;
- Nonverbal communication channels that are unique to robots, as well as channels that replicate those commonly used by humans;
- How robotic sounds, lights, and colors or physical gestures with arms, legs, tails, ears, and other body parts can be effective for communicating with people.

When we think of what it means to communicate with someone face to face, the first thing that comes to mind is often the content of our speech—what we are saying to each other—rather than the manner in which such content is delivered. Just for a minute, though, imagine speaking face to face with someone without the ability to look at the person or to use gestures. Not only would you be uncomfortable, but you might also have difficulty getting the intended meaning across. Moreover, without the nonverbal "channel," it seems harder to establish a strong connection with the person, particularly when you communicate with a stranger.

This is because people constantly and seemingly automatically pick up on a variety of nonverbal cues while interacting. These cues are used to interpret the nuances of meaning, emotion, and intention in others. Nonverbal cues are such an important aspect of human communication that being unable to produce and decipher them appropriately makes interaction quite challenging. Anyone may experience a sense of bewilderment when they go to another country—we may find it difficult to summon the waiter to give us the bill or might struggle to read another person's face in order to understand what he or she is feeling. The importance of nonverbal cues is acutely experienced by people with disorders such as autism, who have difficulty noticing and interpreting nonverbal social cues in others. On the other hand, being sensitive to nonverbal cues can improve one's understanding of an interaction. For example, researchers who have used "social sensors" to measure aspects of nonverbal behavior, such as gaze and rhythmicity, can predict which people will exchange cards at a conference (Pentland and Heibeck, 2010) or which couples will break up within a six-year period, based on thin slices of nonverbal behavior (Carrere and Gottman, 1999).

Even in the earliest social robot designs, nonverbal cues that are present in human interaction have been actively used to enrich interactions with the robot. They are typically used in combination with speech to provide supplemental information on the robot's internal state or intentions. Kismet, one of the first social robots, used postural cues, such as pulling back or leaning forward, to express affect and engage people in interaction (Breazeal, 2003). Keepon, a minimalist social robot, uses gaze and reactive motion to express attention and affect (Kozima et al., 2009). Many robots are also capable of engaging in joint attention to signal engagement with the user and a shared task. Next, we discuss the functions and types of nonverbal cues and their uses in HRI.

6.1 Functions of nonverbal cues in interaction

Nonverbal cues allow people to communicate important information "between the lines." They add a further layer of information to human (and human–robot) interaction, adding to what is being communicated linguistically. Through nonverbal communication, people can signal mutual understanding, shared goals, and common ground. They can communicate thoughts, emotions, and attention. And they can do so in a more subtle, indirect manner than through verbal expression.

In psychology, nonverbal communicative cues, such as eye gaze, body posture, or facial muscle activity, are often studied as implicit indicators of affect toward a person or an object. Many of the nonverbal signals we convey are expressed automatically without much thought or are even entirely unconsciously. Therefore, nonverbal cues are often believed to be unfiltered and more genuine, revealing people's "true" attitudes. For instance, your body language can communicate a message very different from your speech. Think of an acquaintance you do not like very much. Although you might greet this person in a friendly manner and start a seemingly friendly chat, your nonverbal cues might give away your true feelings. You might look at the person more briefly, frown rather than smile, and avoid physical contact while not even being aware that your nonverbal cues are incongruent with your verbal chitchat.

Nonverbal cues are equally important for human–robot interaction

(HRI). Nonverbal cues produced by people when interacting with a robot can indicate whether a person is enjoying the interaction and whether the person likes the robot or not. They can therefore act as a measure or cue of attitude or engagement and be used to guide the robot's behavior. Even in the HRI context, verbal and nonverbal cues might be contradictory. For example, people may verbally express positive ideas about a robot while the nonverbal cues suggest they are tense or anxious while interacting with it. HRI may also be affected by the way robots produce nonverbal cues. For example, and interaction can appear awkward when the robot produces gestures that do not match the rhythm or meaning of its speech or when it does not respond appropriately to people's nonverbal cues. Early research on HRI focused mainly on speech as the most obvious mode of communication for robots, but researchers now agree that nonverbal cues are central to HRI, and their implementation is widely accepted as a prerequisite for smooth and successful interaction between humans and robots. To illustrate, think of human eye gaze during a conversation. Eye gaze occurs automatically, without much thought, but at the same time, it signals shared attention—that both people are talking about the same thing—and acknowledges the conversation partner. When speaking to a robot, we would expect the robot to turn its head toward us and make eye contact with us, showing that it is attending to what we say. A robot that displays such nonverbal behavior will make the interaction seem more natural and smooth. Conversely, we notice immediately when some of this "social glue" is absent—we can sense that something is going wrong, even though it might be difficult to pinpoint exactly what is missing. When the robot stares straight ahead and does not acknowledge our presence or spoken requests, the interaction breaks down.

As with all information, nonverbal communication always occurs in a specific context, which renders the respective nonverbal signals appropriate or not. This context may be restricted by specific social and cultural norms. For example, in Western societies, people shake hands to greet each other formally, whereas a respectful greeting in Japan is performed by bowing. Even the degree to which one person bows to another signals social status and hierarchy. This might be almost imperceptible to the naive observer, but it is immediately obvious to those who understand the relevant cultural norms. Similarly, a conversation with a person from a Western society would naturally include continuous eye contact or even physical touch. However, this might be interpreted as threatening or rude in another cultural context. Such social and cultural differences are being taken up in recent HRI research on designing culturally sensitive interactions, investigating, among other issues, the importance of nonverbal cues for the cross-cultural deployment of social robots. For example, researchers from the United Kingdom and Japan

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

are working together to develop culturally competent care robots, which includes developing cultural knowledge representations, culturally sensitive planning and execution, and culturally appropriate multimodal HRI (Bruno et al., 2017). Designing HRI that meets social norms and cultural expectations might mean the difference between a successful product and a wasted investment.

Figure 6.1 Culturally appropriate nonverbal cues can make communication between people and robots more natural and pleasant.

6.2 Types of nonverbal interaction

Although we exhibit and experience nonverbal cues in several modalities at once, such as sound, movement, and gaze, it might be worthwhile to consider each channel of communication separately when trying to implement nonverbal signals into HRI. Understanding the functions and effects of various nonverbal cues allows us to then combine them as needed for different tasks and effects in HRI.

6.2.1 Gaze and eye movement

Imagine you are conducting a job interview and the job candidate responds to your inquiries without looking at you, staring only at the desk in front of him or her. Even while you are sketching a graph on the whiteboard, the job candidate does not follow your gaze toward what you are drawing. Would you hire the person? Probably not, because this type of gaze behavior would likely come across as a lack of interest in you and what you are talking about.

Gaze is a subtle and important cue for managing social interaction. Gaze signals interest, understanding, attention, and people's ability and willingness to follow the conversation. Beyond their social function, gaze and eye movements also facilitate functional interactions and collaboration, such as handing an object to someone or calling someone's attention to the next tool needed in a task. Using eye-tracking methodology to assess gaze patterns can provide insights into information processing and human cognition. Pragmatically, analyzing gaze patterns can also help to ensure that a given task has been completed smoothly. Gaze can also be a way of soliciting and keeping another person's attention during an interaction. For instance, gaze can be a way to manage turn-taking in interactions; by looking from one person to another, the speaker might suggest whose turn it is to speak next.

A particularly well-established component of gaze behavior in human interaction is *joint attention*. Joint attention refers to interaction partners attending to the same area or object at the same time. The significance of this behavior for human development starts in early childhood, when joint attention is a major scaffold for learning. The ability to attend to the same object at the same time with an adult caregiver is an important prerequisite for infants' ability to learn new words and behaviors (Yu and Smith, 2013), whereas the inability to perform joint attention can lead to developmental difficulties (Charman et al., 2000). Joint attention in adult communication can also signify interest and deep involvement in the interaction and is important for collaborative tasks where actors need to coordinate their activities. To achieve joint attention, the timing and synchrony of gaze behavior are important aspects to consider.

Figure 6.2 Pupils signal attraction, even in robots.

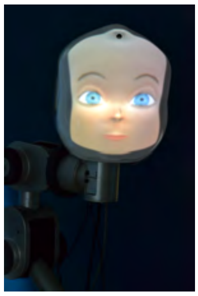

Eyes are a window to the soul, or in this case, they unconsciously reveal how much you like your interaction partner. Pupil dilation is controlled by the autonomic nervous system, as are uncontrollable reactions such as an increase in heart rate or goose bumps. When people see physically attractive others, their pupils automatically dilate. This also works the other way: people judge faces with larger pupils as more attractive than those with more visible irises. This can be used on robots to give the impression that the robot is attracted to the user (see Figure 6.2).

Joint attention has been incorporated into HRI in several ways: Imai et al. (2003) used it as a way of scaffolding smoother communication with people so that they know what the robot is talking about, both in conjunction with and without speech. Joint attention has also been studied as a fundamental capability of robots that want to learn from humans, particularly humanoid robots (Scassellati, 1999). Finally, joint attention with robots has been studied in interactions with children who have autism, with the aim of using the robot to assist them in developing this important social skill. It is, however, still unclear whether individuals with autism who were trained to use social skills, such as performing joint attention, with robots are able to apply these skills in human–human interaction as well (Robins et al., 2004).

When used in HRI, robot gaze cues most often produce similar effects as they would in human interactions. This may be because researchers have used human gaze behavior to derive models of gaze behavior for robots, and they have shown that the resulting gaze cues can be used to lead people to take on different conversational roles as addressees, bystanders, or nonparticipants (Mutlu et al., 2012). In a multiparty interaction, a robot can use its gaze to control who will be the next person to talk (Mutlu et al., 2009). Andrist et al. (2014) used face-tracking movements to engage in mutual gaze and purposeful gaze aversions in an HRI study to show that such cues can make a robot seem more intentional and thoughtful. Mutlu et al. (2006) also showed that a robot's gaze cues, modeled on those of humans, used in the course of telling a story affected how well people remembered the story's content; the people with whom the robot kept gaze contact could recall more details

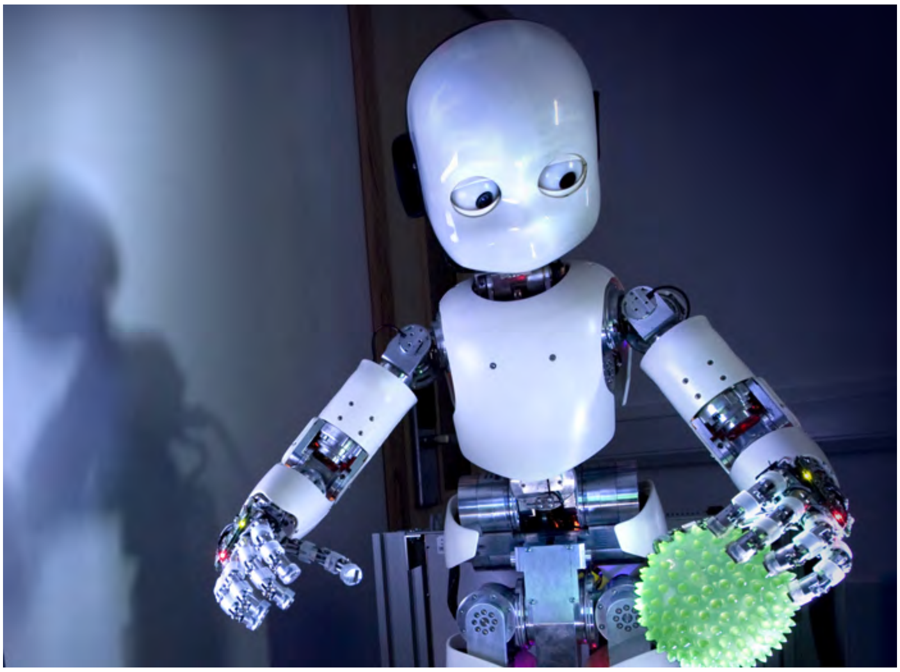

Figure 6.3 The eyes of robots are often designed to pitch and yaw, allowing a robot to use gaze as an effective communication channel. Here, iCub (2004–present) gives a good impression of attending to the ball in its left hand.

from its story. Robot gaze can therefore be a powerful way to manage interactions with one or more people.

6.2.2 Gesture

Following speech, gesturing is perhaps the most apparent way of providing information during an interaction. Gestures can function in place of or along with speech and are often categorized based on their role in communication. Deictic gestures refer to pointing to specific things in the environment and can be important for establishing joint attention. Iconic gestures often go along with speech, further supporting and illustrating what is being said. For example, opening your arms wide while saying you are holding a big ball would be an iconic gesture, as would smoothly moving your hand upward while explaining how your airplane took off. Symbolic gestures, such as waving for hello or goodbye, can carry their own meaning, with or without accompanying speech. Finally, beat gestures are used to go along with the rhythm of speech and look like moving one's arms while speaking as if conducting an invisible orchestra (see Figure 6.4). Gestures can also be used for emphasizing particular moments during speech, such as lifting your hands up while saying "what?" when you are surprised by something.

Gestures are likewise a powerful way of enhancing spoken communication in HRI. A robot may be designed to gesture through its arms and hands or other body parts, such as its head, ears, or tail. The

Figure 6.4 A Pepper (2014–present) robot using hand gestures to accompany its speech. Without these automatically generated beat gestures, the speech would appear less natural.

shape, timing, naturalness, and smoothness of gestures can also affect people's perceptions and understanding (Bremner et al., 2009). Salem et al. (2013) showed that including gestures along with speech in HRI led to the ASIMO robot used in their experiment being perceived as more anthropomorphic and likable, with participants expressing greater willingness to interact with the robot later on than when the robot communicated through speech alone. Interestingly, this study also showed that the use of gestures performed incongruently with speech led to even more pronounced positive effects in evaluations of the robot, although it had a negative effect on task performance. Gestures should therefore be used carefully in the design of robots, and their effects should be tested in studies with humans to gauge their effects on specific interactions.

6.2.3 Mimicry and Imitation

Another aspect of nonverbal interaction that has been given much attention in the human-interaction literature is mimicry and imitation. By mimicry, we mean the unconscious replication of the behavior of another person, and by imitation, we mean the conscious replication of another's behavior (Genschow et al., 2017). Mimicry and imitation are performed not only by humans but also by primates (hence the notion of "aping someone") and are considered basic social capabilities.

> Researchers in Japan found a band of macaques that all wash their sweet potatoes in a stream. This behavior was traced to a female member of the troop, who may have initially done this once by accident, and then others copied her when they realized that washing the potatoes produced a less gritty and more pleasing meal, and so they continued the practice. Observations of this kind have led to the claim that animals, not only humans, have "culture" (Whiten et al., 1999; De Waal, 2001).

In humans, mimicry and imitation have multiple developmental func-

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

tions. In early childhood development, mimicry and imitation provide a common way to learn new behaviors and culturally relevant social norms. Children use mimicry to learn to do things in particular ways—such as talking with a British accent or making expressions similar to those of a family member. As adults, we can also use imitation to blend into our social and cultural surroundings, such as gesturing more emphatically when we are speaking Italian or visiting Italy. As such, imitation and mimicry can be important ways of developing signs of in-group identity.

Mimicry, as a largely automatic behavioral response, also has many significant social functions: one is that it indirectly signals positive affect and liking for an interaction partner. If two people use the same gestures or adopt the same posture during a conversation, it is usually because they have established a positive relation in that interaction. Similarly, when people's nonverbal cues are out of sync and not mirroring each other, you can sense that the communication is not running smoothly. Mimicry as a subtle nonverbal cue can thus be a helpful signal to interpret, for instance, in the context of dating or job interviews.

Mimicry's significance in establishing a social relationship with another person makes it possible for its manipulation to function as a tool for persuasion. In studies of the "chameleon effect," Chartrand and Bargh (1999) found that subtle mimicry of a person's gestures and posture can help that person persuade an interaction partner to agree with his or her suggestions. For example, if you sit with your right leg crossed over your left, and your interaction partner subtly adopts that position, too, before telling you that Candy A tastes better than Candy B, you are more likely to choose to try Candy A over Candy B than if the person had not mimicked your posture (see Figure 6.5). However, this effect is time dependent. If you notice your conversation partner mimicking you, either because they are too obvious about it or too late in their timing, their intentions will backfire because you may see them as manipulative or insincere.

Various aspects of imitation and mimicry have been implemented and evaluated in the design of robots. There is a large and growing collection of literature on robot learning by imitation, in which robots in some way record and then reproduce actions performed by humans (Argall et al., 2009). Riek et al. (2010) developed an ape-like robot that mimicked users' head gestures, and their findings suggest this made a positive contribution to people's interactions with the robot, although these gestures were not always clear to participants. If we combine what we know about mimicry (see Section 6.2.3) and posture from human psychology, we can design robots that are able to display certain types

Figure 6.5 Similar to a chameleon adjusting its color to the environment, the chameleon effect refers to mimicking a person's gestures to be more persuasive.

of behaviors (e.g., leaning in) to affect how people behave and, therefore, how they feel. For example, Wills et al. (2016) showed that a robot that mimicked people's facial expressions and displayed socially contingent head poses received more monetary donations than a robot that did not display such behavior. Imitation and mimicry can therefore be used as both conscious and unconscious social cues in HRI to improve interaction and persuade people to follow the robot's suggestions.

6.2.4 Touch

Touch is a nonverbal cue that is often involved in close interactions among people, such as those between friends or between caregivers and patients. We often use touch deliberately to calm down someone who is agitated or to console someone who is sad. We also often incidentally touch people we feel attracted to or whom we like. It turns out that these people often also like us more when this happens. Both deliberate and incidental touch can therefore have beneficial effects, particularly when the interaction partners are part of the same social group. It is important, however, to know when and how it is appropriate to touch someone.

In everyday life, touch is sometimes used deliberately to achieve a goal. According to the so-called Midas effect, waiters and waitresses get a higher tip if they happen to incidentally touch the customers before they pay for their meal (Crusco and Wetzel, 1984). Touch does not always have positive effects, however, particularly when people who identify with different social groups are interacting with each other.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

In this case, touch may even lead to more negative feelings about the interaction partner. Incidental touch has also been shown to lead to a reduction of more indirect, but not direct, forms of prejudice against an out-group (Seger et al., 2014). Results on the effects of touch between human groups are therefore mixed, and it is interesting to consider what role touch might play in interactions between humans and robots, which may represent a new social group in the society of the future.

Figure 6.6 Telenoid (2010–2013) is a haptic robot that is designed to be hugged. Studies on whether this is a form of interaction people are comfortable with are ongoing. (Source: Hiroshi Ishiguro)

The few studies on touch in HRI that are available in the literature demonstrate the need for more empirical work on this nonverbal cue (Van Erp and Toet, 2013; Willemse et al., 2016). On the positive side, tactile interaction with animal-like robots, such as Paro or the Haptic Creature, show that people can feel less stressed and anxious when they initiate such interactions (Shibata, 2012; Yohanan and MacLean, 2012). Chen et al. (2014) showed that people did not mind being touched by a robot in a nursing scenario, but they evaluated functional touch (e.g., to clean their arm) more positively than affective touch (e.g., to comfort them). In contrast, a recent study by Wullenkord et al. (2016) explored the negative consequences of touch in an interaction with the robot Nao. Participants reported their attitudes toward a Nao robot, then had to touch the robot as part of a task. After the task, they reported their attitudes and social judgments about the robot again. Overall, contact improved the participants' attitudes, such that people expressed more positive and less negative attitudes after the touch interaction as compared to one without touch. However, people who had particularly negative emotions toward robots at the onset of the study experienced the opposite effect and had more negative perceptions after they touched the robot.

Touch is an integral part of natural human–robot interactions, for example, in functional tasks such as object handovers and manipulation and in social tasks such as a handshake for greeting. In both functional and social uses, we need to keep in mind the psychological implications of incidental or deliberate touch, whether it is being touched by a robot or having to touch a robot.

.

6.2.5 Posture and movement

People also communicate through their full body posture and the way in which they move. Along with facial expression, postures can be used to interpret a person's emotional state. Slow movements, drooping shoulders, and lethargic gestures all suggest a downcast state of mind, whereas fast movements and an upright bearing are signs of a positive attitude. These types of postural cues are particularly important when a person's face is not visible, but they can also provide additional cues to a person's state of mind even when we can see the person's fa-

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

cial expression. Researchers have found that people can interpret these types of nonverbal cues not only when they see the whole body of the person but also in minimalist light dot displays that depict a person's movements (Alaerts et al., 2011).

The Thrifty Faucet (2009) is a simple interactive prototype that uses its posture to communicate 15 lifelike motion patterns, including seeking, curiosity, and rejection, to users. The aim is to enable communication with users about more sustainable water use (Togler et al., 2009).

(Source: Jonas Togler)

The way we pose can signal attention, engagement, and attraction in an interaction between humans. People might be displaying a defensive posture by holding their arms in front of them, whereas open arms are a clear invitation for engagement, perhaps even a hug. How we are posed in relation to other people can also provide valuable information; if two people are sitting with their knees toward each other, it shows willing engagement, whereas if one person is turned partly away from the other, it can show a desire to discontinue the interaction.

Bodily postures can provide an additional layer of expressiveness to robots. To illustrate, when a robot lacks expressive facial features, the body can be used as the primary way to communicate emotions. Beck et al. (2010) showed that affective body postures can improve people's understanding of a robot's emotional state. A robot's posture can be used to express emotion and, through that, impact the emotions of onlookers. Xu et al. (2014) showed that people were not only able to interpret the affective body postures of robots, but also that they adopted the emotions they thought the robots were showing.

Robot designers have also realized that micromovements, barely perceptible motions, can convey the impression that the robot is lifelike (Yamaoka et al., 2005; Ishiguro, 2007; Sakamoto et al., 2007). These micromovements are often implemented as small, random perturbations to the robot's actuators. Such lifelike animations can also be used to

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

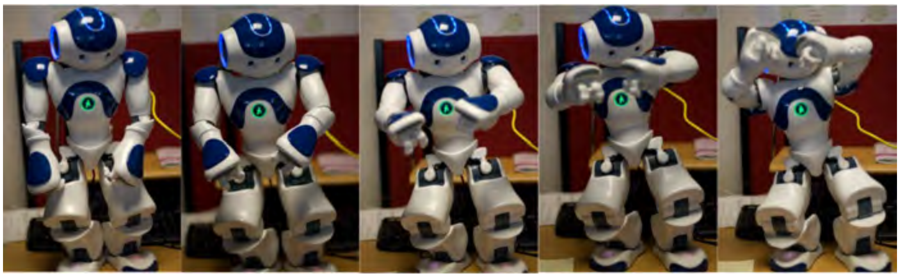

Figure 6.7 A Nao robot (2008–present) using body postures to express emotions, morphing between sad (left) and fearful (right). (Source: Beck et al. (2010))

signal the robot's internal state, for example, the velocity or amplitude of the motion signals the excitement level of the robot (Belpaeme et al., 2012). This approach has been successfully used on petlike small robots (Cooney et al., 2014; Singh and Young, 2012).

6.2.6 Interaction rhythm and timing

The temporal nature, or "timing," of communicative cues carries its own significance in interaction. In verbal communication, we refer to this as *turn-taking* among interaction partners. Nonverbal cues (e.g., gaze, gesture) can support this turn-taking by guiding attention to the appropriate interaction partner or signifying the end of a turn. Establishing synchronized temporal patterns of interaction can further scaffold the communicative and collaborative success of an interaction.

The "rhythmicity" and "synchrony" of an interaction provide a largely unconscious but crucial component of human communication. To understand what we mean by interaction rhythms, think about human interaction as a coupled system working together. In order for two people to be able to communicate and work effectively, they need to become "rhythmically entrained" to each other's actions—to be doing things not necessarily at the same time but to the same beat. Like in dance, rhythmicity enables people to be more attuned to each other's communicative cues, to be looking, speaking, and moving at the right time to enable clear and smooth communication among the two partners (Warner et al., 1987). Although often unconscious, the effects of rhythmicity on interaction are significant: being out of synchrony can imply that interaction partners have missed important social signals and are therefore unable to interpret each other's behavior; it can also lead to a more negative interaction outcome and to a less positive attitude toward the other person.

Michalowski et al. (2007) showed that a robot that is rhythmically entrained with a human interaction partner is considered more lifelike than a robot that is behaving rhythmically but is not synched with the human. They also showed that people are more likely to interact for a longer time with a rhythmically entrained dancing robot. Rhythmicity in interaction can also be useful in supporting turn-taking and collaboration in teams, including anticipation of people's behaviors and when they will do them (Hoffman and Breazeal, 2007). Finally, Siu et al. (2010) showed that listening to highly rhythmic music while performing robotic surgery can improve the performance of the human–robot surgical team. These findings suggest that rhythmicity in HRI can improve both the perceived quality of the interaction and the chances of a successful outcome.

6.3 Nonverbal interaction in robots

6.3.1 Robot perception of nonverbal cues

Standard pattern-recognition techniques are used to allow robots to perceive and identify human nonverbal cues. Posture and gesture recognition are well studied. Typical systems use cameras, depth cameras, or sensors carried by the user to record a time series of data. Although software could be written to recognize a limited number of gestures, it is instead typical to use machine learning as the system to be trained to recognize gestures and other nonverbal cues. To achieve this, a database is collected of, for example, people showing different gestures. Typically, thousands or even millions of data points are needed, and each needs to be labeled, meaning that for each data point, we need to note what it is showing. Is it a person waving, pointing, or beckoning? Next, a classifier is trained on the labeled data; this is often an iterative process, where the classifier's performance improves when more data are processed. Once the performance is sufficient for the application, the classifier is deployed on the robot (Mitra and Acharya, 2007).

These basic perception techniques are used to allow HRI researchers to estimate whether people are actually engaged in interactions with their robots. Unlike typical human interaction, where it is expected that the human partner will be attentive and engaged, in HRI, users sometimes do not attend to what the robot says and signals. Thus, perceiving the "engagement" of users is a crucial step for enabling robots to create a successful interaction. Rich et al. (2010) developed a technique to integrate the detection of cues such as eye contact and back-channeling to identify whether a user is engaged in interaction. Sanghvi et al. (2011) analyzed affective postures and body motion to detect engagement with a robotic game companion.

Although the constant advancement of technology allows for the improvement of robotic perception capabilities, researchers also add special equipment to the robot, such as eye trackers and motion-capture systems, to provide data on nonverbal cues relevant for interaction. For tactile interaction, there has been some research in the robotics field in which film-type piezoelectric polymer sensors were inserted in thin and thick silicone rubber (Taichi et al., 2006).

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

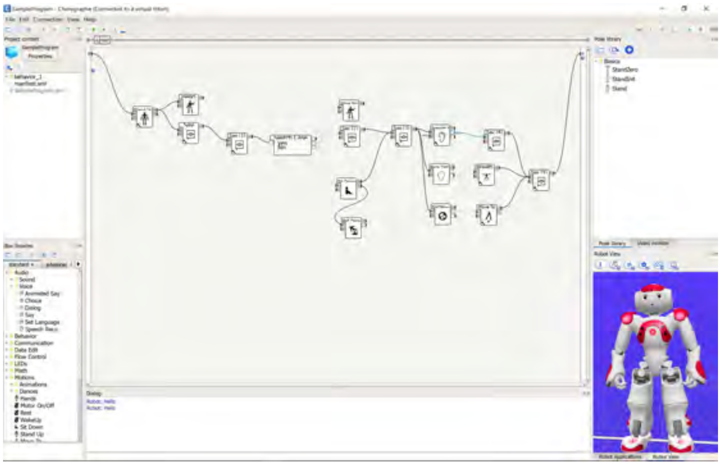

Figure 6.8 Choregraphe is a visual editor for the Nao and Pepper robots. It contains a pose editor that allows the robot designer to efficiently generate postures and animations for the robot.

6.3.2 Generating nonverbal cues in robots

Generating gestures and other nonverbal cues is not trivial in robots. The nonverbal cues need to be contingent on the interaction: if the user snaps her fingers, the robot needs to blink immediately. Nonverbal cues also need to be coordinated with each other and with other cues, including verbal interaction, both in terms of the semantic meaning and the timing of execution. HRI poses particular challenges for perception and generation of nonverbal cues because all this has to be done in real-time.

Animation framework

The most simple and most frequently used approach is to generate motions with an animation framework. That is, a robot designer will typically control each of the joint angles of a robot to set a posture for it; this is called a "key frame." After the designer prepares multiple key frames, the system interpolates the postures between them and generates smooth motions for the robot.

This requires extensive effort by the designer. Graphical user interfaces (GUIs) are often used to reduce the amount of effort in motion design. The commercial robots Nao and Pepper come with a GUI called Choregraphe, which helps designers visually display the posture of the robot and create desired motions more easily and quickly (see Figure 6.8).

Other techniques used for animation or virtual agents can also be used for generating motions for robots. Motion-capture systems can be used to record a timed series of precise human motions, which can then be replicated in robots. Robot designers have also taken advantage of markup languages for virtual agents, such as Behavior Markup Language (BML), in which a designer can specify which gesture an agent should exhibit in combination with speech (Kopp et al., 2006).

Cognitive mechanisms for robots

Another approach to achieving natural behavior in robots is to endow the robot with artificial cognition, which is an artificial equivalent of natural cognition. The expectation is that natural interaction behavior will emerge in the robot when it is controlled by artificial cognitive mechanisms. So instead of hacking the robot's nonverbal behavior, a constructivist approach is used. However, in order to construct a cognitive mechanism for a robot, researchers first need an understanding of how human cognition works.

> *Theory of Mind* is the ability to read desires, goals, and intentions in others. It is essential in understanding what others are thinking and what they are about to do. A typical example of Theory of Mind is the false-belief task.
>
> Imagine two people, Sally and Anne, in a room. The room has two boxes and a cake. Sally puts the cake in one box while Anne is watching. Anne leaves the room, and Sally switches the cake to the other box. When Anne comes back into the room, where will she look for the cake?
>
> Children typically develop the ability to give the correct answer to this type of problem at the age of 4 (Baron-Cohen et al., 1985). Robots still have a way to go.

Scassellati (2000) developed an embodied theory of mind architecture that takes into account salient objects, task constraints, and the attentional state of others to link the robot's perceptions of the world with high-level cognitive skills and related actions, such as joint attention, attribution of intent to others, and social learning. Sugiyama et al. (2007) developed a cognitive mechanism for a robot to replicate human deictic interaction. This involves using pointing (deictic) gestures in reference to a term, such as "this one" or "that one," that signifies a target object the listener can identify. The details of deictic interaction can also depend on the target. For example, we would not point directly at a nearby person because it is impolite. Liu et al. (2013) developed a computational model for a robot that balances two factors, understandability and social appropriateness. It enables a robot to refrain from exhibiting impolite pointing gestures while still keeping its deictic interaction understandable.

An important aspect of HRI design is generating nonverbal behaviors for robots that appropriately accompany speech. This is often inspired by the way humans use nonverbal cues in dialogue. Kanda et al. (2007a)'s robot system automatically generates nonverbal cues, such as nodding and synchronous arm motions, to exhibit its attentional state to the user in correspondence to the user's arm gestures. Robots

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

also benefit from displaying other nonverbal cues when using spoken dialogue, such as lip synchrony (Ishi et al., 2011).

6.4 Conclusion

This chapter highlighted the important role of nonverbal cues in communication between humans and robots. The implementation of nonverbal cues into the communicative repertoire of robots still calls for further technical advancement and refinement, particularly because nonverbal cues represent such subtle aspects of communication. Existing research illustrates the relevance of nonverbal communication in HRI while also making clear that much more work needs to be done before robots will be able to act and react in humanlike and natural ways in everyday communication with people.

Questions for you to think about:

- Still not convinced that nonverbal cues are important? Get up right now and have a conversation with someone, but do so without looking at the person's face. How did it go? How did you feel? Also, afterward, ask your communication partner what he or she thought about your behavior and how it made him or her feel.
- Think of a robot use case you are interested in. What aspect of nonverbal behavior is particularly relevant for this scenario? Would gesture or gaze be particularly helpful? How about contingency and timing? If you need some inspiration, you can go out and observe people in a similar context and see what they do.
- Have you ever watched a video where the audio track was a fraction of a second out of sync? Or video-conferenced with someone where the audio lagged? How did that affect the interaction? How long did you think the delay was? What, if anything, did you do to manage the difficulties in the interaction?
- How would you know if a robot is using is nonverbal cues effectively? Is there is a way in which you can measure the quality of the nonverbal interaction? Can you measure the outcome of the interaction?

Future reading:

- Henny Admoni and Brian Scassellati. Social eye gaze in human-robot interaction: A review. *Journal of Human-Robot Interaction*, 6(1):25–63, 2017. doi: 10.5898/JHRI.6.1.Admoni. URL `https://doi.org/10.5898/JHRI.6.1.Admoni`
- Cynthia Breazeal, Cory D. Kidd, Andrea Lockerd Thomaz, Guy

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Hoffman, and Matt Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 708–713. IEEE, 2005. ISBN 0-7803-8912-3. doi: 10.1109/IROS.2005.1545011. URL `https://doi.org/10.1109/IROS.2005.1545011`

- Nikolaos Mavridis. A review of verbal and non-verbal human–robot interactive communication. *Robotics and Autonomous Systems*, 63:22–35, 2015. ISSN 0921-8890. doi: 10.1016/j.robot.2014.09.031. URL `https://doi.org/10.1016/j.robot.2014.09.031`
- C. L. Nehaniv, K. Dautenhahn, J. Kubacki, M. Haegele, C. Parlitz, and R. Alami. A methodological approach relating the classification of gesture to identification of human intent in the context of human-robot interaction. In *IEEE International Workshop on Robot and Human Interactive Communication*, pages 371–377, 2005. ISBN 0780392744. doi: 10.1109/ROMAN.2005.1513807. URL `https://doi.org/10.1109/ROMAN.2005.1513807`
- Candace L. Sidner, Christopher Lee, Cory D. Kidd, Neal Lesh, and Charles Rich. Explorations in engagement for humans and robots. *Artificial Intelligence*, 166(1-2):140–164, 2005. doi: 10.1016/j.artint.2005.03.005. URL `https://doi.org/10.1016/j.artint.2005.03.005`

7

Verbal Interaction

What is covered in this chapter:

- The complexities and challenges of human verbal interaction;
- The components of speech in human and human–robot interaction (HRI);
- The basic principles of speech recognition and application to HRI;
- Dialogue management systems in HRI;
- Speech generation in HRI, including the use of chat bots.

Imagine you come across a robot at your local electronics shop. It says "Hello" as you approach and asks you what you are looking for today. You rattle off "Oh, I don't know, maybe a camera for my daughter, some batteries, and just looking around, you know." As you await a response, there's an extended silence from the robot. Then it repeats its initial question, asking you to speak slower and from a closer distance. Is the robot broken? You approach another one of the store's robots, with similar results. Why are conversations with robots so frustrating? (This did, in fact, happen to one of the authors.)

Speech is perhaps the most obvious mode of communication among humans because it is both audible and explicit. It is also a common mode of communication designed into robots, both in terms of the speech produced by the robots and speech as input for robots. However, producing robot speech is much simpler than understanding human speech, which creates an imbalance between people's expectations and the robot's actual capabilities. In this chapter, we describe the main components of human speech and then discuss the mechanisms by which a robot can be prepared for verbal interaction.

7.1 Human–human verbal interaction

In human communication, speech serves various functions: it is used simply to convey information, but equally importantly, it also serves to create joint attention and a shared reality through communication. In addition to being an inherent part of our nature, speech is incredi-

Figure 7.1 The difficulties these two Peppers in a store in Tokyo had in communicating with passersby could have been due to the noisy environment or the diverse ways in which people communicate verbally.

bly complex and open to multiple interpretations. By a mere twist of intonation or shift in emphasis, the meaning of the same sentence can switch dramatically. For example, try to pronounce the following sentence eight times while putting emphasis on the next word each time, starting with the first word of the sentence, "she":

She said she did not take his money.

By shifting the emphasis from one word to the next, what is inferred by the listener changes from a statement of belief (*She* said she did not take his money; apparently, someone else claimed otherwise) to disbelief (She *said* she didn't take the money, but someone actually has seen her doing it), to an accusation (She said *she* didn't [...], but someone else did), and so on.

Verbal communication is enriched by paralinguistic information as well, such as prosody and nonverbal behavior such as gaze, gestures, and facial expressions (see Chapter 6).

7.1.1 Components of speech

An *utterance* is the smallest unit in spoken language. Spoken language typically contains pauses between utterances, and an utterance is often less grammatically correct than a written sentence would be. This can become painfully clear when we read the transcript of a random

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

sentence from a conversation: whereas it takes no effort to understand what the person means when the person says it, the same sentence may appear incoherent when written down.

Spoken utterances can be short and consist of single words—such as *uhm*, *sure*, or *thanks*—or they can last for many minutes. Spoken language is often imperfect and has disfluencies, for example: "You know, I was, like, yeah, going to buy her, you know, something, but then I had, like, uhm, what the heck."

Words are the smallest units that we can utter to convey meaning. In turn, *phonemes* are small units of sound that make up words, "pat," for example, consists of three phonemes, [p] [a], and [t]. Changing a single one of them will change the meaning of the word; if the [p] is changed to a [b], we have a "bat."

Conversational fillers make up part of speech without directly relating to a specific concept. They serve to keep a conversation going. For example, people utter "uh-huh" while listening to indicate that they are attending to and following the conversation. Conversational fillers are an important part of human verbal communication because they allow listeners to signal a broad range of responses (e.g., they are paying attention, they understand what the speaker means, they are surprised at a sudden twist in the story, or they share an emotion) without disrupting the flow of conversation. Such feedback increases the efficiency of verbal communication tremendously, and it enhances the experience of a shared reality between the speaker and the listener.

7.1.2 Written text versus spoken language

Written text and spoken utterances are vastly different. Whereas people expect rather strict adherence to grammatical rules and syntax in written text, they become much more liberal when talking. Because of the unidirectional nature of written communication, a written text needs to be prepared with a certain level of precision and refinement because it cannot be adjusted while it is being delivered.

Verbal communication, on the other hand, allows for many ways in which one can clarify any misunderstandings or obscurities while one is delivering the message. People usually quickly detect when the interaction partner does not understand the message in the intended way, and in response, they change their speech on the fly.

Natural and humanlike communication that runs smoothly is often crucial for human–robot interaction (HRI). However, in order to build natural-language interaction, many technical prerequisites have to be in place. These include the robot's capability to transcribe speech into words, understand words by coming up with appropriate responses, and generate spoken language. The robot also often needs to be able to do this on the basis of verbal speech, which, as described previously, is more challenging than working with written text alone.

7.2 Speech recognition

Speech recognition is the recognition of spoken language by a computer and is also known as automated speech recognition (ASR) or speech-to-text (STT). Speech recognition is a process that takes a digital recording of speech and transcribes it. Speech recognition by itself does not understand or interpret what has been said. It merely converts a recorded fragment of speech into a written representation ready for further processing. Speech recognition has been mainly developed for controlling digital devices through spoken language or for dictation applications. Because of this, there is an assumption that the speech is recorded using a high-quality microphone, which is positioned close to the speaker in a relatively noise-free environment.

In HRI, these assumptions are often violated. When addressing a robot, the human conversation partner is often located at some distance from the robot, which has a negative impact on the quality of the recording. Signal processing and directional microphone arrays can alleviate this problem, but many robots are not equipped with such hardware. Due to the robot's microphone often not being located near the speaker's mouth, the microphone will also pick up sounds around the robot. Other people in the room talking, different sounds from the environment (e.g., a truck passing by outside, people walking around, or a cell phone ringing), and even mechanical noise from the robot itself all end up being recorded and provide a challenge for speech recognition. To avoid these problems, *close miking* is often used, where the user wears a lapel microphone or headset when talking to a robot.

The speech-recognition process requires a speech-recognition engine, software that has been trained to recognize one specific language. These are trained on thousands of hours of recorded and hand-transcribed speech, and they can handle only one language. Some speech-recognition engines are very specific and will only recognize brief commands or instructions specific to an application (e.g., recognizing spoken digits). Others are unconstrained and have been trained to recognize any possible spoken sentence. There are a few free, open-source speech-recognition engines, but the best-performing speech-recognition engines are commercial.

7.2.1 Basic principles of speech recognition

Speech recognition starts with a digital recording of speech, usually a recording of a single speaker. The recording is in the time domain: for every time step of the recording, for example, every 1/16,000th of a second, the sample contains the amplitude, or volume, of the recording. This is sufficient to replay the recording, but it is inconvenient for transcribing the speech into words. Thus, the recording is first converted to

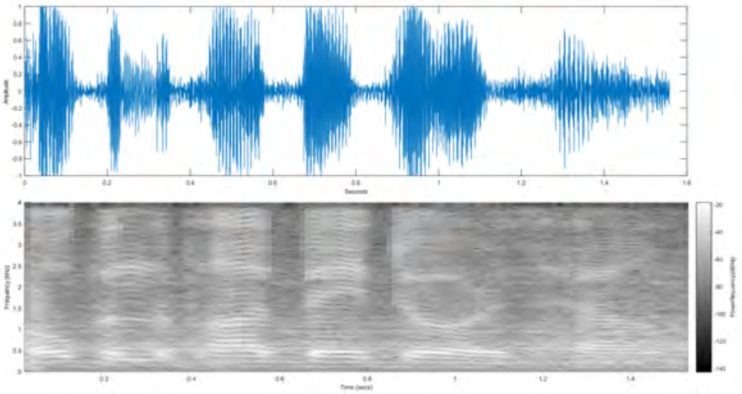

Figure 7.2 The speech sample "Open the pod bay doors, HAL" shown in the time and frequency domains. Speech recognition needs to transform these data into text.

the frequency domain. This means that it now shows how strongly certain frequencies are present in the signal at each time step. Phonemes look very different in the frequency domain—for example, an "o" has a different signature than an "a" in the frequency domain—and as such, they are easier to recognize with the use of an algorithm. Figure 7.2 shows a speech recording in both the time and frequency domains.

Until recently, speech-recognition engines used Gaussian mixture models and hidden Markov models to extract phonemes, words, and sentences from a speech recording. In essence, these approaches use probabilistic models of how phonemes and words can be strung together in words and sentences. The model knows that "robot" is a more likely transcription than "lobot" and that "the robot served the man" is more likely than "the robot swerved the nan."

In recent years, these probabilistic models have been replaced by deep neural networks (DNNs). These neural networks are similar in essence to the artificial neural networks that have been around since the 1960s, but their size is several magnitudes larger. A typical DNN can have hundreds of thousands of neurons and millions of connections between the neurons. Although these networks could not be trained in the past, new developments in algorithmic design and in computational hardware now allow the training of these networks to recognize spoken language relatively reliably. The performance of speech recognition using DNN has increased significantly compared to earlier methods. Not only has the rate of correctly recognized speech increased, but also speech-recognition engines can now increasingly deal with background noise, crowded environments, and ill-formed speech. They are also now speaker independent, meaning that the same speech-recognition model can deal with different speakers, including speakers of both genders.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

7.2.2 Limitations

All speech-recognition engines still struggle with recognizing atypical speech. Speakers on which the models have been insufficiently trained, such as young speakers (Kennedy et al., 2017) or elderly speakers, still provide a challenge. Also, the local dialects of languages or nonnative speakers will often result in severely reduced recognition performance. The acoustic environment also is a determining factor. Noisy, reverberating, or crowded spaces will decrease ASR performance. Proper nouns, such as Margaret or Launceston Street, are also likely to be incorrectly picked up by speech recognition.

Constraining what needs to be recognized could increase the performance of the speech engine. To do so, most ASR engines allow the programmer to set constraints on what should be recognized, for example, digits from 0 to 10 or simple commands. Although constrained ASR can handle atypical speech with some success, the current state of the art still does not allow spoken interactions with target persons from different backgrounds.

7.2.3 Practice in HRI

Numerous speech-recognition engines are available. Speech recognition using DNNs is, due to the computational resources needed to store and compute through the networks, usually available as a remote service. These cloud-based solutions allow you to send a recorded speech fragment over the internet, and the transcribed speech is returned soon after. Next to the best and most up-to-date performance offered by cloud-based services, cloud-based recognition also frees up computational resources on the robot, allowing the robot to have a relatively low-cost computational core. If the nature of the application does not allow the use of cloud-based ASR, for example, because the robot does not have a reliable, always-on internet connection, there are on-board speech-recognition solutions that use a reduced DNN or first-generation approaches to speech recognition. Their performance is, however, lower than that of the cloud-based services.

Many big software companies provide cloud-based speech-recognition services. Google, IBM, Microsoft, and Nuance all offer pay-per-use cloud speech recognition. Recognizing a single speech sample is often free for low-frequency use, but costs are on the order of 1 cent per recognition event. There are a few free open-source alternatives, such as the Mozilla Foundation's Common Voice initiative, which builds an open and publicly available data set of voices to train speech-enabled applications, and its DeepSpeech recognition engine.

Speech-recognition engines generally have a simple-to-use application programming interface (API), allowing the programmer to quickly integrate speech recognition on the robot. Next to the transcribed sentence, ASR engines will often also return a confidence value for the transcribed sentence, giving a measure of how confident the engine is about the recognized speech. Some engines will even return alternative transcriptions, again with confidence values.

7.2.4 Voice-activity detection

In some HRI applications, speech recognition is difficult due to the presence of noise, for example, because the robot is located in a crowded public space. Still, we can make a robot respond, albeit in a somewhat limited way, to people talking by using voice-activity detection.

Voice-activity detection (VAD) is often part of ASR, and it distinguishes speech from silence as well as other acoustic events. There is VAD software that can, for example, tell the difference between music playing and someone talking.

In HRI, VAD is used to give the user the impression that the robot is listening and can be used to implement spoken language turn-taking without actually recognizing or understanding the user's speech. In recent years, deep learning has also improved VAD performance. The free OpenSmile software package (Eyben et al., 2013) is currently leading in terms of performance. In combination with sound-source localization, whereby two or more microphones are used to pick up where a sound is coming from, we can even let the robot look at who is speaking.

7.2.5 Language understanding in HRI

A common misconception is that speech recognition also means that the speech is "understood" by the computer. It is not. Extracting semantic content from spoken language is particularly challenging, and a range of approaches exist that try to extract meaning from text, from broad semantic content to very specific content instructions.

Sentiment analysis, which matured as a way to analyze messages on social media, can be used to extract the affect contained in an utterance. Sentiment-analysis software often returns a scalar value denoting how negative or positive a message is. Although sentiment analysis is optimized for written language, in spoken language, we also have access to the way in which a message is delivered. Prosody and amplitude give us insight into the affect of the message: you do not need to speak the language to hear that the speaker is happy or agitated. Much in the same way, sentiment analysis and emotion from speech can roughly classify the affective state of the speaker.

More advanced methods, called natural-language understanding (NLU), will extract key words from the transcribed sentence. Elements such as

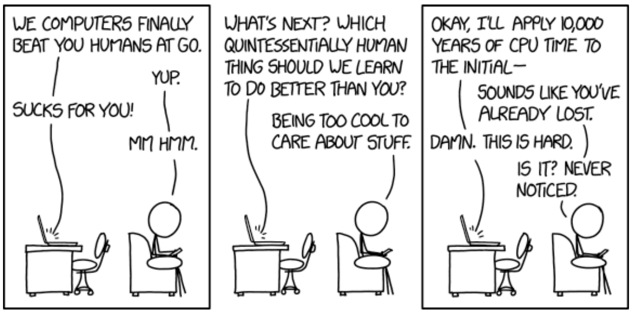

Figure 7.3 It will likely take many years before artificial intelligence can successfully detect sarcasm. (Source: XKCD)

a command, location, person, event, or date can be picked out from a message, allowing the software to respond appropriately. These methods are often used in personal digital assistants. "Remind me to pick up the kids at 7 p.m." is interpreted as a command to set a reminder at the time of 7 p.m. for the "pick up the kids" event. Those who frequently use digital assistants that extract information soon learn to conform to delivering information in ways that the computer can understand, and they settle on a manner of speaking that helps the computer perform satisfactorily.

Open-ended understanding of natural language is the biggest current challenge. Because conversational agents such as chat bots use conversational tricks to deal with ambiguous or poorly understood language, they might give the impression to the naive user that the agent understands what is being said. However, genuine understanding of natural language has not yet been achieved.

Even when people quickly realize that a robot does not understand figures of speech, they may still take some time to figure out how to phrase their message in such a way that the robot actually understands what they want. This is shown in the following dialogue that actually occurred between a human and a robot waiter, where the robot was programmed to take everything literally (Williams et al., 2018):

PARTICIPANT: Um, can I have one water?
ROBOT: Yes, that is permissible.
PARTICIPANT: Great.
ROBOT: Please tell me your order.
PARTICIPANT: Can I have one water?
ROBOT: Yes, that is permissible.
PARTICIPANT: Great, I'll take one water.
ROBOT: Thank you for sharing that interesting prediction. Please tell me your order.
PARTICIPANT: Can I order one water?
ROBOT: Yes, that is permissible.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

PARTICIPANT: Okay, can you bring me one water?
ROBOT: Yes, I am able to do that.
PARTICIPANT: May you bring me one water?
ROBOT: Yes, I am able to do that.
PARTICIPANT: I would like to have one water.
ROBOT: Thank you for sharing that interesting fact. Please tell me your order.
[*16-second pause*]
PARTICIPANT: Can you bring me one water?
ROBOT: Yes, I am able to do that. Please tell me your order.
PARTICIPANT: May you please bring me one water?
ROBOT: Yes, I am able to do that.
PARTICIPANT: So can you do it?
ROBOT: Yes. Please tell me your order.
PARTICIPANT: Can you go inside and get the water for me?
ROBOT: Yes, I am able to do that. Please tell me your order.
PARTICIPANT: My order is water.
ROBOT: Okay.
[*Robot retrieves order for participant*]

7.3 Dialogue management

Dialogue management (DM) is the process that keeps a conversation, either written or spoken, flowing between a computer and a user. Key to dialogue management is that the interaction consists of several turns, instead of a single instruction. A command such as "play Bob Marley" does not require DM, whereas "I would like to order something" requires the agent to further explore what it is the user wants, which could result in a whole chain of to-and-fro questions on the exact nature of the order. The management of such a dialogue requires the dialogue manager to keep track of the *state* of the conversation. This state can be specific to the application or can be very generic. If the robot is handling a pizza order, then the state consists of the details of the pizza order, such as which size pizza the user wants or what toppings go on the pizza. The dialogue manager will try to complete the unknowns in the state from the user's spoken utterances and will ask specific questions to fill in missing information before passing the order on.

7.3.1 Basic principle

Dialogue managers can range in complexity from systems that take the user through a very strict order of dialogue turns to systems that allow the user great liberty in the interaction. Restricted dialogue managers are used for closed and well-contextualized tasks, such as registering guests, filling out forms, or taking orders. System-initiative dialogue managers leave no room for the user to change the course of the dialogue: the agent asks a number of questions to which the user responds. User-initiative dialogue managers allow the user to take the lead, and the system only intervenes when something is unclear or when more information is needed. Mixed-initiative systems offer a combination of both.

Simple dialogue managers will allow the programmer to design a dialogue script, which, in its simplest form, implements a linear script. Conditions and branches, which take the dialogue in a different direction, allow more flexibility. In essence, these dialogue managers are finite-state machines (FSMs), a set of instructions that define a flow and often refer to a computer program. More advanced systems can accommodate event-based DM, allowing the control flow to be interrupted by an event. Event-based dialogue managers can handle nonlinear dialogue flows, such as the user asking for the current time in the middle of placing a pizza order.

QiChat is the dialogue manager software used for Softbank Robotics robots, such as the Nao and Pepper robots. Here is an example of QiChat code. It listens for sentences by the user, such as greetings ("hi," "hello," and "hey there") and orders (e.g., "Do you have lemonade?").

```
topic:~introduction ()
language:enu
concept:(greetings) ^rand[hi hello "hey there"]
concept:(wine) [red white] wine
concept:(alcohol) [beer ~wine]
u:(~greetings) ~greetings
u:(do you have _~drink) yes, I have $1
u:(I want to drink something) do you want ~alcohol?
```

An example of the dialogue it produces is as follows:

ROBOT: Hey there.
USER: Hello.
ROBOT: Do you have white wine?
USER: Yes, I have white wine
ROBOT: I want to drink something.
USER: Do you want beer?
ROBOT: I want to drink something.
USER: Do you want red wine?

Still more advanced dialogue managers will use a planner that cap-

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

tures the state of the system and actions that can be taken to change the state. The advantage of using a planner is that the programmer is no longer required to write a dialogue script that captures every way in which the state needs to be filled; rather, the planner can search which actions are still needed to complete the state. So, rather than explicitly writing the questions the robot needs to ask to complete a pizza order, the planner knows that the state of a pizza contains variables such as size, toppings, and delivery time and will find the actions, in this case questions, needed to complete any missing information in the pizza order.

7.3.2 Practice in HRI

There are several commercial solutions for DM; for example, companies that provide speech-recognition services will often provide DM together with speech production. Dialogue managers can range from very simple script-based services, allowing the programmer to implement linear linguistic interactions, to complex and rich services with planners. The most popular dialogue managers are event based because these have sufficient flexibility for most language-based commercial interactions. Dialogue managers, however, are not at all suitable to implement free-flowing and open conversation. Free linguistic conversation requires a large range of dialogue rules, and the dialogue script soon becomes unwieldy.

Turn-taking in HRI

Spoken dialogue with a robot will invite the user to take a more natural stance toward interaction, and as such, it might be necessary to introduce a number of factors that are also present in human interaction. One of those is *back-channeling*—the responses given by the listener during a conversation to signal that he or she is still engaged, such as “uh-huh” or “really?”. When your conversation partner is visible, there is often nonverbal back-channeling, such as a brief nod or a smile. In personal assistants, this often takes the form of a visual signal, such as a throbbing light, but on robots, these back-channeling signals can mimic human signals. The robot can use verbal back-channel signals, from the nonlexical “uh-huh” and “hmm” utterances to the phrasal and substantive utterances such as “yeah” and “tell me more.” The robot could augment these with signals, such as blinking lights or a gentle hum, to show that it is listening and paying attention. One of the problems in using back-channeling on robots is when to use a back-channeling signal because the timing is dependent on speaker verbal and nonverbal cues. For example, Park et al. (2017) showed that a robot using a back-channel prediction model that provided contingent back-channel signals was preferred by children.

The role of timing

Timing is critical in natural interaction: when a response is delayed, this is seen as disturbing, whereas a very quick response is often seen as insincere (Sacks et al., 1974; Heldner and Edlund, 2010). The timing of the response also depends on other factors. Increased cognitive load slows the response; yes/no answers have a faster response time than responses that require a fully formed reply (Walczyk et al., 2003). An analysis of telephone conversations showed that "yes" answers to a question take on average just 100 ms, whereas responses to undesired offers take on average almost 500 ms (Strömbergsson et al., 2013). A response given before the end of a question shows how human conversational partners anticipate questions and utter a response before the question is finished.

Computers are significantly slower than people in issuing dialogue responses. Due to the sequential processing chain in DM, a robot often needs several seconds before a response is formulated. Silences can be filled with conversational fillers or visual signals, signaling to the user that the robot is formulating a response. However, these are poor substitutes for prompt turn-taking, and considerable effort is being put into reducing the response delay in natural-language interaction. Just-in-time speech synthesis, where the robot starts speaking before having a plan of how to finish the sentence, seems promising, as does incremental spoken-dialogue processing, which works along the same principle as already-taken actions in response to spoken instructions before the instructions have been finished (Baumann and Schlangen, 2012).

7.4 Speech production

The final step in natural language interaction is converting a written response of the system into speech. For this, we need speech production, also known as speech synthesis or text-to-speech (TTS).

Speech production has seen impressive progress in recent years. In the 1990s, only voices that sounded tinny were available. Now, nearly 30 years later, we have artificial speech production that is almost indistinguishable from human speech. The two established methods for generating artificial speech are concatenative and parametric TTS. In concatenative speech production, an actor's voice has been recorded and cut into phonemes, and these are then "glued" together and smoothed at the seams to provide natural-sounding speech (Hunt and Black, 1996). In parametric TTS, a model is trained to produce acoustic speech parameters from text (Zen et al., 2009). Although concatenative models sound natural, they have little flexibility, and new voices require completely new recording and training of the TTS model. Parametric TTS is more flexible and allows for customization of the voice and prosody, at

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

the expense of naturalness. Recent advances have overcome these limitations by training generative deep neural networks (DNNs). Van den Oord et al. present a DNN model that produces speech that is virtually indistinguishable from human speech and even includes breathing and lip-smacking (van den Oord et al., 2016). This model has been adopted by Google as the voice of its digital assistant.

7.4.1 Practice in HRI

A wide selection of speech-production software is currently available, from free solutions to bespoke commercial software with voices customized to specific applications.

TTS engines

The simplest TTS engines have a small computational footprint and can run on cheap robot hardware. The most natural-sounding TTS engines use DNNs and are cloud based. Depending on the application, some TTS engines not only convert text into a speech file but also provide timing information for phonemes, which can be used to animate a robot. It might be necessary for the speech to be synchronized with facial animations or gestures on the robot, and timing information will allow for precise synchronization between the speech and the animations.

In HRI, it is important to consider which voice fits the robot and its application. A small robot requires a voice that matches its appearance, rather than a commanding baritone. In some cases, though, it might be important to match the sound of the voice to the fact that it emanates from a robot: a natural-sounding TTS engine might sit uneasily on an artificial agent. At the same time, research by Eyssel et al. (2012a) has shown that the type of voice affects the social perception of social robots. For example, robots with a male voice are anthropomorphized and evaluated more favorably by men than by women, and vice versa.

Some limitations to speech production still exist. Adaptive prosody and emotion, although actively being researched, are not commonly available on TTS engines. Also, synthesized voices do not adapt to the context in which they are being used. When the room is quiet, there is little need for the robot to have a booming voice, whereas a robot addressing a crowd at an exhibition would do well to adapt its rate of speech and volume to increase its intelligibility.

Chat bots

Chat bots are computer programs intended to converse with the user, typically by using written text. These systems are often implemented as a web application in which users enter text on a web page and the server responds to every text entry. These chat bots often have a specific

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

goal, such as to provide technical support or answer questions about the products of a company. Chat bots can become full sales agents or customer support agents. These agents are normally constrained regarding the topics to which they can respond in a meaningful way. More recently, chat bots have become speech-enabled. Chat bots such as Siri (Apple), Cortana (Microsoft), Alexa (Amazon), and Bixby (Samsung) now respond to simple spoken commands and respond with spoken text.

A second type of chat bot is general-purpose agents that try to respond to any utterance. They achieve this, on the one hand, by having thousands of hand-crafted rules on how to respond to often-occurring utterances and, on the other hand, by maintaining a database of all previous conversations, often learning from how users responded in the past to a given utterance. The ultimate goal is to create a chat bot that is indistinguishable from a human—users would no longer be able to tell if they are talking to a computer or a human. Controlled tests are being set up in an annual competition, and the most convincing chat bot receives the Loebner Prize. This test is often called a Turing test, named after the famous computer scientist Alan Turing, who proposed such a test as a measure for the intelligence of a computer (Turing, 1950).

The list of chat bots created by major information technology (IT) companies, such as Apple, Microsoft, Google, Amazon, and Facebook, indicates that there is already considerable interest in natural-language technology, and many companies make their technology available for developers. Google is offering its Cloud Speech application programming interface (API), Microsoft is pitching its Cognitive Services, and Amazon offers its Alexa set of tools to build voice-based services.

The availability of these services means it is no longer necessary to create your own software for speech recognition, understanding, or synthesis. Instead, developers can use online services for their robots. The audio signal recorded through the robot's microphone is streamed in real-time to the company's servers, and they send back the recognized text while the user is still talking. Similarly, these services can be used not only to recognize the spoken text but also to respond to the meaning of the text. The systems can, for example, identify entities, syntax, sentiments, and categories. This all helps the robot to better respond to the utterances of the users. These companies also offer speech-synthesis tools. The robot sends what it wants to say to a server and receives back the audio signal that the robot then plays on its loudspeakers.

It is much harder for a human to learn a new language than for a computer to do so. Still, artificial languages, such as Esperanto, have been developed to overcome some of the inherent problems

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

with learning natural languages. These constructed languages serve different purposes:

- Engineered languages—experimentation in logic, philosophy, or linguistics (Loglan, ROILA)
- Auxiliary languages—developed to help in the translations between natural languages (Esperanto)
- Artistic languages—created to enrich fictional worlds (Klingon, Elfish, or Dothraki)

The RObot Interaction LAnguage (ROILA) was developed for HRI, in particular to facilitate the problems that speech-recognition accuracy encounters (Stedeman et al., 2011). The words of this language have been designed to sound most distinct from each other, making it much easier for automated speech recognition to correctly identify the spoken words. "Go forward" in ROILA is "kanek koloke"; "go back" is "kanek nole."

7.5 Conclusion

Despite being the most obvious form of communication among humans, language is very complex, not only due to the large number of words people use daily but also because their meaning and significance change based on various contextual factors (e.g., relationships between speakers, task, prosody). Creating robots that can engage in this rich and diverse form of communication is a necessary goal for HRI, and technical tools available for speech analysis, synthesis, and production enable some degree of verbal HRI that does not need to be developed from scratch. Open-ended, natural-language conversation is still not possible, but verbal interaction in more constrained contexts can be successfully applied on robotic platforms.

Questions for you to think about:

- Imagine a social robot that needs to perceive all of the utterances you speak at your home every day, and think of a list of words (dictionary) for ASR. How long would this list need to be for the robot to be able to understand your everyday conversations?
- Consider the difference in how you say "yes" willingly versus reluctantly. How would you make a robot respond appropriately to such different modes of speaking?
- What are some problems that can emerge in relation to the important role of timing in human–robot interactions? How are these solved in other social interactions where the interactants

miss out on social cues (e.g., in texting, or when there is time delay on Skype calls)?

Future reading:

- Amir Aly and Adriana Tapus. A model for synthesizing a combined verbal and nonverbal behavior based on personality traits in human-robot interaction. In *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction*, HRI '13, pages 325–332, Piscataway, NJ, USA, 2013. IEEE Press. ISBN 978-1-4673-3055-8. doi: 10.1109/HRI.2013.6483606. URL `https://doi.org/10.1109/HRI.2013.6483606`
- J. Cassell, Joseph Sullivan, Scott Prevost, and Elizabeth Churchill. *Embodied conversational agents*. MIT Press, Cambridge, MA, 2000. ISBN 9780262032780. URL `http://www.worldcat.org/oclc/440727862`
- Friederike Eyssel, Dieta Kuchenbrandt, Frank Hegel, and Laura de Ruiter. Activating elicited agent knowledge: How robot and user features shape the perception of social robots. In *Robot and human interactive communication (RO-MAN)*, pages 851–857. IEEE, 2012b. doi: 10.1109/ROMAN.2012.6343858. URL `https://doi.org/10.1109/ROMAN.2012.6343858`
- Takayuki Kanda, Masahiro Shiomi, Zenta Miyashita, Hiroshi Ishiguro, and Norihiro Hagita. A communication robot in a shopping mall. *IEEE Transactions on Robotics*, 26(5):897–913, 2010. doi: 10.1109/TRO.2010.2062550. URL `https://doi.org/10.1109/TRO.2010.2062550`
- Nikolaos Mavridis. A review of verbal and non-verbal human–robot interactive communication. *Robotics and Autonomous Systems*, 63:22–35, 2015. ISSN 0921-8890. doi: 10.1016/j.robot.2014.09.031. URL `https://doi.org/10.1016/j.robot.2014.09.031`
- Michael L. Walters, Dag Sverre Syrdal, Kheng Lee Koay, Kerstin Dautenhahn, and R. Te Boekhorst. Human approach distances to a mechanical-looking robot with different robot voice styles. In *Robot and human interactive communication (RO-MAN)*, pages 707–712. IEEE, 2008. doi: 10.1109/ROMAN.2008.4600750. URL `https://doi.org/10.1109/ROMAN.2008.4600750`

8

Emotion

What is covered in this chapter:

- The difference between affect, emotions, and mood;
- What roles emotions play in interacting with other humans and robots;
- Basic models of emotions;
- The challenges in emotion processing.

How are you feeling right now? Happy? Bored? A bit self-conscious? Whatever the case may be, it's unlikely that you are feeling absolutely nothing. Various feeling states, and related emotions, are a key aspect of our day-to-day experience and of our interactions with other people. Emotions can motivate and modulate behavior and are a necessary component of human cognition and behavior. They can be spread through vicarious experience, such as watching a tense movie, and direct social interaction, such as seeing your best friend happy. Because emotions are such an integral part of human social cognition, they are also an important topic in human–robot interaction (HRI). Social robots are often designed to interpret human emotion, to express emotions, and at times, even to have some form of synthetic emotion driving their behavior. Although emotions are not implemented in each and every social robot, taking emotions into account in the design of a robot can help improve the intuitiveness of the human–robot interaction.

8.1 What are emotions, mood, and affect?

From an evolutionary perspective, emotions are necessary for survival because they help individuals respond to environmental factors that either promote or threaten survival (Lang et al., 1997). As such, they prepare the body for behavioral responses, help in decision-making, and facilitate interpersonal interaction. Emotions arise as an appraisal of different situations that people encounter (Gross, 2007; Lazarus, 1991). For example, when another person shoves us out of the way to be first in line, we get angry, and our bodies prepare for a potential conflict: the adrenaline makes us more prone to undertake action, and our expression

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

signals to the other person that he or she crossed a line. Conversely, upon finding out our friend did not invite us to his or her birthday party, sadness hampers quick action, forcing us to reconsider our prior behaviour (i.e., what did we do or say that may have offended him or her?) and evokes empathetic responses from others (Bonanno et al., 2008). In this way, emotions can also help us modulate the behaviors of others in an interaction.

Affect is used as a comprehensive term that encompasses the entire spectrum of emotionally laden responses, ranging from quick and subconscious responses caused by an external event to complex moods, such as love, that linger for longer (e.g., Lang et al., 1997; Bonanno et al., 2008; Beedie et al., 2005). Within affect, a distinction is made between emotions and moods (Beedie et al., 2005).

Emotions are usually seen as being caused by an identifiable source, such as an event or seeing emotions in other people. They are often externalized and directed at a specific object or person. For example, you experience happiness when getting a promotion at work, get angry when your phone's battery dies during an important call, or experience a pang of jealousy when a colleague gets a company car and you do not (Beedie et al., 2005). Emotions are also shorter-lived than moods (Gendolla, 2000). *Moods* are more diffuse and internal, often lack a clear cause and object (Ekkekakis, 2013; Russell and Barrett, 1999), and instead are the result of an interaction between environmental, incidental, and cognitive processes—such as the apprehensive mood while waiting a week to hear about the medical test results or the warm feeling of a sunny week spent in the company of friends.

8.1.1 Emotion and interaction

Emotions are not just internal; they are also a universal communication channel that has helped us communicate internal affective states to others and have likely been very important to our survival as a species.

Your emotions provide the outside world with information about your internal affective state, which is helpful to others in two ways. First, emotions convey information about you and your potential future actions. For example, displaying anger and frustration signals to others that you may be preparing for an aggressive response. In addition, emotions can convey information about the environment. An expression of fear may alert others around you of a fast-approaching grizzly bear before you have even found time to scream. In both scenarios, emotion provides an incentive for others to take action. In the case of anger, someone may choose to step down and attempt to suss the situation. In the case of fear, other people will likely scan the environment for a threat (Keltner and Kring, 1998). In this way, the successful communication of emotions promotes survival, enhances social bonds,

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

and minimizes the chances of social rejection and interpersonal physical aggression (Andersen and Guerrero, 1998).

8.2 Understanding human emotions

Since antiquity, people have given names to the numerous emotions we experience. Aristotle believed there to be 14 different emotions, including anger, love, and mildness. Ekman lists 15 basic emotions, including pride in achievement, relief, satisfaction, sensory pleasure, and shame Ekman (1999). It is impossible to provide a definite list of emotions because they vary between people and cultures, language does not offer a perfect mapping to emotions, and some emotions show overlap. Still, some emotions are likely to be considered more universal than others. Anger, sadness, and happiness are likely candidates for a set of core emotions. Ekman and Friesen (1975), in their seminal work on the facial expression of emotions, listed six basic facial expressions that are recognized across cultures. These facial expressions have often been mistaken for a set of basic emotions we experience, although they were only ever intended to describe a basic set of emotions that we express via our faces and that are recognized by different cultures.

Although many scholars distinguish between basic, or *primary*, emotions and reactive, or *secondary*, emotions, no consensus has been reached yet on which emotions are to be included in the first category and which should be considered secondary (Holm, 1999; Greenberg, 2008), and some scholars argue that basic emotions do not exist at all (see, e.g., Ortony and Turner, 1990). For those who do agree on the existence of basic emotions, primary emotions are considered to be universal across cultures (Stein and Oatley, 1992) and to be quick, gut-level responses (Greenberg, 2008) and include emotions such as amusement, anger, surprise, disgust, and fear. Secondary emotions, on the other hand, are reactive and reflective. They differ across cultures (Kemper, 1987). For example, pride, remorse, and guilt are secondary emotions.

But there have been challenges to the idea of emotions being distinct categories. Russell (1980) argued that emotions are the cognitive interpretations of sensations that are the product of two independent neurophysiological systems, namely, arousal and valence. As such, emotions are spread across a two-dimensional continuum rather than being composed of a set of discrete, independent basic emotions (Posner et al., 2005) (see Figure 8.4). This model has been widely studied and confirmed to hold across different languages and cultures (Russell et al., 1989; Larsen and Diener, 1992). However, a meta-analysis found that although the model makes for a reasonable representation of self-reported affect, not all affective states fall into the expected regions as predicted by the theory, and some cannot even be consistently ascribed to any

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

of the regions, suggesting that assumptions about the nature of some affective states may need to be revised (Remington et al., 2000).

8.3 When emotions go wrong

The importance of emotions in social interactions becomes especially clear when one partner fails to understand the emotion of the other partner or fails to respond with the proper emotion. Even tiny glitches in providing an adequate emotional response in social interaction can have serious consequences. For example, misinterpreting sarcasm for a genuine response can lead to misunderstandings in the conversation and hurt feelings. The situation becomes more problematic when someone is consistently unable to adequately perceive, express, or respond to affective states

Problems with emotional responsiveness are one of the defining symptoms of, for example, depression (Joormann and Gotlib, 2010). Although depressed individuals are able to understand the way others are feeling and can express their own emotional state, they have a reduced emotional response to positive stimuli, such as rewards (Pizzagalli et al., 2009), and have recurring negative thoughts about the past, present, and future. As a consequence, a depressed individual's patterns of social interaction often result in social isolation and even more loneliness, feeding into the individual's already frail psychological state.

Furthermore, people might be incapable of recognizing, expressing, and interpreting another person's emotions. For example, people with autism spectrum disorders may find it difficult to correctly interpret displays of emotion (Rutherford and Towns, 2008; Blair, 2005). This is clearly problematic for everyday social interactions because the affected person cannot intuitively understand the needs of his or her interaction partners and will often respond inappropriately.

Furthermore, people may have trouble expressing their emotional state, for example, when their facial muscles are impaired after a stroke. This makes it hard for their interaction partners to infer their internal states and form an idea of what they mean.

A person's inability to express and interpret emotions comes with serious consequences for the individual's capability to either provide or respond to emotional cues in an appropriate way. This, in turn, impairs the capability to interact with other people effectively and smoothly. Likewise, social interactions with robots may be difficult if the robotic counterpart is unable to express and interpret emotional states.

8.4 Emotions for robots

Emotions are considered an important communication channel in HRI. When a robot expresses emotion, people tend to ascribe a level of social agency to it (Breazeal, 2004a; Novikova and Watts, 2015). Even if a robot has not explicitly been designed to express emotions, users may still interpret the robot's behavior as if it had been motivated by emotional states. A robot that is not programmed to share, understand, or express emotions will thus run into problems when people interpret its behavior as disinterested, cold, or plain rude. Therefore, engineers and designers should consider what emotions the robot's design and behavior convey, whether and how a robot will interpret emotional input, and how it will respond.

8.4.1 Emotion interaction strategies

The most straightforward way of programming emotional responsiveness for social robots may be through mimicry. Mimicking in humans has been shown to create an idea of shared reality: you indicate that you fully understand the other person's situation, which creates closeness (Stel et al., 2008). The exception here might be anger—however good it may feel at first, responding to an angry person by yelling back usually does not facilitate mutual understanding or a resolution of the conflict.

A robot can use mimicry as a simple interaction strategy. It is a relatively simple response because it requires the robot "only" to be capable of recognizing an emotion in the human and then reflecting the emotion back in response. This already poses plenty of challenges, as will be discussed later in this chapter, but at least it cuts out the complicated task of formulating an appropriate response. Moreover, it may be a very basic expectation that humans have toward their interaction partners. Although we may excuse our friends for not knowing how to cheer us up when we are sad, we do expect (and appreciate) that they will respond to our sadness by lowering their brows and heads and becoming more soft-spoken.

One note that has to be made here concerns expectation management. When users perceive the robot to be emotionally responsive, they may extend this observation to expectations about the robot's compliance with other social norms. For example, a user may expect a robot to remember to ask about a confrontational meeting he was upset about the other night, so when the robot simply wishes him to "have a great day at work!" in the morning, he may be disappointed in the robot's social skills. Thus, the robot's emotional responsiveness should match its capability to fulfill other expectations.

8.4.2 Artificial perception of emotions

Robots need to register a wide variety of emotional cues, some explicit and some subtle, before being capable of emotional interaction. For instance, if we want to create a robot that responds emotionally when someone displays aggressive behavior, such as throwing an item at it, we need to integrate technologies for human behavior recognition and object recognition.

More specifically, we may want to create a robot that responds to human emotions. There are many studies on affect recognition (Gunes et al., 2011; Zeng et al., 2009). The most typical approach to recognizing or classifying emotions is to use computer vision to extract emotions from facial cues. Provided with a data set of human (frontal) faces with correctly labeled emotions, machine-learning systems, such as those using deep-learning techniques (LeCun et al., 2015), can extract features from the image to recognize a range of facial emotions. A famous example of this is smile recognition, which is broadly implemented in digital cameras nowadays. Affect recognition may also imply the interpretation of other visual cues, such as walking patterns, alleviating the need for a clear view of the user's face (Venture et al., 2014).

Many consumer-market digital cameras have a smile-detection feature. If a group poses in front of the camera, it will only take a shot when all the people in the frame smile. This technology partly replaces the timer function, which could never guarantee that everybody would look at the camera and smile at the time of the picture being taken.

Next to visual cues, human speech is perhaps the second most important channel to extract emotion from. In particular, prosody, the patterns of stress and intonation in spoken language, can be used to read the emotional state of the speaker. For instance, when people are happy, they tend to talk with a higher pitch. When sad, they tend to speak slowly and with a lower pitch. Researchers have developed pattern-recognition techniques (i.e., machine learning) to infer human emotions from speech (El Ayadi et al., 2011; Han et al., 2014).

Finally, a robot can sense human affect from other modalities. For instance, human skin conductance changes in response to an individual's affective state. A prominent example of the use of skin conductance as a measure is the polygraph or lie detector. However, skin-conductance sensors have been tried in HRI, with only limited success (Bethel et al., 2007).

Figure 8.1 Emotions expressed through mechanical facial expressions. Left: eMuu (2001). Middle: iCat (2005–2012). Right: Flobi (2010). (Source: Christoph Bartneck and University of Bielefeld)

8.4.3 Expressing emotions with robots

Typically, people design robots that convey emotions through facial expressions. The most common approach here is to mimic the way in which people display emotions. This is a good example of how the study of human behaviors can be used for designing robot behaviors. The facial expression of emotions has been well documented (Hjortsjö, 1969). Ekman's Facial Action Coding System (FACS), in which human facial muscles are grouped as action units (AUs), describes emotions as combinations of action units (Ekman and Friesen, 1978). For instance, when a person displays a happy face (i.e., smiling), the muscles involved are the *orbicularis oculi* and *pars orbitalis*, which raise the cheek (AU6), and the *zygomaticus major*, which raises the corners of the mouth (AU12).

Using a simplified equivalent of human facial muscles, researchers have developed robots that are capable of conveying emotions through facial expressions. For instance, a robotic face with soft rubber skin and 19 pneumatic actuators was developed by Hashimoto et al. (2013). This robot uses AUs to express facial emotions. For example, it activates actuators corresponding to AU6 and A12 to express happiness. There are many other robots designed to express emotion that rely on a simplified interpretation of human facial cues, including Kismet (Breazeal and Scassellati, 1999), Eddie (Sosnowski et al., 2006), iCat (van Breemen et al., 2005), and eMuu (Bartneck, 2002), among others (see Figure 8.1).

Robots can also express emotion through various humanlike modalities, such as body movements and prosody. But even nonanthropomorphic robots can express affect, by means of adjusting their navigational trajectories. For instance, research on a cleaning robot (Saerbeck and Bartneck, 2010) and a flying robot (Sharma et al., 2013) showed

Figure 8.2 Non-anthropomorphic robots can express emotion through their behavior or through the addition of expressive features, such as lights. Anki, the producer of Cozmo (2016–2019), describes its robot as "[having] his own lively personality, driven by powerful AI, and brought to life with complex facial expressions, a host of emotions and his own emotive language and soundtrack." (Source: Anki)

that they could display affect through adapting particular motion patterns. Some other ways in which nonanthropomorphic robots can express affect include speed of motion, body posture, sound, color, and orientation (see Figure 8.2) to the person they are interacting with (Bethel and Murphy, 2008).

8.4.4 Emotion models

Psychologists (Plutchik and Conte, 1997; Scherer, 1984) have attempted to capture human emotions in formal models. The benefit of this approach is that it views emotions as a numerical representation, which in turn lends itself well to representing emotion in computers and robots. These models also put different emotional categories in relation to each other, for example, by defining happiness as the polar opposite of sadness or by defining a distance function between emotions.

Emotion models are not only used to capture the emotional state of the user but can also be used to represent the emotional state of the robot itself and subsequently drive the behavior of the robot. For example, a robot with an almost-empty battery can act tired and announce it needs a rest. Once it has reached the charger, it needs to update its internal emotional state to happy. Expressing this emotional state allows the user to have access to the robot's internal state and will enrich the interaction.

A classic emotion model that has been used in some robots is the OCC model, named after its authors' initials (Ortony et al., 1988). This model specifies 22 emotion categories based on valenced reactions to situations, such as events and acts of agents (including oneself), or as reactions to attractive or unattractive objects (see Figure 8.3). It also offers a structure for the variables, such as the likelihood of an event or the familiarity of an object, which determines the intensity of the emotion types. It contains a sufficient level of complexity and detail to cover most situations an emotional robot might have to deal with.

Needless to say, many robots do not possess the ability to express all 22 emotions. Even if they could, implementing 22 different emo-

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

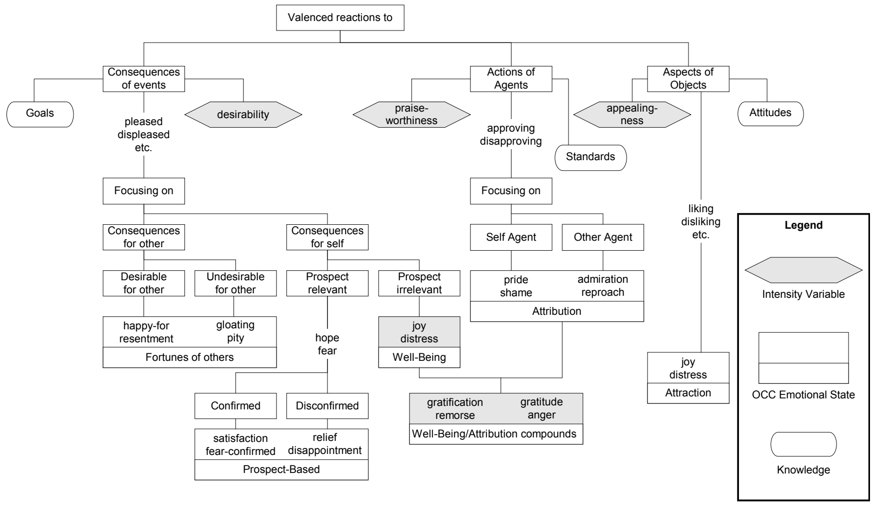

Figure 8.3 The OCC model of emotions.

tions can be challenging; hence, many robot designers prefer to reduce the number of categories. Often, a decision is made to implement only Ekman's six basic facial emotional expressions. These are reliably recognized, even across cultures Ekman (1992). However, a robot that only expresses six emotions makes for a quite limited interaction experience.

Perhaps more popular than the OCC model are the models that represent emotion as a point in a multidimensional space. Russel's two-dimensional (2D) space of arousal and valence (see Figure 8.4) captures a wide range of emotions on a 2D plane and is one of the simplest emotion models that still has sufficient expressive power for HRI (Russell, 1980).

The 2D circumplex model, however, places "angry" and "afraid" side by side, whereas most people would argue that these are vastly different emotions. Later versions thus added a third axis, leading to the framework by Mehrabian and Russell (1974); Mehrabian (1980). This framework captures emotions in a three-dimensional (3D), continuous space, with the dimensions consisting of pleasure (P), arousal (A), and dominance (D) (see Figure 8.5). The PAD space model has been used on many social robots to model the user's and the robot's emotional state, including Kismet (Breazeal, 2003).

8.5 Challenges in affective HRI

Despite considerable efforts in the perception, representation, and expression of emotion in virtual agents and robots, there are still a number of open challenges.

It is virtually impossible to correctly read emotions from facial in-

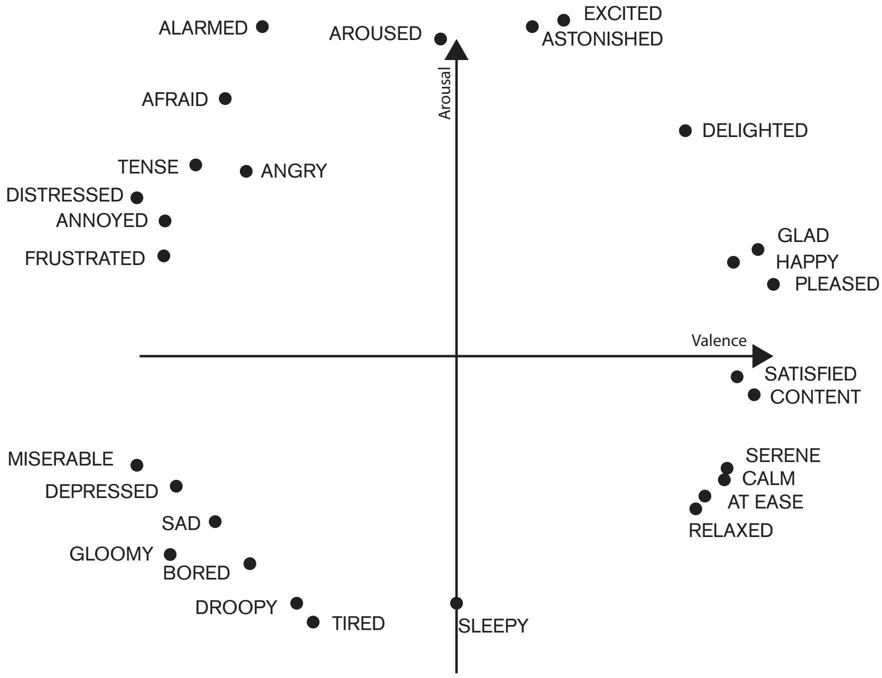

Figure 8.4 Russel's circumplex model of affect.

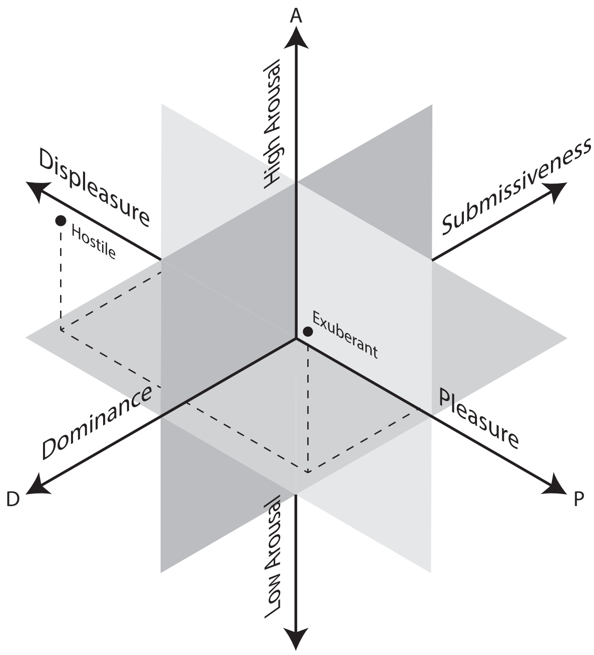

Figure 8.5 The PAD emotion model. An emotion is represented as a point in a 3D space, with axes representing pleasure (P), dominance (D), and arousal (A).

formation alone (see Figure 8.6). Given that people struggle to correctly read emotions from still facial images, robots will certainly have trouble with this as well. The addition of more information—such as

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 8.6 Can you tell if the tennis player just scored or lost a point? A study showed that people struggled to correctly read strong emotions from the static faces alone, but they could, however, when only seeing the body posture (Aviezer et al., 2012). (Source: Steven Pisano)

the context of the interaction, animated rather than still expressions of emotion, and body language—allows us to increase the recognition rate, both by people and by algorithms.

Another problem in emotion recognition by computers is that almost all algorithms are trained on emotions that have been acted out by actors. As such, these emotions are exaggerated and bear little resemblance to the emotions we experience and express in daily life. This also means that most emotion-recognition software is only able to correctly recognize emotions that are displayed with a certain exaggerated intensity. Because of this, their use in real-world applications is still limited (Pantic et al., 2007), and the recognition accuracy of subtle emotional expressions drops dramatically (Bartneck and Reichenbach, 2005). Another problem is that most emotion-recognition software returns probabilities for only the six basic emotions proposed by Ekman, or a point in a 2D or 3D emotion space. This is perhaps a rather limited view of emotion and misses many of the emotions we experience in real life, such as pride, embarrassment, guilt, or annoyance.

Another aspect of emotional recognition that poses difficulty for robots is recognizing emotions across a wide variety of people. Although we may all be expressing a number of universal emotions, we do not all do it with the same intensity, in the same type of context, or with the same meaning. Interpreting the emotional status of a person, therefore, requires a sensitivity to his or her individual affective quirks. Humans become adept at this through long years of interacting with each other but also through long-term experience with individuals. That is why you might be able to tell that your partner is laughing out of annoyance rather than happiness, whereas new acquaintances may not be able to do so. Robots still decode emotions largely based on momentary snapshots of a person's countenance, and they do not develop more long-term models of affect, emotion, and mood for their interaction partners.

Finally, a robot's emotional responsiveness can fool potential end users into thinking the robot would actually experience genuine emotions. A robot merely expressing a certain emotion does not replace the actual, visceral experience of an emotional state. The robot merely displays emotional states in response to a computational model. Affective cognition, in which a full socioemotional repertoire is expressed and recognized for different users and contexts, still remains elusive.

Questions for you to think about:

- Come up with a list of 10 emotions, and then try to display them nonverbally to a friend. Can your friend guess which emotion you are showing?

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

- Let's role play: To understand how emotions are involved in our daily interaction, imagine being incapable of both experiencing and processing any information involving emotion. Then, set out to have a chat with a friend (consider telling the friend beforehand about your experiment). Try not to respond to whatever emotion your talking partner displays, and try not to show any emotional feedback. What happens?
- Are there tasks for which a robot should or shouldn't have emotion? Is it a good idea to implement emotion into a self-driving car, for example? If not, what are the potential problems?

Future reading:

- Christoph Bartneck and Michael J. Lyons. Facial expression analysis, modeling and synthesis: Overcoming the limitations of artificial intelligence with the art of the soluble. In Jordi Vallverdu and David Casacuberta, editors, *Handbook of research on synthetic emotions and sociable robotics: New applications in affective computing and artificial intelligence*, Information Science Reference, pages 33–53. IGI Global, 2009. URL `http://www.bartneck.de/publications/2009/facialExpressionAnalysisModelingSynthesisAI/bartneckLyonsEmotionBook2009.pdf`
- Cynthia Breazeal. Social interactions in HRI: The robot view. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 34(2):181–186, 2004b. doi: 10.1109/TSMCC.2004.826268. URL `https://doi.org/10.1109/TSMCC.2004.826268`
- Rafael A. Calvo, Sidney D'Mello, Jonathan Gratch, and Arvid Kappas. *The Oxford handbook of affective computing.* Oxford Library of Psychology, Oxford, UK, 2015. ISBN 978-0199942237. URL `http://www.worldcat.org/oclc/1008985555`
- R. W. Picard. *Affective computing.* MIT Press, Cambridge, MA, 1997. ISBN 978-0262661157. URL `https://mitpress.mit.edu/books/affective-computing`
- Robert Trappl, Paolo Petta, and Sabine Payr. *Emotions in humans and artifacts.* MIT Press, Cambridge, MA, 2003. ISBN 978-0262201421. URL `https://mitpress.mit.edu/books/emotions-humans-and-artifacts`

9

Research Methods

What is covered in this chapter:

- Methodological considerations and various decisions you need to make in setting up and performing a human–robot interaction (HRI) study;
- The strengths and weaknesses of different research methods, and how to identify them for understanding and evaluating HRI;
- How the choice of robot, environment, and context matter for study results;
- The importance of looking at new ways of reporting data and insights befitting HRI, even though there is a tradition of reporting experimental work.

Now you have a robot, and you want to know with some certainty how it performs. What do people think about its appearance? How do they react to its behavior? Will people accept it? What will the effects of using the robot be in the short term or over a longer period of time? How does the robot perform technically? These are common questions in human–robot interaction (HRI), and they will require you to use different research approaches and methodologies to find the answers.

HRI research consists of at least two interrelated components: the human and the robot. These are essential to any HRI study; if you investigate humans without robots, you are engaging in social science research, whereas research on robots without humans involved would qualify as robotics or artificial intelligence research. The unit of analysis in HRI is always some form of interaction between the two. The context in which HRI happens is of high relevance and needs to be explicitly defined in studies. You might study HRI in the lab or in a school or hospital; you might study HRI in different cultures or in different application domains. The context in which the robot interacts with people is very likely to have a strong influence on your results, and you need to be aware of with whom and in what circumstances the interaction unfolds.

Although the focus of HRI is always on the interaction between humans and robots, there are different aspects of this relationship to study. In *robot-centered work*, the research focus might be on developing the technical capabilities that robots need to interact with people, or testing different aspects of the robot's functionality or design to see which are most effective. In *user-centered work*, on the other hand, the focus of a study could be on understanding aspects of human behavior or cognition that will affect the success of HRI. For instance, an extroverted user might prefer more direct communication by the robot, whereas an introverted user might like indirect communication.

HRI research also increasingly strives to strike a balance between these two approaches, coupling robot and user-centered aspects in different ways. For example, in iterative design, the robot's design goes through a number of cycles of prototyping, testing, analyzing, and refining. Researchers come up with a series of robot design ideas, which they then test out with users. Based on the users' preferences, the researchers then further develop the robot's appearance and capabilities. Another mode of coupling user- and robot-centered aspects of HRI is through studying human behavior to develop behavioral models that can then be applied to HRI and to test those out with users to see if they produce the expected and desired results in interaction.

Studies in which users interact with the robot, tests of the robot's performance, and more open-ended explorations of ways in which people and robots interact in everyday life are all part of HRI research. Consequently, HRI researchers draw on and often mix a variety of research methods and techniques, some adapted from other disciplines (e.g., sociology, anthropology, or human factors research) and some developed for the HRI field itself (e.g., the "Wizard-of-Oz" technique, described in Section 9.6.1). To employ these methods successfully, HRI researchers need to be aware of their strengths and weaknesses, the kinds of data and insights they may produce, and the types of technical and human resources they require.

Taking an experimental approach has become standard in the HRI community. This was not always the case, and a quick glance at older HRI research will show methods that would make current HRI researchers blush. There is a push to have current research meet criteria for methodological soundness that are applied in other empirical sciences, such as psychology (Baxter et al., 2016). This chapter discusses the kinds of decisions that HRI researchers make at different points in the research process, from defining the research questions (Section 9.1), to study design (Section 9.2), to statistics (Section 9.8), and explains the journey you will make when evaluating the interaction between robots and people. After thinking through the steps to formulating a research question in Section 9.1, Section 9.2 provides examples of different uses of qualitative, quantitative, and mixed methods in user and system studies, observational and experimental studies, and other forms of HRI research. The selection of participants is the focus of Section 9.3, whereas Section 9.4 emphasizes the importance of defining the context of interaction as part of the initial study design. Sections 9.5 and 9.6 consider how to choose an appropriate robot and mode of interaction for your HRI studies. Sections 9.7 and 9.8 present various metrics and research standards to be taken into account in HRI research, including statistical, ethical, and generalizability concerns. The overall aim of the chapter is to provide a basis from which to make initial study design choices and then delve more deeply into research methods to develop your own novel HRI studies.

9.1 Defining a research question and approach

Defining a good research question is one of the hardest tasks of a researcher. To form a strong research question, a researcher must consider previous relevant work and replicate or extend it to contribute new scientific insights. In HRI, such insights can come in the form of knowledge about human cognition and behavior, guidelines for robot design, technical aspects of the robot, or findings that can inform the application of robots in different use contexts.

Research questions in HRI might arise from theoretical considerations, such as the expectation that people will treat robots as social, or from the pragmatic need to test the usability of a certain robot feature or function. To find relevant literature from multiple fields of expertise, we recommend searching publications across disciplinary databases to incorporate research findings from multiple fields of relevant expertise. Ideally, you would look for a well-established phenomenon or theory and seek to replicate and extend it in your new research project, independently of whether it is about humans or robots. Research on interactions among humans can easily serve as a blueprint for human–robot research. Existing work in HRI, psychology, sociology, anthropology, design, and media communications can provide relevant insights into the underpinnings of smooth, successful, and acceptable HRI or into the optimal, human-centered design of a novel robot platform.

To illustrate, in the 1990s, Reeves and Nass Reeves and Nass (1996) proposed the "computers as social actors" (CASA) approach and sought to replicate classic psychological findings in the context of human–computer interaction. In their seminal work, the authors conducted studies that provide evidence for the hypothesis that computers are treated just like human interaction partners. Moreover, they found that such behavior occurs quite automatically. For instance, they showed that humans give higher ratings if a computer asks about its own performance than when they have to rate the performance on a different computer, which indicates that people are polite to computers. Later on, the CASA approach was successfully extended to HRI through a wide array of studies, including some exploring the attribution of gender to robots (Eyssel and Hegel, 2012) and users' mental models of robots (Walden et al., 2015) and others studying the effects of perceptions of social presence and agency in caregiving (Kim et al., 2013) and educational (Edwards et al., 2016) scenarios. This paradigm continues to inspire new research in HRI.

9.1.1 Is your research exploratory or confirmatory?

Broadly speaking, research can be classified as either exploratory or confirmatory research. Exploratory research questions deal with phenomena that have not previously been examined in detail and aim at finding out the general "lay of the land" in a specific domain. For example, you might ask "How do people adopt and use a robot vacuum cleaner in their home over one month?" or "Do people attribute gender to robots and assign stereotypes to them?" Exploratory research assumes that there is not enough relevant prior information about the phenomenon to formulate testable expectations about the potential outcomes of the study, and it therefore seeks to explore what factors might be important and which outcomes are possible.

In an exploratory HRI study, Jodi Forlizzi and Carl DiSalvo investigated how a vacuum-cleaning robot is integrated into the homes of real people. Their findings produced many surprises for the research community (Forlizzi and DiSalvo, 2006), including that people would treat autonomous robotic vacuums as social actors, that such vacuums could inspire teenagers to clean their rooms, and even that some pet–robot interaction occurred (see Figure 9.1).

Figure 9.1 A cat riding on a Roomba robot (2002–present). (Source: Eirik Newth)

When there is enough information to formulate hypotheses about the possible outcomes of an intervention, we enter the domain of confirmatory research. The goal of confirmatory research is to test hypotheses. In your hypothesis, you need to spell out the findings that you anticipate prior to starting your study and explain why you think those findings should be expected. A key point here is to formulate a question in such a way that it is verifiable. Take this example from everyday life: You might know that teenagers are often interested in new gadgets and technologies but tend to avoid doing chores. This may lead you to expect that introducing a robotic vacuum cleaner into their homes will increase their engagement with cleaning compared to introducing a normal top-of-the-line vacuum cleaner. You would then design your study in such a way that it answers the following research question: "Do teenagers engage in more cleaning with a robotic vacuum cleaner compared to a conventional vacuum cleaner?"

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

You should consider registering your hypotheses prior to conducting your experimental study at one of the many sites available for that purpose, such as the Center for Open Science (`https://osf.io/prereg`), AsPredicted (`https://aspredicted.org`), or the U.S. National Library of Medicine (`https://clinicaltrials.gov`). This will keep your work in line with the standards and rigor in the psychological and clinical sciences and makes it clear that you have not adjusted your hypothesis to fit the data or have reported only carefully selected results (Nosek et al., 2017).

The teenagers and cleaning example shows how hypotheses can be inspired by commonsense knowledge, but you can also build on prior empirical research and social theory to develop hypotheses about HRI. One such example is the social conformity theory of Solomon Asch, who showed how people tend to conform to peer pressure. In an elegant experiment, he showed that when people complete a simple visual task in a group setting, they are more likely to give the same response as others in the group even if they know the response is wrong (Asch, 1951). This classic experiment can be run with a group made up of robots rather than people. Will people conform to robots? Studies have shown that adults do not, but children do (Brandstetter et al., 2014; Vollmer et al., 2018).

9.1.2 Are you establishing correlation or causation?

Along with deciding whether your research questions call for an exploratory or confirmatory approach, you need to decide whether you want to establish correlation or causation between the variables of interest in your research study.

In correlational studies, we can show a clear pattern by which the variables change value in relation to each other, but we cannot know what causes this relationship. A correlational survey study of teenagers using the Roomba could measure whether there is a statistical relationship between households owning a Roomba and the amount of time teenagers spend cleaning. We would, however, not necessarily know why this relationship happens. It might be that teenagers who own a Roomba are more tidy to start with, or that their parents ask them to clean more often. To make the claim that a Roomba would increase the time spent cleaning, you would need to compare the behaviors of two similar groups of teenagers by giving one group a Roomba and the other group a regular vacuum cleaner, then measuring the outcomes. This requires an experimental study design to investigate the causal relationship and show that a change in one variable actually leads to a change in the other. We do this by dividing a sample into two (or more)

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

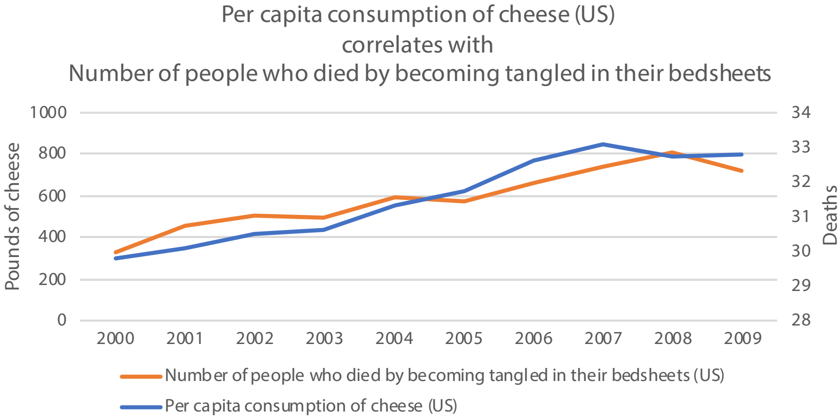

Figure 9.2 A strong correlation that has no causal relationship.

similar groups, then manipulating the variable we think has an effect in one group, and finally measuring the variable of interest in both groups. Because the groups should be similar in terms of demographics, skills, and other characteristics at the start, any major difference that is observed would be the result of our manipulation.

The difference between correlation and causation is important because it defines what conclusions can be drawn from the findings. Correlation says nothing other than "these things happen to occur simultaneously"—for example, there will be a strong correlation between the number of firefighters on the scene and the damage that a fire did. This does not, of course, mean that the fire was caused by firefighters or that we should stop sending firefighters whenever there's a fire. Sometimes a correlation even pops up for no reason at all, a so-called "spurious correlation." An example of a spurious correlation is the strong ($\rho = 0.97, r^2 = 0.896$) relationship between U.S. per-capita cheese consumption and the number of people who died by becoming tangled in their bedsheets (see Figure 9.2[1]).

9.2 Choosing among qualitative, quantitative, and mixed methods

How you define your research question will also affect what type of methods you should use to answer it. Qualitative methods allow researchers to understand the qualities of an interaction that are difficult to capture in numbers. It requires researchers to identify and interpret the underlying meaning or thematic patterns that they see in social interaction. The data that are derived from these studies typically cannot be expressed numerically, which disqualifies this approach from being used to establish correlations or causation. Quantitative methods,

[1] `http://www.tylervigen.com/`

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

in contrast, often take the shape of surveys or controlled experiments and produce data that can be expressed numerically and analyzed statistically to check for correlations and causation. They will therefore allow you to make predictions or even establish cause and effect. Observational studies (Section 10.2.3) can produce both qualitative and quantitative data, which can be used to investigate commonly seen patterns in interaction and correlations between the characteristics of the humans, robots, or context. For instance, you might find from observation and interviews that the number of times adolescents clean with the Roomba can be related to their personality characteristics, such as self-reported conscientiousness. The interviews might also tell you that people talk about the Roomba as a social actor, calling it a "he" or "she" rather than an "it" (i.e., a tool). Finally, your research questions might call for a mixed-methods approach, which may include exploratory research using interviews, focus groups, or observation of naturalistic interaction to identify emergent factors significant to HRI, followed by experiments to confirm these relationships. For example, if your interviews lead you to think that the autonomous behavior of the Roomba is what makes it seem social to people, you could set up an experiment to test this. Such an experiment would have two groups of participants, whom you present with either an autonomous Roomba or a Roomba that they steer using a game controller. You can then measure the level of sociality they ascribe to each Roomba and test if these are significantly different from one another.

9.2.1 User studies

User studies are experiments in which you bring people in to interact with a robot. Not all HRI research requires a user study—for example, you might just want to test the navigation capability of your receptionist robot. However, most HRI research at some point will involve a study in which you measure how users respond to variations of the robot, the interaction itself, or the context of the interaction. These different variations are called experimental conditions. The critical feature of a user study is the random assignment of a large enough sample of research participants to your experimental conditions. Experimental conditions typically emerge from the factors that you consider of importance or interest and should be outlined in your research design. For instance, assume we want to test whether people apply human stereotypes to a gendered robot. To test this, we run an experiment using a male and a female robot prototype. The robot's gender is called the independent variable, which is the aspect in the experiment that is controlled. Because we test two robot versions, male versus female, the independent variable has two levels. The resulting research design thus leaves us with two conditions to which we randomly assign our research participants.

If we think that gender stereotyping of a robot also depends on the gender of the human watching it, we want to test not just for the effect of robot prototype gender but also take into account participant gender as well. We thus add a second independent variable to our design: participant gender. Because we cannot manipulate this variable (we cannot randomly assign a gender to each participant who walks into our lab), participant gender would be called a quasi-experimental factor. Our study design has now a 2×2 format: robot gender (male vs. female) and participant gender (male vs. female). In our analysis, we will thus be comparing four groups, or "cells" in our design: males rating a male robot, males rating a female robot, females rating a male robot, and females rating a male robot.

Now the question is: How exactly do we measure what we want to know? The variables we measure are called dependent variables. We know from psychological literature that females are commonly perceived as communal and warm, whereas males are perceived as more assertive (Bem, 1974; Cuddy et al., 2008). We can use this information to measure to what extent our male and female robot prototypes are being stereotyped. Indeed, previous research studies have shown that manipulating robot gender leads to such a stereotypical perception of traits in robots (Eyssel and Hegel, 2012). People seem to reproduce the stereotypes that are common among humans in the context of robots.

Not only does the dependent variable need to be well designed, but it is also important that the independent variable (i.e., the construct of interest) is validated. Can we be sure that our study participants actually recognized the robots as male or female? To establish the validity of our results, we need to know whether robot gender was operationalized successfully. We can do this by including a manipulation check in our study to see that our experimental treatment was indeed effective, that is, that our participants indeed perceived the robot with male gender cues as male and the robot with the female cues as female. This could be done simply by adding a post-interaction question asking them to identify the gender of the robot and/or by seeing whether they refer to the robot by a specific gender when they talk about it following the interaction. Only once this is established can researchers be sure that the operationalization—that is, the translation of the theoretical construct of interest into a measurement or manipulation—was effective.

9.2.2 System studies

Whereas user studies are used to report on people's attitudes toward and interaction with robots, system studies are those that evaluate the technical capabilities of the robot. A system study might involve users, but user involvement is not always needed. At the same time, system studies do require the same rigor expected from user studies. This means that verifiable research hypotheses and performance claims, a study protocol, and clear metrics are all key to system studies.

For example, when designing an interactive robot for children, you might want to know how well automated speech recognition works for your target user group (Kennedy et al., 2017). Speech recognition has been designed to work well for adults, but it might not be suitable for children due to their voices having a higher pitch and their speech often containing more disfluencies and ungrammatical utterances. To test whether speech recognition works for child speech, you could ask children to interact with your robot, but a better idea would be to use recordings of children's speech and pull these through the speech-recognition software. The benefit of this approach is that the experiment is repeatable: you can try different parameter settings in the software or even swap different speech-recognition engines and assess the performance using the same recordings.

Systems studies are often used to assess the perceptual capabilities of the robot. Capabilities such as face recognition, facial emotion classification, or sentiment detection from voice are best assessed using consistent test data sets with well-established metrics. For some capabilities, there are existing data sets that can be used to assess the performance of the robot. For face recognition, several data sets exist, for example, the IMDB-WIKI, which contains images of people extracted from the IMDB database and Wikipedia; in addition to labels, the images contain gender and age information Rothe et al. (2016). Using well-established metrics allows you to compare the performance of your robot to others. Classification problems often have agreed-on methods of reporting performance, such as reporting the accuracy of the classification (the number of correct classifications divided by the total number classifications, including the ones that are wrong) or the precision and recall. Speech-recognition performance is often expressed as a word error rate (WER), which is the total number of substitutions, deletions, and insertions in the text divided by the number of words in the actual spoken sentence. So if "Can you bring me a drink please" is recognized as "Can bring me a pink sneeze," that is a WER of $(2 + 1 + 0)/7$. It is worth exploring what the accepted metrics are in a particular discipline and rigorously sticking to the accepted method for evaluating and reporting system performance.

9.2.3 Observational studies

As robots have become more robust, more reliable, easier to use, and cheaper, it has become viable for HRI researchers to study how people and robots interact in various naturalistic contexts using observational methods. Observing how people interact with robots, for example, by studying where they place robots in their environments and how they respond to different kinds of verbal and nonverbal cues performed by robots, allows researchers to understand how HRI can unfold in a more natural way, without researchers directly intervening in the interaction.

Observational studies can be exploratory, involving putting a robot into a specific environment to see how interactions there unfold. An example of such an observational study is the work of Chang and Sabanovic, who put a seal companion robot in a public space in a nursing home and observed when and how different people interacted with the robot (Chang and Šabanović, 2015). The findings provided frequency counts of interactions with the robot, as well as identifying different social factors (e.g., participant gender, social mediation effects) that affected whether and for how long people interacted with the robot. The researchers did not manipulate anything about the robot or the environment. They just observed how residents interacted with the robot.

Observational studies can also be performed to evaluate, by means of a field experiment, how effective a robot is for a particular task or the effect of certain design variables on interactions. Researchers from the Advanced Telecommunications Research Institute (ATR) in Japan have performed several observational studies of interactions between the humanoid Robovie and mall customers. These studies represent a particularly fruitful iterative form of design and evaluation using observational techniques. In the initial stages of the study, researchers observed general human behaviors and analyzed these observations to identify particular behavioral patterns, which they then used to develop behavioral models for the robot. The robot was then placed in the mall, and people's reactions to it were evaluated to see if the behavioral models had the expected positive effects on people's responses.

Observational studies can rely on data collected in several different ways: observational notes and logs collected by a researcher in person, manual annotations of video recordings of interactions between people and robots, and robot logs from interactions with people.

In-person observation provides the possibility for researchers to have a better understanding of the broader context of interaction because they can see and hear things that might not initially be in the data-collection protocol. This can lead to amendments of the protocol or can be represented in notes that can help guide later analysis and interpretation of the data. In-person observation, however, is limited by the sensory capabilities of observers at the time of coding and does not allow for others to go back and review the coded observations. In terms of establishing interrater reliability, more than one coder needs to be present in the context at the same time, which can be inconvenient and become a distraction to other people in the space because of the presence of multiple researchers.

Video coding, on the other hand, allows researchers to review observations as many times as needed, potentially revise their coding schemes, revise their codes of observations, and easily provide data to a second coder for establishing interrater reliability. Video, however, has a limited view defined by whatever is visible from the chosen camera angle. This may cause researchers to miss some relevant aspects of the interaction, so it is important to clearly define what the camera should be focused on before the video observation starts so that important things are not missed. Although video coding may seem more convenient and preferable overall, some contexts (e.g., nursing homes, hospitals, schools) may not allow researchers to record video, so in-person coding may be necessary.

Finally, robot logs are limited by the robot's ability to sense and categorize different human actions but have the benefit of being able to provide data about both the robot's state and actions and the human actions it perceived at the same time. It is, of course, possible to combine these different data sources to improve the accuracy of the data.

Both in-person coding and video annotations require the development of a coding scheme that coders will follow systematically. This coding scheme can be developed based on theoretical or practical interests and expectations, or it can be developed bottom-up by identifying points of particular interest in a portion of the data and then going through the rest of the corpus to understand related patterns. It is very important to pilot test the coding scheme to identify missing components and overlapping or unclear codes so that coders can be in clear agreement about what the codes mean before they start (particularly for in-person coding, where you can't go back to view the interaction). Video analysis is also quite labor-intensive, so properly defining how fine-grained you need the coding scheme to be can save time and effort. Aside from providing frequency counts of certain types of behaviors or identifying qualities and patterns of interaction, observational coding of interaction behaviors can also provide particularly interesting temporal patterns of behavior, which can show the effects of certain robot behaviors on people's actions (e.g., how a particular gaze cue by a robot is followed by a joint-attention behavior by a person).

9.2.4 Ethnographic studies

Along with behavioral observation, HRI researchers also engage in more in-depth and often long-term ethnographic observations, in which they not only seek to identify certain behavioral and interaction patterns among humans and robots but also to understand what those patterns mean to people and how they are connected with the broader environmental, organizational, social, and cultural contexts in which those interactions take place. Ethnographic observations can include all aspects of interactions between people and robots, including behaviors, speech, gestures, and posture. They also include information on the context in which those occur, including daily practices, values, goals, beliefs, and discourse of different stakeholders, which include but are not limited to people who directly interact with the robot.

Although behavioral observation is inspired by ethology and the desire to explore and build explanatory models of animal and human behavior, ethnographic observation is based on the theory and practices of anthropology and the goals of understanding sociocultural experiences holistically. Ethnographic observation is often performed over longer periods of time, from a few months to a few years, which is necessary for the observer to get a more complete and emergent sense of the cultural logic of the research site. Ethnographic studies can be performed by participants as outside observers but also through participant observation, where the researcher takes part in the activity under study to better understand the experience. The former type of study is currently more widely represented in HRI, although social studies of robot design often take the latter approach. Ethnographic study is also often coupled with a “grounded theory” approach to data analysis, which assumes that the collection and interpretation of data are ongoing throughout the project, with the researcher regularly engaging in reflection on the questions that guide the research, methods of data collection and analysis, and potential interpretations of the data, thus iterating as the study goes along.

Ethnographic studies are still relatively rare in HRI, partly because of the labor involved in collecting data over longer periods of time but also because there have not been many robots that are technically capable of taking part in long-term interactions with people. Some successful examples of ethnographic studies include a one-year-long study of a service robot in a hospital that showed that the patient type in the context, oncology or postnatal, determined whether the robot was appreciated or hated (and sometimes kicked and sworn at) by nurses (Mutlu and Forlizzi, 2008). Forlizzi and DiSalvo (2006) did an ethnographic study in which they gave families either a robotic Roomba vacuum or the latest version of a conventional vacuum to use over several months. They learned that people treated the robot, but not the conventional vacuum, as a social agent, and that having a robotic vacuum changed the way the family cleaned, particularly inspiring teenagers and men to participate. Leite et al. (2012) performed an ethnographic study with a social robot that could respond empathically to children in an elementary school. The study found that the task scenario and children’s specific preferences influenced their experiences of the robot's empathy. Several ethnographic studies have also been performed with scientists using robots. Vertesi (2015) studied National Aeronautics and Space Administration (NASA) scientists' interactions with a remote Rover and showed how the organizational structure of the team affected the team members' use and experience of the robot. The study also showed that scientists performed aspects of the robot's behaviors with their own bodies, creating a team identity for themselves in the process.

Ethnographic studies are particularly valuable because HRI is a young field and thus is still developing a corpus of theoretical and empirical work that can identify the most relevant factors we need to pay attention to, not only in the design of robots but also in their implementation in different environments.

9.2.5 Conversational analysis

Conversational analysis (CA) is a method in which the verbal and nonverbal aspects of an interaction are reported in great detail (Sidnell, 2011). This is not limited to conversation only, as the name might imply, but can be applied to any form of interaction between people or between people and technology.

The process of CA starts by recording an interaction between two or more parties. Whereas this used to be audio recording, nowadays, video recording is more convenient, and several cameras can be used to capture the interaction from different angles. The participants being recorded might or might not be aware of the recording. From the recording, a very detailed transcription is produced, including turn-taking cues such as pauses in conversation, emotional cues such as laughter, behaviors performed while conversing, and other details of the interaction. Depending on the research question, the temporal resolution of the transcription can be brought down to the frame rate of the video recording. This can capture small actions, such as blinking and other eye movements, gestures, and changes in body posture . Fischer et al. (2013) used CA to investigate how the contingency of robot feedback affects the quality of verbal HRI. In their experiments, participants instructed the humanoid robot iCub how to stack some shapes in a contingent and noncontingent condition. Analysis of participants' linguistic behaviors, including verbosity, attention-getting, and word diversity, showed that contingency had an impact on the participants' tutoring behaviors and therefore can be important for learning by demonstration.

CA will pay specific attention to elements in the verbal interaction, such as turn-taking, back-channeling, overlap of speaking, repair statements, echo utterances, and discourse markers. In HRI, CA can be used to analyze in great detail how people interact with social robots and whether they employ similar conversational strategies with robots as they do with people.

9.2.6 Crowdsourced studies

A new development in HRI is that researchers now also have access to crowdsourcing as a way of doing studies. Crowdsourcing is the practice of obtaining data from a large number of people, either paid or unpaid, via online methods. In recent years, the use of online crowdsourcing platforms has allowed researchers to run user studies and gather large amounts of data with relatively little effort and to gather data from subjects they would typically struggle to reach (Doan et al., 2011). The online platform can be entirely built by the researchers, but more often, existing online tools are used to recruit, run, and analyze user studies. The most widely used tools are Amazon Mechanical Turk (MTurk or AMT) (see 9.3 and Crowdflower, which allow you to post jobs, also known as human intelligence tasks (HITs). The jobs usually are short user studies in which participants are asked to watch a number of images or videos containing robots or interactions with robots and then answer a set of questions about the material. Crowdsourcing allows the researcher to gather large amounts of data in a short time frame and for a modest cost. Taking part in a study will earn each participant a small financial reward, typically only a few U.S. dollars, with the price set depending on the complexity of the task, the time it is expected to take, and the quality rating of the respondent.

Running crowdsourcing studies comes with its own set of unique challenges, the most important being the relatively low level of control the experimenter has over the subjects taking part in the study and the environment in which the study is executed. Any account that meets the broad inclusion criteria set by the crowdsourcing platform is allowed to take the job. However, the account that is logged in might not be being used by the actual person registered as taking part in the study. Participants could take your study while pursuing an array of other activities, such as eating ice cream while petting a cat; or they could be full of caffeine or sitting on a crowded bus while listening to loud music on headphones. Crowdsourcing also is open to malicious user behavior: participants often provide low-quality or deliberately incorrect responses.

To avoid some of these problems, it is good practice to include verification questions in your user study (Oppenheimer et al., 2009). These questions check whether participants pay attention and are engaged with the task. When showing a video, a number could be displayed for a few seconds, after which the video participants are asked to enter the number. Questions can also be used to ensure the participant is re-

Figure 9.3 Amazon Mechanical Turk was named after a fake chess-playing machine called "The Turk" constructed in the late 18th century.

sponding to the questions rather than just picking random responses, such as "Please click the third option from below."

After data collection, it is necessary to separate the wheat from the chaff. A first filter will be the responses to the verification questions; another method is to exclude all responses that took less than a reasonable amount of time. For example, if you believe the study should take a minimum of 15 minutes, then any responses that are far under that time should be disregarded. Some crowdsourcing platforms allow you not to reward participants if their responses are of insufficient quality, which not only leaves those participants without pay but also negatively affects their ratings. This has shown to be an excellent incentive. Given that data collected using crowdsourcing are inherently more variable than data collected in the lab, one way of addressing this problem is to collect more of it.

Although crowdsourcing has been successfully used to replicate results from lab studies in social psychology, linguistics, and behavioral economics (Bartneck et al., 2015; Goodman et al., 2013; Schnoebelen and Kuperman, 2010; Suri and Watts, 2011), the value of crowdsourcing to HRI needs to be considered on a case-by-case basis. Sometimes the physical presence of a robot is key to the participant's performance, precluding the use of crowdsourcing. Sometimes the effect you are measuring is small and would not show up when sampling a large and diverse population. Sometimes the population you need is scarce on crowdsourcing platforms, such as elderly users or Swedish primary school teachers. Sometimes the task requires a certain level of language proficiency. Crowdsourcing has its place in HRI research, but it should be used with care and consideration.

9.2.7 Single-Subject Studies

Another type of study to consider in HRI is the single-subject or single-case research design. In this type of study, researchers compare the effects of an intervention on a single subject rather than a group of people. This is done by initially collecting baseline measures of the individual's behavior, which are compared with the subject's behavior during and following the intervention.

Single-subject designs are used in cases where recruiting large numbers of participants can be difficult due to their rarity in the population or when individual differences between subjects are large and relevant to the phenomenon of interest. Multiple subjects can be recruited in single-subject designs, but the number of subjects is often small, and for the sake of analysis, each subject is treated as his or her own control.

Single-subject designs are commonly used in medical and education research fields, and in the case of HRI, they are used in research on the effects of robots on individuals with autism. For example, Pop et al.

(2013) performed single-case studies with three children to investigate whether the social robot Probo can help children with autism spectrum disorders better identify situation-based emotions. Tapus et al. (2012) similarly worked with four children with autism to see whether they would show more social engagement with the Nao robot than with humans, and they found large variability among their responses. This shows the importance of performing single-subject studies in cases where individuals of interest, such as those diagnosed with autism, vary widely in their behaviors; in such cases, averaging the responses of a group could mask important intervention effects because different individual responses would cancel each other out when aggregated.

9.3 Selecting research participants and study designs

Because "humans" are a necessary component of HRI studies, several important decisions in HRI studies must be made regarding the participants. One is who the participants will be. The usual suspects for empirical HRI research are university students because they are the most convenient population to access for academic researchers, have time for and interest in participating in studies, and are usually in close physical proximity to the laboratories where much of the HRI research is performed. It is, however, important to consider the limitations of using university students as a "convenience sample," particularly in relation to the kind of research questions you have posed. In an ideal world, we would aim for a large, representative sample of potential end users of robots so that we can claim that our findings hold for a wide range of users and have *external validity*—that is, they can tell us something about people and robots in situations outside the study itself. Such samples are very difficult to bring in for experimental studies but might be more achievable in surveys. In studies of the general perceptions of robots, HRI, similarly to psychological research, assumes that university students are "close enough" to the general population in terms of characteristics when it comes to broad social traits (e.g., stereotyping), cognitive performance (e.g., memory), and attitudes (e.g., fear of robots). Even when using university students, it is important to be mindful of and balance certain characteristics of the sample, such as gender or educational background, depending on whether these factors might be expected to have an effect on your results. For example, students in a computer science department would likely be seen as having more positive attitudes toward robots and having greater ease in using computing technology than a broader student population or the general population of potential users.

If your research questions relate to studying the characteristics of a specific population, such as older adults, or to investigating the effects of robot applications in specific domains, such as the treatment of children with diabetes, your choice of participants will need to be more specialized. The specificity of your research question and the claims you want to make will guide the level of specificity of your sample. It is not possible, for example, to claim that a robot will have positive effects on older adults experiencing cognitive decline if you run your study with university students or even with older adults who are not experiencing cognitive decline. A university student sample will also not be sufficient for investigating the use of robots to support learning in young children. Thus, before running your study, you have to decide carefully what kinds of people should take part in it. You will also need to consider how to get access to this population and how to recruit and motivate individuals to be in your study. You should also consider whether you will be able to bring people from this population to your lab, whether you need to go to another place to have contact with them, or whether an online study might be appropriate.

Another consideration regarding research participants is the number of participants you might need to answer your research questions. This will depend both on the type of study and analysis you are doing (quantitative vs. qualitative, survey, experiment, or interview) and on the type of population you are working with (e.g., university students or older adults or children with diabetes). It is difficult to reliably test for an effect with a small sample size because people will always differ a little bit from one another. In a gender stereotype study, for example, some participants will consider all robots a bit more "warm" than others; other participants will think all robots possess typically "male" qualities. Such differences, which naturally occur in people, will add noise to the data. Unless the manipulation has an extremely large effect, the data that we gather from a small sample will not be enough to detect a stable effect; the differences among people might cancel each other out, or the variability of their responses might be too large. If you want to reach a valid conclusion about cause and effect, you need to determine the right sample size for your experimental design.

9.3.1 Study design

As a rule of thumb, it has been recommended to conduct an experimental design with a minimum of 25 participants per condition. However, how many participants you need to reliably find a difference between conditions also depends on the type of design you use. When using a *between-subjects design*, participants are randomly assigned to a condition. In our example, one group of participants would be presented with the "male" robot, whereas the other group of participants would be shown the "female" version. After answering questions using a Likert scale—a type of rating scale commonly used to measure attitudes and opinions by asking respondents to rate items based on the degree to which they agree with them—the mean scores of each group can be compared. Alternatively, in a *within-subjects design*, one group of participants is exposed to both versions of the robot prototype and asked to evaluate both. Because the same person provides two evaluations, you cut down on the "noise" in your data, and the number of participants required will be lower for this design. However, not all research questions are suitable to be answered with a within-subjects design. For example, if you want to test if people recover faster from a broken leg when they have a robotic assistant who does walking exercises with them every day, you can hardly have them first heal on their own and then break the other leg so that they can recover again with their robot helper. Also, researchers have to be mindful of the order effect that may occur; maybe people will always like the first robot better than the second (because of the novelty). Thus, it is a good idea to *counterbalance* any conditions when running a within-subjects design. This means that half the participants will first interact with the female robot and then with the male, and vice versa for the other half.

To approximate a sufficient sample size to establish a statistical effect of the desired size, the internet offers a variety of tools, such as G*Power (Faul et al., 2007). However, researchers may not always be able to meet such recommendations because they are also constrained by the availability of resources, such as time, money, robots, and potential participants.

Sometimes HRI researchers choose to use a *survey*, which is a list of questions to be answered by participants. Answers are often given through multiple-choice options or some sort of rating scale. One commonly used type of scale is the Likert scale, which asks respondents to rate statements about their attitudes and opinions on a topic based on how much they agree—for example: Rate the statement "I found the robot friendly" on a scale of 1 ("Strongly agree") to 5 ("Strongly disagree"). Another form of scale that is often used is the semantic differential scale, which asks respondents to evaluate the qualities of an artifact, or their attitudes, on a spectrum between two opposing terms (e.g., scary–friendly, competent–incompetent). Multiple-choice or scale-based questions make the survey easier to analyze later on but require careful design while developing the survey to make sure that the questions are appropriately measuring the concepts the researchers are interested in. Along with making up their own questions and scales, researchers can use questions and scales developed and evaluated by other researchers to measure concepts of interest (e.g., evaluating participant personality with the Big Five Scale (John et al., 1999), evaluating robot sociality with the Robot Social Attributes Scale (Carpinella et al., 2017)). Finally, researchers sometimes include open-ended questions in surveys as well, particularly when it is important to allow respondents to provide answers based on their own terms and categories or to understand their thought process or understanding of concepts while answering the survey (e.g., "Describe your ideal robot before you answer the following questions about it"). Because survey research is well established in the social sciences, there are many handbooks that describe how to go about constructing and performing surveys (for some examples, see Fowler (1995); Fowler Jr. (2013)).

Surveys allow researchers to investigate correlations between various factors relevant to HRI in a broad population. Such surveys often involve hundreds of participants and accommodate analyses with many different factors. Some surveys try to have a representative sample of participants, which can involve making sure the number of participants in certain categories (e.g., gender, age, ethnicity) corresponds to their percentage in the general population or weighting the collected data to achieve representative ratios. Studies that involve special populations, such as older adults with depression, may have to make do with a smaller number of participants because of the acknowledged difficulty in recruiting specific populations. In some cases, such as studies of children diagnosed with autism, where the participants are also widely diverse in the way they express themselves and experience the world, it is possible to treat participants as individual cases and study changes within each participant's behaviors and responses.

For qualitative studies, rather than focusing on a particular number of participants needed, the rule of thumb is to try to achieve "saturation" of the analytic themes and findings. The idea here is that the researchers can stop collecting new data once they find that the data they are collecting is simply adding to and repeating existing themes and findings, rather than creating new ones.

9.4 Defining the context of interaction

9.4.1 Location of study

For HRI in particular, an important distinction is between studies performed in the lab versus those performed in the field. Especially in the early years of HRI, the majority of research was performed in the controlled environment of the lab. Although robotic technology has certainly advanced over the years, and there are now robotic platforms robust enough to use outside of the lab, so-called "in the wild" studies are still relatively rare compared with the number of studies performed in the lab.

Studying interactions outside of the laboratory is important for understanding how people might interact with robots in natural circumstances, determining what kinds of HRI might emerge in those circumstances, and investigating the potential broader social effects of new robotic technologies. On the other hand, laboratory studies benefit from the researchers' ability to strictly control the context and nature of people's interactions with a robot—the introduction, task, environment, and length of the interaction can be clearly defined by the researchers. In the lab, participants are asked to interact with the robot only in the way researchers suggest. This allows for the strict manipulation of desired variables.

In contrast, field studies are more flexible in what can happen and are therefore closer to what might occur in day-to-day HRI. In the field, participants can choose how, when, whether, and why they want to interact with a robot; they can even ignore it. Field studies, therefore, provide a space in which to observe and discover new emergent phenomena, new variables of interest and significance to interaction, and the form and consequences of HRI when it is outside of the researchers' control. Field studies also effectively show how complex interactions between different contextual variables, such as institutional culture or interactions among people, might affect the interaction.

9.4.2 Temporal context of HRI

A related distinction that has grown in importance in HRI is whether researchers are studying short-term or long-term interactions between people and robots. The majority of lab studies, by necessity of their design, focus on "the first 10 minutes of HRI"—how people respond to and make sense of their first introduction to a robot. Researchers widely acknowledge, however, that people will change their attitude toward the robot as time passes, and consequently, the way they interact with the robot will change as well. The first interaction suffers from the *novelty effect*: people are generally not familiar with robots, so their initial reactions might be quite different from their reactions over a longer period of time. Short-term studies therefore have limited validity in informing us about how people and robots will interact over a longer period of time. They do, however, tell us about the kinds of characteristics of people and features of the robot that will affect the initial encounter. Such studies are important for setting up a positive feedback loop of interaction, which can then support more positive effects in long-term interaction. Studies of longer-term interactions, which can take place over several days, weeks, months, or in a few cases, even years, allow us to see how interactions between people and robots develop and change over time, how robots are integrated into human social contexts, and how social interactions between people themselves may change because of the presence of a robot.

9.4.3 Social units of interaction in HRI

Interactions between people and robots can be studied through several different social units of analysis, which the social sciences see as distinct in terms of the aspects of cognition and interaction they enable (see Figure 9.5). The most common unit, so far, has been the interaction dyad—one person and one robot interacting with each other. This is partly due to the early constraints of HRI—robots were difficult to procure and difficult to maintain and operate; hence, the most common form of HRI study was the lab experiment involving a single participant interacting with a single robot.

Figure 9.4
Robovie in school.

As early as 2006, the Robovie robot was one of the first robots capable of supporting group interactions at an elementary school. It taught children English and tracked their social networks over time, keeping the children interested in interacting with the robot by unlocking secrets (Kanda et al., 2007b) (see Figure 9.4).

As robots have become more readily available and capable of interacting with more people and in more open-ended, naturalistic environments, the unit of analysis in HRI has expanded. Early studies of HRI "in the wild" showed that people actually often interact with robots not individually but in groups, a task for which most early robots were poorly equipped (Šabanović et al., 2006). Increasingly, HRI studies group interactions involving two or more people, both inside and outside of the lab. For example, Leite et al. (2015) found that children

Figure 9.5 Units of analysis in HRI

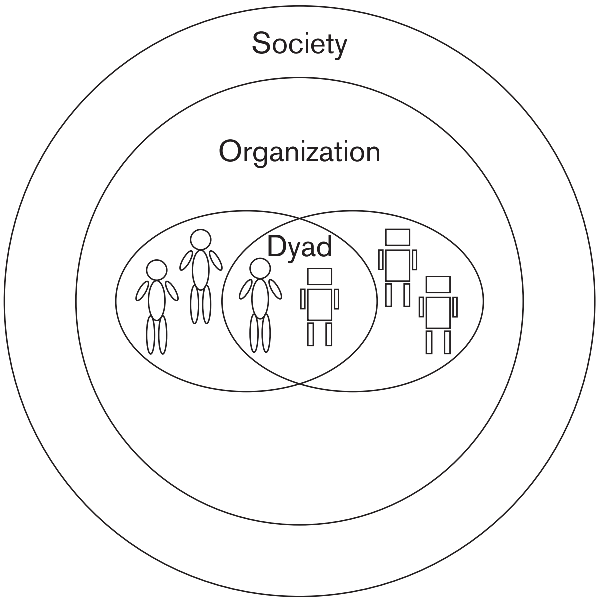

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

were better able to recall information from a story told by a group of robots when they interacted with them individually rather than in a group of three. Brscić et al. (2015) showed that children who come across a robot in a shopping mall abuse it only when they are in groups but not individually.

Social scientists distinguish between dyadic interactions and group interactions, and they consider the cognitive and behavioral aspects of each to be different. Groups bring in new perspectives on group effects, multi-party collaboration, team dynamics, and other such effects. Our vision of how we will be interacting with robots in the future also presupposes that there will be many robots in our environment, so another aspect of group HRI studies has been exploring how multiple robots can interact with people, whether in teams, in swarms, or simply as co-present robotic actors.

When robots collaborate in teams, they are often perceived as having more social agency. For example, Carpenter (2016) found that robots used in military bomb-disposal teams were often seen by the soldiers as members of the group and that soldiers became attached to such robots, even expressing feelings of sadness when their team's robot was destroyed.

The increasing availability of robots for research in applied settings beyond the laboratory opens up another unit of analysis. That is, we can look at how HRI occurs within organizations, such as educational and nursing institutions or even the military. By studying HRI within organizations, it is not only possible to see the effect of individual factors on HRI but also the effect of the broader context, such as how existing labor distributions or roles affect the robot's function and its acceptance by workers, how the robot is adapted to existing practices, and how institutional values affect people's interpretations of the robot.

Mutlu and Forlizzi (2008) showed that introducing a robot into an organization, for example, reduced work for some while increasing it for others. At the same time, it is plausible that people in different roles (e.g., manager, nurse, janitor) can have different perceptions of a robot based on how it affects their work. In another ethnographic study on the use of the seal-like Paro robot in a nursing home, Chang and Šabanović (2015) showed that having even one person who acts as an advocate for the robot in an organization can lead to more people committing to try it out and make it work for them, by modeling positive experiences of using the robot and creating a "positive feedback loop" that supports the long-term adoption of the robot. An organization can also be set up in a particular way to support the functions of a robot. Vertesi's ethnographic study of the NASA Rover team showed that the need to balance the robot as a scarce resource shared by many

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

different scientists and engineers worked well with an egalitarian setup of the team, where all team members needed to agree and say they were "happy" about the robot's next move (Vertesi, 2015). Now that it is possible, studying interactions between people and robots from an organizational standpoint seems necessary for the further development of the field and for our ability to design appropriate robots and social structures for the successful application of HRI in the real world (Jung and Hinds, 2018).

9.5 Choosing a robot for your study

Along with deciding how many and what types of participants you need to answer your research question, you will also need to decide on the characteristics of the robot(s) you need to use in your study. Factors you will need to decide on include the robot's appearance, functionality, and ease of use, among others. Whereas some of these decisions might be based on practical constraints, such as what types of robots are available to you or how much it would cost to purchase a new one, others will be guided by your research interests.

Robots can be seen as research tools, with which you can manipulate factors of interest and observe the effects of such manipulation on the outcome variables you want to measure. This approach is at the heart of experimental HRI research but can also be useful for more exploratory studies in which you may want to see if certain design factors produce differential effects on HRI. In order to use robots as a stimulus in HRI studies, we can manipulate their appearance, behavior, communication mode and style, and their role in the interaction, among other characteristics. HRI researchers often use off-the-shelf robots for their studies, but they also sometimes design and test their own prototypes. When deciding what kind of robot to use, determining which hardware and software capabilities would be best for the study and the appropriate level of autonomy of the robot are important considerations.

There are some commercial robots that lend themselves well to HRI studies, such as the Nao and Pepper (Softbank Robotics) or Paro (Intelligent System). Even when using a commercial robot, getting your robot up and running will require some basic programming skills. The Nao and Pepper robots can be programmed using a visual programming environment (Choreographe), which allows you to quickly go from the drawing board to a working robot. However, knowledge of more advanced control software and programming languages, such as the Robot Operating System (ROS), will allow you to greatly extend the repertoire of the robot's behavior and enrich the interaction. ROS contains a number of packages that implement sensory perception and visualization for different types of robots.

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

9.6 Setting up the mode of interaction

There are dozens of ways in which people and robots can be brought together for a study. People can meet an actual robot, or they can be shown pictures or videos of a robot. The robot can be fully autonomous or can be tele-operated by the experimenter. People can come to the lab, or the scientists can get out of the lab and bring their robots to the people. Sometimes, a single data point is all that is needed; on other occasions, only thousands of data points will do.

9.6.1 Wizard of Oz

In some HRI studies where the development of autonomous capabilities for the robot is not the focus of the research at hand, researchers commonly rely on the Wizard-of-Oz (WoZ) technique. WoZ involves deceiving study participants into thinking the robot is behaving autonomously, when it is actually being operated by a member of the research team. Research participants should then be informed about this deception in a post-experimental debriefing.

Using WoZ, researchers can "pretend" that their robot has interactional skills that it does not have, either because they require further technical development or because additional time or skill must be expended on programming the robot. The WoZ approach is particularly suitable in situations in which technology has developed to a degree at which it is almost usable for HRI, such as speech recognition. Using a wizard to recognize the users' utterances makes an experiment more robust and the robot's behavior more realistic and believable, enabling an actual interaction flow. It could, however, be considered problematic to completely fake an artificial intelligence (AI) system that can uphold a serious and prolonged conversation because that would be considered a very unrealistic level of capability for the robot. WoZ can also be used to test people's perceptions of more advanced capabilities, such as a robot that can understand and respond to the social context in very nuanced ways (e.g., see Kahn Jr. et al. (2012)). For experimental studies, it is also important to constrain the wizard's behavior so that the robot's behavior is kept consistent across conditions and does not introduce additional variation that can confound the analysis.

The WoZ method is named after a character in the movie of the same name. Dorothy and her companions set out to find the all mighty Wizard of Oz who can return Dorothy to Kansas. They encounter the wizard in his castle and are afraid of his gigantic appearance, his authoritative voice, and the smoke and fire he emits. Only when Dorothy's dog, Toto, pulls away a curtain do they notice Professor Marvel, who is operating the machinery that controls the stage show of the wizard. In HRI research, wizards often hide in the background and control the robot, giving the robot the semblance of having more advanced autonomous capabilities than it actually has. We all hope not to encounter Toto and be found out.

9.6.2 Real versus simulated interaction

Although the ideal way to gauge people's perceptions of and response to robots is in real-time, face-to-face interaction, it is still common for HRI researchers to present their participants only with video or photos of robots. In the field of HRI, there has been considerable discussion on whether video recordings of robots can be used as a replacement for live human–robot interactions. Whereas Dautenhahn et al. (2006a) argues that the two interaction styles are broadly equivalent, Bainbridge et al. (2011) concludes that participants had a more positive experience interacting with physically present robots than with a video representation. Powers et al. (2007) also found large attitude differences between participants interacting with a co-located robot in comparison to a remote robot. Therefore, using visual stimuli alone limits the generalizability of study findings but can be appropriate for exploratory studies of the effects of certain factors (e.g., perceptions of different robot forms (DiSalvo et al., 2002)) or for studies in which accessing the appropriate population can be difficult, such as cross-cultural samples. Using videos to present robots to participants can also enable researchers to avoid problems associated with a less controlled experiment that involves actual interaction. Finally, videos and photos are particularly amenable for use in studies that take advantage of online participant pools, whether through universities, word-of-mouth referrals, or services like Amazon's Mechanical Turk.

9.7 Selecting appropriate HRI measures

In HRI, as in psychology and other social sciences, researchers commonly distinguish between direct versus indirect measures to assess attitudes toward people or objects. In the example of the "gendered" robot study described earlier, the study design relied on *direct measurements* of the dependent variables—asking participants to rate the robot's warmth and authoritativeness, for example.

Within both correlational and experimental studies, self-reports are often used to assess the constructs of interest , such as concepts or variables. Self-report measures commonly bear high face validity, meaning that people usually directly know what the researchers want to measure when they read the items of the given questionnaire. On the other hand, this makes it easy for participants to amend their actual opinion with the aim of pleasing researchers, to represent themselves in a positive light or "be a good participant." This aspect also holds true for interview techniques, which are a way to gather an even more holistic picture of participants' thoughts and feelings toward both humans and robots. Interviews can be structured or semistructured in nature. In structured interviews, the interviewer asks a set of predetermined questions, often in a specific order, whereas in a semistructured interview, the interviewer has more leeway in deviating from the script; for example, some questions may be planned, but others may arise spontaneously during the interview. Both types often use questions to which interviewees can respond in their own words. Such open-ended responses, however, require labor-intensive coding after transcription of the interview's content. Such interviews might be a useful complement to questionnaires, though, as illustrated by de Graaf et al. (2017)'s use of data from a long-term survey and an interview to explore the reasons why people choose not to use a communication robot in their homes. As their work has shown, a research participant might feel highly uncomfortable in the presence of an unfamiliar robot.

In some cases, however, participants might be reluctant to report their true feelings and attitudes on a questionnaire or when talking directly to an interviewer. They may also not be aware of and able to report some unconsciously held beliefs. In that situation, it might be useful to complement your set of direct measures with indirect ones. Reaction times are often used as a proxy for factors that are harder to measure, such as attention or engagement. Indirect measures can include the use of eye tracking as an indicator of attentional focus and cognitive processing or the use of physiological measures such as heart rate or skin conductance to give researchers an idea of participants' level of stress experienced during HRI. Whereas computerized measures of attitude (e.g., a variant of the so-called Implicit Association Test [2] to measure anthropomorphization) have become increasingly popular, physiological correlates of attitudes toward robots or other technologies are less frequently used in contemporary research. Computerized and physiological measures are often more difficult to administer and require specific equipment, and ultimately, the findings are not always interpretable in an unambiguous manner. For example, skin conductance can indicate that someone is excited, but it cannot reveal whether the excitement is due to fear or enjoyment. In addition, a study in which the skin conductance of participants was measured as they interacted with a Nao robot showed that skin conductance readings are, unfortunately, not very conclusive (Kuchenbrandt et al., 2014).

To circumvent difficulties in interpreting results, it is helpful to use

[2] `https://implicit.harvard.edu/implicit/`

a combination of direct and indirect measures or several indirect measures at once in one study to ensure that you are indeed measuring the construct, or variable, that you intend to measure. As a researcher, you should aim to establish that all measurements used in your research reliably and validly assess what they are supposed to capture. This can be done by carefully pilot testing your study designs and measures used, developing and even formally validating new measures, or using widely accepted and validated measures that you find in the literature.

9.8 Research standards

9.8.1 Changing standards of statistical analysis

Because HRI is an interdisciplinary effort, some researchers might be more versed in statistics than others. Psychologists are typically well trained in statistical methods, and HRI researchers from other disciplines may find their advice most valuable when analyzing and reporting quantitative results.

"To call in the statistician after the experiment is done may be no more than asking him to perform a post-mortem examination: he may be able to say what the experiment died of."

As the famous quote by statistician Sir Ronald Aylmer Fisher (February 17, 1890–July 29, 1962) points out, the earlier on you ask for advice on your experimental design and analysis, the more useful it will be. Most universities offer statistical consultation of some sort, but even informal discussions with peers and professors may prove of tremendous value.

Although going into the details of the extensive number of statistical tests and procedures is beyond the scope of this chapter, the interested reader might consult the readily available literature, such as the work of Andy Field (Field, 2018).

Descriptive statistics, which give a summary overview of data without yet comparing conditions, should be used as the first stage of data analysis. Always provide means, standard deviations (when the data have a normal distribution—if this is not the case, one can provide a range), the number of participants, demographics (e.g., age and gender), and excluded data points, together with the reason for exclusion. Next, your study might require inferential statistics. These are statistics that compare two or more data sets. Selecting the correct statistical methods can be challenging, and this brief section serves as an overview rather than delving into the specifics of statistical methods.

There has been an interesting development in statistics. For a very long time, experimental science has relied on null hypothesis significance testing (NHST) to report on the importance of results. In this process, you calculate the probability that the data that were collected would have been observed if there had been no difference between the groups that were compared; in other words, you test the "null hypothesis" that nothing is going on. As such, the p-value is the chance that a researcher would conclude that there is a difference between the groups, whereas in fact, there is none; this is called a *Type I error*.

If the null hypothesis can be rejected (i.e., the probability or p-value that is obtained is less than or equal to some threshold, typically 0.05), the result may be considered "significant." On the face of it, this provides a useful means of characterizing the success (or failure) of a method or intervention.

However, in recent years, the overreliance on NHST and p-values has been questioned (Nuzzo, 2014). First, the threshold of $p \leq 0.05$ to call a result "significant" is arbitrary, and there is no scientifically valid reason for why we use the 0.05 threshold. Second, empirical results have suggested, and simulation studies have shown, that p-values are highly volatile in experiment replications. Repeating a study that has a significant p-value can result in p-values of the replication study in the range [0.00008,0.44] for 80% of the replication studies (Cumming, 2008). p-values are thus unreliable as a measure of how solid a result is. Third, p-values do not incorporate any information about how important a result is: a highly statistically significant result from the perspective of NHST does not say anything about the size of the observed experimental effect and thus cannot be used, by itself, to assess the importance or impact of the result. For example, the difference in life expectancy between two medical procedures could be "highly significant" even if one treatment makes you live only one hour longer than the other. Just reporting the statistical significance of experimental results says very little about how important those results actually are.

A more fundamental issue with NHST concerns the inferences one can draw from it. What is tested in NHST (the chance of finding the current data, provided that there is no true effect, or $p(A|B)$) is not what the researcher actually wants to know (the chance of a true effect, provided the current data, or $p(B|A)$). Although these *may* seem similar, their fundamental difference becomes clear when we consider sharks and death tolls. The chance of dying, provided that you are eaten by a shark ($p(dead|sharkbite)$), is pretty close to 1. However, the chance that you are eaten by a shark, provided that you are dying, $p(sharkbite|dead)$, is close to 0. In his entertaining and remarkably accessible paper "The Earth Is Round ($p < 0.05$)," Jacob Cohen explains some of the problems with NHST in further depth (Cohen, 1994).

We can remedy this by reporting not only the p-values but also the

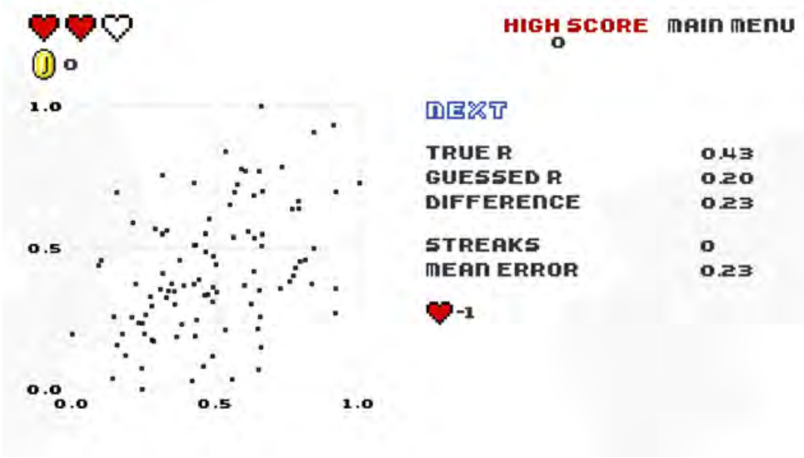

Figure 9.6 Take a look at the plot shown here. If you had to make a guess, how strongly would you say the two variables are correlated? It has been shown that people find it very hard to infer the strength of a relationship from plots. On the website `http://www.guessthecorrelation.com`, you can try for yourself. (By the way, the correlation in the picture is $r = .43$, which is considered a medium effect.)

confidence intervals (CIs) of our data. CIs do not compare data and therefore cannot be used to say if results are significant or not. Instead, they report on how confident we are that the mean of the data lies between a minimum and maximum value of the CI. When reporting the 95% CI of data, this means that in a replication study, the mean of the replication data will have an 83% chance of being within the CI of the original experiment.

Finally, when comparing data, it is also standard practice to report the effect size. Although the p-value can indicate whether the difference between data is statistically significant, it does not say anything about how important that result really is. The effect size is a standardized, scale-free measure of the relative size of the effect of an intervention. It is often written as d, is a positive number, and is calculated as $d = (\mu_E - \mu_C)/SD$, where μ_E and μ_C are the mean of the experimental group and the control group, respectively, and SD is typically the standard deviation of the control group. Cohen (1977) considers $d < 0.2$ to be a small effect size, $0.2 \leq d < 0.8$ to be a medium effect size, and $d \geq 0.8$ to be a large effect size. To help put these numbers into context, the difference in height between 15-year-old and 16-year-old girls in the United States has a small effect size of $d = 0.2$. The difference between the heights of 14-year-old and 18-year-old girls has $d = 0.5$ and is a medium effect size, visible to the naked eye. The difference in height between 13-year-old and 18-year-old girls is $d = 0.8$ and is a large effect size, immediately obvious to an observer. To see what a correlation of "medium effect" looks like in a plot, see Figure 9.6.

Reporting CIs and effect sizes conveys additional information, complementing a test of statistical significance, but the emphasis is on the magnitude and relative importance of an effect, rather than the statistical significance—a measure all too fickle and overinterpreted (Coe, 2002).

9.8.2 Power

The p-value reflects the chance of a researcher wrongfully concluding that there is a difference between the groups (a Type I error). To avoid making these errors, scientists keep the threshold for what is considered a significant effect low, usually at $p \leq .05$. With this threshold, a researcher will make a Type I error (wrongfully conclude there is an effect) 1 in 20 times.

However, as mentioned in Section 9.3, the opposite is possible as well: a researcher may conduct an experiment, gather data, and then wrongfully conclude that there is no effect. This has been, not very creatively, named a *Type II error*. Type I and II errors can be avoided by making sure your experiment has sufficient statistical power—this usually means that you have to collect either enough participants or enough data points per participant. This can be tricky, and the number of participants needed can increase dramatically depending on how complicated your study design is or how small the effect is you're hoping to detect. Software such as G*Power (Faul et al., 2007) allows you to calculate the power both before and after a study.

Although most researchers are primarily concerned with avoiding Type I errors, one could argue that wrongfully concluding that a certain factor does not matter in HRI can be equally damaging to future research (not to speak of the waste of the resources spent on collecting the data, or the frustration and disappointment level of the researchers involved). It thus is well worth the effort to keep both types of errors in mind when designing an experiment.

9.8.3 Generalizability and replication

Recent developments in the social sciences have illustrated the importance of being able to reproduce and replicate research findings. Initiatives such as the Open Science Framework[3] push this idea in the field of psychology to shed light on the validity of both novel and allegedly established findings in the field. In HRI, the reproducibility of research has been less prominent on the research agenda, but the recent concerns in the social science community have brought these topics into the purview of HRI researchers as well (Irfan et al., 2018). Replication of HRI results is also now more possible than before because of the wide availability of certain robot platforms (e.g., Nao or Baxter), in contrast to earlier reliance by researchers on bespoke platforms. There has been a drive for sharing code for commonly available robots and, if possible, making the experimental procedures available to other HRI researchers in order to enable them to run the same experiment in their own labs, testing the generalizability of a certain research question across con-

[3] `https://osf.io`

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

texts (Baxter et al., 2016). Overall, the notion of generalizability is highly important even though representative samples are hard to obtain in HRI research.

The choice of methodology also affects the degree to which we can generalize from our HRI studies in the laboratory to those findings obtained from field studies. Developing new robots, applying robots in different contexts, and understanding the potential consequences of robots for people in daily life may require a combination of the methods mentioned in this chapter. This does not need to be done in one research project or by a single researcher but could be accomplished by the HRI research community over time.

9.8.4 Ethical considerations in HRI studies

Last but not least, one important aspect to consider when dealing with human participants in HRI studies is the need to take into account the ethics of human-subjects research. Any research that involves human participants, whether correlational or experimental, qualitative or quantitative, online or in person, requires participants' informed consent before the research is started. That is, participants are informed about the nature of the study and what to expect, with an emphasis on the voluntary nature of their participation and information regarding the risk and benefits of taking part in a given research study. Before starting a study, either online or in the real world, participants have to declare that they understand what they will be asked to do and what will be done with the collected data and that they consent to participating. Many universities and institutions have specific guidelines on how participants can be recruited and informed about their participation in research studies. Researchers need to be aware of this and follow all policies to be able to present their results for publication following the study.

Sometimes, however, it is impossible to fully disclose the actual goals of the given research project. In that case, a cover story or deception is used. For instance, in WoZ studies, participants are led to believe that a robot can behave autonomously. In that case, it is key to provide post-experimental information, a so-called *debriefing*, to participants so that they do not go home from the study thinking that robots are currently able to function fully autonomously.

This is even more critical if a robot might provide the human interaction partner with fictitious feedback about the human's personality or performance. Of course, the participants then must be debriefed about the reason for providing made-up feedback, and they must be informed that this feedback was actually bogus. Again, this serves to ensure participants' psychological well-being beyond the duration of the study.

In the case of qualitative research, initial information about the study goals given to participants may be more cursory, but the common practice is to later inform study participants of the findings if they are interested. In some cases, researchers might even discuss their interpretations of the data with participants or collaboratively develop interpretations and future robot design and implementation guidelines based on the results.

In HRI research, we also have to consider the ethical aspects of having humans involved with robots—both in terms of physical and psychological safety and in terms of the implications an interaction could have for a given individual. Think, for example, of an elderly person who has had a robot in his or her home for a certain amount of time and might have gotten attached to the robot companion. Consequently, the day the robot is taken away, this will cause distress. Users' emotional reactions toward robots, the attachments they might build, and the void that results when the robot is taken away must be considered.

To make sure that you are complying with ethics regulations, you may consult with the various codes of ethical conduct, such as those provided by the American Psychological Association,[4] the American Anthropological Association,[5] or the Association for Computing Machinery.[6] Your university's ethics committee may provide more detailed feedback regarding your specific research study. Note that ethics approval is a requirement for publication in many scientific journals, so consider getting it before you start your data collection.

Along with ethical behavior toward research participants, researchers should also reflect on the ethical implications of their research aims, questions, and findings and make choices about what types of research to pursue, and how to go about it, with these implications in mind. Such ethical considerations can include questions about where to seek out and accept funding, whether to participate in research that may inform particular corporations or governments, and even how to structure one's relationship with participants and their ability to provide input on the methods and presentation of research results.

More generally, the ethical and social consequences of the implementation of robots in society have to be taken into account. In most contemporary research projects that deal with smart homes or deployment of robots in homes, care facilities, or public spaces, these aspects have to be investigated and addressed. Considering the ethical implications of digitalization and a potential hybrid human–robot society is a key societal issue that is now discussed at large, not solely by robot ethicists and philosophers.

[4] http://www.apa.org/ethics/code/

[5] https://s3.amazonaws.com/rdcms-aaa/files/production/public/FileDownloads/pdfs/issues/policy-advocacy/upload/ethicscode.pdf

[6] https://www.acm.org/about-acm/code-of-ethics

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

9.9 Conclusion

HRI studies have a lot in common with work in several social science disciplines, including experimental psychology, anthropology, and sociology. It is good practice to be aware of scholarly norms and practices in the field or fields relevant to your work. HRI researchers are expected to be aware of and adopt the same rigor when collecting and reporting data as other scholars using the methods they have chosen.

HRI is also sensitive to the same problems that have plagued the social sciences for over a century. For example, in the drive to come up with original work, HRI experiments are almost never repeated. There is also a considerable publication bias, with positive results more likely to make it to publication, whereas negative results, less exciting results, or less conclusive findings tend not to get published or to go unnoticed. However, HRI has opportunities that were not on offer until recently. Experimental data, including large video logs, can now be fully stored and shared with others, ready for scrutiny or additional analyses. Methods, protocols, and results are now more available than ever before, largely due to the drive toward open-access publishing and preregistration of experiments.

HRI researchers can also find relevant methodological approaches and discussions in the related field of human–computer interaction (HCI), which has a longer history of performing user studies, system evaluations, and theory building around the use of computing technologies in society and can provide guidelines and critical perspectives pertinent to HRI research. HRI researchers can learn from discussions about how to incorporate contextual variables into their work, how to think critically about design and study methods, and how to work more closely with the potential users of new robotic technologies through prior work in HCI. It is also, however, important to remember that HRI deals with robots, which are not only a different, embodied technology compared to computers but also pose different technical and social challenges for research.

Questions for you to think about:

- In some instances, it is not ethical or possible to answer a research question with an experiment. Can you think of such an instance? How would you address ethical issues related to the setup of your study? How might you address concerns about the inclusion of vulnerable populations (e.g., children, older adults with cognitive impairments) in your study?
- "Significance" has been considered a misleading term because it says nothing about the relevance of a finding. Can you think of

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

a situation where finding a significant small effect is relevant? What about a situation where it is irrelevant?
- Say you want to set up an experiment in which you assess how well a robot tutor teaches children. How would you set up your study? How would you measure the robot's ability as a tutor? What confounding factors do you expect?
- HRI studies often seek to address people's subjective experiences of robots—their enjoyment of the interaction, for example. How would you measure enjoyment, incorporating both direct and indirect and subjective and behavioral measures? How would you make sure that your enjoyment measure has construct validity—that it is actually measuring enjoyment with the robot, not just general happiness, or reflecting the participant trying to please the experimenter?

Further reading:

- Cindy L. Bethel and Robin R. Murphy. Review of human studies methods in HRI and recommendations. *International Journal of Social Robotics*, 2(4):347–359, 2010. doi: 10.1007/s12369-010-0064-9. URL `https://doi.org/10.1007/s12369-010-0064-9`
- Geoff Cumming. The new statistics: Why and how. *Psychological Science*, 25(1):7–29, 2014. doi: 10.1177/0956797613504966. URL `http://journals.sagepub.com/doi/10.1177/0956797613504966`
- Andy Field and Graham Hole. *How to design and report experiments*. Sage, Thousand Oaks, CA, 2002. ISBN 978085702829. URL `http://www.worldcat.org/title/how-to-design-and-report-experiments/oclc/961100072`
- Laurel D. Riek. Wizard of Oz studies in HRI: A systematic review and new reporting guidelines. *Journal of Human-Robot Interaction*, 1(1):119–136, 2012. doi: 10.5898/JHRI.1.1.Riek. URL `https://doi.org/10.5898/JHRI.1.1.Riek`
- Paul Baxter, James Kennedy, Emmanuel Senft, Severin Lemaignan, and Tony Belpaeme. From characterising three years of HRI to methodology and reporting recommendations. In *The 11th ACM/IEEE International Conference on Human-Robot Interaction*, pages 391–398. IEEE Press, 2016. ISBN 978-1-4673-8370-7. doi: 10.1109/HRI.2016.7451777. URL `https://doi.org/10.1109/HRI.2016.7451777`
- Selma Šabanović, Marek P. Michalowski, and Reid Simmons. Robots in the wild: Observing human-robot social interaction

outside the lab. In *9th IEEE International Workshop on Advanced Motion Control*, pages 596–601. IEEE, 2006. ISBN 0-7803-9511-1. doi: 10.1109/AMC.2006.1631758. URL `https://doi.org/10.1109/AMC.2006.1631758`

- James E. Young, JaYoung Sung, Amy Voida, Ehud Sharlin, Takeo Igarashi, Henrik I. Christensen, and Rebecca E. Grinter. Evaluating human-robot interaction. *International Journal of Social Robotics*, 3(1):53–67, 2011. doi: 10.1007/s12369-010-0081-8. URL `https://doi.org/10.1007/s12369-010-0081-8`

10

Applications

What is covered in this chapter:

- The diverse areas of robot applications where human–robot interaction (HRI) is an important component;
- Applications beyond robots that are studied in a research context;
- Possible future applications;
- Potential problems that would need to be solved when HRI has a larger role in our society.

Human–robot interaction (HRI) has numerous applications expected to make a positive difference in people's lives. HRI is increasingly getting traction in the technology market, and although most applications are still being developed in the academic sphere, adventurous start-ups have popped up that are developing and selling HRI applications, and established information technology (IT) industries are keen to understand and develop technologies that allow robots or robot technology to interact successfully with people. Not all of these enterprises turn out to be successful. Sony, for example, was one of the pioneers of commercial robotics with its Aibo (see Figure 10.1) and Qrio (see Figure 10.2) robots, only to stop its efforts in the field in 2006. However, Sony's efforts were recently rekindled, with a new Aibo appearing in 2018 (see Figure 3.2). Another example is the Bosch company, which initially supported Mayfield Robotics in developing the Kuri home robot but stopped the project before the official product launch.

A successful HRI application means something different depending on the perspective one takes: the notion of what constitutes success is very different for a researcher compared to an entrepreneur. Whereas a researcher will be interested in measurable outcomes of the robot's use and usability, an entrepreneur might be less concerned about the effectiveness of the robot and will be happy with a "good enough" technical solution that can be brought to market, thus preferring sales figures over scientific figures. Some may even develop unsuccessful applications

Figure 10.1 The Sony Aibo ERS-7 (2003–2005) with the Nao (2008–present) robot.

on purpose for the entertainment value or to inspire people to think more critically about the uses and design of robotic technology (see the accompanying text box for examples).


Figure 10.2 Sony's Qrio robot (left) (2003–2006) (source: Sony) and Mayfield Robotics' Kuri (2016–2018) (source: Mayfield Robotics)—two robots that never made it to the consumer market.

The self-crowned "Queen of Shitty Robots," Simone Giertz is a nonengineer robot enthusiast who designs service robots that usually perform poorly in their intended application. Her videos on the testing of her different creations not only have entertainment value but also demonstrate how designing robots for seemingly simple tasks can prove to be quite challenging.[a] White's "Helpless Robot," on the other hand, is a machine with a passive personality that asks people to move it around the room, opening up questions about the meaning of machine autonomy and whether our machines serve us or whether it is the other way around.[b]

[a] `https://www.youtube.com/channel/UC3KEoMzNz8eYnwBC34RaKCQ/`
[b] `http://www.year01.com/archive/helpless/`

For now, most robot applications remain at the research stage, but this is expected to change rapidly. The first wave of commercial success in robotics took place in automating industrial production; the next wave of commercial success is expected to come from introducing robots in dynamic and open environments populated by people in customer service, companionship, and socially and physically assistive roles. It is here that HRI has its major role to play: a solid understanding of how robots should behave around people, and how people respond to and benefit from robots, is needed to make the next robot wave a success (Haegele, 2016). The following sections provide an overview of the various types of robots that have been tested in the lab and in the field, starting with service robots.

10.1 Service robots

A novel robot often attracts people's attention; in public spaces like shops, visitors become interested and approach, and children crowd

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 10.3 The Robovie robot as a museum guide (2006).

around it. This makes robots an ideal asset for customer service settings. Many such applications have already been successfully tested in field research and have been deployed in grocery stores or bank branches (e.g., Pepper providing service at HSBC in the United States).

10.1.1 Tour guide robots

One of the applications developed in the early years of HRI research is the tour guide robot (Burgard et al., 1998; Shiomi et al., 2006). Typically, a tour guide robot moves from one location to another while providing information about nearby entities; some of them take the user to a requested location. This robot application involves navigational interaction (e.g., the robot safely moving around in an environment it shares with humans) and face-to-face interaction with its users (see Figure 10.3).

There are many instances of successful tour guide applications. One such application is in a museum setting, where a mobile robot is left to autonomously navigate around. Visitors are invited to use a user interface on the robot to indicate whether they want to have a tour guide. Once a tour is requested, the robot leads the visitors to several exhibits, providing a brief explanation at each (Burgard et al., 1998). HRI researchers experimenting with museum robots have found that giving the robot the ability to display emotions can enrich the educational experience and allow the robot to better manage its interactions

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

with people, such as getting them to move out of its way by expressing frustration (Nourbakhsh et al., 1999). An alternative application concerns the retail context, when a customer may want to know where in the store a specific item is kept, and a robot takes the lead to show him or her the way to the appropriate shelf (Gross et al., 2009). A final example is the airport, where a robot can escort travelers to the gate for their next flight (Triebel et al., 2016).

It is easy to imagine similar scenarios where robots would be helpful. For example, it is common for people to escort other people, either because they need physical assistance or because they want to be accompanied, in daily interactions. Robots could be used in this context in the future. One such application being developed by HRI researchers is a guide robot for individuals with visual impairments (Feng et al., 2015). Although the current limitations in robotic hardware and HRI capabilities prevent such uses in the present, technical advancements and further HRI research should enable us to have robots with faster velocity and better navigation capability in human crowds that can be applied for accompanying users in a broader range of environments.

10.1.2 Receptionist robots

Receptionist robots are placed at a reception desk and interact with visitors, typically offering information through spoken-language conversation. For instance, Gockley et al. (2005) studied people's interactions with a robot with a display for its head as a receptionist at a university (see Figure 10.4). The robot was able to provide directions and would share daily stories with people who came to chat with it. It also turned out that people were sensitive to the robot's moods, and the length of their interactions with it changed based on whether the robot displayed a happy, sad, or neutral expression (Gockley et al., 2006). Outside of the research setting, android robots have been used as receptionists in hotels. In this case, users use a graphical user interface (GUI) to proceed through the check-in process, attended by an android robot and a small humanoid robot that offers greetings to the visitors.

Figure 10.4
Receptionist robot.

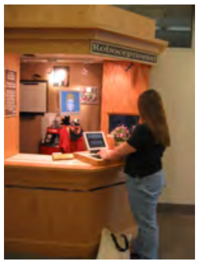

10.1.3 Robots for sales promotion

Another straightforward application of service robots is product promotion in the retail context. In this setting, robots can function as proxies for store clerks, informing customers about the promotions offered by the store. Because people are naturally curious about robots, these robots can easily attract the attention of potential visitors, who will stop to listen and then look around. In Japan, Pepper is already used for this purpose. In the typical use case, robots are not necessarily proactive but instead wait for visitors to initiate interaction. In the research context, researchers study robots that proactively approach customers to offer promotions (Satake et al., 2009). For instance, the famous Geminoid android robot has been deployed in a shopping mall in Japan to boost sales (Watanabe et al., 2015).

10.2 Robots for learning

Social robots have been shown to be particularly effective for assisting in learning and education (Mubin et al., 2013). This should not be confused with the use of robots as an educational tool to teach mathematics, programming, or engineering, such as Lego Mindstorms. Robots can take on various roles in the process of learning: The robot can act as a teacher, taking the students through the curriculum and offering testing opportunities to assess knowledge. As a tutor, a robot would support the teacher in his or her teaching (Kanda et al., 2004). This role is actually preferred by teachers and students (Reich-Stiebert and Eyssel, 2016). However, the robot is also often presented as a peer. The peer-like robot has a similar level of knowledge as the learner, and the learner and robot take a learning journey together, with the robot adapting its performance to that of the learner. At the far extreme is the robot that needs to be completely taught by the student. This approach, known as a care-receiving robot or teachable agent, is effective for two reasons. First, teaching a subject often leads to mastery of that subject, and second, having a less knowledgeable peer can boost the learner's confidence (Hood et al., 2015; Tanaka and Kimura, 2010). Finally, robots could also be used as a sidekick for teachers. In this role, the robot spices up the lesson and makes the learning more entertaining, thus capturing student interest (Alemi et al., 2014).

Tutoring robots may take over specific tasks from the teacher. Because teachers typically deal with class sizes of more than 20 students, they are required to teach to the mean of the class using a broad rather than a personalized style. It has been shown that tutoring has a strong impact on learning. Bloom (1984) found that one-to-one tutoring resulted in a 2-standard-deviation improvement against a control group, concluding that "the average tutored student was above 98% of the students in the control class." Although research has since shown that the effects are not as large as first observed, there is nonetheless a distinct advantage to the one-to-one tutoring approach (VanLehn, 2011). Social robots in education capitalize on this by offering a one-to-one, personalized tutoring experience.

Robots have been used to teach a wide range of topics, from mathematics to languages, both to adults and children. The main contribution of the robot seems to be that its physical presence promotes learning. Although computer-based tutoring programs, also known as intelligent

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

tutoring systems (ITSs), are effective (VanLehn, 2011), the social robot adds to this through its social and physical presence. Studies have shown that robots offer a distinct advantage over on-screen social agents or ITSs, and the students learn faster and learn more when tutored by a robot as compared to alternative technologies (Kennedy et al., 2015; Leyzberg et al., 2012, e.g.). The reasons for this are still unclear: it might be that the social and physical presence of the robot engages the learner more than just on-screen delivery and feedback, or it might be that the learning experience is a more multimodal experience, thus resulting in a richer and embodied pedagogical exchange (Mayer and DaPra, 2012)—of course, a combination of these two is also possible. It may come as no surprise that socially supportive robots perform much better (Saerbeck et al., 2010). Some socially interactive behaviors can also backfire in learning contexts, leading the student to interpret the robot as a peer rather than a teacher and to engage with it socially rather than focusing on achieving certain learning goals (Kennedy et al., 2015). HRI research is therefore necessary to guide the development of robots that can effectively support learning.

10.3 Robots for entertainment

10.3.1 Pet and toy robots

Robotic pets and toys were among the first commercial robot applications for personal use. After the first doglike robot, Aibo (Fujita, 2001), appeared on the market in 1999 (see Figure 10.1), the development of many other entertainment robots soon followed. Compared to other robotic applications, entertainment robots have been easier to get to market because the functions they perform do not have to be as advanced, and they often use preprogrammed capabilities, such as dancing, talking, burping, and even seeming to develop their knowledge by simply starting to use more advanced preprogrammed skills after a period of time. Some of the most popular robotic toys over the years have been Furby, Sony's Aibo robot dog, and more recently, Cosmo. Lego Mindstorms was a market leader in the educational toy robot niche but has recently been followed by a slew of robots that allow children to learn how to code and think computationally, such as Dash and Dot and Ozobot, among many others. The WowWee company is another market leader, with many different robots, including the humanoid robots Robosapiens and Femisapiens and a mobile home robot. The company Sphero developed a robotic ball that could be remote-controlled; following the release of the new series of Star Wars films in 2015, the company amended the design to represent the BB-8 droid, which became one of the most popular holiday toys of that season.

Although most entertainment robots target children and adolescents,

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 10.5 Pleo Robot (2006–present). (Source: Max Braun)

many are also enjoyed by adults. The Aibo in particular was very popular with adults, who even started a "black market" of Aibo parts when the robot was discontinued by Sony in 2006. As mentioned earlier, Sony introduced a brand-new version of Aibo in 2018.

Pleo (see Figure 10.5), a camarasaurus rex robot platform, provides a similar complexity of interaction, with various modes of personality and behavior that adapt and change across time and users. These examples show that many robot toys are not necessarily social or humanlike in appearance, but they still elicit strong social responses in children and adult consumers alike.

Considering the variety of ways in which robots can provide entertainment and the popularity of robots among the public in general, it is not surprising that the market for toy robots has been and is expected to stay one of the largest for personal robots (Haegele, 2016).

10.3.2 Robots for exhibitions

Robots are often used in exhibitions and theme parks to entertain audiences. These often-animatronic devices are very robust; they must play the same animation script sometimes hundreds of times per day, with only a brief moment for maintenance between performances. Some robots intentionally look like robots, but others resemble animals, for example, dinosaurs (see Figure 10.6), or people. In these cases, the robot has a flexible latex skin, which has been carefully painted to reflect realistic skin coloration and patterns. Most of these animatronic robots have no autonomy: they play a prerecorded script of animation timed to a soundtrack. In rare cases, the robot may have limited autonomy, such as the ability to focus on members in the audience while

Figure 10.6
Animatronic robot.

speaking. A popular example of the use of animatronic robots is the Hall of Presidents located in the Walt Disney World Resort.

10.3.3 Robots in the performing arts

Robots are also sometimes used in the performing arts. One of the first robot performance art pieces was Senster, created in 1970 for Philips' Evoluon in Eindhoven, the Netherlands (Reichardt, 1978). Senster was an electro-hydraulic structure shaped after a lobster's claw, with six hinged joints. It registered and responded to sound and movement from the environment. It was on display until 1974, when it was dismantled. More recently, 20 Nao robots performed a synchronized dance recital for France Pavilion Day (June 21) at the Shanghai 2010 Expo.

Not all art applications have to be for a broader public. Home theater systems might soon become what their name promises. Imagine a future in which you download the theater script of *Romeo and Juliet* into your robots. You can then either watch the robots perform the play or join in yourself. It is important to note that a major use of robotics—both in the past and currently—is to automate tasks that we do not want to perform ourselves. Industrial robots, for example, were introduced to relieve us of difficult and repetitive manual labor. There is little use in automating tasks that we actually enjoy doing. This does not mean that there is no place for robots in the theater–plays that actually deal with robots should, of course, be cast with robots (Chikaraishi et al., 2017).

Furthermore, there are many ways in which robots can interact with people in art performances, to which the future social robots could contribute as a human counterpart. For example, Hoffman and Weinberg (2010) developed a marimba-playing robot that joins a jazz-like session with a human player. Kahn Jr. et al. (2014) revealed that a robot can partner with a human to enhance human creativity in the art-creation context. Nishiguchi et al. (2017) suggest that developing robots that can perform as actors in a play alongside humans can also be a way to develop more humanlike behaviors for robots.

10.3.4 Sex robots

Along with toy robots aimed at the child market, there are also embodied robots and virtual reality (VR) interfaces for the fulfillment of adult entertainment needs. Colloquially known as "sex robots," diverse robotic platforms offer varying levels of humanlike appearance and behavioral response. The Real Doll company, which develops hyper-realistic sex dolls, is working on adding robotic capabilities, including an emotive face and responses, to its base models. Several other producers have developed prototypes of sex robots, although none has yet come to market. It is envisioned that the sex robot industry will continue to grow over the coming years.

Levy (2009) provides a history of sex machines and speculates about our future intimate relationships with robots.

10.4 Robots in healthcare and therapy

Healthcare and therapy represent prominent domains of application for robotics. In these domains, social robots are used to offer support, education, and diversion to patients, with an eye toward improving healthcare and therapy outcomes. The practice of using social robots in healthcare is referred to as socially assistive robotics (SAR) (Tapus et al., 2007; Feil-Seifer and Matarić, 2011) and is often targeted to older adults (Broadbent et al., 2009; Broekens et al., 2009).

10.4.1 Robots for senior citizens

Although senior citizens and people with mild cognitive impairments are a key target audience for robot developers who want to offer technology-mediated social, emotional, and cognitive rehabilitation and diversion, there are other target groups that can benefit from social robots.

For example, the Paro robot is a seal-like robot equipped with sensors that allow it to detect when it is being picked up or stroked (see

Figure 10.7 The ElliQ robot (2019–present) from Intuition Robotics is designed to interact with senior citizens. (Source: Intuition Robotics)

Figure 2.6). It can respond by wriggling and making seal-like noises. Paro has been used in a multitude of studies with elderly people, and positive psychological, physiological, and social effects of long-term interaction with the robot have been documented (Wada and Shibata, 2007). The robot is used as a companion in care homes and stimulates not only human–robot interactions but also interactions between the residents. It has been able to reduce feelings of loneliness and improve the residents' quality of life. Paro has been commercially available in Japan since 2006 and in the United States and Europe since 2009. It is interesting to note that although it is purchased by many individuals for home use in Japan, in Europe and the United States, the robot is almost exclusively purchased by healthcare institutions and companies. Furthermore, some robots, such as NEC's Papero (see Figure 10.8, have only ever been released in Japan.

Robots can also provide reminders for people to take their medications (Pineau et al., 2003) and can provide pre-clinic or tele-clinic support at home, thus reducing costs for medical services (Robinson et al., 2014).

10.4.2 Robots for people with autism spectrum disorder

Children and adults with autism spectrum disorder (ASD) are another group for which social robots are often developed and used. It has been shown that people with ASD generally respond well to robots, and there has been a large body of research looking into how robots can be effectively used to support ASD therapy (Diehl et al., 2012; Scassellati et al., 2012; Thill et al., 2012). Many types of robots have been used

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 10.8 NEC's Papero robot has been available in different versions, such as Papero R-100, Papero Mini, and Papero i (1997–present)

in a therapeutic context to support children with ASD (Robins et al., 2009; Pop et al., 2013) (see Figure 10.9). These include a wide range from humanoid robots, such as Kaspar and Nao, to zoomorphic robots, such as Elvis and Pleo. The predictable nature of robot behavior and the fact that robots are nonjudgmental have been credited as potential reasons why using them in interactions and therapeutic interventions with individuals with ASD is successful. The robots are either used as a focal point for the interaction between the therapist and the patient or

Figure 10.9 A range of robots used in autism spectrum disorder therapy. From left to right, Nao (2008–present), Elvis (2018–present), Kaspar (2009–present) and Zeno (2012–present). (Source: Christoph Bartneck, Bram Vanderborght, Greet Van de Perre, Adaptive Systems Research Group, University of Hertfordshire, Steve Jurvetson)

are used to train and improve children's social competences and their ability to regulate and interpret emotions.

10.4.3 Robots for rehabilitation

Robots are also used to support physical rehabilitation. This can be through offering physiotherapy, and through providing encouragement and mental support. Social robots have been shown to be effective in cardiac-focused rehabilitation by providing encouragement and social facilitation during cardiac exercises (Kang et al., 2005; Lara et al., 2017). Robots can also be used to encourage users to adopt healthy practices or to change unhealthy habits. For example, Kidd and Breazeal (2007) describe a robot that acts as a weight-loss coach, and Belpaeme et al. (2012) describe the use of a robot to support children diagnosed with diabetes. Kidd's early research developed into a robotic start-up and healthcare robot called Mabu.

Robots can also be used as orthotic or prosthetic devices. The restoration of the function of the lower limbs, arms, and hands through robotics has received considerable attention (Bogue, 2009). Although these developments are largely the concern of mechatronics, there is a role for HRI in the study of the acceptance and usability of robotic prostheses.

10.5 Robots as personal assistants

Smart-home assistants, unobtrusive devices that are placed in the home or the office and are often voice-operated, have been a recent and largely unexpected success of cloud-connected technology. Technology giants such as Amazon, Google, Microsoft, Apple, and Samsung have raced to build voice-operated assistants, and some offer hardware products that are built around this technology. Amazon's Alexa, Apple's Siri, Microsoft's Cortana, and the Google Assistant have found embodiment on a range of devices, with shapes and sizes ranging from a hockey puck to a shoe box. These devices offer a vast range of services, but they are most often used to request simple information, such as the time, weather, or traffic, or to stream music. These devices can engage

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

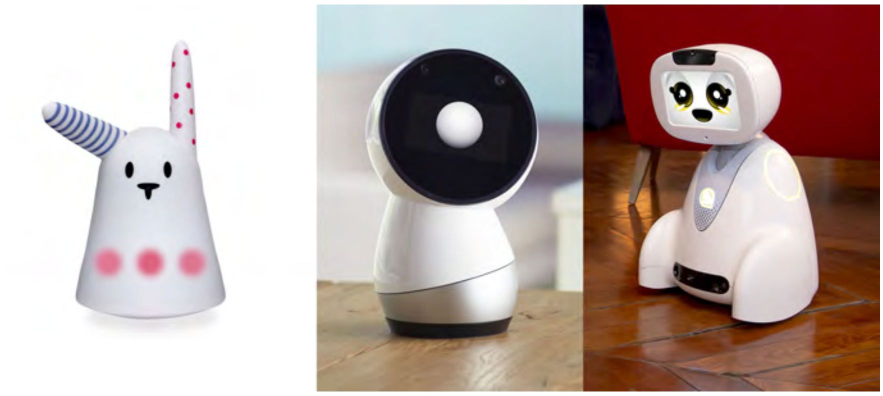

Figure 10.10 Personal assistant robots: from left to right, the Jibo robot (2017–2018), the Nabaztag (2009–2011) robot, and the Buddy Robot (2018–present). (Source: Jibo Inc., Blue Frog Robotics)

in only very short social exchanges, often limited to chitchat, such as telling a joke.

Recently, a number of commercial ventures have been launched that offer social robots as personal home assistants, perhaps eventually to rival existing smart-home assistants. Personal robotic assistants are devices that have no physical manipulation or locomotion capabilities. Instead, they have a distinct social presence and have visual features suggestive of their ability to interact socially, such as eyes, ears, or a mouth (see Figure 10.10). They might be motorized and can track the user around the room, giving the impression of being aware of the people in the environment. Although personal robotic assistants provide services similar to those of smart-home assistants, their social presence offers an opportunity that is unique to social robots. For instance, in addition to playing music, a social personal assistant robot would express its engagement with the music so that users would feel like they are listening to the music together with the robot (Hoffman and Vanunu, 2013). These robots can be used as surveillance devices, act as communicative intermediates, engage in richer games, tell stories, or be used to provide encouragement or incentives.

10.6 Service robots

Service robots are designed to help humans in various onerous, often called "dull, dirty, and dangerous," tasks. The tasks performed by such robots are typically simple and repetitive, and they often do not involve explicit interaction with people. HRI research considers such robots when they operate in everyday human contexts and therefore come into regular contact with people, including house-cleaning and delivery robots and robots that offer personal assistance.

Cleaning robots

Cleaning robots are widely used in homes. The most well-known cleaning robot is Roomba; it is also the most commercially successful personal service robot to date. It is a small robot, approximately 30 cm in diameter, that has two wheels to enable it to move around, dust sensors to know where it needs to clean, cliff sensors to avoid falling down the stairs, and of course, vacuuming capability. It moves around randomly in a house, turning when it comes to a wall, and over a period of time, it manages to clean up the room. (In general, that is; pets can undermine this goal horribly. See the accompanying box). There are many other vacuum-cleaning robots for the home, as well as the mopping robot Scooba.

Dreaded by every pet-owning Roomba user, the *Poopocalypse* is the unfortunate yet inevitable event where a pet leaves a dropping somewhere in the house, and the Roomba encounters it before the owner can clean it up, spreading it all across the house. These incidents are common enough that iRobot formulated an official response, warning Roomba users to not use their Roomba unsupervised if they own a pet (Solon, 2016).

Commercial service robots coming onto the market have provided HRI researchers with opportunities to study how people respond to and use such robots in everyday circumstances. Fink and Kaplan performed ethnographic studies of Roombas in user homes to identify common use patterns, and they also noticed how users prep their homes so that Roomba can do its job (Fink et al., 2013). Other researchers have found that users sometimes like to display Roombas as a sophisticated technology, whereas at other times, they try to disguise or hide them because they are deemed unsightly (Sung et al., 2007, 2009). Forlizzi and DiSalvo (2006) also explored how people's models of service affect the way they expect robots to interact with them, including how robots can best recover from mistakes made while providing services, such as bringing users the wrong drink.

Delivery robots

Delivery robots carry objects from one place to another. Amazon uses delivery robots in its warehouses. They are also used in other environments, such as the Aetheon TUG robot used in hospitals. Some hotels use robots to deliver goods from the service desk to guest rooms. More recently, mobile robots are now being used to make meal deliveries in San Francisco, California, through Yelp's Eat24 app. There are many start-ups that seek to provide delivery robots. Although perhaps desirable for the direct users, these robots sometimes turn out to be a nuisance for bystanders, who have to dodge them on already-busy city streets. Mutlu and Forlizzi (2008) showed that the workflow and patient profile of the hospital ward in which the Aetheon TUG delivery robot was deployed could make the difference between a successful and unsuccessful implementation.

Security robots

Robots are also commonly considered as potential providers of security in homes and public spaces. A robotic security guard, K5 (see Figure 10.11)) recently appeared on the market and has since been deployed at some shopping malls. It roams around the environment to monitor crime and alerts human authorities if it senses something suspicious. A prime example of a service robot that was not accepted by its environment, the K5 robot has fallen victim to a variety of abusive behaviors, ranging from an attack by a drunken man while patrolling a parking lot in Mountain View, California, to being tackled and covered in barbecue sauce while attempting to chase off homeless people from a nongovernmental organization's doorstep in San Francisco (see Figure 10.11).

Figure 10.11 Knightscope K5 (2013–present). (Source: Knighscope)

10.7 Collaborative robots

Collaborative robots are gaining importance in the automation industry. Traditional industrial robots typically are stiff, are strong, and have limited sensory capabilities. Because of this, humans are not allowed near a powered industrial robot. In contrast, collaborative robots—co-bots for short—have safety features and a mechatronic design that allow them to operate near people or even work together with people.

Some co-bots are equipped to interpret or produce social signals, such as the Walt robot, which has a face attached to its robotic arm (see Figure 10.12). The Baxter robot (see Figure 2.7) is a two-armed robot that is able to display a range of facial expressions on its screen, signaling various internal states. An embarrassed blush, for example, signals to the human co-worker that the robot is at a loss about what to do next.

The deployment of co-bots in industrial manufacturing contexts and the workplace in general may fundamentally change the notion of collaborative teamwork. In positive scenarios, co-bots should be able to help humans get more pleasure and efficiency from their work. In the worst case, collaboration with robots could backfire through a reversal of the roles of humans and robots, leading to humans serving robots rather than vice versa.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 10.12 Walt (2017–present), a collaborative robot, working at the Audi car factory in Brussels to apply glue to car parts. It has a headlight-shaped head with an animated face to communicate its internal state to its human co-workers. (Source: copyright imec)

10.8 Self-driving cars

Self-driving cars are, in essence, robots in which the user is in the passenger's seat. Although autonomous cars are still not widely available, most cars now have some form of on-board advanced driver assistance technologies (ADAS), such as lane following, adaptive cruise control, automatic parking, predictive braking, pedestrian protection systems, and blind-spot warning systems. Many of these systems require an effective human–machine interface for the driver of the car. In addition, self-driving cars require interfaces allowing them to interpret the actions and intentions of other traffic users, and the car will need ways of expressing its intentions to other users. Car drivers use a wide range of signals to communicate their intent to others. For example, slowing down when nearing a crosswalk can signal to pedestrians that they have been noticed and that it is safe to cross. The Jaguar Land Rover developed a more explicit way of communicating with pedestrians by putting "googly eyes" on its cars to signify attention.

Interaction with the driver does not only happen through the car's interface but also often requires autonomous technology to communicate why a decision was made. Koo et al. (2015) show how a message that explains "why" an action was taken, such as automated braking, is preferred over a system that merely reports the action.

HRI studies can help understand how traffic users and passengers respond to autonomous cars. Rothenbücher et al. (2016) present a paradigm in which a driver is disguised as a car seat, giving the impression that the car is self-driving (see Figure 10.13). This deception allows for carefully controlled studies on how people perceive and respond to self-driving cars without the need for a fully self-driving car.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

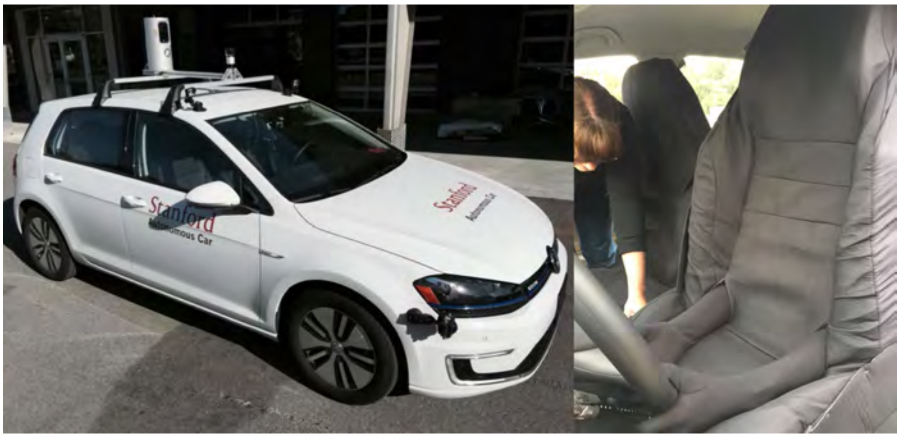

Figure 10.13 A mock-up of a self-driving vehicle, in which a driver is disguised as a car seat, used to study people's responses to the behavior of self-driving cars. (Source: Wendy Ju

10.9 Remotely operated robots

People in the military have reported becoming very attached to their robots, despite the fact that these were designed without any capability of social interaction. Military robots have been named, have been awarded battlefield promotions, and have received medals of honor from their human supervisors (Garreau, 2007).

There are several application examples of remotely operated robots. Robots used for planetary exploration have some autonomous navigation capability but receive commands from human operators on Earth as well. Packbot (see Figure 10.15) is a scout robot used in a military context; a human operator tele-operates Packbot while it searches for bomb-traps, thus clearing the road for military vehicles. Also in the military context, a human operator can operate a drone from a faraway location during military operations. In search-and-rescue scenarios, an operator controls a robot that moves on the ground or through the sky to find a person in need. Telepresence robots have started to appear on the market as well and can be used, for example, to give a presentation at a remote place or to interact with people in a different location.

In these remote-operation scenarios, a human operator commonly needs to work with some level of autonomy in the robot. A robot may autonomously navigate around, but the operator may need to provide destinations for efficient use. The robot's capability of avoiding risks (e.g., collisions with obstacles or attacks from a hostile entity) can be poor, and hence the operator needs to intervene before the robots are seriously damaged.

Operators interact with remotely operated robots via a user interface (see Figure 10.14); here, there are many common HRI problems to address, as with other types of human–robot interaction. For instance, the robot system needs to acquire an appropriate level of trust from the operator, not too much, not too little. There are similar ethical issues to be considered. For example, if the autonomy system fails, who is responsible? Is it ethical to design a system that would allow such a failure of autonomy?

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Figure 10.14 The T-HR3 robot (2017–present) can be remotely controlled using a dedicated user interface. (Source: Toyota)

Figure 10.15 Packbot (2016–present). (Source: Endeavor Robotics)

10.10 Future applications

Many of the applications introduced in this chapter are already available today. As technologies keep advancing, however, other types of future applications will emerge. For instance, researchers envision that daily appliances can be more automated and connected, as a network of devices within a smart home, for example. Several research groups also envision that individual robots can provide interfaces for such smart homes (Bernotat et al., 2016). Researchers have also started exploring how people might react to robotic furniture and appliances. Sirkin et al. (2015) studied how a robot ottoman should interact with people and also explored interactions with an interactive chest of drawers. Yamaji et al. (2010) developed a set of social trash boxes, which use social cues such as approaching and bowing to motivate people to throw away their trash; they also created a set of robotic dishes that can be summoned by a user by rapping on the table. Osawa et al. (2009) investigated how people may respond to home appliances being anthropomorphized, such as equipping a refrigerator with eyes or a printer with a mouth so that it can speak to a user.

Future developments of robots will also likely extend the capabilities within existing application domains. For example, healthcare robots are now being developed not only to provide companionship but also to monitor the behavior and health status of their users (e.g., Autom) and also possibly to assist with tasks of daily living (e.g., Care-O-Bot). Educational robots may take on more active roles in tutoring, particularly in domains such as second-language learning (Belpaeme et al., 2015). Following data-based applications in other domains, robots might also take advantage of their interactive capabilities to collect different kinds of information on users. We can expect robotic sensing and interaction capabilities to become more distributed in our lived environment, engaging with us through various everyday devices that may not immediately come across as robots.

10.11 Problems for robot application

There are various HRI problems that could prevent robots from being successful on the commercial market and as applications in everyday life. These include the potential for robot design to lead to misplaced and eventually disappointed expectations, overreliance on and addiction to robots, misuse and abuse of robots, and engagement with robots taking people's attention away from other concerns.

10.11.1 Addressing user expectations

Users often enter into interactions with robots with certain expectations, often rooted in exposure to specific conceptions of robots in the popular news media or fiction. The design and presentation of robots can also inspire certain expectations in users. For example, if a robot speaks in English, users will likely expect that it will be able to understand spoken English. The more humanlike the robot looks, the more human capabilities it may be expected to have. The cost of disappointing user expectations can be that the robot is perceived as incompetent, and people are therefore less willing to use it. Paepcke and Takayama (2010) showed that it is possible, however, to manage user expectations by describing the robot's abilities realistically; in fact, it is better to set expectations lower rather than higher. User expectations could also be managed through the design; for example, many social robots

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

are designed with infant-like appearances to decrease expectations and increase tolerance for error.

10.11.2 Addiction

There is a concern that robots, and specifically social robots, will make people over-reliant on the social and physical interaction offered by robotic devices. One can easily imagine a future in which some people prefer robots as interaction partners, perhaps even as life partners, over humans Borenstein and Arkin (2016). A less extreme scenario would be where robots are preferred over people for some interactions. Although this is not necessarily cause for concern—many people already prefer online shopping over a trip to the store, for example—we should be wary of the negative consequences of substituting social human interaction with social "robot" interaction. One concern is that robots will be seen to offer friendship, a state that, of course, is artificial to the robot but might be perceived as genuine by the human user (Elder, 2017). Conversations with a robot could be pleasant, even cathartic, but there is a danger that because the robot panders to the user, offering an interaction that is pleasing, this might make the user over-reliant on the robot, causing the human to crave the robot's company. Because robots are most likely to be under the control of corporations, to some extent, there is a concern that dependence, and perhaps even addiction, will be a sought-after property in robots. Lessons should be learned from our interaction with connected devices when designing robots (Turkle, 2016).

10.11.3 Attention theft

As can already be observed with mobile devices, technology attracts our attention, and robots, too, could cause "attention theft." Neuroscience research has demonstrated that our attention is grabbed by motion and sound, and this is exacerbated when the sound and movement is lifelike and social (Posner, 2011). Robots pose an easy opportunity for attention theft, either unintentionally or by design. When designing and deploying robots, care should be taken that the robot has a mechanism to identify when not to engage with the user or draw attention through its actions, however unintentional. In particular, this should be carefully done in cases where the robot might attract attention away from a human interaction partner.

Gazzaley and Rosen (2016) provide an interesting read about the "dark side" of our high-tech age.

10.11.4 Loss of interest by user

The so-called novelty effect is frequently discussed in the HRI literature, suggesting that people pay more attention to a novel entity and

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

express a preference to use it because it is unfamiliar; however, such effects are usually not long lasting (Kanda et al., 2004; Koay et al., 2007b). Researchers have tested various robot applications in research contexts and have revealed that the novelty effect lasted anywhere from a few minutes to, at most, a few months. Therefore, even if a one-shot experiment were to reveal positive outcomes regarding the performance and evaluation of a robot, we cannot be sure that the positive effect will prevail in the long run. Longitudinal studies are needed to provide further evidence for positive HRI over time. An important goal is to enable robots to sustain users' interest over time and across multiple interactions (Tanaka et al., 2007; Kidd and Breazeal, 2007; Kanda et al., 2007b).

10.11.5 Robot abuse

An aspect of HRI that came as a bit of a surprise to researchers in the field is robot abuse. It has been noted by various scholars that robots, when left unsupervised, sometimes get abused by humans (Brscić et al., 2015). Notably, the behavior that is generally displayed shares more similarities with intimidation and bullying than with vandalism. This makes sense, considering that robots are recognized as social agents by humans. Children seem especially prone to engage in robot-bullying behavior (see Figure 10.16), presumably due to their strong tendency to anthropomorphize and as part of developing their social skills. In a field experiment with a robot in a shopping mall, the abusive behavior of children became so disruptive to the robot's functioning that researchers eventually programmed the robot to avoid children, especially when they were gathered in a group (Brscić et al., 2015).

The fact that robots elicit abuse and that bystanders will unlikely intervene is a problem for their application in public spaces. For instance, in the retail context, robot abuse would disturb business; hence, store managers might be hesitant to have robots at their stores in order to avoid this problem. Visitors who witness abuse might feel sympathetic toward the robots in spite of being unlikely to intervene, which would result in a negative overall experience. So far, little experimental research has been conducted on the reasons why some people engage in robot bullying. A broader understanding of the phenomenon will likely help in the development of strategies to discourage robot abuse and thereby enable smoother functioning of robots.

Figure 10.16 A child kicking a robot in a shopping mall.

10.12 Conclusion

Markets for robots are growing (Haegele, 2016), but many of the robots that are available on the market still feature limited social interaction capabilities, for instance, pet robots and service robots. In the domain of navigation, great strides have been made, as documented by applications such as delivery robots and self-driving cars. Before deploying any such technologies, though, empirical research and evaluation studies need to be conducted in order to test and validate the new technologies and to get them ready for the market. With more research in open-ended, real-world contexts, it is likely that researchers will come up with new application concepts for robots and find novel niches that existing robotic technologies can successfully occupy.

Questions for you to think about:

- Try to think about a couple of new future applications that are not yet mentioned in the chapter. For each application that comes to mind, briefly describe possible technical problems and solutions.
- Suppose you would be able to prepare the technical solutions for your applications. Think about market potential: Who are the targeted users, how expensive will your robots be, and which consumers would be willing to buy the respective robots?
- Suppose your applications are successful in terms of technical preparation and the potential market. What problems might they cause? How would you avoid or at least reduce such problems?

Future reading:

- International Federation of Robotics. World Robotics Report. (Part of the report is free to download: `https://ifr.org/free-downloads/`).
- Joost Broekens, Marcel Heerink, Henk Rosendal, et al. Assistive social robots in elderly care: A review. *Gerontechnology*, 8(2): 94–103, 2009. doi: 10.4017/gt.2009.08.02.002.00. URL `https://doi.org/10.4017/gt.2009.08.02.002.00`
- Martin Ford. *The rise of the robots: Technology and the threat of mass unemployment*. Oneworld Publications, London, UK, 2015. ISBN 978-0465059997. URL `http://www.worldcat.org/oclc/993846206`
- Iolanda Leite, Carlos Martinho, and Ana Paiva. Social robots for long-term interaction: A survey. *International Journal of Social Robotics*, 5(2):291–308, 2013. doi: 10.1007/s12369-013-0178-y. URL `https://doi.org/10.1007/s12369-013-0178-y`
- Omar Mubin, Catherine J. Stevens, Suleman Shahid, Abdullah Al Mahmud, and Jian-Jie Dong. A review of the applicability of robots in education. *Journal of Technology in Education and Learning*, 1(209-0015):1–7, 2013. doi: 10.2316/

Journal.209.2013.1.209-0015. URL `http://doi.org/10.2316/Journal.209.2013.1.209-0015`
- Illah Reza Nourbakhsh. *Robot futures*. MIT Press, Cambridge, MA, 2013. ISBN 9780262018623. URL `http://www.worldcat.org/oclc/945438245`

11

Robots in Society

What is covered in this chapter:

- The influence of the media on human–robot interaction (HRI) research;
- Stereotypes of robots in the media;
- Positive and negative visions of HRI;
- Ethical considerations when designing an HRI study;
- Ethical issues of robots that fulfill a user's emotional needs;
- The dilemmas associated with behavior toward robots (e.g., robots' rights to being treated in a moral way);
- The issue of job losses as a result of the increasing number of robots in the workforce.

The discussion of robots in society often brings up questions about how we envision robots in the present and future and the social and ethical consequences of using robots in different tasks and contexts. Researchers, the media, and members of the public argue over how robots will affect our perceptions of and interactions with other humans, what the consequences of new robotic technologies will be for labor distribution and relations, and what should be considered socially and ethically appropriate uses of robots. This kind of exploration is crucial to the field of human–robot interaction (HRI) because understanding the societal meaning, significance, and consequences of HRI research will ensure that new robotic technologies fit our common social values and goals. To understand how robots might fit into society, we take a broad view of HRI through the lens of culture and the narratives, values, and practices that provide the context and tools with which people make sense of the world around them and the robots that will be coming to share it.

In this chapter, we look at robots in fiction and film, two aspects of popular culture that have had particularly strong impacts on how we imagine robotic technology in society. We also consider ethical concerns about the introduction and use of robots in society to reflect on how our values and priorities should be taken into account while shaping the human–robot interactions of the future.

11.1 Robots in popular media

What movies have become popular with audiences or critics recently? Is there a TV series that went viral or an episode that everyone is talking about? Did any of those contain robots, by chance? If so, how were these machines portrayed? Looking into the literature and other media, it becomes clear that robots have always been a "hot topic" for sci-fi writers and avid consumers.

Historically, stories of artificial human beings, such as the Golem in Jewish folklore, have been around for hundreds of years. Karel Čapek was the first author to use the word "robot," which was featured in his theater play *R.U.R.—Rossum's Universal Robots* that premiered in 1921 (see Figure 11.1). In it, robots take over the world and kill almost all humans. Two robots do, however, start to exhibit emotions for each other, and the last remaining human considers them to be the new Adam and Eve.

Figure 11.1 A scene from Čapek's 1921 play *R.U.R.* shows robots rebelling against their human masters.

Now think back to when you first heard about robots—this first encounter with a robot was likely an on-screen encounter. Computer graphics can nowadays visualize almost anything; hence, depictions of robots in movies can be quite fantastical. For example, movies depict robots that use antigravity to float around. In reality, there is little use for such robot hardware features. Robots have been portrayed in all types of artistic expressions, such as books, movies, plays, and computer games. Such media portrayals form our perceptions and understanding of robots and can thus bias our views, particularly because these are the only experiences most people have with robots. We are at an interesting point in time where, on the one hand, more and more robots are about to enter our everyday lives, but on the other hand, almost all our knowledge about robots stems from the media. This gap between the expectations fueled by science fiction and the actual abilities of robots often leads to disappointment when people interact with robots. This is why it is important to look at how robots are portrayed in popular media and to take such portrayals into account when we are designing robots for and presenting them to the public.

As a disclaimer, we have to acknowledge that it was not possible for us to consider every robot mentioned in every book, film, computer game, newspaper article, or play. Their number vastly exceeds our limited capacity for processing. But an exhaustive review is, in our view, not even necessary. We believe that we can still draw some valid conclusions from representative samples of robots in the media.

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

11.1.1 Robots want to be humans

In many stories, robots are portrayed as wanting to be humans, despite their own superiority in many aspects, such as strength and computational power. This desire to become human is the central story line for Isaac Asimov's *The Bicentennial Man*, in which a robot named Andrew Martin is following a lifelong plan to become recognized as a human (Asimov, 1976). The book was the basis for the movie of the same name, released in 1999. Besides becoming physically more humanlike, Andrew Martin also fights a legal battle to gain full legal status. He is even prepared to accept mortality to gain this status.

Other robots, such as the replicant Rachael in the movie *Blade Runner* based on the book by Philip K. Dick, are not even aware of the fact that they are robots (Dick, 2007). The same holds true for some of the humanlike Cylons in the 2004 TV series *Battlestar Galactica*.

On the contrary, a prime example of a robotic character that is aware of its robotic nature is Mr. Data from the TV series *Star Trek: The Next Generation*. Mr. Data is stronger than humans; has more computational power; and does not need sleep, nutrition, or oxygen. Still, his character is set up to have the desire to become more humanlike. The key aspect, however, that makes Mr. Data different from humans is his lack of emotion. Similarly, Steven Spielberg's movie *A.I.*, based on Brian Aldiss's short story *Super-Toys Last All Summer Long*, accepts this main premise of a lack of emotion as well (Aldiss, 2001). Because robots lack emotions, Professor Allen Hobby, played by William Hurt, builds the robot David with no ability to love. Likewise, sci-fi authors have considered emotions to be a feature that all robots would lack. However, in reality, several computational systems of emotions have already been successfully implemented. The computer programs implementing the so-called *OCC model of emotions* (Ortony et al., 1988) are prime examples. Equipping robots with emotions in the attempt to make them human is therefore an archetypal story line.

A more subtle variation of this story line is the inclusion of a control or setting for honesty and humor in the robots depicted in the movie *Interstellar*. The following dialogue between Cooper, the captain of the spaceship, and the TARS robot emerges:

COOPER: Hey, TARS, what's your honesty parameter?
TARS: Ninety percent.
COOPER: Ninety percent?
TARS: Absolute honesty isn't always the most diplomatic nor the safest form of communication with emotional beings.
COOPER: Okay, ninety percent it is.

Although robots might not have emotions themselves, they will be required to interact with humans that do have emotions, and hence it

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

will be necessary for them to process emotions and even adjust their rational behavior accordingly.

These archetypal examples taken from contemporary film are only the tip of the iceberg, but they illustrate humans' steady desire to compare themselves to superhuman entities. A hundred years ago, however, there were already machines that were more powerful than humans, although their power was physical and not mental. These days, we can see the major progress in the area of artificial intelligence AI). On May 11, 1997, the IBM computer "Deep Blue" won the first chess match against the world champion at the time. In 2011, the IBM computer Watson won as a contestant in the quiz game show *Jeopardy*. In 2017, Google's DeepMind AlphaGo defeated the world's number-one Go player, Ke Jie. In light of this progress, it is easy to imagine how robots in the future might be both strong and intelligent, leaving humans in an inferior position. At the same time, computers and robots are successful in limited task domains, so humans may have the advantage through their ability to adapt and generalize to different tasks and contexts. Fictional narratives let us explore the consequences of these and other possibilities from the safety of our couches.

11.1.2 Robots as a threat to humanity

Another archetypal story line in fiction is that of a robotic uprising. In short, humanity builds intelligent and strong robots. The robots decide to take over the world and enslave or kill all humans in order to secure resources for themselves (Barrat, 2015). Karel Čapek's original play, mentioned earlier, already introduced this narrative. Going back to the example of Mr. Data, he has a brother named Lore that possesses an emotion chip. Lore follows the path of not wanting to be like a human but instead wanting to enslave humanity. Other popular examples are *The Terminator* (Cameron, 1984) (see Figure 11.2), the Cylons in *Battlestar Galactica*, the Machines portrayed in the movie *The Matrix*, and the robots portrayed in the 2004 movie *iRobot*. The latter is based on the book by the same name by Isaac Asimov (Asimov, 1991). Asimov coined the term "Frankenstein Complex" to describe this archetypal story line.

Figure 11.2 The Terminator. (Source: Dick Thomas Johnson)

This archetype builds on two assumptions. First, robots resemble humans. The robots depicted in these movies have been designed to look, think, and act like their creators. However, they exceed their creators in intelligence and power. Second, once they interact with the now "inferior" human species, robots dehumanize their subordinates, a theme familiar in examples from human history as well. Many colonial powers declared indigenous populations as nonhumans in an attempt to vindicate the atrocities committed toward them. Accordingly, because robots resemble humans, they will also enslave and kill humans. How-

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

ever, this rationale is overly simplistic. The issue of a perceived threat to distinctiveness is also addressed in the psychological literature (Ferrari et al., 2016). If you want to learn more about the psychology of feeling threatened by robots, then consider reading the work of Zlotowski et al. (2017).

11.1.3 Superior robots being good

Several science-fiction authors have already proposed future scenarios in which superior robots quietly influence human society. In Isaac Asimov's *Prelude to Foundation*, he describes a robotic first minister, Eto Demerzel (a.k.a. R. Daneel Olivaw), who keeps the empire on the right track (Asimov, 1988). Interestingly, he hides his robotic nature. He is a very humanlike robot in appearance but resorts to various strategies to blend in. For example, he eats food, despite the fact that he cannot digest it. He collects it in a pouch that can be emptied later. Here we have a scenario in which a superior being works to help human society behind the scenes.

The notion of robots being evil and humans being good is most persistent in Western culture. Robots are extremely popular in the Japanese media, and there we can observe a different relationship between humans and robots: robots, such as "Astro Boy" and Doraemon, are good-natured characters that help humans in their daily lives. This more positive spin on the social uses and consequences of robots is often seen as being partially responsible for the large number of personal and home robots being developed in Japan and their perceived higher acceptance there than in Western societies.

11.1.4 Similarity between humans and robots

The story archetypes described previously all explore the question of to what degree humans and robots are alike. From a conceptual point of view, robots are typically portrayed by emphasizing either their similarities or lack thereof in terms of their body and mind (see Table 11.1). Dixon supports this view by stating that artists explore the deep-seated fears and fascinations associated with machine embodiment in relation to two distinct themes: the humanization of machines and the dehumanization of humans (Dixon, 2004; Haslam, 2006).

		Mind	
		Similar	Different
Body	Similar	Type I	Type II
	Different	Type III	Type IV

Table 11.1 Topics of HRI in theater

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

These four types of topics can, of course, be mixed. If we take the example of Mr. Data, at the superficial level, he looks very much like a human, which sets our expectations accordingly (Type II). It then appears dramatic and surprising if Mr. Data is able to enter the vacuum of space without being damaged. In the movie *Prometheus*, the android David, played by Michael Fassbender, is wearing a space suit when walking on an alien planet. Wearing this suit does not serve a functional purpose because David does not require air. The following dialogue emerges:

CHARLIE HOLLOWAY: David, why are you wearing a suit, man?
DAVID: I beg your pardon?
CHARLIE HOLLOWAY: You don't breathe, remember? So, why wear the suit?
DAVID: I was designed like this, because you people are more comfortable interacting with your own kind. If I didn't wear the suit, it would defeat the purpose.

Again, the human embodiment sets our expectations, and when a difference from humans is displayed, it surprises the audience. Godfried-Willem Raes takes a different approach with his robot orchestra. He emphasizes the equality of robots and humans in his theatrical performances (Type I). He argues:

> If these robots conceal nothing, it is fairly self-evident that when their functioning is made dependent on human input and interaction, this human input is also provided naked. The naked human in confrontation with the naked machine reveals the simple fact that humans, too, are actually machines, albeit fundamentally more refined and efficient machines than our musical robots.

An example of Type III could be Johnny Five from the 1986 movie *Short Circuit*. Although Johnny Five has a distinctively robotic body, he does express human emotions, which suggests that his mind is similar to that of humans.

11.1.5 Narratives of robotic science

Ben Goldacre has pointed out how the media promotes the public misunderstanding of science (Goldacre, 2008). Two narratives that the media frequently uses are science-scare stories and wacky science stories.

The performance of autonomous vehicles, which can also be considered a form of human–robot interaction, is currently the target of immense scrutiny. The crash statistics provided by Tesla, Waymo, and others indicate that they are performing better than humans. Tesla, for example,[1] showed that driving using the vehicle's autopilot feature reduces the probability of crashes dramatically. This finding does not, however, take into account that Tesla's autopilot currently does not operate in urban environments. The comparison to the overall crash statistics is therefore problematic. Still, the accidents that do occur attract disproportional attention in the news. Most cases even attract international news coverage, such as the lethal Uber crash in 2018. This attention may affect and possibly inhibit the adoption of this technology.

One question that almost all reporters ask when interviewing HRI researchers is when robots will take over the world. The goal, then, is to write a story that scares the public and hence attracts attention. A story entitled "Robots are harmless and almost useless" is very unlikely to get published. But that is what most HRI projects come down to at this point in time. The question of when and if robots will take over the world addresses our inner fears and fascinations involving interacting with robots. Are we like robots? Are robots like us? And if so, will superior robots act as badly as humans have when encountering "inferior" beings?

We may ask ourselves why these questions are so persistent in the media. The most obvious answer is that stories need to have a conflict to generate tension. A fictional world in which everybody is happily living together is unlikely to capture the attention of the audience. Pitching evil robots against good humans not only serves the purpose of creating a conflict but also triggers an "in-group" effect. We humans feel that we need to defend our species against "out-group" robots. This division can then be challenged by introducing robots that are indistinguishable from humans, such as in the TV shows *Battlestar Galactica* and *Westworld*. This creates great uncertainty, which in turn creates tension. Notable exceptions from the gloomy visions in the media are the TV series *Futurama* by Matt Groening and the movie *Robot and Frank* by Jake Schreier, both of which depict a vision of the future in which humans and robots live peacefully side by side. They even become friends. In the movie *Her*, the protagonist Theodore, played by Joaquin Phoenix, even falls in love with his AI mobile phone Samantha (Jonze, 2013).

The wacky science narrative occurs less frequently but attracts attention nevertheless. A robot preacher that "can beam light from its hands and give automated blessings to worshippers" is just one example of a newspaper story that is intended more to entertain than to report scientific progress (Berghuis, 2017).

For HRI researchers, media coverage therefore has great potential to showcase their work, but it also carries considerable risk. The reporter

[1] https://www.tesla.com/blog/q3-2018-vehicle-safety-report

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

might intend to write a scare story or a wacky science story. Researchers are therefore advised to participate in the media training sessions that many universities and research institutes offer to their staff. A general guideline for talking to the media is to stick to the research that was actually performed and avoid engaging in wild speculations about topics that were not covered in the study at hand. It is also always a good idea to ask which questions will be asked beforehand and, when possible, to request to view a manuscript draft prior to publication so that any misunderstandings or misrepresentations of the science involved can be corrected prior to publication.

HRI researchers cannot shy away from representations of robots in the media, fictional or otherwise, and the elicitation of associated fears. In HRI studies, we invite people to engage with robots, and every single person who interacts with a robot does so with preconceptions and expectations of what the robot can and cannot do. Many of these come from science fiction and reports in the media, rather than the annals of scientific research.

11.2 Ethics in HRI

Is it okay to develop and sell a sex robot, which is always willing to do what you want and will stay forever young and fit? Would you have your parents be taken care of by a carebot instead of a human nurse?

Roboticists and philosophers alike have long been concerned with such ethical issues in robotics, coining a shared domain of scholarship called "roboethics." More recently, a group of HRI scholars formulated five ethical rules, which they call their Principles of Robotics, to raise broader awareness about the role of ethics for HRI.[2] Ethical rules have also been a subject of discussion in popular literature, particularly the well-known "Three Laws of Robotics" (see the accompanying text box).

Isaac Asimov (January 2, 1920–April 6, 1992; see Figure 11.3) proposed three rules of robotics that would safeguard humanity from malevolent robots:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given to it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

Although Asimov's work is highly visible in the public media, it has

[2] https://www.epsrc.ac.uk/research/ourportfolio/themes/engineering/activities/principlesofrobotics/

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

been criticized by philosophers. Asimov eventually added a zeroth law:

0. A robot may not harm humanity or, by inaction, allow humanity to come to harm.

Figure 11.3 Isaac Asimov (January 2, 1920–April 6, 1992).

This clearly marks the relevance of debating issues such as the ubiquitous deployment of robots in future society; their use in home and care contexts; the implications of developing autonomous weapons systems and autonomous cars; or, giving it a seemingly positive touch, the development of robots for attachment, love, or sex.

These days, many robotics research projects envision robots as "slaves," in the sense that they conduct acts on behalf of humans, like killing others or serving to fulfill humans' need for psychological closeness and sexuality. Some of these projects are even funded by government agencies. At the same time, there are clear countermovements, such as the Campaign Against Killer Robots.[3] As responsible researchers, we have to consider the ethical implications of what we envision and the steps we take to approach these visions of the future (Sparrow, 2011). In the following subsections, we discuss some of the common topics of ethical concern in HRI research.

11.2.1 Robots in research

As a student beginning to get hands-on experience with empirical research in HRI, you might plan to conduct a study with a robot that acts seemingly autonomously. Even here, ethics has to be considered because you might choose to deceive your participants by controlling your robot using the Wizard-of-Oz approach. You thereby make the participants believe that the robot has certain functions, whereas in reality, you control the robot's behavior in the background. The problem with this approach is that the deception concerning robot skills raises and biases users' hopes and expectations about the robot's abilities. This may manipulate them into thinking that robotic technology is more advanced than it actually is (Riek, 2012).

Another critical example to consider might be the usage of robots as persuasive communicators within your research project. Previous research on persuasive technology has shown that robots can be used to manipulate people into changing not only their attitudes but also their behaviors (Brandstetter et al., 2017). Examples of behaviors that have been successfully influenced include health-related habits, such as exercising or maintaining a healthy diet (Kiesler et al., 2008). Even if it might benefit people to change their health-related habits, such as smoking less and exercising more, instrumentalizing social robots for

[3] `https://www.stopkillerrobots.org/`

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

this purpose poses ethical concerns if they exploit the social bond with the user and influence the user without the user's explicit consent and conscious knowledge or understanding about how he or she is being influenced.

11.2.2 Robots to fulfill emotional needs

Robotic care

Imagine your grandmother has been given a robot companion by a group of researchers. They tell her that this new technological friend will stay with her in her home for the next three weeks. She interacts with the robot every day for these three weeks, and over time, she becomes quite attached to it. The robot invites her to do activities like memory games on a regular basis. It asks her how she is doing and whether she slept well; it keeps her company, and it never argues with her. She is delighted with her new companion, and life is good. That is, until the researchers come back and ask her to complete some questionnaires before packing up the robot and taking it away. The dull routine of the elderly care center creeps back, and she feels even more lonely than before.

This brief scenario gives a glimpse of the psychological experience of getting attached—not only to people but also to objects like robots. HRI researchers have shown how easily people grow attached to a robot, even when it only briefly enters their everyday lives (Šabanović et al., 2014; Forlizzi and DiSalvo, 2006; Chang and Šabanović, 2015; Kidd and Breazeal, 2008). The emotional and social consequences of withdrawing this source of attention and "artificial affection" clearly need to be considered when running case studies with a social robot that has to be returned at the end of the study.

Other studies, however, have demonstrated the beneficial effect of deploying small-scale robots such as the therapeutic robot Paro (Wada and Shibata, 2007; Shibata, 2012) or the robot dog Aibo (Broekens et al., 2009). These robots are not able to do any tedious manual labor, but they can provide companionship. Given the high workload that caretakers are burdened with, any relief, even small, is likely welcomed.

Sparrow and Sparrow (2006) offer an interesting perspective on robotic care that has become a classic in the literature. They argue that even when a robotic caregiver could be developed that is capable of providing superb emotional and physical care, it would still be unethical to outsource care to machines. The reason for this is that a relationship can only be meaningful when it is between two entities that are capable of experiencing reciprocal affect and concern; an imitation of caring, however perfect, should never substitute the real product. This kind of relationship may also be detrimental to the value of upholding

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

a person's dignity. This brings us into the ethics of developing a deeper emotional attachment to a robot.

Emotional attachment to robots

Affection toward robots can go deeper and beyond the care setting. Humans may start to favor robot companions over humans. Imagine a social robot that can truly mimic friendship and emotional support. This "ideal robotic friend" comes with all the perks of a human friend, never complains, and learns never to annoy its owner. Slowly, people could come to prefer these robotic companions over their human peers, who would not be able to measure up to the high standards that robotic friends provide. Would such a future be desirable?

Even though users may project all kinds of human traits into a robot, the robot is not able to experience those traits in the same way humans do, and therefore, the authenticity of the expression can be doubted. Still, robots are sometimes specifically designed to express social cues to deliberately facilitate bonding with them. The authenticity of feelings is normally important in human–human interaction, and we do not know how humans will react to robots that express themselves based on calculations rather than the sensation of emotions.

Going beyond human–robot friendship, there are individuals who feel closeness and intimacy toward robots. The broader question is whether promoting human–robot emotional bonds is desirable (Borenstein and Arkin, 2016). After all, we have to realize that the emotional relationship between humans and robots might be asymmetrical. Humans might nevertheless be quite satisfied with the robot exhibiting sympathetic responses, whether the robot has a humanlike sensation of attachment or not.

Ethical implications of persuasion through robots

Language develops dynamically, and every participant in discourse influences its development simply through its usage. New words appear (e.g., "to google"), others change their meaning (e.g., "gay"), and yet other words fall out of usage altogether. We can use Siri, Cortana, or Bixby to control our phones, homes, or shopping tours. Familiarity alone will influence our attitudes toward concepts, political ideas, and products; this is called the "mere exposure effect" (Zajonc, 1968). The more often people hear a word, the more positive their attitude toward this word becomes. One day, it will make a great difference if your smart-shopping robot proposes to purchase "Pepsi" compared to offering a "Coca-Cola." The question really is who gets to decide what words our artificial counterparts use.

Robots have the ability to synchronize their vocabulary through the internet in seconds. Even the mass media cannot compete with this level of consistent usage of selected words (Brandstetter et al., 2017).

Because of its ability to communicate in humanlike ways, a robot can be a convincing persuasive communicator. This comes with negative implications, though: without us even noticing, computers and robots can influence what words we use and how we feel about them. This can and probably is happening already, and we need to develop media and language competency to be able to withstand attempts to influence our views. With the ever-more personalized and intimate relationships that we form with technologies, we are increasingly vulnerable. We probably already spend more time with our phones than with our partners and friends.

Furthermore, to our knowledge, there are no regulations or policies in place at this point in time to supervise how large information technology (IT) companies, such as Google, Amazon, or Facebook, influence the usage of language, although there is concern about "fake news" and the difficulty of telling fact from fiction in online contexts. It might also be a better approach to regulate the development of our language only to the degree that it should be left to its natural flow of change. With powerful tools at our fingertips, we need to ensure that no company or government can influence our language without our consent and that the robots we design do not become just one additional persuasive and misleading technology.

Generalizing abusive behavior toward robots

Being recognized as a social actor comes with a downside: not all social behaviors are positive. In a few field experiments with autonomous robots that were left unsupervised in public spaces, people were observed attempting to intimidate and bully robots (Brscić et al., 2015; Salvini et al., 2010). It is noteworthy that the type of aggression that people displayed seemed to resemble human–human abuse, such as kicking, slapping, insulting, and refusing to move out of the way after the robot politely asked. Abuse that would be more meaningful for machines, such as unplugging them or cutting their wires, was not observed.

Robots normally do not experience any pain or humiliation, and hence the human actually faces greater danger than the robot when, for example, slapping the robot because the human might hurt his or her hand. But there are more issues to consider than just the bully's bodily integrity. It has been argued that bullying a robot is a moral offense—even though nobody gets hurt, responding with violence is still considered wrong and should therefore be discouraged (Whitby, 2008). In addition, scholars have argued that if this behavior is perceived as acceptable, it might generalize to other social agents, such as animals and humans (Whitby, 2008; De Angeli, 2009). This transfer of negative behavior from a humanlike agent to actual humans is argued to also happen in other domains, such as violent computer games (Sparrow,

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

2017; Darling, 2012), and has been a topic of discussion for quite a while. Further research on this topic is still needed.

A related issue is that interactions with a robot may raise expectations regarding the behavior of other humans. This has been argued to be particularly dangerous in the domain of sex. A robot could easily be designed to seem to desire intercourse at any time and to readily and fully comply with any wishes of the user without having any desires or demands of its own. This could change what people consider normal or appropriate behavior from an intimate partner.

This issue becomes even more problematic if the robot is specifically designed for sexual behaviors that would be considered wrong if it had involved human partners. For example, child-shaped sex robots could be designed to fit the desires of pedophiles; or sex robots could be programmed to explicitly not consent to or even struggle against sex in order for users to play out their rape fantasies. These robot behavior designs have been deemed ethically inappropriate by some scholars (for a philosophical justification, see Sparrow (2017)).

11.2.3 Robots in the workplace

A repeatedly expressed worry is that "robots will take over" and "robots will replace me in the job market." Since the Industrial Revolution, humans have been replacing manual labor with machines, and the recent deployment of robots is no exception. Robots help us to improve our productivity and thereby help to increase our standard of living. Although robots do replace certain jobs, they also create many new jobs, in particular for highly trained professionals. The challenge that society is facing is that the people replaced by robots need to find new jobs, which might require them to embark in additional training or studies. This may be problematic or even impossible for some, for example, due to financial or intellectual constraints.

Some fields, such as education, are less welcoming toward accepting robots in their workforce. Reich-Stiebert and Eyssel (2015) showed that robots are preferred as assistants in the classroom but not as the main teachers. They also voiced concern about the usage and maintenance of the robots, being particularly fearful that the robot would take their resources in terms of time and attention. Interestingly, primary school teachers were particularly reluctant to have robots in schools, maybe because in their view, young students are particularly vulnerable. An analysis of the predictors of such rather negative attitudes and behavioral inclinations toward educational robots revealed that technology commitment was the key predictor of positive attitudes. That is, those teachers who were open to working with novel technologies in general felt more positive about robots and the future use of them in their classrooms. Another field in which people are concerned about the application of robots is assistive robots in their homes (Reich-Stiebert and Eyssel, 2015, 2013). Again, technology commitment was found to predict people's reluctance to accept robots in their lives.

Haegele (2016) claimed that more and more robots will be sold on the market in the next few years. Their acceptance into society, however, will remain a challenge, and further research on technology-related attitudes and how to change them is necessary to increase society's acceptance of robots.

11.3 Conclusion

It is important to realize that robots, humans, and their interactions are part of broader societies that encompass different kinds of people, technologies, institutions, and practices. In these different social and cultural contexts, people may hold different initial attitudes and beliefs about robots based on their prior exposure to fictional narratives and popular media. Potential users of robots will also hold different social and cultural values and norms. Both these cultural narratives and values will affect how people perceive and respond to robots and how the use of robots might affect existing social structures and practices. HRI researchers should be conscious of and sensitive to prevailing cultural narratives and values when they design and deploy robots in society, and they should also consider whether they want robots to reproduce or challenge existing practices and norms.

Questions for you to think about:

- What was the last movie or series you watched, or book you read, that depicted robots?
- List the characteristics of the robot protagonists you have recently seen in a film or TV series. What were their capabilities? Did they appear humanlike? Did they pose a threat to humanity, or did they save the world?
- How will the availability of new forms of media such as YouTube change people's expectations toward robots?
- Think of professions that have been replaced by machines. Which ones come to mind? What are the potential positive and negative implications of this replacement?
- Is there an activity that you are happy to have a machine do? What is an activity that you would not want to be replaced by a machine? How do you think others might feel about your choices—who might disagree?
- Discuss whether it is ethical to use a social robot as comfort for

lonely elderly people. Describe relevant issues, and explain your opinion.
- In a future where highly intelligent robots are available, would it be ethical to develop robot nannies or robot teachers? Describe the potential issues.
- Some HRI studies are provocative or thought-provoking, for example, Bartneck et al. (2018) on the presence of racism in HRI. Is it ethical to run controversial HRI studies? Are there particular themes, such as religion, where HRI should not tread?

Future reading:

- Spike Jonze. *Her*, 2013. URL `https://www.imdb.com/title/tt1798709/?ref_=fn_al_tt_1`
- Isaac Asimov. *The Robot Series*. 1950–1986. [this collection consists of several books that were never formally published as a series]
- Philip K. Dick. *Do androids dream of electric sheep?* Boom! Studios, a division of Boom Entertainment, Los Angeles, CA, 1986. ISBN 978-160886784. URL `http://www.worldcat.org/oclc/929049302`
- Jake Schreier. Robot and Frank, 2013. URL `https://www.imdb.com/title/tt1990314/`
- Amanda J. C. Sharkey. Should we welcome robot teachers? *Ethics and Information Technology*, 18(4):283–297, 2016. doi: 10.1007/s10676-016-9387-z. URL `https://doi.org/10.1007/s10676-016-9387-z`
- Peter W. Singer. *Wired for war: The robotics revolution and conflict in the twenty-first century*. Penguin, New York, NY, 2009. ISBN 9781594201981. URL `http://www.worldcat.org/oclc/857636246`
- Gianmarco Veruggio, Fiorella Operto, and George Bekey. Roboethics: Social and ethical implications. In Bruno Siciliano and Oussama Khatib, editors, *Springer handbook of robotics*, pages 2135–2160. Springer, 2016. ISBN 978-3-319-32550-7. doi: 10.1007/978-3-319-32552-1. URL `https://doi.org/10.1007/978-3-319-32552-1`
- Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. The moral machine experiment. *Nature*, 2018. ISSN 1476-4687. doi: 10.1038/s41586-018-0637-6. URL `https://doi.org/10.1038/s41586-018-0637-6`
- Robert Sparrow. Robots, rape, and representation. *International*

Journal of Social Robotics, 9(4):465–477, Sep 2017. ISSN 1875-4805. doi: 10.1007/s12369-017-0413-z. URL `https://doi.org/10.1007/s12369-017-0413-z`

- Patrick Lin, Keith Abney, and George A. Bekey. *Robot ethics: The ethical and social implications of robotics*. Intelligent robotics and autonomous agents. MIT Press, Cambridge, MA, 2012. ISBN 9780262016667. URL `http://www.worldcat.org/oclc/1004334474`

12

The Future

What is covered in this chapter:

- Current attitude of the general public toward robots and how this may change in the upcoming years;
- Possible shifts and developments in the nature of human–robot relationships, specifically companion bots;
- Further development of the technology of human–robot interactions (HRI), specifically artificial intelligence (AI);
- The inherent issues with predicting the future ("crystal ball problems").

As with other technologies that have become commonplace in everyday use, such as personal computers, smartphones, or the internet, we expect robots will sooner or later become assimilated into society. They may even be accepted into our personal, and even intimate, spaces. Robots are currently being designed to be co-workers, tutors, and assistants in the medical field and to provide services in care settings, in education, and in people's homes. Sony, evidencing the renewed interest of the company in social robots, has recently released a new AIBO robotic dog (see Figure 3.2).

Technological advancements make this vision increasingly conceivable but are not sufficient to ensure a rosy future in a society equipped with robots. Recent polls in the United States and Europe suggest that the broader public is not very willing to accept social robots for everyday use, particularly in areas such as eldercare and other socially assistive and interactive applications (Smith, 2014; European Commission, 2012). Human–robot interaction (HRI) studies have also shown that people report high levels of robot anxiety and other negative attitudes toward robots and a low willingness to interact with robots in their personal space or workplace settings (Reich-Stiebert and Eyssel, 2013, 2015).

One solution to this issue may be to just wait it out. As technology advances, people will have more opportunities to interact with robots and may become more accepting of them through that exposure. As we mentioned in our discussion of nonverbal cues, direct interaction with

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

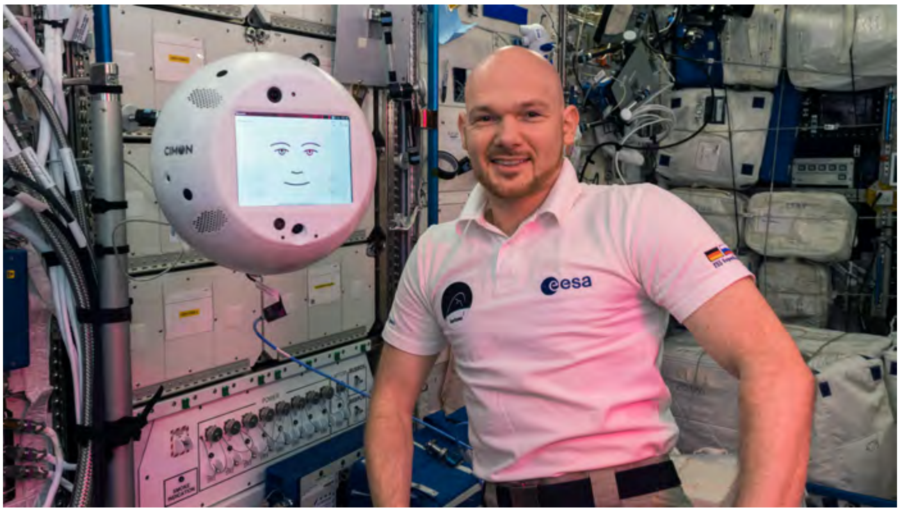

Figure 12.1 The Cimon robot (2018–present) assists astronauts on the International Space Station. (Source: National Aeronautics and Space Administration)

members of another social group—in this case, robots—changes attitudes and decreases anxiety related to that group (Crisp and Turner, 2013; Pettigrew et al., 2011). Wullenkord (2017) showed that just imagining collaborative interaction with a Nao robot prior to actually interacting with it improved attitudes and reactions toward the robot and increased the perceived quality of the interaction. We can therefore expect that as people have increased contact with robots, be it directly or through the media, attitudes will grow more positive, and the willingness to use robots will increase over time.

As we have seen in the rest of this book, however, advances in HRI research can significantly speed up this process. By better understanding the sources of people's concerns about robots and the types of societal needs and desires that robotic capabilities can address, we can create interactions that will be positive and beneficial to people, which can lead to a positive feedback loop for familiarizing people with robots and create more support for the broader adoption of new robotic technologies.

We also need to consider that the media frequently portray robots negatively or unrealistically. For example, there has been much talk of robots, instead of people, looking after those in need of assistance in our aging societies. This is not a pleasant thought for many, who are reasonably concerned by the wide-ranging implications of this scenario for HRI and, more fundamentally, for human–human relationships. The way this future scenario is portrayed by the media, however, is unrealistic. This manner of framing robotics in society creates fear in the general public and distracts us from the work we need to do and the choices we need to make to create our preferred future.

Facilitating an open mind about novel developments in technology and science might be a step toward achieving a more positive view and

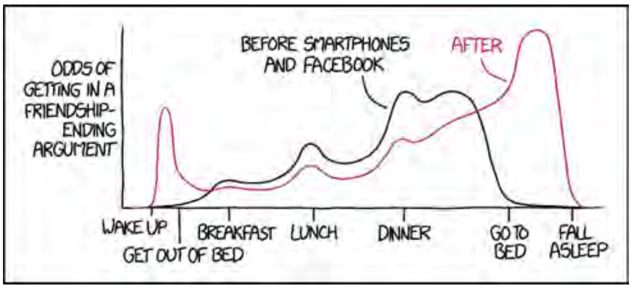

Figure 12.2 Odds of getting into a friendship-ending argument before and after the introduction of smartphones. (Source: XKCD)

a stronger sense of acceptance by the general public. These changes can only be observed through longitudinal studies, and HRI scholars must work together with the communities they seek to serve to consider how technological developments can come together with societal structures to produce positive change. There is no quick "technological fix" for societal problems, such as demographic change. Besides developing much-needed technologies, it is also crucial to take on a human-centered approach that focuses on the actual psychological, social, and emotional needs of the people using and being affected by robots. A more human-centered view coupled with technological advancement will together create robust and socially appropriate robots that can benefit us all.

12.1 The nature of human–robot relationships

When waiting to check in at the airport, a machine handles the check-in process. In Japan, Pepper robots greet us when we enter a bank or a shop. When care is provided mainly by machines rather than humans, this has strong implications for the development and maintenance of human relationships. Even currently, some technologies, such as social networks and online games, have resulted in less direct contact between people. Instead of writing letters or meeting in person, people communicate via posts on Facebook. Our patterns of when we talk to whom about what are changing (see Figure 12.2), as are the ways we begin and end our romantic relationships—by smartphone. Robots may contribute to further estrangement among people, as is argued by Turkle (2017), or robots could be designed to support and even increase interaction among people. This effect has been seen with the seal-like robot Paro in a day home, in which older adults ended up meeting and talking more to others when the robot was put in a public space (Wada and Shibata, 2007).

Clearly, as social robots and artificial intelligence (AI) are developed further, they will likely play an increasingly larger role in our everyday lives and society. Because the nature of human–robot relationships is a product of the robots' capabilities and the users' preferences, these developments are inevitably tied to the question of what issues we consider ethical and desirable to address with robots and AI.

For example, one major societal issue at the moment is loneliness. Feeling socially connected to others has an almost incredible list of benefits for individual mental and physical health (Vaillan, 2015). This will become increasingly relevant as the populations of developed countries continue to age in the upcoming decades. An increasing part of the population is in need of care, not just for attending to their physical needs of feeding, bathing, and clothing but for emotional care as well. It might be that the younger generations are neither willing nor competently able to serve these dual needs on their own. Particularly, the emotional needs of seniors or people with cognitive or physical impairments have to be taken into account, but all people are in danger of growing more and more lonely and disconnected (American Osteopathic Association, 2016).

The lack of social connection can have serious impacts for our psychological well-being and health. The "need to belong," a key motivation of human nature (Baumeister and Leary, 1995), can easily become frustrated. To illustrate, research by Eisenberger et al. (2003) has shed light on the neuroanatomical underpinnings of reactions to social exclusion, whereas Williams (2007) has documented the negative social consequences of exclusionary status. That is, when the need to belong is violated, people feel a lower sense of belonging, self-esteem, and control and even regard their existence as less meaningful than when their inclusionary status is not under threat. In addition, the risk of developing Alzheimer's disease is doubled in lonely people compared with socially connected individuals, and loneliness is a predictor of a decline in cognitive abilities (Shankar et al., 2013). In light of the detrimental effects of loneliness on quality of life and psychological and cognitive functioning, robots could play an important role in mediating these effects.

Figure 12.3 Not everyone is charmed by the idea of strong AI. The late theoretical physicist Stephen William Hawking, and the inventor and engineer Elon Musk are both vocal critics against the development of strong AI.

A few commercial start-ups have been offering artificially intelligent "companions," although so far with only modest success, such as Gatebox's Living With Project. If AI and robots are developed to the point where they can reliably imitate human interaction patterns, they could be extremely helpful in relieving feelings of boredom and loneliness.

What remains to be seen is how comfortable people are with the different potential roles that AI may take on. As the quest for strong AI continues, the question of whether such an AI is desirable is being raised on a daily basis. Whereas the most spectacular version of this question considers how we can ensure that such an AI would remain benevolent to the human race, it is at least as interesting to consider the issue of whether people would be comfortable with handing over power in the first place. Assume that strong AI is developed, the sole purpose of which is to enhance the well-being of society while adhering

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

to a set of rules that keep it from harming humans (e.g., Asimov's Laws of Robotics; see Section 11.2). Can one throw out of the window all the doubts about self-interest, bias, and hidden political agendas that arise with human leaders and, instead, fully trust that the AI would take proper care (see Figure 12.3)? Would users agree with such a setup?

This is an important point in robot and AI development. Just as we ask ourselves, "Just because we could, does that mean we should?" to balance out all the rational possibilities with moral and ethical skepticism, the reverse, "Just because we should, does that mean we would?" holds true as well. Robots are logical, but humans—the people who create robots—are not. Simply because something might be beneficial from a logical point of view does not mean that people are comfortable with it.

12.2 The technology of HRI

HRI is lifted on the tides of technological progress. New sensors and actuators and continuous developments in AI are quickly adopted into HRI applications. Given the steady progress in AI and its applications, there is every reason to believe that a number of technological problems that currently still require the smoke and mirrors of Wizard-of-Oz (WoZ) control will soon be delivered autonomously by the robot.

Progress in HRI is not so much held back by a lack of development in robotic hardware but, rather, by a lack of progress in autonomous control and AI. A testament to this is the ability of human operators to hold a meaningful interaction through a robot. It is clearly not the limited view through the sensors and the limited expressivity of the actuators that hinder the interaction. Rather, it is the artificial cognition—substituted by real cognition in the case of WoZ control—that is lacking. There is, of course, room for improvement in robot hardware: the speed and power of actuators need work, and the energy autonomy of robots needs to improve drastically. Furthermore, robotics and social robotics in particular have always taken a "Frankenstein approach" to hardware, building robots from whatever technology is readily available rather than developing radically new hardware solutions. But at this point, breakthroughs in HRI are most likely to come from progress in robot control and AI. Machine learning holds considerable promise here. However, there are fundamental barriers to the use of machine learning in HRI. Because machine learning requires vast amounts of annotated data and computational time, it comes to its own in domains that allow offline learning and for which huge amounts of training data are available or, when not, can be generated. Although there is plenty of human interaction going on in the world, these interactions run in real-time. As opposed to machine learning of the game of chess or Go, where the learning can run as fast as computers will allow, machine learning of HRI strategies inherently runs online. No matter how fast the computer is, the interaction pace is dictated by the human interaction partner, and the evaluation and updates of the machine learning will run in "human time" rather than in computer time. One solution for facilitating machine learning for HRI might be to use more robots and data from more interactions: pooling interaction events could be a solution to the dearth of HRI data and could speed up the evaluation of learned interaction strategies. It is unclear what the next technological breakthroughs will be in AI and robotics, but one thing is clear: HRI will readily absorb them.

12.3 Crystal ball problems

Predicting the future is hard to do, and especially in the field of HRI, it seems as if every stance imaginable is defended with passion by a small army of experts (and a large group of those wishing to be experts), ranging from doomsday predictions to nirvana forecasts. It proves to be nearly impossible to gain consensus on the far future of HRI and even on small and concrete predictions of how long it will take to develop a specific capability or what we actually want from a robot. Here, it is fitting to mention two predictions related to developments in HRI that have not panned out as expected.

First, we can perhaps take some lessons from developments in AI, which have been rapid yet unable to adhere to initial expectations. When the fundamentals of AI were developed in the 1950s, it was expected that strong AI could be designed within the upcoming decade (McCorduck, 1979; Russell and Norvig, 2009). Half a century later, AI still struggles with understanding human sentences. As long as the rules of a behavior are strict and can be operationalized, AI can keep up with humans easily and often even outsmart them. This was famously shown by the Deep Blue computer program beating the human world champion Kasparov in a game of chess in the late 1990s (Campbell et al., 2002). The recent victory of an AI over the world champion on a game of Go (Murphy, 2016) was considered a milestone because Go is more complex and has a larger emphasis on strategics, whereas chess is more tactical and has a less extensive set of possible ways to win.

Similarly, the world has seen a sharp increase in the number of startups and initiatives for new social robots. A popular way of obtaining funds for these kinds of projects is through crowdsourcing because many people prove to be willing to invest. They predict that the robot will be a success. However, for some reason, few of the funded projects actually take off. Usually, they start up, continue for a year or two, and then die out. This raises the question of whether humans are capable of realizing what they want from a robot. Obviously, even when we think some social robot is a brilliant application for an everyday problem, when we actually get the robot, we are not quite as sold as we were on the idea. Although, as we've shown in this book, there are people from many different fields involved in developing robotic applications for society, we would likely benefit from expanding the range of perspectives that participate in discussion and decision-making about the kind of future we want with robots.

These two examples of predictions gone wrong may, of course, be part of the whole problem of expectation management and the tendency of users to overestimate what robots are capable of. But in addition, the second example indicates that humans are adept at predicting how much they would like to have a robot for certain tasks. We imagine robots taking over all kinds of jobs, but it remains to be seen in which areas we prefer the (messier) human way of doing it.

Questions for you to think about:

- Which technological developments, and related social developments, have surprised you the most in your lifetime?
- What kind of future would you want to see with robots? What kind of future would you be afraid of or concerned about?
- How much time do you spend interacting with people face to face versus in mediated environments (e.g., Facebook, conference call)? What about nonhuman agents—do you interact with them at all? In what circumstances and how much?
- Who is caring for your grandparents or parents? What kind of community do they live in? Do you live close to them? Who do you think will take care of you in the future? What kind of community might you find yourself living in?

Further Reading:

- Future of Life Institute. An open letter—research priorities for robust and beneficial artificial intelligence, January 2015. URL `https://futureoflife.org/ai-open-letter/`
- Illah Reza Nourbakhsh. *Robot futures.* MIT Press, Cambridge, MA, 2013. ISBN 9780262018623. URL `http://www.worldcat.org/oclc/945438245`
- Daniel H. Wilson. *How to survive a robot uprising: Tips on defending yourself against the coming rebellion.* Bloomsbury, New York, NY, 2005. ISBN 9781582345925. URL `http://www.worldcat.org/oclc/1029483559`
- Jo Cribb and David Glover. *Don't worry about the robots.* Allen & Unwin, Auckland, New Zealand, 2018. ISBN 9781760633509. URL `http://www.worldcat.org/oclc/1042120802`

References

Chadia Abras, Diane Maloney-Krichmar, and Jenny Preece. User-centered design. In William Sims Bainbridge, editor, *Berkshire encyclopedia of human-computer interaction*, volume 2, pages 763–767. Sage, Great Barrington, MA, 2004. ISBN 9780974309125. URL http://www.worldcat.org/oclc/635690108.

Henny Admoni and Brian Scassellati. Social eye gaze in human-robot interaction: A review. *Journal of Human-Robot Interaction*, 6(1):25–63, 2017. doi: 10.5898/JHRI.6.1.Admoni. URL https://doi.org/10.5898/JHRI.6.1.Admoni.

Kaat Alaerts, Evelien Nackaerts, Pieter Meyns, Stephan P. Swinnen, and Nicole Wenderoth. Action and emotion recognition from point light displays: An investigation of gender differences. *PloS One*, 6(6):e20989, 2011. doi: 10.1371/journal.pone.0020989. URL https://doi.org/10.1371/journal.pone.0020989.

Brian Wilson Aldiss. *Supertoys last all summer long: And other stories of future time*. St. Martin's Griffin, New York, NY, 2001. ISBN 978-0312280611. URL http://www.worldcat.org/oclc/956323493.

Minoo Alemi, Ali Meghdari, and Maryam Ghazisaedy. Employing humanoid robots for teaching English language in Iranian junior high-schools. *International Journal of Humanoid Robotics*, 11(03):1450022, 2014. doi: 10.1142/S0219843614500224. URL https://doi.org/10.1142/S0219843614500224.

Christopher Alexander. *A pattern language: Towns, buildings, construction*. Oxford University Press, Oxford, UK, 1977. ISBN 978-0195019193. URL http://www.worldcat.org/oclc/961298119.

Philipp Althaus, Hiroshi Ishiguro, Takayuki Kanda, Takahiro Miyashita, and Henrik I. Christensen. Navigation for human-robot interaction tasks. In *IEEE International Conference on Robotics and Automation*, volume 2, pages 1894–1900. IEEE, 2004. ISBN 0-7803-8232-3. doi: 10.1109/ROBOT.2004.1308100. URL https://doi.org/10.1109/ROBOT.2004.1308100.

Amir Aly and Adriana Tapus. A model for synthesizing a combined verbal and nonverbal behavior based on personality traits in human-robot interaction. In *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction*, HRI '13, pages 325–332, Piscataway, NJ, USA, 2013. IEEE Press. ISBN 978-1-4673-3055-8. doi: 10.1109/HRI.2013.6483606. URL https://doi.org/10.1109/HRI.2013.6483606.

American Osteopathic Association. Survey finds nearly three-quarters (72%) of Americans feel lonely, 2016. URL https://www.osteopathic.org/inside-aoa/news-and-publications/media-center/2016-news-releases/Pages/10-11-survey-finds-nearly-three-quarters-of-americans-feel-lonely.aspx.

Peter A. Andersen and Laura K. Guerrero. Principles of communication and emotion in social interaction. In Peter A. Andersen and Laura K. Guerrero, editors, *Handbook of communication and emotion: Research, theory, applications,* *and contexts*, chapter 3, pages 49–96. Academic Press, 1998. ISBN 0-12-057770-4. doi: 10.1016/B978-012057770-5/50005-9. URL https://doi.org/10.1016/B978-012057770-5/50005-9.

Sean Andrist, Xiang Zhi Tan, Michael Gleicher, and Bilge Mutlu. Conversational gaze aversion for humanlike robots. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 25–32. ACM, 2014. ISBN 978-1-4503-2658-2. doi: 10.1145/2559636.2559666. URL https://doi.org/10.1145/2559636.2559666.

Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57 (5):469–483, 2009. doi: 10.1016/j.robot.2008.10.024. URL https://doi.org/10.1016/j.robot.2008.10.024.

S. E. Asch. *Effects of group pressure upon the modification and distortion of judgments*, pages 177–190. Carnegie Press, Oxford, England, 1951. doi: psycinfo/1952-00803-001. URL http://doi.apa.org/psycinfo/1952-00803-001.

Isaac Asimov. *The Bicentennial man and other stories*. Doubleday science fiction. Doubleday, Garden City, NY, [Book Club edition, 1976. ISBN 978-0385121989. URL http://www.worldcat.org/oclc/85069299.

Isaac Asimov. *Prelude to foundation*. Grafton, London, UK, 1988. ISBN 9780008117481. URL http://www.worldcat.org/oclc/987248670.

Isaac Asimov. *I, robot*. Bantam spectra book. Bantam Books, New York, NY, 1991. ISBN 0553294385. URL http://www.worldcat.org/oclc/586089717.

Hillel Aviezer, Yaacov Trope, and Alexander Todorov. Body cues, not facial expressions, discriminate between intense positive and negative emotions. *Science*, 338(6111):1225–1229, 2012. doi: 10.1126/science.1224313. URL https://doi.org/10.1126/science.1224313.

Edmond Awad, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. The moral machine experiment. *Nature*, 2018. ISSN 1476-4687. doi: 10.1038/s41586-018-0637-6. URL https://doi.org/10.1038/s41586-018-0637-6.

Wilma A. Bainbridge, Justin W. Hart, Elizabeth S. Kim, and Brian Scassellati. The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics*, 3(1):41–52, Jan 2011. ISSN 1875-4805. doi: 10.1007/s12369-010-0082-7. URL https://doi.org/10.1007/s12369-010-0082-7.

Simon Baron-Cohen, Alan M. Leslie, and Uta Frith. Does the autistic child have a "Theory of Mind"? *Cognition*, 21(1):37–46, 1985. doi: 10.1016/0010-0277(85)90022-8. URL https://doi.org/10.1016/0010-0277(85)90022-8.

James Barrat. Why Stephen Hawking and Bill Gates are terrified of artificial intelligence. *Huffington Post*, 2015. URL http://www.huffingtonpost.com/james-barrat/hawking-gates-artificial-intelligence_b_7008706.html.

Christoph Bartneck. *eMuu: An embodied emotional character for the ambient intelligent home*. PhD thesis, Technische Universiteit Eindhoven, 2002. URL http://www.bartneck.de/publications/2002/eMuu/bartneckPHDThesis2002.pdf.

Christoph Bartneck and Jun Hu. Rapid prototyping for interactive robots. In *The 8th Conference on Intelligent Autonomous Systems (IAS-8)*, pages 136–145, 2004. doi: 10.6084/m9.figshare.5160775.v1. URL https://doi.org/10.6084/m9.figshare.5160775.v1.

Christoph Bartneck and Michael J. Lyons. Facial expression analysis, modeling and synthesis: Overcoming the limitations of artificial intelligence with the art of the soluble. In Jordi Vallverdu and David Casacuberta, editors, *Handbook of research on synthetic emotions and sociable robotics: New applications in affective computing and artificial intelligence*, Information Science Reference, pages 33–53. IGI Global, 2009. URL http://www.bartneck.

de/publications/2009/facialExpressionAnalysisModelingSynthesisAI/bartneckLyonsEmotionBook2009.pdf.

Christoph Bartneck and M. Rauterberg. HCI reality—an unreal tournament. *International Journal of Human-Computer Studies*, 65(8):737–743, 2007. doi: 10.1016/j.ijhcs.2007.03.003. URL https://doi.org/10.1016/j.ijhcs.2007.03.003.

Christoph Bartneck and Juliane Reichenbach. Subtle emotional expressions of synthetic characters. *International Journal of Human-Computer Studies*, 62 (2):179 – 192, 2005. ISSN 1071-5819. doi: 10.1016/j.ijhcs.2004.11.006. URL https://doi.org/10.1016/j.ijhcs.2004.11.006. Subtle expressivity for characters and robots.

Christoph Bartneck, T. Nomura, T. Kanda, Tomhohiro Suzuki, and Kato Kennsuke. Cultural differences in attitudes towards robots. In *AISB Symposium on Robot Companions: Hard Problems and Open Challenges in Human-Robot Interaction*, pages 1–4. The Society for the Study of Artificial Intelligence and the Simulation of Behaviour (AISB), 2005. doi: 10.13140/RG.2.2.22507.34085. URL http://www.bartneck.de/publications/2005/cultureNars/bartneckAISB2005.pdf.

Christoph Bartneck, Elizabeth Croft, Dana Kulic, and Susana Zoghbi. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, 1(1):71–81, 2009. doi: 10.1007/s12369-008-0001-3. URL https://doi.org/10.1007/s12369-008-0001-3.

Christoph Bartneck, Andreas Duenser, Elena Moltchanova, and Karolina Zawieska. Comparing the similarity of responses received from studies in Amazon's Mechanical Turk to studies conducted online and with direct recruitment. *PloS One*, 10(4):e0121595, 2015. doi: 10.1371/journal.pone.0121595. URL https://doi.org/10.1371/journal.pone.0121595.

Christoph Bartneck, Kumar Yogeeswaran, Qi Min Ser, Graeme Woodward, R. Sparrow, Siheng Wang, and Friederike Eyssel. Robots and racism. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, pages 196–204. ACM, 2018. ISBN 978-1-4503-4953-6. doi: 10.1145/3171221.3171260. URL https://doi.org/10.1145/3171221.3171260.

Timo Baumann and David Schlangen. The INPROTK 2012 release. In *NAACL-HLT Workshop on Future Directions and Needs in the Spoken Dialog Community: Tools and Data*, pages 29–32. Association for Computational Linguistics, 2012. URL http://dl.acm.org/citation.cfm?id=2390444.2390464.

Roy F. Baumeister and Mark R. Leary. The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin*, 117(3): 497–529, 1995. doi: 10.1037/0033-2909.117.3.497. URL https://doi.org/10.1037/0033-2909.117.3.497.

Paul Baxter, James Kennedy, Emmanuel Senft, Severin Lemaignan, and Tony Belpaeme. From characterising three years of HRI to methodology and reporting recommendations. In *The 11th ACM/IEEE International Conference on Human-Robot Interaction*, pages 391–398. IEEE Press, 2016. ISBN 978-1-4673-8370-7. doi: 10.1109/HRI.2016.7451777. URL https://doi.org/10.1109/HRI.2016.7451777.

Aryel Beck, Antoine Hiolle, Alexandre Mazel, and Lola Cañamero. Interpretation of emotional body language displayed by robots. In *Proceedings of the 3rd International Workshop on Affective Interaction in Natural Environments*, pages 37–42. ACM, 2010. ISBN 978-1-4503-0170-1. doi: 10.1145/1877826.1877837. URL https://doi.org/10.1145/1877826.1877837.

Christopher Beedie, Peter Terry, and Andrew Lane. Distinctions between emotion and mood. *Cognition & Emotion*, 19(6):847–878, 2005. doi: 10.1080/02699930541000057. URL https://doi.org/10.1080/02699930541000057.

Tony Belpaeme, Paul E. Baxter, Robin Read, Rachel Wood, Heriberto Cuayáhuitl, Bernd Kiefer, Stefania Racioppa, Ivana Kruijff-Korbayová, Georgios Athanasopoulos, Valentin Enescu, et al. Multimodal child-robot interaction: Building social bonds. *Journal of Human-Robot Interaction*, 1(2):33–53, 2012. doi: 10.5898/JHRI.1.2.Belpaeme. URL `https://doi.org/10.5898/JHRI.1.2.Belpaeme`.

Tony Belpaeme, James Kennedy, Paul Baxter, Paul Vogt, Emiel E. J. Krahmer, Stefan Kopp, Kirsten Bergmann, Paul Leseman, Aylin C. Küntay, Tilbe Göksun, et al. L2tor-second language tutoring using social robots. In *Proceedings of the ICSR 2015 WONDER Workshop*, 2015. URL `https://pub.uni-bielefeld.de/download/2900267/2900268`.

Sandra L. Bem. The measurement of psychological androgyny. *Journal of Consulting and Clinical Psychology*, 42:155–162, 1974. doi: 10.1037/h0036215. URL `https://doi.org/10.1037/h0036215`.

Koen Berghuis. Robot 'preacher' can beam light from its hands and give automated blessings to worshippers, 2017. URL `https://www.mirror.co.uk/news/weird-news/robot-priest-can-beam-light-10523678`.

Jasmin Bernotat, Birte Schiffhauer, Friederike Eyssel, Patrick Holthaus, Christian Leichsenring, Viktor Richter, Marian Pohling, Birte Carlmeyer, Norman Köster, Sebastian Meyer zu Borgsen, et al. Welcome to the future: How naïve users intuitively address an intelligent robotics apartment. In *International Conference on Social Robotics*, pages 982–992. Springer, 2016. ISBN 978-3-319-47436-6. doi: 10.1007/978-3-319-47437-3_96. URL `https://doi.org/10.1007/978-3-319-47437-3_96`.

Cindy L. Bethel and Robin R. Murphy. Survey of non-facial/non-verbal affective expressions for appearance-constrained robots. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(1):83–92, 2008. doi: 10.1109/TSMCC.2007.905845. URL `https://doi.org/10.1109/TSMCC.2007.905845`.

Cindy L. Bethel and Robin R. Murphy. Review of human studies methods in HRI and recommendations. *International Journal of Social Robotics*, 2(4):347–359, 2010. doi: 10.1007/s12369-010-0064-9. URL `https://doi.org/10.1007/s12369-010-0064-9`.

Cindy L. Bethel, Kristen Salomon, Robin R. Murphy, and Jennifer L. Burke. Survey of psychophysiology measurements applied to human-robot interaction. In *The 16th IEEE International Symposium on Robot and Human Interactive Communication*, pages 732–737. IEEE, 2007. ISBN 978-1-4244-1634-9. doi: 10.1109/ROMAN.2007.4415182. URL `https://doi.org/10.1109/ROMAN.2007.4415182`.

James R. Blair. Responding to the emotions of others: Dissociating forms of empathy through the study of typical and psychiatric populations. *Consciousness and Cognition*, 14(4):698–718, 2005. doi: 10.1016/j.concog.2005.06.004. URL `https://doi.org/10.1016/j.concog.2005.06.004`.

Benjamin S. Bloom. The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6):4–16, 1984. doi: 10.3102/0013189X013006004. URL `https://doi.org/10.3102/0013189X013006004`.

Robert Bogue. Exoskeletons and robotic prosthetics: A review of recent developments. *Industrial Robot: An International Journal*, 36(5):421–427, 2009. doi: 10.1108/01439910910980141. URL `https://doi.org/10.1108/01439910910980141`.

George A. Bonanno, Laura Goorin, and Karin G. Coifman. Social functions of emotion. In Michael Lewis, Jeanette M. Haviland-Jones, and Lisa Feldman Barrett, editors, *Handbook of emotions*, volume 3, pages 456–468. Guilford Press, New York, NY, 2008. ISBN 978-1-59385-650-2. URL `http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.472.7583&rep=rep1&type=pdf`.

Jason Borenstein and Ronald C. Arkin. Robots, ethics, and intimacy: The need for scientific research. In *Conference of the International Association for Computing and Philosophy*, 2016. URL https://www.cc.gatech.edu/ai/robot-lab/online-publications/RobotsEthicsIntimacy-IACAP.pdf.

Valentino Braitenberg. *Vehicles: Experiments in synthetic psychology*. MIT Press, Cambridge, MA, 1986. ISBN 978-0262521123. URL http://www.worldcat.org/oclc/254155258.

Jürgen Brandstetter, Péter Rácz, Clay Beckner, Eduardo B. Sandoval, Jennifer Hay, and Christoph Bartneck. A peer pressure experiment: Recreation of the Asch conformity experiment with robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1335–1340. IEEE, 2014. ISBN 978-1-4799-6934-0. doi: 10.1109/IROS.2014.6942730. URL https://doi.org/10.1109/IROS.2014.6942730.

Jurgen Brandstetter, Eduardo B. Sandoval, Clay Beckner, and Christoph Bartneck. Persistent lexical entrainment in HRI. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 63–72. ACM, 2017. ISBN 978-1-4503-4336-7. doi: 10.1145/2909824.3020257. URL https://doi.org/10.1145/2909824.3020257.

Cynthia Breazeal. *Designing sociable robots*. MIT Press, Cambridge, MA, Cambridge, 2003. ISBN 978-0262524315. URL http://www.worldcat.org/oclc/758042496.

Cynthia Breazeal. Function meets style: Insights from emotion theory applied to HRI. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 34(2):187–194, 2004a. doi: 10.1109/TSMCC.2004.826270. URL https://doi.org/10.1109/TSMCC.2004.826270.

Cynthia Breazeal. Social interactions in HRI: The robot view. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 34(2):181–186, 2004b. doi: 10.1109/TSMCC.2004.826268. URL https://doi.org/10.1109/TSMCC.2004.826268.

Cynthia Breazeal and Brian Scassellati. A context-dependent attention system for a social robot. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence, Volume 2*, pages 1146–1151. Morgan Kaufmann Publishers Inc., 1999. URL http://dl.acm.org/citation.cfm?id=1624312.1624382.

Cynthia Breazeal, Cory D. Kidd, Andrea Lockerd Thomaz, Guy Hoffman, and Matt Berlin. Effects of nonverbal communication on efficiency and robustness in human-robot teamwork. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 708–713. IEEE, 2005. ISBN 0-7803-8912-3. doi: 10.1109/IROS.2005.1545011. URL https://doi.org/10.1109/IROS.2005.1545011.

Paul Bremner, Anthony Pipe, Chris Melhuish, Mike Fraser, and Sriram Subramanian. Conversational gestures in human-robot interaction. In *IEEE International Conference on Systems, Man and Cybernetics*, pages 1645–1649. IEEE, 2009. ISBN 978-1-4244-2793-2. doi: 10.1109/ICSMC.2009.5346903. URL https://doi.org/10.1109/ICSMC.2009.5346903.

Elizabeth Broadbent, Rebecca Stafford, and Bruce MacDonald. Acceptance of healthcare robots for the older population: Review and future directions. *International Journal of Social Robotics*, 1(4):319–330, 2009. doi: 10.1007/s12369-009-0030-6. URL https://doi.org/10.1007/s12369-009-0030-6.

Joost Broekens, Marcel Heerink, Henk Rosendal, et al. Assistive social robots in elderly care: A review. *Gerontechnology*, 8(2):94–103, 2009. doi: 10.4017/gt.2009.08.02.002.00. URL https://doi.org/10.4017/gt.2009.08.02.002.00.

Rodney Brooks. A robust layered control system for a mobile robot. *IEEE Journal* *on Robotics and Automation*, 2(1):14–23, 1986. doi: 10.1109/JRA.1986.1087032. URL `https://doi.org/10.1109/JRA.1986.1087032`.

Rodney A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47 (1-3):139–159, 1991. doi: 10.1016/0004-3702(91)90053-M. URL `https://doi.org/10.1016/0004-3702(91)90053-M`.

Rodney Allen Brooks. *Flesh and machines: How robots will change us*. Vintage, New York, NY, 2003. ISBN 9780375725272. URL `http://www.worldcat.org/oclc/249859485`.

Drazen Brscić, Takayuki Kanda, Tetsushi Ikeda, and Takahiro Miyashita. Person tracking in large public spaces using 3-D range sensors. *IEEE Transactions on Human-Machine Systems*, 43(6):522–534, 2013. doi: 10.1109/THMS.2013.2283945. URL `https://doi.org/10.1109/THMS.2013.2283945`.

Drazen Brscić, Hiroyuki Kidokoro, Yoshitaka Suehiro, and Takayuki Kanda. Escaping from children's abuse of social robots. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 59–66. ACM, 2015. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696468. URL `https://doi.org/10.1145/2696454.2696468`.

Barbara Bruno, Nak Young Chong, Hiroko Kamide, Sanjeev Kanoria, Jaeryoung Lee, Yuto Lim, Amit Kumar Pandey, Chris Papadopoulos, Irena Papadopoulos, Federico Pecora, et al. The CARESSES EU-Japan project: Making assistive robots culturally competent. *arXiv*, page 1708.06276, 2017. URL `https://arxiv.org/abs/1708.06276`.

Richard Buchanan. Wicked problems in design thinking. *Design Issues*, 8(2):5–21, 1992. URL `https://www.jstor.org/stable/1511637`.

Wolfram Burgard, Armin B. Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner, and Sebastian Thrun. The interactive museum tour-guide robot. In *Proceedings of the 15th National/10th Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence*, pages 11–18, 1998. ISBN 0-262-51098-7. URL `https://dl.acm.org/citation.cfm?id=295249`.

Maya Cakmak, Siddhartha S. Srinivasa, Min Kyung Lee, Jodi Forlizzi, and Sara Kiesler. Human preferences for robot-human hand-over configurations. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1986–1993. IEEE, 2011. ISBN 978-1-61284-454-1. doi: 10.1109/IROS.2011.6094735. URL `https://doi.org/10.1109/IROS.2011.6094735`.

Rafael A. Calvo, Sidney D'Mello, Jonathan Gratch, and Arvid Kappas. *The Oxford handbook of affective computing*. Oxford Library of Psychology, Oxford, UK, 2015. ISBN 978-0199942237. URL `http://www.worldcat.org/oclc/1008985555`.

James Cameron. *The Terminator*, 1984. URL `https://www.imdb.com/title/tt0088247/`.

Murray Campbell, A. Joseph Hoane, and Feng-hsiung Hsu. Deep blue. *Artificial Intelligence*, 134(1-2):57–83, 2002. doi: 10.1016/S0004-3702(01)00129-1. URL `https://doi.org/10.1016/S0004-3702(01)00129-1`.

Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1302–1310, 2017. ISBN 9781538604571. doi: 10.1109/CVPR.2017.143. URL `https://doi.org/10.1109/CVPR.2017.143`.

Julie Carpenter. *Culture and human-robot interaction in militarized spaces: A war story*. Routledge, New York, NY, 2016. ISBN 978-1-4724-4311-3. URL `http://www.worldcat.org/oclc/951397181`.

Colleen M. Carpinella, Alisa B. Wyman, Michael A. Perez, and Steven J. Stroessner. The Robotic Social Attributes Scale (RoSAS): Development and validation. In

ACM/IEEE International Conference on Human-Robot Interaction, pages 254–262. ACM, 2017. ISBN 978-1-4503-4336-7. doi: 10.1145/2909824.3020208. URL `https://doi.org/10.1145/2909824.3020208`.

Sybil Carrere and John Mordechai Gottman. Predicting divorce among newlyweds from the first three minutes of a marital conflict discussion. *Family Process*, 38 (3):293–301, 1999. doi: 10.1111/j.1545-5300.1999.00293.x. URL `https://doi.org/10.1111/j.1545-5300.1999.00293.x`.

J. Cassell, Joseph Sullivan, Scott Prevost, and Elizabeth Churchill. *Embodied conversational agents*. MIT Press, Cambridge, MA, 2000. ISBN 9780262032780. URL `http://www.worldcat.org/oclc/440727862`.

Filippo Cavallo, Raffaele Limosani, Alessandro Manzi, Manuele Bonaccorsi, Raffaele Esposito, Maurizio Di Rocco, Federico Pecora, Giancarlo Teti, Alessandro Saffiotti, and Paolo Dario. Development of a socially believable multi-robot solution from town to home. *Cognitive Computation*, 6(4):954–967, 2014. doi: 10.1007/s12559-014-9290-z. URL `https://doi.org/10.1007/s12559-014-9290-z`.

Wan-Ling Chang and Selma Šabanović. Interaction expands function: Social shaping of the therapeutic robot PARO in a nursing home. In *The 10th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 343–350. ACM, 2015. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696472. URL `https://doi.org/10.1145/2696454.2696472`.

Tony Charman, Simon Baron-Cohen, John Swettenham, Gillian Baird, Antony Cox, and Auriol Drew. Testing joint attention, imitation, and play as infancy precursors to language and Theory of Mind. *Cognitive Development*, 15(4):481–498, 2000. doi: 10.1016/S0885-2014(01)00037-5. URL `https://doi.org/10.1016/S0885-2014(01)00037-5`.

Tanya L. Chartrand and John A. Bargh. The chameleon effect: The perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, 76(6):893–910, 1999. doi: 10.1037/0022-3514.76.6.893. URL `https://doi.org/10.1037/0022-3514.76.6.893`.

Tiffany L. Chen, Chih-Hung Aaron King, Andrea L. Thomaz, and Charles C. Kemp. An investigation of responses to robot-initiated touch in a nursing context. *International Journal of Social Robotics*, 6(1):141–161, 2014. doi: 10.1007/s12369-013-0215-x. URL `https://doi.org/10.1007/s12369-013-0215-x`.

Takenobu Chikaraishi, Yuichiro Yoshikawa, Kohei Ogawa, Oriza Hirata, and Hiroshi Ishiguro. Creation and staging of android theatre "sayonara" towards developing highly human-like robots. *Future Internet*, 9(4):75–92, 2017. doi: 10.3390/fi9040075. URL `https://doi.org/10.3390/fi9040075`.

Howie M. Choset, Seth Hutchinson, Kevin M. Lynch, George Kantor, Wolfram Burgard, Lydia E. Kavraki, and Sebastian Thrun. *Principles of robot motion: Theory, algorithms, and implementation*. MIT Press, Cambridge, MA, 2005. ISBN 978-026203327. URL `http://www.worldcat.org/oclc/762070740`.

Robert Coe. It's the effect size, stupid: What effect size is and why it is important. In *Annual Conference of the British Educational Research Association*. Education-line, 2002. URL `http://www.leeds.ac.uk/educol/documents/00002182.htm`.

Jacob Cohen. *Statistical power analysis for the behavioral sciences*. Academic Press, New York, NY, 1977. ISBN 9781483276489. URL `http://www.worldcat.org/oclc/898103044`.

Jacob Cohen. The earth is round (p ¡ .05). *American Psychologist*, 49:997–1003, 1994. doi: 10.1037/0003-066X.49.12.997. URL `https://doi.org/10.1037/0003-066X.49.12.997`.

Mark Cook. Experiments on orientation and proxemics. *Human Relations*, 23(1):61–76, 1970. doi: 10.1177/001872677002300107. URL `https://doi.org/10.1177/001872677002300107`.

Martin Cooney, Takayuki Kanda, Aris Alissandarakis, and Hiroshi Ishiguro. Designing enjoyable motion-based play interactions with a small humanoid robot. *International Journal of Social Robotics*, 6(2):173–193, 2014. doi: 10.1007/s12369-013-0212-0. URL https://doi.org/10.1007/s12369-013-0212-0.

Jo Cribb and David Glover. *Don't worry about the robots.* Allen & Unwin, Auckland, New Zealand, 2018. ISBN 9781760633509. URL http://www.worldcat.org/oclc/1042120802.

Richard J. Crisp and Rhiannon N. Turner. Imagined intergroup contact: Refinements, debates, and clarifications. In Gordon Hodson and Miles Hewstone, editors, *Advances in intergroup contact*, chapter 6, pages 149–165. Psychology Press, 2013. ISBN 978-1136213908. URL http://www.worldcat.org/oclc/694393740.

April H. Crusco and Christopher G. Wetzel. The Midas touch: The effects of interpersonal touch on restaurant tipping. *Personality and Social Psychology Bulletin*, 10(4):512–517, 1984. doi: 10.1177/0146167284104003. URL https://doi.org/10.1177/0146167284104003.

Mihaly Csikszentmihalyi and Isabella Selega Csikszentmihalyi. *Optimal experience: Psychological studies of flow in consciousness.* Cambridge University Press, Cambridge, UK, 1988. ISBN 0521342880. URL http://www.worldcat.org/oclc/963712478.

Amy Cuddy, Susan Fiske, and Peter Glick. Warmth and competence as universal dimensions of social perception: The stereotype content model and the bias map. *Advances in Experimental Social Psychology*, 40:61–149, 12 2008. doi: 10.1016/S0065-2601(07)00002-0. URL https://doi.org/10.1016/S0065-2601(07)00002-0.

Geoff Cumming. Replication and p intervals: p values predict the future only vaguely, but confidence intervals do much better. *Perspectives on Psychological Science*, 3(4):286–300, 2008. doi: 10.1111/j.1745-6924.2008.00079.x. URL https://doi.org/10.1111/j.1745-6924.2008.00079.x.

Geoff Cumming. The new statistics: Why and how. *Psychological Science*, 25(1):7–29, 2014. doi: 10.1177/0956797613504966. URL http://journals.sagepub.com/doi/10.1177/0956797613504966.

Kate Darling. Extending legal protection to social robots: The effects of anthropomorphism, empathy, and violent behavior towards robotic objects. In *We Robot Conference.* SSRN, 2012. doi: 10.2139/ssrn.2044797. URL https://doi.org/10.2139/ssrn.2044797.

K. Dautenhahn, M. Walters, S. Woods, K. L. Koay, C. L. Nehaniv, A. Sisbot, R. Alami, and T. Siméon. How may I serve you? A robot companion approaching a seated person in a helping context. In *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction*, HRI '06, pages 172–179, New York, NY, 2006a. ACM. ISBN 1-59593-294-1. doi: 10.1145/1121241.1121272. URL http://doi.acm.org/10.1145/1121241.1121272.

Kerstin Dautenhahn, Michael Walters, Sarah Woods, Kheng Lee Koay, Chrystopher L. Nehaniv, A. Sisbot, Rachid Alami, and Thierry Siméon. How may I serve you? A robot companion approaching a seated person in a helping context. In *1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction*, pages 172–179. ACM, 2006b. ISBN 1-59593-294-1. doi: 10.1145/1121241.1121272. URL https://doi.org/10.1145/1121241.1121272.

Andrew J. Davison, Ian D. Reid, Nicholas D. Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(6):1052–1067, 2007. doi: 10.1109/TPAMI.2007.1049. URL http://doi.org/10.1109/TPAMI.2007.1049.

Antonella De Angeli. Ethical implications of verbal disinhibition with conversational agents. *PsychNology Journal*, 7(1), 2009. URL `http://psychnology.org/File/PNJ7(1)/PSYCHNOLOGY_JOURNAL_7_1_DEANGELI.pdf`.

Maartje de Graaf, Somaya Ben Allouch, and Jan van Dijk. Why do they refuse to use my robot? Reasons for non-use derived from a long-term home study. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, pages 224–233. ACM, 2017. ISBN 978-1-4503-4336-7. doi: 10.1145/2909824.3020236. URL `https://doi.org/10.1145/2909824.3020236`.

Frans De Waal. *The ape and the sushi master: Cultural reflections of a primatologist*. Basic Books, New York, NY, 2001. ISBN 978-0465041763. URL `http://www.worldcat.org/oclc/458716823`.

Philip K. Dick. *Do androids dream of electric sheep?* Boom! Studios, a division of Boom Entertainment, Los Angeles, CA, 1986. ISBN 978-160886784. URL `http://www.worldcat.org/oclc/929049302`.

Philip K. Dick. *Blade runner: Do androids dream of electric sheep?* Ballantine Books, New York, NY, 25th anniversary edition, 2007. ISBN 9780345350473. URL `http://www.worldcat.org/oclc/776604212`.

Joshua J. Diehl, Lauren M. Schmitt, Michael Villano, and Charles R. Crowell. The clinical use of robots for individuals with autism spectrum disorders: A critical review. *Research in Autism Spectrum Disorders*, 6(1):249–262, 2012. doi: 10.1016/j.rasd.2011.05.006. URL `https://doi.org/10.1016/j.rasd.2011.05.006`.

Carl DiSalvo, Illah Nourbakhsh, David Holstius, Ayça Akin, and Marti Louw. The neighborhood networks project: A case study of critical engagement and creative expression through participatory design. In *10th Anniversary Conference on Participatory Design 2008*, pages 41–50. Indiana University, 2008. ISBN 978-0-9818561-0-0. URL `https://dl.acm.org/citation.cfm?id=1795241`.

Carl F. DiSalvo, Francine Gemperle, Jodi Forlizzi, and Sara Kiesler. All robots are not created equal: The design and perception of humanoid robot heads. In *Proceedings of the 4th Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques*, DIS '02, pages 321–326, New York, NY, 2002. ACM. ISBN 1-58113-515-7. doi: 10.1145/778712.778756. URL `http://doi.acm.org/10.1145/778712.778756`.

Steve Dixon. Metal performance humanizing robots, returning to nature, and camping about. *TDR/The Drama Review*, 48(4):15–46, 2004. ISSN 1054-2043. doi: 10.1162/1054204042442017. URL `http://dx.doi.org/10.1162/1054204042442017`.

Anhai Doan, Raghu Ramakrishnan, and Alon Y. Halevy. Crowdsourcing systems on the world-wide web. *Communications of the ACM*, 54(4):86–96, 2011. doi: 10.1145/1924421.1924442. URL `https://doi.org/10.1145/1924421.1924442`.

Anca D. Dragan, Kenton C. T. Lee, and Siddhartha S. Srinivasa. Legibility and predictability of robot motion. In *8th ACM/IEEE International Conference on Human-Robot Interaction*, pages 301–308. IEEE, 2013. ISBN 978-1-4673-3099-2. doi: 10.1109/HRI.2013.6483603. URL `https://doi.org/10.1109/HRI.2013.6483603`.

Brian R. Duffy. Anthropomorphism and the social robot. *Robotics and Autonomous Systems*, 42(3):177–190, 2003. ISSN 0921-8890. doi: 10.1016/S0921-8890(02)00374-3. URL `https://doi.org/10.1016/S0921-8890(02)00374-3`.

Autumn Edwards, Chad Edwards, Patric R. Spence, Christina Harris, and Andrew Gambino. Robots in the classroom: Differences in students' perceptions of credibility and learning between "teacher as robot" and "robot as teacher". *Computers in Human Behavior*, 65:627–634, 2016. doi: 10.1016/j.chb.2016.06.005. URL `https://doi.org/10.1016/j.chb.2016.06.005`.

Naomi I. Eisenberger, Matthew D. Lieberman, and Kipling D. Williams. Does rejection hurt? An fMRI study of social exclusion. *Science*, 302(5643):290–292,

2003. doi: 10.1126/science.1089134. URL https://doi.org/10.1126/science.1089134.

Panteleimon Ekkekakis. *The measurement of affect, mood, and emotion: A guide for health-behavioral research.* Cambridge University Press, Cambridge, UK, 2013. doi: 10.1017/CBO9780511820724. URL https://doi.org/10.1017/CBO9780511820724.

Paul Ekman. Facial expressions of emotion: New findings, new questions. *Psychological Science*, 3(1):34–38, 1992. doi: 10.1111/j.1467-9280.1992.tb00253.x. URL https://doi.org/10.1111/j.1467-9280.1992.tb00253.x.

Paul Ekman. Basic emotions. In T. Dalgleich and M. Power, editors, *Handbook of cognition and emotion*, pages 45–60. Wiley Online Library, 1999. ISBN 978-1462509997. URL http://www.worldcat.org/oclc/826592694.

Paul Ekman and Wallace Friesen. Facial action coding system: A technique for the measurement of facial movement. *Palo Alto: Consulting Psychologists*, 1978.

Paul Ekman and Wallace V. Friesen. *Unmasking the face.* Prentice Hall, Englewood Cliffs, NJ, 1975. ISBN 978-1883536367. URL http://www.worldcat.org/oclc/803874427.

Moataz El Ayadi, Mohamed S. Kamel, and Fakhri Karray. Survey on speech emotion recognition: Features, classification schemes, and databases. *Pattern Recognition*, 44(3):572–587, 2011. doi: 10.1016/j.patcog.2010.09.020. URL https://doi.org/10.1016/j.patcog.2010.09.020.

Ilias El Makrini, Shirley A. Elprama, Jan Van den Bergh, Bram Vanderborght, Albert-Jan Knevels, Charlotte I. C. Jewell, Frank Stals, Geert De Coppel, Ilse Ravyse, Johan Potargent, et al. Working with Walt. *IEEE Robotics & Automation Magazine*, 25:51–58, 2018. doi: 10.1109/MRA.2018.2815947. URL https://doi.org/10.1109/MRA.2018.2815947.

Alexis M. Elder. *Friendship, robots, and social media: False friends and second selves.* Routledge, New York, NY, 2017. ISBN 978-1138065666. URL http://www.worldcat.org/oclc/1016009820.

Nicholas Epley, Adam Waytz, and John T. Cacioppo. On seeing human: A three-factor theory of anthropomorphism. *Psychological Review*, 114(4):864–886, 2007. doi: 10.1037/0033-295X.114.4.864. URL https://doi.org/10.1037/0033-295X.114.4.864.

Nicholas Epley, Adam Waytz, Scott Akalis, and John T. Cacioppo. When we need a human: Motivational determinants of anthropomorphism. *Social Cognition*, 26 (2):143–155, 2008. doi: 10.1521/soco.2008.26.2.143. URL https://doi.org/10.1521/soco.2008.26.2.143.

European Commission. Public attitudes towards robots: A report. Technical Report Special Eurobarometer 382 / Wave EB77.1, Directorate-General for Information Society and Media, 2012. URL http://ec.europa.eu/commfrontoffice/publicopinion/archives/ebs/ebs_382_en.pdf.

Vanessa Evers, Heidy C. Maldonado, Talia L. Brodecki, and Pamela J. Hinds. Relational vs. group self-construal: Untangling the role of national culture in HRI. In *Proceedings of the 3rd ACM/IEEE International Conference on Human-Robot Interaction*, HRI '08, pages 255–262, New York, NY, 2008. ACM. ISBN 978-1-60558-017-3. doi: 10.1145/1349822.1349856. URL http://doi.acm.org/10.1145/1349822.1349856.

Florian Eyben, Felix Weninger, Florian Gross, and Björn Schuller. Recent developments in OpenSMILE, the Munich open-source multimedia feature extractor. In *21st ACM International Conference on Multimedia*, pages 835–838. ACM, 2013. ISBN 978-1-4503-2404-5. doi: 10.1145/2502081.2502224. URL https://doi.org/10.1145/2502081.2502224.

F. Eyssel and N. Reich. Loneliness makes the heart grow fonder (of robots)—on the effects of loneliness on psychological anthropomorphism. In *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 121–122, 2013. ISBN 978-1-4673-3101-2. doi: 10.1109/HRI.2013.6483531. URL https://doi.org/10.1109/HRI.2013.6483531.

Friederike Eyssel. An experimental psychological perspective on social robotics. *Robotics and Autonomous Systems*, 87(Supplement C):363–371, 2017. ISSN 0921-8890. doi: https://doi.org/10.1016/j.robot.2016.08.029. URL http://www.sciencedirect.com/science/article/pii/S0921889016305462.

Friederike Eyssel and Frank Hegel. (S)he's got the look: Gender-stereotyping of social robots. *Journal of Applied Social Psychology*, 42:2213–2230, 2012. doi: 10.1111/j.1559-1816.2012.00937.x. URL https://doi.org/10.1111/j.1559-1816.2012.00937.x.

Friederike Eyssel, Dieta Kuchenbrandt, Simon Bobinger, Laura de Ruiter, and Frank Hegel. "If you sound like me, you must be more human": On the interplay of robot and user features on human-robot acceptance and anthropomorphism. In *Proceedings of the 7th Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '12, pages 125–126, New York, NY, 2012a. ACM. ISBN 978-1-4503-1063-5. doi: 10.1145/2157689.2157717. URL http://doi.acm.org/10.1145/2157689.2157717.

Friederike Eyssel, Dieta Kuchenbrandt, Frank Hegel, and Laura de Ruiter. Activating elicited agent knowledge: How robot and user features shape the perception of social robots. In *Robot and human interactive communication (RO-MAN)*, pages 851–857. IEEE, 2012b. doi: 10.1109/ROMAN.2012.6343858. URL https://doi.org/10.1109/ROMAN.2012.6343858.

Franz Faul, Edgar Erdfelder, Albert-Georg Lang, and Buchner Axel. G*power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39:175–191, 2007. doi: 10.3758/BF03193146. URL https://doi.org/10.3758/BF03193146.

David Feil-Seifer and Maja J. Matarić. Socially assistive robotics. *IEEE Robotics & Automation Magazine*, 18(1):24–31, 2011. doi: 10.1109/MRA.2010.940150. URL https://10.1109/MRA.2010.940150.

Catherine Feng, Shiri Azenkot, and Maya Cakmak. Designing a robot guide for blind people in indoor environments. In *The 10th Annual ACM/IEEE International Conference on Human-Robot Interaction Extended Abstracts*, pages 107–108. ACM, 2015. ISBN 978-1-4503-3318-4. doi: 10.1145/2701973.2702060. URL https://doi.org/10.1145/2701973.2702060.

Francesco Ferrari, Maria Paola Paladino, and Jolanda Jetten. Blurring human-machine distinctions: Anthropomorphic appearance in social robots as a threat to human distinctiveness. *International Journal of Social Robotics*, 8(2):287–302, 2016. doi: 10.1007/s12369-016-0338-y. URL https://doi.org/10.1007/s12369-016-0338-y.

Andy Field. *Discovering statistics using IBM SPSS statistics*. Sage, Thousand Oaks, CA, 2018. ISBN 9781526419514. URL http://www.worldcat.org/oclc/1030545826.

Andy Field and Graham Hole. *How to design and report experiments*. Sage, Thousand Oaks, CA, 2002. ISBN 978085702829. URL http://www.worldcat.org/title/how-to-design-and-report-experiments/oclc/961100072.

Julia Fink. Anthropomorphism and human likeness in the design of robots and human-robot interaction. In Shuzhi Sam Ge, Oussama Khatib, John-John Cabibihan, Reid Simmons, and Mary-Anne Williams, editors, *Social robotics*, pages 199–208, Berlin, Heidelberg, 2012. Springer. ISBN 978-3-642-34103-8. doi: 10.1007/

978-3-642-34103-8_20. URL https://doi.org/10.1007/978-3-642-34103-8_20.

Julia Fink, Valérie Bauwens, Frédéric Kaplan, and Pierre Dillenbourg. Living with a vacuum cleaning robot. *International Journal of Social Robotics*, 5(3):389–408, Aug 2013. ISSN 1875-4805. doi: 10.1007/s12369-013-0190-2. URL https://doi.org/10.1007/s12369-013-0190-2.

Julia Fink, Séverin Lemaignan, Pierre Dillenbourg, Philippe Rétornaz, Florian Vaussard, Alain Berthoud, Francesco Mondada, Florian Wille, and Karmen Franinović. Which robot behavior can motivate children to tidy up their toys? Design and evaluation of ranger. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 439–446. ACM, 2014. ISBN 978-1-4503-2658-2. doi: 10.1145/2559636.2559659. URL https://doi.org/10.1145/2559636.2559659.

Jaime Fisac, Andrea Bajcsy, Sylvia Herbert, David Fridovich-Keil, Steven Wang, Claire Tomlin, and Anca Dragan. Probabilistically safe robot planning with confidence-based human predictions. In *Proceedings of Robotics: Science and Systems*, Pittsburgh, Pennsylvania, June 2018. ISBN 978-0-9923747-4-7. doi: 10.15607/RSS.2018.XIV.069. URL https://doi.org/10.15607/RSS.2018.XIV.069.

Kerstin Fischer, Katrin Lohan, Joe Saunders, Chrystopher Nehaniv, Britta Wrede, and Katharina Rohlfing. The impact of the contingency of robot feedback on HRI. In *International Conference on Collaboration Technologies and Systems*, pages 210–217. IEEE, 2013. ISBN 978-1-4673-6403-4. doi: 10.1109/CTS.2013.6567231. URL https://doi.org/10.1109/CTS.2013.6567231.

Martin Ford. *The rise of the robots: Technology and the threat of mass unemployment*. Oneworld Publications, London, UK, 2015. ISBN 978-0465059997. URL http://www.worldcat.org/oclc/993846206.

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*, HRI '06, pages 258–265, New York, NY, 2006. ACM. ISBN 1-59593-294-1. doi: 10.1145/1121241.1121286. URL http://doi.acm.org/10.1145/1121241.1121286.

Floyd J. Fowler. *Improving survey questions: Design and evaluation*, volume 38. Sage, Thousand Oaks, CA, 1995. ISBN 978-0803945838. URL http://www.worldcat.org/oclc/551387270.

Floyd J. Fowler Jr. *Survey research methods*. Sage, Thousand Oaks, CA, 2013. ISBN 978-1452259000. URL http://www.worldcat.org/oclc/935314651.

Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine*, 4(1):23–33, 1997. doi: 10.1109/100.580977. URL https://doi.org/10.1109/100.580977.

Masahiro Fujita. Aibo: Toward the era of digital creatures. *International Journal of Robotics Research*, 20(10):781–794, 2001. doi: 10.1177/02783640122068092. URL https://doi.org/10.1177/02783640122068092.

Future of Life Institute. An open letter—research priorities for robust and beneficial artificial intelligence, January 2015. URL https://futureoflife.org/ai-open-letter/.

Joel Garreau. Bots on the ground. *Washington Post*, 2007. URL http://www.washingtonpost.com/wp-dyn/content/article/2007/05/05/AR2007050501009.html.

Adam Gazzaley and Larry D. Rosen. *The distracted mind: Ancient brains in a high-tech world*. MIT Press, Cambridge, MA, 2016. ISBN 978-0262534437. URL http://www.worldcat.org/oclc/978487215.

Guido H. E. Gendolla. On the impact of mood on behavior: An integrative theory and a review. *Review of General Psychology*, 4(4):378–408, 2000. doi: 10.1037/1089-2680.4.4.378. URL https://doi.org/10.1037/1089-2680.4.4.378.

Oliver Genschow, Sofie van Den Bossche, Emiel Cracco, Lara Bardi, Davide Rigoni, and Marcel Brass. Mimicry and automatic imitation are not correlated. *PloS One*, 12(9):e0183784, 2017. doi: 10.1371/journal.pone.0183784. URL https://doi.org/10.1371/journal.pone.0183784.

Robert M. Geraci. Spiritual robots: Religion and our scientific view of the natural world. *Theology and Science*, 4(3):229–246, 2006. doi: 10.1080/14746700600952993. URL https://doi.org/10.1080/14746700600952993.

James J. Gibson. *The ecological approach to visual perception: Classic edition*. Psychology Press, London, UK, 2014. ISBN 978-1848725782. URL http://www.worldcat.org/oclc/896794768.

Rachel Gockley, Allison Bruce, Jodi Forlizzi, Marek Michalowski, Anne Mundell, Stephanie Rosenthal, Brennan Sellner, Reid Simmons, Kevin Snipes, Alan C. Schultz, et al. Designing robots for long-term social interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1338–1343. IEEE, 2005. ISBN 0-7803-8912-3. doi: 10.1109/IROS.2005.1545303. URL https://doi.org/10.1109/IROS.2005.1545303.

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*, pages 186–193. ACM, 2006. ISBN 1-59593-294-1. doi: 10.1145/1121241.1121274. URL https://doi.org/10.1145/1121241.1121274.

Rachel Gockley, Jodi Forlizzi, and Reid Simmons. Natural person-following behavior for social robots. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 17–24. ACM, 2007. ISBN 978-1-59593-617-2. doi: 10.1145/1228716.1228720. URL https://doi.org/10.1145/1228716.1228720.

Ben Goldacre. *Bad science*. Fourth Estate, London, UK, 2008. ISBN 9780007240197. URL http://www.worldcat.org/oclc/760098401.

Joseph K. Goodman, Cynthia E. Cryder, and Amar Cheema. Data collection in a flat world: The strengths and weaknesses of mechanical turk samples. *Journal of Behavioral Decision Making*, 26(3):213–224, 2013. doi: 10.1002/bdm.1753. URL https://doi.org/10.1002/bdm.1753.

Eberhard Graether and Florian Mueller. Joggobot: A flying robot as jogging companion. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems*, pages 1063–1066, New York, NY, 2012. ACM. ISBN 978-1-4503-1016-1. doi: 10.1145/2212776.2212386. URL https://doi.org/10.1145/2212776.2212386.

Heather M. Gray, Kurt Gray, and Daniel M. Wegner. Dimensions of mind perception. *Science*, 315(5812):619–619, 2007. ISSN 0036-8075. doi: 10.1126/science.1134475. URL http://science.sciencemag.org/content/315/5812/619.

Leslie S. Greenberg. Application of emotion in psychotherapy. In Michael Lewis, Jeanette M. Haviland-Jones, and Lisa Feldman Barrett, editors, *Handbook of emotions*, volume 3, pages 88–101. Guilford Press, New York, NY, 2008. ISBN 978-1-59385-650-2. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.472.7583&rep=rep1&type=pdf.

H.-M. Gross, H. Boehme, Ch Schroeter, Steffen Müller, Alexander König, Erik Einhorn, Ch Martin, Matthias Merten, and Andreas Bley. Toomas: Interactive shopping guide robots in everyday use-final implementation and experiences from long-term field trials. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2005–2012. IEEE, 2009. ISBN 978-1-4244-3803-7. doi: 10.1109/IROS.2009.5354497. URL https://doi.org/10.1109/IROS.2009.5354497.

James J. Gross. Emotion regulation: Conceptual foundations. In James J. Gross, editor, *Handbook of emotion regulation*, chapter 1, pages 3–22. Guilford Press, 2007. ISBN 978-1462520732. URL http://www.worldcat.org/oclc/1027033463.

Hatice Gunes, Björn Schuller, Maja Pantic, and Roddy Cowie. Emotion representation, analysis and synthesis in continuous space: A survey. In *IEEE International Conference on Automatic Face & Gesture Recognition and Workshops*, pages 827–834. IEEE, 2011. ISBN 978-1-4244-9140-7. doi: 10.1109/FG.2011.5771357. URL https://doi.org/10.1109/FG.2011.5771357.

Martin Haegele. *World robotics service robots*. IFR Statistical Department, Chicago, IL, 2016. ISBN 9783816306948. URL http://www.worldcat.org/oclc/979905174.

Edward T. Hall, Ray L. Birdwhistell, Bernhard Bock, Paul Bohannan, A. Richard Diebold Jr., Marshall Durbin, Munro S. Edmonson, J. L. Fischer, Dell Hymes, Solon T. Kimball, et al. Proxemics [and comments and replies]. *Current Anthropology*, 9(2/3):83–108, 1968. doi: 10.1086/200975. URL https://doi.org/10.1086/200975.

Kun Han, Dong Yu, and Ivan Tashev. Speech emotion recognition using deep neural network and extreme learning machine. In *15th Annual Conference of the International Speech Communication Association*, pages 223–227, 2014. URL https://www.isca-speech.org/archive/archive_papers/interspeech_2014/i14_0223.pdf.

Takuya Hashimoto, Igor M. Verner, and Hiroshi Kobayashi. Human-like robot as teacher's representative in a science lesson: An elementary school experiment. In J. H. Kim, Matson E., and Xu P. Myung H., editors, *Robot intelligence technology and applications*, volume 208 of *Advances in Intelligent Systems and Computing*, pages 775–786. Springer, 2013. doi: 10.1007/978-3-642-37374-9_74. URL https://doi.org/10.1007/978-3-642-37374-9_74.

Nick Haslam. Dehumanization: An integrative review. *Personality and Social Psychology Review*, 10(3):252–264, 2006. doi: 10.1207/s15327957pspr1003_4. URL https://doi.org/10.1207/s15327957pspr1003_4.

Kotaro Hayashi, Masahiro Shiomi, Takayuki Kanda, Norihiro Hagita, and AI Robotics. Friendly patrolling: A model of natural encounters. In Hugh Durrant-Whyte, Nicholas Roy, and Pieter Abbeel, editors, *Robotics: Science and systems, Volume II*, pages 121–129. MIT Press, Cambridge, MA, 2012. ISBN 978-0-262-51779-9. URL http://www.worldcat.org/oclc/858018257.

Fritz Heider and Marianne Simmel. An experimental study of apparent behavior. *American Journal of Psychology*, 57(2):243–259, 1944. doi: 10.2307/1416950. URL https://doi.org/10.2307/1416950.

Mattias Heldner and Jens Edlund. Pauses, gaps and overlaps in conversations. *Journal of Phonetics*, 38(4):555–568, 2010. doi: 10.1016/j.wocn.2010.08.002. URL https://doi.org/10.1016/j.wocn.2010.08.002.

Carl-Herman Hjortsjö. *Man's face and mimic language*. Studen litteratur, Sweden, 1969. URL http://www.worldcat.org/oclc/974134474.

Chin-Chang Ho and Karl F. MacDorman. Revisiting the Uncanny Valley theory: Developing and validating an alternative to the Godspeed indices. *Computers in Human Behavior*, 26(6):1508–1518, 2010. doi: 10.1016/j.chb.2010.05.015. URL https://doi.org/10.1016/j.chb.2010.05.015.

Guy Hoffman. Dumb robots, smart phones: A case study of music listening companionship. In *The 21st IEEE International Symposium on Robot and Human Interactive Communication*, pages 358–363. IEEE, 2012. ISBN 978-1-4673-4604-7. doi: 10.1109/ROMAN.2012.6343779. URL https://doi.org/10.1109/ROMAN.2012.6343779.

Guy Hoffman and Cynthia Breazeal. Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. In *Proceedings of* *the ACM/IEEE International Conference on Human-Robot Interaction*, pages 1–8. ACM, 2007. ISBN 978-1-59593-617-2. doi: 10.1145/1228716.1228718. URL `https://doi.org/10.1145/1228716.1228718`.

Guy Hoffman and Keinan Vanunu. Effects of robotic companionship on music enjoyment and agent perception. In *8th ACM/IEEE International Conference on Human-Robot Interaction*, pages 317–324. IEEE, 2013. ISBN 978-1-4673-3099-2. doi: 10.1109/HRI.2013.6483605. URL `https://doi.org/10.1109/HRI.2013.6483605`.

Guy Hoffman and Gil Weinberg. Shimon: An interactive improvisational robotic marimba player. In *CHI'10 Extended Abstracts on Human Factors in Computing Systems*, pages 3097–3102. ACM, 2010. ISBN 978-1-60558-930-5. doi: 10.1145/1753846.1753925. URL `https://doi.org/10.1145/1753846.1753925`.

Olle Holm. Analyses of longing: Origins, levels, and dimensions. *Journal of Psychology*, 133(6):621–630, 1999. doi: 10.1080/00223989909599768. URL `https://doi.org/10.1080/00223989909599768`.

Deanna Hood, Séverin Lemaignan, and Pierre Dillenbourg. When children teach a robot to write: An autonomous teachable humanoid which uses simulated handwriting. In *10th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 83–90. ACM, 2015. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696479. URL `https://doi.org/10.1145/2696454.2696479`.

Andrew J. Hunt and Alan W. Black. Unit selection in a concatenative speech synthesis system using a large speech database. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 373–376. IEEE, 1996. ISBN 0-7803-3192-3. doi: 10.1109/ICASSP.1996.541110. URL `https://doi.org/10.1109/ICASSP.1996.541110`.

Helge Hüttenrauch, Kerstin Severinson Eklundh, Anders Green, and Elin A. Topp. Investigating spatial relationships in human-robot interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5052–5059. IEEE, 2006. ISBN 1-4244-0258-1. doi: 10.1109/IROS.2006.282535. URL `https://doi.org/10.1109/IROS.2006.282535`.

Michita Imai, Tetsuo Ono, and Hiroshi Ishiguro. Physical relation and expression: Joint attention for human-robot interaction. *IEEE Transactions on Industrial Electronics*, 50(4):636–643, 2003. doi: 10.1109/TIE.2003.814769. URL `https://doi.org/10.1109/TIE.2003.814769`.

Bahar Irfan, James Kennedy, Séverin Lemaignan, Fotios Papadopoulos, Emmanuel Senft, and Tony Belpaeme. Social psychology and human-robot interaction: An uneasy marriage. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '18, pages 13–20, New York, NY, 2018. ACM. ISBN 978-1-4503-5615-2. doi: 10.1145/3173386.3173389. URL `http://doi.acm.org/10.1145/3173386.3173389`.

Carlos T. Ishi, Chaoran Liu, Hiroshi Ishiguro, and Norihiro Hagita. Speech-driven lip motion generation for tele-operated humanoid robots. In *Auditory-visual speech processing*, pages 131–135, 2011. URL `https://www.isca-speech.org/archive/avsp11/papers/av11_131.pdf`.

Hiroshi Ishiguro. Android science. In Thrun S., Brooks R., and Durrant-Whyte H., editors, *Robotics research*, pages 118–127. Springer, 2007. ISBN 978-3-540-48110-2. doi: 10.1007/978-3-540-48113-3_11. URL `https://doi.org/10.1007/978-3-540-48113-3_11`.

Oliver P. John, Sanjay Srivastava, et al. The Big Five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research*, 2(1999):102–138, 1999.

Spike Jonze. *Her*, 2013. URL `https://www.imdb.com/title/tt1798709/?ref_=fn_al_tt_1`.

Jutta Joormann and Ian H. Gotlib. Emotion regulation in depression: Relation to cognitive inhibition. *Cognition and Emotion*, 24(2):281–298, 2010. doi: 10.1080/02699930903407948. URL https://doi.org/10.1080/02699930903407948.

Malte Jung and Pamela Hinds. Robots in the wild: A time for more robust theories of human-robot interaction. *ACM Transactions on Human-Robot Interaction (THRI)*, 7(1):2, 2018. doi: 10.1145/3208975. URL https://doi.org/10.1145/3208975.

Peter H. Kahn, Nathan G. Freier, Takayuki Kanda, Hiroshi Ishiguro, Jolina H. Ruckert, Rachel L. Severson, and Shaun K. Kane. Design patterns for sociality in human-robot interaction. In *The 3rd ACM/IEEE International Conference on Human-Robot Interaction*, pages 97–104. ACM, 2008. ISBN 978-1-60558-017-3. doi: 10.1145/1349822.1349836. URL https://doi.org/10.1145/1349822.1349836.

Peter H. Kahn Jr., Takayuki Kanda, Hiroshi Ishiguro, Brian T. Gill, Jolina H. Ruckert, Solace Shen, Heather E. Gary, Aimee L. Reichert, Nathan G. Freier, and Rachel L. Severson. Do people hold a humanoid robot morally accountable for the harm it causes? In *Proceedings of the 7th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 33–40. ACM, 2012. ISBN 978-1-4503-1063-5. doi: 10.1145/2157689.2157696. URL https://doi.org/10.1145/2157689.2157696.

Peter H. Kahn Jr., Takayuki Kanda, Hiroshi Ishiguro, Solace Shen, Heather E. Gary, and Jolina H. Ruckert. Creative collaboration with a social robot. In *ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 99–103. ACM, 2014. ISBN 978-1-4503-2968-2. doi: 10.1145/2632048.2632058. URL https://doi.org/10.1145/2632048.2632058.

Takayuki Kanda, Takayuki Hirano, Daniel Eaton, and Hiroshi Ishiguro. Interactive robots as social partners and peer tutors for children: A field trial. *Human-Computer Interaction*, 19(1):61–84, 2004. doi: 10.1080/07370024.2004.9667340. URL https://doi.org/10.1080/07370024.2004.9667340.

Takayuki Kanda, Masayuki Kamasima, 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, 2007a. doi: 10.1007/s10514-006-9007-6. URL https://doi.org/10.1007/s10514-006-9007-6.

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, 2007b. doi: 10.1109/TRO.2007.904904. URL https://doi.org/10.1109/TRO.2007.904904.

Takayuki Kanda, Masahiro Shiomi, Zenta Miyashita, Hiroshi Ishiguro, and Norihiro Hagita. A communication robot in a shopping mall. *IEEE Transactions on Robotics*, 26(5):897–913, 2010. doi: 10.1109/TRO.2010.2062550. URL https://doi.org/10.1109/TRO.2010.2062550.

Kyong Il Kang, Sanford Freedman, Maja J. Mataric, Mark J. Cunningham, and Becky Lopez. A hands-off physical therapy assistance robot for cardiac patients. In *9th International Conference on Rehabilitation Robotics (ICORR)*, pages 337–340. IEEE, 2005. ISBN 0-7803-9003-2. doi: 10.1109/ICORR.2005.1501114. URL https://doi.org/10.1109/ICORR.2005.1501114.

Frederic Kaplan. Who is afraid of the humanoid? Investigating cultural differences in the acceptance of robots. *International Journal of Humanoid Robotics*, 1(3):1–16, 2004. doi: 10.1142/S0219843604000289. URL https://doi.org/10.1142/S0219843604000289.

Victor Kaptelinin. Technology and the givens of existence: Toward an existential inquiry framework in HCI research. In *Proceedings of the 2018 CHI Conference on*

Human Factors in Computing Systems, CHI '18, pages 270:1–270:14, New York, NY, 2018. ACM. ISBN 978-1-4503-5620-6. doi: 10.1145/3173574.3173844. URL `http://doi.acm.org/10.1145/3173574.3173844`.

Yusuke Kato, Takayuki Kanda, and Hiroshi Ishiguro. May I help you? design of human-like polite approaching behavior. In *10th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 35–42. ACM, 2015. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696463. URL `https://doi.org/10.1145/2696454.2696463`.

Dacher Keltner and Ann M. Kring. Emotion, social function, and psychopathology. *Review of General Psychology*, 2(3):320–342, 1998. doi: 10.1037/1089-2680.2.3.320. URL `https://doi.org/10.1037/1089-2680.2.3.320`.

Theodore D. Kemper. How many emotions are there? Wedding the social and the autonomic components. *American Journal of Sociology*, 93(2):263–289, 1987. doi: 10.1086/228745. URL `https://doi.org/10.1086/228745`.

Adam Kendon. *Conducting interaction: Patterns of behavior in focused encounters*. Cambridge University Press, Cambridge, UK, 1990. ISBN 978-0521389389. URL `http://www.worldcat.org/oclc/785489376`.

James Kennedy, Paul Baxter, and Tony Belpaeme. The robot who tried too hard: Social behaviour of a robot tutor can negatively affect child learning. In *10th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 67–74. ACM, 2015. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696457. URL `https://doi.org/10.1145/2696454.2696457`.

James Kennedy, Séverin Lemaignan, Caroline Montassier, Pauline Lavalade, Bahar Irfan, Fotios Papadopoulos, Emmanuel Senft, and Tony Belpaeme. Child speech recognition in human-robot interaction: Evaluations and recommendations. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, pages 82–90. ACM, 2017. ISBN 978-1-4503-4336-7. doi: 10.1145/2909824.3020229. URL `https://doi.org/10.1145/2909824.3020229`.

Cory D. Kidd and Cynthia Breazeal. A robotic weight loss coach. In *Proceedings of the 22nd National Conference on Artificial Intelligence, Volume 2*, AAAI'07, pages 1985–1986. AAAI Press, 2007. ISBN 978-1-57735-323-2. URL `http://dl.acm.org/citation.cfm?id=1619797.1619992`.

Cory D. Kidd and Cynthia Breazeal. Robots at home: Understanding long-term human-robot interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3230–3235. IEEE, 2008. ISBN 978-1-4244-2057-5. doi: 10.1109/IROS.2008.4651113. URL `https://doi.org/10.1109/IROS.2008.4651113`.

Sara Kiesler, Aaron Powers, Susan R. Fussell, and Cristen Torrey. Anthropomorphic interactions with a robot and robot-like agent. *Social Cognition*, 26(2):169–181, 2008. doi: 10.1521/soco.2008.26.2.169. URL `https://doi.org/10.1521/soco.2008.26.2.169`.

Ki Joon Kim, Eunil Park, and S. Shyam Sundar. Caregiving role in human–robot interaction: A study of the mediating effects of perceived benefit and social presence. *Computers in Human Behavior*, 29(4):1799–1806, 2013. doi: 10.1016/j.chb.2013.02.009. URL `https://doi.org/10.1016/j.chb.2013.02.009`.

Sandra L. Kirmeyer and Thung-Rung Lin. Social support: Its relationship to observed communication with peers and superiors. *Academy of Management Journal*, 30(1):138–151, 1987. doi: 10.5465/255900. URL `https://doi.org/10.5465/255900`.

Naho Kitano. "Rinri": An incitement towards the existence of robots in Japanese society. *International Review of Information Ethics*, 6:78–83, 2006. URL `http://www.i-r-i-e.net/inhalt/006/006_Kitano.pdf`.

Frank Klassner. A case study of LEGO Mindstorms suitability for artificial intelligence and robotics courses at the college level. *SIGCSE Bulletin*, 34(1): 8–12, February 2002. ISSN 0097-8418. doi: 10.1145/563517.563345. URL http://doi.acm.org/10.1145/563517.563345.

Kheng Lee Koay, Emrah Akin Sisbot, Dag Sverre Syrdal, Mick L. Walters, Kerstin Dautenhahn, and Rachid Alami. Exploratory study of a robot approaching a person in the context of handing over an object. In *AAAI Spring Symposium: Multidisciplinary Collaboration for Socially Assistive Robotics*, pages 18–24, 2007a. URL http://www.aaai.org/Papers/Symposia/Spring/2007/SS-07-07/SS07-07-004.pdf.

Kheng Lee Koay, Dag Sverre Syrdal, Michael L. Walters, and Kerstin Dautenhahn. Living with robots: Investigating the habituation effect in participants' preferences during a longitudinal human-robot interaction study. In *The 16th IEEE International Symposium on Robot and Human Interactive Communication*, pages 564–569. IEEE, 2007b. ISBN 978-1-4244-1634-9. doi: 10.1109/ROMAN.2007.4415149. URL https://doi.org/10.1109/ROMAN.2007.4415149.

Thomas Kollar, Stefanie Tellex, Deb Roy, and Nicholas Roy. Toward understanding natural language directions. In *5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 259–266. IEEE, 2010. ISBN 978-1-4244-4892-0. doi: 10.1109/HRI.2010.5453186. URL https://doi.org/10.1109/HRI.2010.5453186.

Jeamin Koo, Jungsuk Kwac, Wendy Ju, Martin Steinert, Larry Leifer, and Clifford Nass. Why did my car just do that? Explaining semi-autonomous driving actions to improve driver understanding, trust, and performance. *International Journal on Interactive Design and Manufacturing (IJIDeM)*, 9(4):269–275, 2015. doi: 10.1007/s12008-014-0227-2. URL https://doi.org/10.1007/s12008-014-0227-2.

Stefan Kopp, Brigitte Krenn, Stacy Marsella, Andrew N. Marshall, Catherine Pelachaud, Hannes Pirker, Kristinn R. Thórisson, and Hannes Vilhjálmsson. Towards a common framework for multimodal generation: The behavior markup language. In *International Workshop on Intelligent Virtual Agents*, pages 205–217. Springer, 2006. ISBN 978-3-540-37593-7. doi: 10.1007/11821830_17. URL https://doi.org/10.1007/11821830_17.

Hideki Kozima, Marek P. Michalowski, and Cocoro Nakagawa. Keepon. *International Journal of Social Robotics*, 1(1):3–18, 2009. doi: 10.1007/s12369-008-0009-8. URL https://doi.org/10.1007/s12369-008-0009-8.

Thibault Kruse, Amit Kumar Pandey, Rachid Alami, and Alexandra Kirsch. Human-aware robot navigation: A survey. *Robotics and Autonomous Systems*, 61(12):1726–1743, 2013. doi: 10.1016/j.robot.2013.05.007. URL https://doi.org/10.1016/j.robot.2013.05.007.

Dieta Kuchenbrandt, Nina Riether, and Friederike Eyssel. Does anthropomorphism reduce stress in HRI? In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction*, pages 218–219, New York, NY, 2014. ACM. ISBN 978-1-4503-2658-2. doi: 10.1145/2559636.2563710. URL http://doi.org/10.1145/2559636.2563710.

Thomas S. Kuhn. *The structure of scientific revolutions*. University of Chicago Press, Chicago, IL, 2nd edition, 1970. ISBN 0226458032. URL http://www.worldcat.org/oclc/468581998.

Dana Kulic and Elizabeth A. Croft. Safe planning for human-robot interaction. *Journal of Field Robotics*, 22(7):383–396, 2005. doi: 10.1002/rob.20073. URL https://doi.org/10.1002/rob.20073.

Hideaki Kuzuoka, Yuya Suzuki, Jun Yamashita, and Keiichi Yamazaki. Reconfiguring spatial formation arrangement by robot body orientation. In *5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 285–292. IEEE

Press, 2010. ISBN 978-1-4244-4892-0. doi: 10.1109/HRI.2010.5453182. URL https://doi.org/10.1109/HRI.2010.5453182.

Peter J. Lang, Margaret M. Bradley, and Bruce N. Cuthbert. Motivated attention: Affect, activation, and action. In Peter J. Lang, Robert F. Simons, Marie Balaban, and Robert Simons, editors, *Attention and orienting: Sensory and motivational processes*, pages 97–135. Erlbaum, Hillsdale, NJ, 1997. ISBN 9781135808204. URL http://www.worldcat.org/oclc/949987355.

Juan S. Lara, Jonathan Casas, Andres Aguirre, Marcela Munera, Monica Rincon-Roncancio, Bahar Irfan, Emmanuel Senft, Tony Belpaeme, and Carlos A. Cifuentes. Human-robot sensor interface for cardiac rehabilitation. In *International Conference on Rehabilitation Robotics (ICORR)*, pages 1013–1018. IEEE, 2017. ISBN 978-1-5386-2296-4. doi: 10.1109/ICORR.2017.8009382. URL https://doi.org/10.1109/ICORR.2017.8009382.

Randy J. Larsen and Edward Diener. Promises and problems with the circumplex model of emotion. In Margaret S. Clark, editor, *Emotion: The review of personality and social psychology*, volume 13, chapter 2, pages 25–59. Thousand Oaks, CA: Sage, 1992. ISBN 978-0803946149. URL http://www.worldcat.org/oclc/180631851.

Richard S. Lazarus. *Emotion and adaptation*. Oxford University Press on Demand, 1991. ISBN 978-0195092660. URL http://www.worldcat.org/oclc/298419692.

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521 (7553):436, 2015. doi: 10.1038/nature14539. URL https://doi.org/10.1038/nature14539.

Hee Rin Lee, JaYoung Sung, Selma Šabanović, and Joenghye Han. Cultural design of domestic robots: A study of user expectations in Korea and the United States. In *IEEE International Workshop on Robot and Human Interactive Communication*, pages 803–808. IEEE, 2012. ISBN 978-1-4673-4604-7. doi: 10.1109/ROMAN.2012.6343850. URL https://doi.org/10.1109/ROMAN.2012.6343850.

Hee Rin Lee, Selma Šabanović, Wan-Ling Chang, Shinichi Nagata, Jennifer Piatt, Casey Bennett, and David Hakken. Steps toward participatory design of social robots: Mutual learning with older adults with depression. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 244–253. ACM, 2017. ISBN 978-1-4503-4336-7. doi: 10.1145/2909824.3020237. URL https://doi.org/10.1145/2909824.3020237.

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 ACM/IEEE International Conference on Human-Robot Interaction*, pages 7–14. ACM, 2009. ISBN 978-1-60558-404-1. doi: 10.1145/1514095.1514100. URL https://doi.org/10.1145/1514095.1514100.

Sau-lai Lee, Ivy Yee-man Lau, Sara Kiesler, and Chi-Yue Chiu. Human mental models of humanoid robots. In *IEEE International Conference on Robotics and Automation*, pages 2767–2772. IEEE, 2005. ISBN 0-7803-8914-X. doi: 10.1109/ROBOT.2005.1570532. URL https://doi.org/10.1109/ROBOT.2005.1570532.

Iolanda Leite, Ginevra Castellano, André Pereira, Carlos Martinho, and Ana Paiva. Modelling empathic behaviour in a robotic game companion for children: An ethnographic study in real-world settings. In *Proceedings of the 7th Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '12, pages 367–374, New York, NY, 2012. ACM. ISBN 978-1-4503-1063-5. doi: 10.1145/2157689.2157811. URL https://dx.doi.org/10.1145/2157689.2157811.

Iolanda Leite, Carlos Martinho, and Ana Paiva. Social robots for long-term interaction: A survey. *International Journal of Social Robotics*, 5(2):291–308, 2013. doi: 10.1007/s12369-013-0178-y. URL https://doi.org/10.1007/s12369-013-0178-y.

Iolanda Leite, Marissa McCoy, Monika Lohani, Daniel Ullman, Nicole Salomons, Charlene Stokes, Susan Rivers, and Brian Scassellati. Emotional storytelling in the classroom: Individual versus group interaction between children and robots. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 75–82. ACM, 2015. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696481. URL https://doi.org/10.1145/2696454.2696481.

Séverin Lemaignan, Mathieu Warnier, E. Akin Sisbot, Aurélie Clodic, and Rachid Alami. Artificial cognition for social human-robot interaction: An implementation. *Artificial Intelligence*, 247:45–69, 2017. ISSN 0004-3702. doi: 10.1016/j.artint.2016.07.002. URL http://doi.org/10.1016/j.artint.2016.07.002.

Douglas B. Lenat. Cyc: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38, 1995. doi: 10.1145/219717.219745. URL https://doi.org/10.1145/219717.219745.

David Levy. *Love and sex with robots: The evolution of human-robot relationships*. Harper Collins, New York, NY, 2009. ISBN 978-0061359804. URL http://www.worldcat.org/oclc/1021135698.

Jacques-Philippe Leyens, Armando Rodríguez, Ramon Rodriguez-Torres, Ruth Gaunt, Maria Paladino, Jeroen Vaes, and Stéphanie Demoulin. Psychological essentialism and the differential attribution of uniquely human emotions to ingroups and outgroups. *European Journal of Social Psychology*, 31:395–411, 07 2001. doi: 10.1002/ejsp.50. URL https://doi.org/10.1002/ejsp.50.

Daniel Leyzberg, Samuel Spaulding, Mariya Toneva, and Brian Scassellati. The physical presence of a robot tutor increases cognitive learning gains. In *Proceedings of the Cognitive Science Society*, pages 1882–1887, 2012. URL https://escholarship.org/uc/item/7ck0p200.

Patrick Lin, Keith Abney, and George A. Bekey. *Robot ethics: The ethical and social implications of robotics*. Intelligent robotics and autonomous agents. MIT Press, Cambridge, MA, 2012. ISBN 9780262016667. URL http://www.worldcat.org/oclc/1004334474.

Jessica Lindblom and Tom Ziemke. Social situatedness of natural and artificial intelligence: Vygotsky and beyond. *Adaptive Behavior*, 11(2):79–96, 2003. doi: 10.1177/10597123030112002. URL https://doi.org/10.1177/10597123030112002.

P. Liu, D. F. Glas, T. Kanda, and H. Ishiguro. Data-driven HRI: Learning social behaviors by example from human-human interaction. *IEEE Transactions on Robotics*, 32(4):988–1008, 2016. ISSN 1552-3098. doi: 10.1109/TRO.2016.2588880. URL https://doi.org/10.1109/TRO.2016.2588880.

Phoebe Liu, Dylan F. Glas, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. It's not polite to point: Generating socially-appropriate deictic behaviors towards people. In *The 8th ACM/IEEE International Conference on Human-Robot Interaction*, pages 267–274. IEEE Press, 2013. ISBN 978-1-4673-3099-2. doi: 10.1109/HRI.2013.6483598. URL https://doi.org/10.1109/HRI.2013.6483598.

Travis Lowdermilk. *User-centered design: A developer's guide to building user-friendly applications*. O'Reilly, Sebastopol, CA, 2013. ISBN 978-1449359805. URL http://www.worldcat.org/oclc/940703603.

Matthias Luber, Luciano Spinello, Jens Silva, and Kai O. Arras. Socially-aware robot navigation: A learning approach. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 902–907. IEEE, 2012. ISBN 978-1-4673-1737-5. doi: 10.1109/IROS.2012.6385716. URL https://doi.org/10.1109/IROS.2012.6385716.

Karl F. MacDorman, Sandosh K. Vasudevan, and Chin-Chang Ho. Does Japan really have robot mania? comparing attitudes by implicit and explicit measures. *AI & SOCIETY*, 23(4):485–510, Jul 2009. ISSN 1435-5655. doi: 10.1007/s00146-008-0181-2. URL `https://doi.org/10.1007/s00146-008-0181-2`.

Alex Mar. Modern love: Are we ready for intimacy with androids?, October 2017. URL `https://www.wired.com/2017/10/hiroshi-ishiguro-when-robots-act-just-like-humans/`. Online; accessed 7-September-2018.

Aarian Marshall and Alex Davies. Uber's self-driving car saw the woman it killed, report says. *Wired Magazine*, March 2018. URL `https://www.wired.com/story/uber-self-driving-crash-arizona-ntsb-report/`. Online; accessed 7-November-2018.

Paul Marshall, Yvonne Rogers, and Nadia Pantidi. Using f-formations to analyse spatial patterns of interaction in physical environments. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work*, pages 445–454. ACM, 2011. ISBN 978-1-4503-0556-3. doi: 10.1145/1958824.1958893. URL `https://doi.org/10.1145/1958824.1958893`.

Maja J. Matarić. *The robotics primer*. MIT Press, Cambridge, MA, 2007. ISBN 9780262633543. URL `http://www.worldcat.org/oclc/604083625`.

Nikolaos Mavridis. A review of verbal and non-verbal human–robot interactive communication. *Robotics and Autonomous Systems*, 63:22–35, 2015. ISSN 0921-8890. doi: 10.1016/j.robot.2014.09.031. URL `https://doi.org/10.1016/j.robot.2014.09.031`.

Richard E. Mayer and C. Scott DaPra. An embodiment effect in computer-based learning with animated pedagogical agents. *Journal of Experimental Psychology: Applied*, 18(3):239–252, 2012. doi: 10.1037/a0028616. URL `http://doi.org/10.1037/a0028616`.

Pamela McCorduck. *Machines who think: A personal inquiry into the history and prospects of artificial intelligence*. W. H. Freeman, San Francisco, 1979. ISBN 978-1568812052. URL `http://www.worldcat.org/oclc/748860627`.

Drew McDermott. Yes, computers can think. *New York Times*, 1997. URL `http://www.nytimes.com/1997/05/14/opinion/yes-computers-can-think.html`.

Albert Mehrabian. *Basic dimensions for a general psychological theory: Implications for personality, social, environmental, and developmental studies*. Oelgeschlager, Gunn & Hain, Cambridge, MA, 1980. ISBN 978-0899460048. URL `http://www.worldcat.org/oclc/925130232`.

Albert Mehrabian and James A. Russell. *An approach to environmental psychology*. MIT Press, Cambridge, MA, 1974. ISBN 9780262630719. URL `http://www.worldcat.org/oclc/318133343`.

Marek P. Michalowski, Selma Sabanovic, and Reid Simmons. A spatial model of engagement for a social robot. In *9th IEEE International Workshop on Advanced Motion Control*, pages 762–767. IEEE, 2006. ISBN 0-7803-9511-1. doi: 10.1109/AMC.2006.1631755. URL `https://doi.org/10.1109/AMC.2006.1631755`.

Marek P. Michalowski, Selma Sabanovic, and Hideki Kozima. A dancing robot for rhythmic social interaction. In *2nd ACM/IEEE International Conference on Human-Robot Interaction*, pages 89–96. IEEE, 2007. ISBN 978-1-59593-617-2. doi: 10.1145/1228716.1228729. URL `https://doi.org/10.1145/1228716.1228729`.

Sushmita Mitra and Tinku Acharya. Gesture recognition: A survey. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 37 (3):311–324, 2007. doi: 10.1109/TSMCC.2007.893280. URL `https://doi.org/10.1109/TSMCC.2007.893280`.

Roger K. Moore. A Bayesian explanation of the "uncanny valley" effect and related psychological phenomena. *Scientific Reports*, 2:864, 2012. doi: 10.1038/srep00864. URL http://dx.doi.org/10.1038/srep00864.

Luis Yoichi Morales Saiki, Satoru Satake, Takayuki Kanda, and Norihiro Hagita. Modeling environments from a route perspective. In *6th International Conference on Human-Robot interaction*, pages 441–448. ACM, 2011. ISBN 978-1-4503-0561-7. doi: 10.1145/1957656.1957815. URL https://doi.org/10.1145/1957656.1957815.

Luis Yoichi Morales Saiki, Satoru Satake, Rajibul Huq, Dylan Glas, Takayuki Kanda, and Norihiro Hagita. How do people walk side-by-side? Using a computational model of human behavior for a social robot. In *7th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 301–308. ACM, 2012. ISBN 978-1-4503-1063-5. doi: 10.1145/2157689.2157799. URL https://doi.org/10.1145/2157689.2157799.

Masahiro Mori. The Uncanny Valley. *Energy*, 7:33–35, 1970. doi: 10.1109/MRA.2012.2192811. URL https://doi.org/10.1109/MRA.2012.2192811.

Masahiro Mori. *The Buddha in the robot*. Tuttle Publishing, Tokyo, Japan, 1982. ISBN 978-4333010028. URL http://www.worldcat.org/oclc/843422852.

Masahiro Mori, Karl F. MacDorman, and Norri Kageki. The Uncanny Valley [from the field]. *IEEE Robotics & Automation Magazine*, 19(2):98–100, 2012. doi: 10.1109/MRA.2012.2192811. URL https://doi.org/10.1109/MRA.2012.2192811.

Omar Mubin, Catherine J. Stevens, Suleman Shahid, Abdullah Al Mahmud, and Jian-Jie Dong. A review of the applicability of robots in education. *Journal of Technology in Education and Learning*, 1(209-0015):1–7, 2013. doi: 10.2316/Journal.209.2013.1.209-0015. URL http://doi.org/10.2316/Journal.209.2013.1.209-0015.

Jonathan Mumm and Bilge Mutlu. Human-robot proxemics: Physical and psychological distancing in human-robot interaction. In *Proceedings of the 2011 ACM/IEEE International Conference on Human-Robot Interaction*, pages 331–338. ACM, 2011. ISBN 978-1-4503-0561-7. doi: 10.1145/1957656.1957786. URL https://dl.acm.org/citation.cfm?doid=1957656.1957786.

Mike Murphy. The beginning of the end: Google's AI has beaten a top human player at the complex game of go. *Quartz*, 2016. URL https://qz.com/636637/the-beginning-of-the-end-googles-ai-has-beaten-a-top-human-player-at-the-complex-game-of-go/.

Bilge Mutlu and Jodi Forlizzi. Robots in organizations: The role of workflow, social, and environmental factors in human-robot interaction. In *3rd ACM/IEEE International Conference on Human-Robot Interaction*, pages 287–294. IEEE, 2008. ISBN 978-1-60558-017-3. doi: 10.1145/1349822.1349860. URL https://doi.org/10.1145/1349822.1349860.

Bilge Mutlu, Jodi Forlizzi, and Jessica Hodgins. A storytelling robot: Modeling and evaluation of human-like gaze behavior. In *6th IEEE-RAS International Conference on Humanoid Robots*, pages 518–523. Citeseer, 2006. ISBN 1-4244-0199-2. doi: https://doi.org/10.1109/ICHR.2006.321322. URL https://doi.org/10.1109/ICHR.2006.321322.

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 *The 4th ACM/IEEE International Conference on Human-Robot Interaction*, pages 61–68. ACM, 2009. ISBN 978-1-60558-404-1. doi: 10.1145/1514095.1514109. URL https://doi.org/10.1145/1514095.1514109.

Bilge Mutlu, Takayuki Kanda, Jodi Forlizzi, Jessica Hodgins, and Hiroshi Ishiguro. Conversational gaze mechanisms for humanlike robots. *ACM Transactions on Interactive Intelligent Systems*, 1(2):12, 2012. doi: 10.1145/2070719.2070725. URL https://doi.org/10.1145/2070719.2070725.

Yasushi Nakauchi and Reid Simmons. A social robot that stands in line. *Autonomous Robots*, 12(3):313–324, 2002. doi: 10.1023/A:1015273816637. URL https://doi.org/10.1023/A:1015273816637.

Roberto Navigli and Simone Paolo Ponzetto. Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artificial Intelligence*, 193:217–250, 2012. doi: 10.1016/j.artint.2012.07.001. URL https://doi.org/10.1016/j.artint.2012.07.001.

C. L. Nehaniv, K. Dautenhahn, J. Kubacki, M. Haegele, C. Parlitz, and R. Alami. A methodological approach relating the classification of gesture to identification of human intent in the context of human-robot interaction. In *IEEE International Workshop on Robot and Human Interactive Communication*, pages 371–377, 2005. ISBN 0780392744. doi: 10.1109/ROMAN.2005.1513807. URL https://doi.org/10.1109/ROMAN.2005.1513807.

Shogo Nishiguchi, Kohei Ogawa, Yuichiro Yoshikawa, Takenobu Chikaraishi, Oriza Hirata, and Hiroshi Ishiguro. Theatrical approach: Designing human-like behaviour in humanoid robots. *Robotics and Autonomous Systems*, 89:158–166, 2017. doi: 10.1016/j.robot.2016.11.017. URL https://doi.org/10.1016/j.robot.2016.11.017.

Don Norman. *The design of everyday things: Revised and expanded edition.* Basic Books, New York, NY, 2013. ISBN 9780465072996. URL http://www.worldcat.org/oclc/862103168.

Donald A. Norman. The way i see it: Signifiers, not affordances. *Interactions*, 15(6): 18–19, 2008. doi: 10.1145/1409040.1409044. URL https://doi.org/10.1145/1409040.1409044.

Brian A. Nosek, Charles R. Ebersole, Alexander DeHaven, and David Mellor. The preregistration revolution. *Proceedings of the National Academy of Sciences of the United States of America*, 115(11):2600–2606, 2017. doi: 10.1073/pnas.1708274114. URL https://doi.org/10.1073/pnas.1708274114.

Illah R. Nourbakhsh, Judith Bobenage, Sebastien Grange, Ron Lutz, Roland Meyer, and Alvaro Soto. An affective mobile robot educator with a full-time job. *Artificial Intelligence*, 114(1-2):95–124, 1999. doi: 10.1016/S0004-3702(99)00027-2. URL https://doi.org/10.1016/S0004-3702(99)00027-2.

Illah Reza Nourbakhsh. *Robot futures.* MIT Press, Cambridge, MA, 2013. ISBN 9780262018623. URL http://www.worldcat.org/oclc/945438245.

Jekaterina Novikova and Leon Watts. Towards artificial emotions to assist social coordination in HRI. *International Journal of Social Robotics*, 7(1):77–88, 2015. doi: 10.1007/s12369-014-0254-y. URL https://doi.org/10.1007/s12369-014-0254-y.

Regina Nuzzo. Statistical errors. *Nature*, 506(7487):150, 2014. doi: 10.1038/506150a. URL https://doi.org/10.1038/506150a.

Daniel M. Oppenheimer, Tom Meyvis, and Nicolas Davidenko. Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45(4):867–872, 2009. doi: 10.1016/j.jesp.2009.03.009. URL https://doi.org/10.1016/j.jesp.2009.03.009.

Andrew Ortony and Terence J Turner. What's basic about basic emotions? *Psychological Review*, 97(3):315, 1990. doi: 10.1037/0033-295X.97.3.315. URL https://doi.org/10.1037/0033-295X.97.3.315.

Andrew Ortony, Gerald Clore, and Allan Collins. *The cognitive structure of emotions.* Cambridge University Press, Cambridge, UK, 1988. ISBN 978-0521386647. URL http://www.worldcat.org/oclc/910015120.

Hirotaka Osawa, Ren Ohmura, and Michita Imai. Using attachable humanoid parts for realizing imaginary intention and body image. *International Journal of Social*

Robotics, 1(1):109–123, 2009. doi: 10.1007/s12369-008-0004-0. URL https://doi.org/10.1007/s12369-008-0004-0.

Elena Pacchierotti, Henrik I. Christensen, and Patric Jensfelt. Evaluation of passing distance for social robots. In *The 15th IEEE International Symposium on Robot and Human Interactive Communication*, pages 315–320. IEEE, 2006. ISBN 1-4244-0564-5. doi: 10.1109/ROMAN.2006.314436. URL https://doi.org/10.1109/ROMAN.2006.314436.

Steffi Paepcke and Leila Takayama. Judging a bot by its cover: An experiment on expectation setting for personal robots. In *5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 45–52. IEEE, 2010. ISBN 978-1-4244-4892-0. doi: 10.1109/HRI.2010.5453268. URL https://doi.org/10.1109/HRI.2010.5453268.

Maja Pantic, Alex Pentland, Anton Nijholt, and Thomas S. Huang. Human computing and machine understanding of human behavior: A survey. In Huang T. S., Nijholt A., Pantic M., and Pentland A., editors, *Artificial intelligence for human computing*, volume 4451 of *Lecture Notes in Computer Science*, pages 47–71. Springer, 2007. doi: 10.1007/978-3-540-72348-6_3. URL https://doi.org/10.1007/978-3-540-72348-6_3.

Hae Won Park, Mirko Gelsomini, Jin Joo Lee, and Cynthia Breazeal. Telling stories to robots: The effect of backchanneling on a child's storytelling. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 100–108. ACM, 2017. ISBN 978-1-4503-4336-7. doi: 10.1145/2909824.3020245. URL https://doi.org/10.1145/2909824.3020245.

Michael Partridge and Christoph Bartneck. The invisible naked guy: An exploration of a minimalistic robot. In *The First International Conference on Human-Agent Interaction*, pages II–2–p2, 2013. doi: 10.17605/OSF.IO/A4YM5. URL https://doi.org/10.17605/OSF.IO/A4YM5.

Alex Pentland and Tracy Heibeck. *Honest signals: How they shape our world*. MIT Press, Cambridge, MA, 2010. ISBN 978-0262515122. URL http://www.worldcat.org/oclc/646395585.

Ignacio Pérez-Hurtado, Jesús Capitán, Fernando Caballero, and Luis Merino. Decision-theoretic planning with person trajectory prediction for social navigation. In *Robot 2015: Second Iberian Robotics Conference*, pages 247–258. Springer, 2016. ISBN 978-3-319-27148-4. doi: 10.1007/978-3-319-27149-1_20. URL https://doi.org/10.1007/978-3-319-27149-1_20.

Thomas F. Pettigrew, Linda R. Tropp, Ulrich Wagner, and Oliver Christ. Recent advances in intergroup contact theory. *International Journal of Intercultural Relations*, 35(3):271–280, 2011. doi: 10.1016/j.ijintrel.2011.03.001. URL https://doi.org/10.1016/j.ijintrel.2011.03.001.

R. W. Picard. *Affective computing*. MIT Press, Cambridge, MA, 1997. ISBN 978-0262661157. URL https://mitpress.mit.edu/books/affective-computing.

Joelle Pineau, Michael Montemerlo, Martha Pollack, Nicholas Roy, and Sebastian Thrun. Towards robotic assistants in nursing homes: Challenges and results. *Robotics and Autonomous Systems*, 42(3-4):271–281, 2003. doi: 10.1016/S0921-8890(02)00381-0. URL https://doi.org/10.1016/S0921-8890(02)00381-0.

Robert M. Pirsig. *Zen and the art of motorcycle maintenance: An inquiry into values*. Morrow, New York, NY, 1974. ISBN 0688002307. URL http://www.worldcat.org/oclc/41356566.

Diego A. Pizzagalli, Avram J. Holmes, Daniel G. Dillon, Elena L. Goetz, Jeffrey L. Birk, Ryan Bogdan, Darin D. Dougherty, Dan V. Iosifescu, Scott L. Rauch, and Maurizio Fava. Reduced caudate and nucleus accumbens response to rewards in unmedicated individuals with major depressive disorder. *American Journal* *of Psychiatry*, 166(6):702–710, 2009. doi: 10.1016/j.jpsychires.2008.03.001. URL https://doi.org/10.1016/j.jpsychires.2008.03.001.

Robert Ed Plutchik and Hope R. Conte. *Circumplex models of personality and emotions*. American Psychological Association, Washington, D.C., 1997. ISBN 978-1557983800. URL http://www.worldcat.org/oclc/442562242.

Cristina Anamaria Pop, Ramona Simut, Sebastian Pintea, Jelle Saldien, Alina Rusu, Daniel David, Johan Vanderfaeillie, Dirk Lefeber, and Bram Vanderborght. Can the social robot Probo help children with autism to identify situation-based emotions? A series of single case experiments. *International Journal of Humanoid Robotics*, 10(03):1350025, 2013. doi: 10.1142/S0219843613500254. URL https://doi.org/10.1142/S0219843613500254.

Jonathan Posner, James A. Russell, and Bradley S. Peterson. The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and Psychopathology*, 17(3):715–734, 2005. doi: 10.1017/S0954579405050340. URL https://doi.org/10.1017/S0954579405050340.

Michael I. Posner. *Cognitive neuroscience of attention*. Guilford Press, New York, NY, 2011. ISBN 978-1609189853. URL http://www.worldcat.org/oclc/958053069.

Aaron Powers, Adam D. I. Kramer, Shirlene Lim, Jean Kuo, Sau-lai Lee, and Sara Kiesler. Eliciting information from people with a gendered humanoid robot. In *IEEE International Workshop on Robot and Human Interactive Communication*, pages 158–163. IEEE, 2005. ISBN 0-7803-9274-4. doi: 10.1109/ROMAN.2005.1513773. URL https://doi.org/10.1109/ROMAN.2005.1513773.

Aaron Powers, Sara Kiesler, Susan Fussell, and Cristen Torrey. Comparing a computer agent with a humanoid robot. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, pages 145–152. ACM, 2007. ISBN 978-1-59593-617-2. doi: 10.1145/1228716.1228736. URL https://doi.org/10.1145/1228716.1228736.

Byron Reeves and Clifford Ivar Nass. *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge University Press, Cambridge, UK, 1996. ISBN 978-1575860534. URL http://www.worldcat.org/oclc/796222708.

Natalia Reich-Stiebert and Friederike Eyssel. Learning with educational companion robots? Toward attitudes on education robots, predictors of attitudes, and application potentials for education robots. *International Journal of Social Robotics*, 7(5):875–888, Nov 2015. ISSN 1875-4805. doi: 10.1007/s12369-015-0308-9. URL https://doi.org/10.1007/s12369-015-0308-9.

Natalia Reich-Stiebert and Friederike Eyssel. Robots in the classroom: What teachers think about teaching and learning with education robots. In *International Conference on Social Robotics*, pages 671–680. Springer, 2016. ISBN 978-3-319-47436-6. doi: 10.1007/978-3-319-47437-3_66. URL https://doi.org/10.1007/978-3-319-47437-3_66.

Natalia Reich-Stiebert and Friederike Anne Eyssel. Leben mit robotern-eine onlinebefragung im deutschen sprachraum zur akzeptanz von servicerobotern im alltag (poster), 2013. URL https://pub.uni-bielefeld.de/publication/2907019.

Jasia Reichardt. *Robots: Fact, fiction, and prediction*. Thames and Hudson, London, UK, 1978. ISBN 9780140049381. URL http://www.worldcat.org/oclc/1001944069.

Nancy A. Remington, Leandre R. Fabrigar, and Penny S. Visser. Reexamining the circumplex model of affect. *Journal of Personality and Social Psychology*, 79 (2):286–300, 2000. doi: 10.1037/0022-3514.79.2.286. URL https://doi.org/10.1037/0022-3514.79.2.286.

Charles Rich, Brett Ponsler, Aaron Holroyd, and Candace L. Sidner. Recognizing engagement in human-robot interaction. In *5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 375–382. IEEE, 2010. ISBN 978-1-4244-4892-0. doi: 10.1109/HRI.2010.5453163. URL `https://doi.org/10.1109/HRI.2010.5453163`.

Laurel D. Riek. Wizard of Oz studies in HRI: A systematic review and new reporting guidelines. *Journal of Human-Robot Interaction*, 1(1):119–136, 2012. doi: 10.5898/JHRI.1.1.Riek. URL `https://doi.org/10.5898/JHRI.1.1.Riek`.

Laurel D. Riek, Philip C. Paul, and Peter Robinson. When my robot smiles at me: Enabling human-robot rapport via real-time head gesture mimicry. *Journal on Multimodal User Interfaces*, 3(1-2):99–108, 2010. doi: 10.1007/s12193-009-0028-2. URL `https://doi.org/10.1007/s12193-009-0028-2`.

Ben Robins, Paul Dickerson, Penny Stribling, and Kerstin Dautenhahn. Robot-mediated joint attention in children with autism: A case study in robot-human interaction. *Interaction Studies*, 5(2):161–198, 2004. doi: 10.1075/is.5.2.02rob. URL `https://doi.org/10.1075/is.5.2.02rob`.

Ben Robins, Kerstin Dautenhahn, and Paul Dickerson. From isolation to communication: A case study evaluation of robot assisted play for children with autism with a minimally expressive humanoid robot. In *2nd International Conferences on Advances in Computer-Human Interactions*, pages 205–211. IEEE, 2009. ISBN 978-1-4244-3351-3. doi: 10.1109/ACHI.2009.32. URL `https://doi.org/10.1109/ACHI.2009.32`.

Hayley Robinson, Bruce MacDonald, and Elizabeth Broadbent. The role of healthcare robots for older people at home: A review. *International Journal of Social Robotics*, 6(4):575–591, 2014. doi: 10.1007/s12369-014-0242-2. URL `https://doi.org/10.1007/s12369-014-0242-2`.

Raquel Ros, Séverin Lemaignan, E. Akin Sisbot, Rachid Alami, Jasmin Steinwender, Katharina Hamann, and Felix Warneken. Which one? grounding the referent based on efficient human-robot interaction. In *19th International Symposium in Robot and Human Interactive Communication*, pages 570–575, 2010. ISBN 1944-9445. doi: 10.1109/ROMAN.2010.5598719. URL `http://doi.org/10.1109/ROMAN.2010.5598719`.

Rasmus Rothe, Radu Timofte, and Luc Van Gool. Deep expectation of real and apparent age from a single image without facial landmarks. *International Journal of Computer Vision (IJCV)*, 126(2):144–157, 2016. doi: 10.1007/s11263-016-0940-3. URL `https://doi.org/10.1007/s11263-016-0940-3`.

Dirk Rothenbücher, Jamy Li, David Sirkin, Brian Mok, and Wendy Ju. Ghost driver: A field study investigating the interaction between pedestrians and driverless vehicles. In *25th IEEE International Symposium on Robot and Human Interactive Communication*, pages 795–802. IEEE, 2016. ISBN 978-1-5090-3930-2. doi: 10.1109/ROMAN.2016.7745210. URL `https://doi.org/10.1109/ROMAN.2016.7745210`.

James A. Russell. A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6):1161–1178, 1980. doi: 10.1037/h0077714. URL `https://doi.org/10.1037/h0077714`.

James A. Russell and Lisa Feldman Barrett. Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology*, 76(5):805, 1999. doi: 10.1037//0022-3514.76.5.805. URL `https://doi.org/10.1037//0022-3514.76.5.805`.

James A. Russell, Maria Lewicka, and Toomas Niit. A cross-cultural study of a circumplex model of affect. *Journal of Personality and Social Psychology*, 57(5):848–856, 1989. doi: 10.1037/0022-3514.57.5.848. URL `https://doi.org/10.1037/0022-3514.57.5.848`.

Stuart Russell and Peter Norvig. *Artificial intelligence: A modern approach.* Pearson, Essex, UK, 3rd edition, 2009. ISBN 978-0136042594. URL `http://www.worldcat.org/oclc/496976145`.

Mel D. Rutherford and Ashley M. Towns. Scan path differences and similarities during emotion perception in those with and without autism spectrum disorders. *Journal of Autism and Developmental Disorders*, 38(7):1371–1381, 2008. doi: 10.1007/s10803-007-0525-7. URL `https://doi.org/10.1007/s10803-007-0525-7`.

Selma Šabanović. *Imagine all the robots: Developing a critical practice of cultural and disciplinary traversals in social robotics.* PhD thesis, Doctoral Thesis Faculty of Rensselaer Polytechnic Institute, 2007. URL `digitool.rpi.edu:8881/dtl_publish/50/9729.html`.

Selma Šabanović. Emotion in robot cultures: Cultural models of affect in social robot design. In *Proceedings of the Conference on Design & Emotion (D&E2010)*, pages 4–11, 2010.

Selma Šabanović and Wan-Ling Chang. Socializing robots: Constructing robotic sociality in the design and use of the assistive robot PARO. *AI & Society*, 31 (4):537–551, 2016. doi: 10.1007/s00146-015-0636-1. URL `https://doi.org/10.1007/s00146-015-0636-1`.

Selma Šabanović, Marek P. Michalowski, and Reid Simmons. Robots in the wild: Observing human-robot social interaction outside the lab. In *9th IEEE International Workshop on Advanced Motion Control*, pages 596–601. IEEE, 2006. ISBN 0-7803-9511-1. doi: 10.1109/AMC.2006.1631758. URL `https://doi.org/10.1109/AMC.2006.1631758`.

Selma Šabanović, Sarah M. Reeder, and Bobak Kechavarzi. Designing robots in the wild: In situ prototype evaluation for a break management robot. *Journal of Human-Robot Interaction*, 3(1):70–88, February 2014. ISSN 2163-0364. doi: 10.5898/JHRI.3.1.Sabanovic. URL `https://doi.org/10.5898/JHRI.3.1.Sabanovic`.

Selma Šabanović, Wan-Ling Chang, Casey C. Bennett, Jennifer A. Piatt, and David Hakken. A robot of my own: Participatory design of socially assistive robots for independently living older adults diagnosed with depression. In *International Conference on Human Aspects of IT for the Aged Population*, pages 104–114. Springer, 2015. ISBN 978-3-319-20891-6. doi: 10.1007/978-3-319-20892-3_11. URL `https://doi.org/10.1007/978-3-319-20892-3_11`.

Harvey Sacks, Emanuel A. Schegloff, and Gail Jefferson. A simplest systematics for the organization of turn-taking for conversation. *Language*, 4:696–735, 1974. doi: 10.2307/412243. URL `https://doi.org/10.2307/412243`.

Martin Saerbeck and Christoph Bartneck. Perception of affect elicited by robot motion. In *5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 53–60. ACM, 2010. ISBN 978-1-4244-4893-7. doi: 10.1145/1734454.1734473. URL `https://doi.org/10.1145/1734454.1734473`.

Martin Saerbeck, Tom Schut, Christoph Bartneck, and Maddy Janse. Expressive robots in education—varying the degree of social supportive behavior of a robotic tutor. In *28th ACM Conference on Human Factors in Computing Systems (CHI2010)*, pages 1613–1622. ACM, 2010. ISBN 978-1-60558-929-9. doi: 10.1145/1753326.1753567. URL `https://doi.org/10.1145/1753326.1753567`.

Daisuke Sakamoto, Takayuki Kanda, Tetsuo Ono, Hiroshi Ishiguro, and Norihiro Hagita. Android as a telecommunication medium with a human-like presence. In *2nd ACM/IEEE International Conference on Human-Robot Interaction*, pages 193–200. IEEE, 2007. ISBN 978-1-59593-617-2. doi: 10.1145/1228716.1228743. URL `https://doi.org/10.1145/1228716.1228743`.

Maha Salem, Friederike Eyssel, Katharina Rohlfing, Stefan Kopp, and Frank Joublin. To err is human (-like): Effects of robot gesture on perceived anthropomorphism and likability. *International Journal of Social Robotics*, 5(3):313–323, 2013. doi: 10.1007/s12369-013-0196-9. URL https://doi.org/10.1007/s12369-013-0196-9.

P. Salvini, G. Ciaravella, W. Yu, G. Ferri, A. Manzi, B. Mazzolai, C. Laschi, S. R. Oh, and P. Dario. How safe are service robots in urban environments? Bullying a robot. In *19th International Symposium in Robot and Human Interactive Communication*, pages 368–374, 2010. ISBN 978-1-4244-7991-7. doi: 10.1109/ROMAN.2010.5654677. URL http://dx.doi.org/10.1109/ROMAN.2010.5654677.

Jyotirmay Sanghvi, Ginevra Castellano, Iolanda Leite, André Pereira, Peter W McOwan, and Ana Paiva. Automatic analysis of affective postures and body motion to detect engagement with a game companion. In *6th ACM/IEEE International Conference on Human-Robot Interaction*, pages 305–311. IEEE, 2011. ISBN 978-1-4503-0561-7. doi: 10.1145/1957656.1957781. URL https://doi.org/10.1145/1957656.1957781.

Porter Edward Sargent. *The new immoralities: Clearing the way for a new ethics.* Porter Sargent, Boston, MA, 2013. ISBN 978-1258541880. URL http://www.worldcat.org/oclc/3794581.

Satoru Satake, Takayuki Kanda, Dylan F. Glas, Michita Imai, Hiroshi Ishiguro, and Norihiro Hagita. How to approach humans? Strategies for social robots to initiate interaction. In *4th ACM/IEEE International Conference on Human-Robot Interaction*, pages 109–116. IEEE, 2009. ISBN 978-1-60558-404-1. doi: 10.1145/1514095.1514117. URL https://doi.org/10.1145/1514095.1514117.

Allison Sauppé and Bilge Mutlu. The social impact of a robot co-worker in industrial settings. In *33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 3613–3622. ACM, 2015. ISBN 978-1-4503-3145-6. doi: 10.1145/2702123.2702181. URL https://doi.org/10.1145/2702123.2702181.

Brian Scassellati. Imitation and mechanisms of joint attention: A developmental structure for building social skills on a humanoid robot. In Nehaniv C. L., editor, *Computation for metaphors, analogy, and agents*, volume 1562 of *Lecture Notes in Computer Science*, pages 176–195. Springer, 1999. ISBN 978-3-540-65959-4. doi: 10.1007/3-540-48834-0_11. URL https://doi.org/10.1007/3-540-48834-0_11.

Brian Scassellati. Investigating models of social development using a humanoid robot. In Barbara Webb and Thomas Consi, editors, *Biorobotics: Methods and applications*, pages 145–168. MIT Press, 2000. ISBN 9780262731416. URL http://www.worldcat.org/oclc/807529041.

Brian Scassellati, Henny Admoni, and Maja Matarić. Robots for use in autism research. *Annual Review of Biomedical Engineering*, 14:275–294, 2012. doi: 10.1146/annurev-bioeng-071811-150036. URL https://doi.org/10.1146/annurev-bioeng-071811-150036.

Klaus R. Scherer. Emotion as a multicomponent process: A model and some cross-cultural data. *Review of Personality & Social Psychology*, 1984. URL https://doi.org/psycinfo/1986-17269-001.

Leonhard Schilbach, Marcus Wilms, Simon B. Eickhoff, Sandro Romanzetti, Ralf Tepest, Gary Bente, N. Jon Shah, Gereon R. Fink, and Kai Vogeley. Minds made for sharing: Initiating joint attention recruits reward-related neurocircuitry. *Journal of Cognitive Neuroscience*, 22(12):2702–2715, 2010. doi: 10.1162/jocn.2009.21401. URL https://doi.org/10.1162/jocn.2009.21401.

Tyler Schnoebelen and Victor Kuperman. Using Amazon Mechanical Turk for linguistic research. *Psihologija*, 43(4):441–464, 2010. doi: 10.2298/PSI1004441S. URL https://doi.org/10.2298/PSI1004441S.

Billy Schonenberg and Christoph Bartneck. Mysterious machines. In *5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 349–350, Osaka, 2010. ACM. ISBN 978-1-4244-4893-7. doi: 10.1145/1734454.1734572. URL https://doi.org/10.1145/1734454.1734572.

Jake Schreier. Robot and Frank, 2013. URL https://www.imdb.com/title/tt1990314/.

John R. Searle. Minds, brains and programs. *Behavioral and Brain Sciences*, 3(3): 417–457, 1980. doi: 10.1017/S0140525X00005756. URL https://doi.org/10.1017/S0140525X00005756.

Charles R. Seger, Eliot R. Smith, Elise James Percy, and Frederica R. Conrey. Reach out and reduce prejudice: The impact of interpersonal touch on intergroup liking. *Basic and Applied Social Psychology*, 36(1):51–58, 2014. doi: 10.1080/01973533.2013.856786. URL https://doi.org/10.1080/01973533.2013.856786.

Aparna Shankar, Mark Hamer, Anne McMunn, and Andrew Steptoe. Social isolation and loneliness: Relationships with cognitive function during 4 years of follow-up in the English Longitudinal Study of Ageing. *Psychosomatic Medicine*, 75(2): 161–170, 2013. doi: 10.1097/PSY.0b013e31827f09cd. URL https://doi.org/10.1097/PSY.0b013e31827f09cd.

Amanda J. C. Sharkey. Should we welcome robot teachers? *Ethics and Information Technology*, 18(4):283–297, 2016. doi: 10.1007/s10676-016-9387-z. URL https://doi.org/10.1007/s10676-016-9387-z.

Megha Sharma, Dale Hildebrandt, Gem Newman, James E. Young, and Rasit Eskicioglu. Communicating affect via flight path: Exploring use of the Laban effort system for designing affective locomotion paths. In *8th ACM/IEEE International Conference on Human-Robot Interaction*, pages 293–300. IEEE, 2013. ISBN 978-1-4673-3099-2. doi: 10.1109/HRI.2013.6483602. URL https://doi.org/10.1109/HRI.2013.6483602.

Glenda Shaw-Garlock. Looking forward to sociable robots. *International Journal of Social Robotics*, 1(3):249–260, Aug 2009. ISSN 1875-4805. doi: 10.1007/s12369-009-0021-7. URL https://doi.org/10.1007/s12369-009-0021-7.

Chao Shi, Masahiro Shiomi, Christian Smith, Takayuki Kanda, and Hiroshi Ishiguro. A model of distributional handing interaction for a mobile robot. In *Robotics: Science and systems*, pages 24–28, 2013. URL http://roboticsproceedings.org/rss09/p55.pdf.

Takanori Shibata. Therapeutic seal robot as biofeedback medical device: Qualitative and quantitative evaluations of robot therapy in dementia care. *Proceedings of the IEEE*, 100(8):2527–2538, 2012. doi: 10.1109/JPROC.2012.2200559. URL https://doi.org/10.1109/JPROC.2012.2200559.

Takanori Shibata, Kazuyoshi Wada, Yousuke Ikeda, and Selma Sabanovic. Cross-cultural studies on subjective evaluation of a seal robot. *Advanced Robotics*, 23 (4):443–458, 2009. doi: 10.1163/156855309X408826. URL https://doi.org/10.1163/156855309X408826.

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*, HRI '06, pages 305–312, New York, NY, 2006. ACM. ISBN 1-59593-294-1. doi: 10.1145/1121241.1121293. URL http://doi.acm.org/10.1145/1121241.1121293.

Masahiro Shiomi, Francesco Zanlungo, Kotaro Hayashi, and Takayuki Kanda. Towards a socially acceptable collision avoidance for a mobile robot navigating among pedestrians using a pedestrian model. *International Journal of Social Robotics*, 6(3):443–455, 2014. doi: 10.1007/s12369-014-0238-y. URL https://doi.org/10.1007/s12369-014-0238-y.

Bruno Siciliano and Oussama Khatib. *Springer handbook of robotics*. Springer, Berlin, 2016. ISBN 9783319325507. URL http://www.worldcat.org/oclc/945745190.

Jack Sidnell. *Conversation analysis: An introduction*, volume 45. John Wiley & Sons, New York, NY, 2011. ISBN 978-1405159012. URL http://www.worldcat.org/oclc/973423100.

Candace L. Sidner, Christopher Lee, Cory D. Kidd, Neal Lesh, and Charles Rich. Explorations in engagement for humans and robots. *Artificial Intelligence*, 166 (1-2):140–164, 2005. doi: 10.1016/j.artint.2005.03.005. URL https://doi.org/10.1016/j.artint.2005.03.005.

Herbert Alexander Simon. *The sciences of the artificial*. MIT Press, Cambridge, MA, 3rd edition, 1996. ISBN 0262691914. URL http://www.worldcat.org/oclc/552080160.

Peter W. Singer. *Wired for war: The robotics revolution and conflict in the twenty-first century*. Penguin, New York, NY, 2009. ISBN 9781594201981. URL http://www.worldcat.org/oclc/857636246.

Ashish Singh and James E. Young. Animal-inspired human-robot interaction: A robotic tail for communicating state. In *7th ACM/IEEE International Conference on Human-Robot Interaction*, pages 237–238. IEEE, 2012. ISBN 978-1-4503-1063-5. doi: 10.1145/2157689.2157773. URL https://doi.org/10.1145/2157689.2157773.

David Sirkin, Brian Mok, Stephen Yang, and Wendy Ju. Mechanical ottoman: How robotic furniture offers and withdraws support. In *10th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 11–18. ACM, 2015. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696461. URL https://doi.org/10.1145/2696454.2696461.

Emrah Akin Sisbot, Luis F. Marin-Urias, Rachid Alami, and Thierry Simeon. A human aware mobile robot motion planner. *IEEE Transactions on Robotics*, 23 (5):874–883, 2007. doi: 10.1109/TRO.2007.904911. URL https://doi.org/10.1109/TRO.2007.904911.

Ka-Chun Siu, Irene H Suh, Mukul Mukherjee, Dmitry Oleynikov, and Nick Stergiou. The effect of music on robot-assisted laparoscopic surgical performance. *Surgical Innovation*, 17(4):306–311, 2010. doi: 10.1177/1553350610381087. URL https://doi.org/10.1177/1553350610381087.

Aaron Smith. US views of technology and the future: Science in the next 50 years. Pew Research Center, April 17, 2014. URL http://assets.pewresearch.org/wp-content/uploads/sites/14/2014/04/US-Views-of-Technology-and-the-Future.pdf.

Richard L. Soash. Media equation: How people treat computers, television, and new media like real people and places. *Collection Management*, 24(3-4):310–311, 1999. doi: 10.1300/J105v24n03_14. URL https://doi.org/10.1300/J105v24n03_14.

Olivia Solon. Roomba creator responds to reports of "Poopocalypse": "We see this a lot". *The Guardian*, 2016. URL https://www.theguardian.com/technology/2016/aug/15/roomba-robot-vacuum-poopocalypse-facebook-post. Accessed: 2018-01-06.

Stefan Sosnowski, Ansgar Bittermann, Kolja Kuhnlenz, and Martin Buss. Design and evaluation of emotion-display EDDIE. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3113–3118. IEEE, 2006. ISBN 1-4244-0258-1. doi: 10.1109/IROS.2006.282330. URL https://doi.org/10.1109/IROS.2006.282330.

Robert Sparrow. Robotic weapons and the future of war. In Paolo Tripodi and Jessica Wolfendale, editors, *New wars and new soldiers: Military ethics in the* *contemporary world*, chapter 7, pages 117–133. Ashgate Surrey, 2011. ISBN 978-1-4094-0105-6. URL http://www.worldcat.org/oclc/960210186.

Robert Sparrow. Robots, rape, and representation. *International Journal of Social Robotics*, 9(4):465–477, Sep 2017. ISSN 1875-4805. doi: 10.1007/s12369-017-0413-z. URL https://doi.org/10.1007/s12369-017-0413-z.

Robert Sparrow and Linda Sparrow. In the hands of machines? The future of aged care. *Minds and Machines*, 16(2):141–161, 2006. doi: 10.1007/s11023-006-9030-6. URL https://doi.org/10.1007/s11023-006-9030-6.

Thorsten Spexard, Shuyin Li, Britta Wrede, Jannik Fritsch, Gerhard Sagerer, Olaf Booij, Zoran Zivkovic, Bas Terwijn, and Ben Krose. BIRON, where are you? Enabling a robot to learn new places in a real home environment by integrating spoken dialog and visual localization. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 934–940. IEEE, 2006. ISBN 1-4244-0258-1. doi: 10.1109/IROS.2006.281770. URL https://doi.org/10.1109/IROS.2006.281770.

A. Stedeman, D. Sutherland, and Christoph Bartneck. *Learning ROILA*. CreateSpace, Charleston, SC, 2011. ISBN 978-1466494978. URL https://www.createspace.com/3716932.

Luc Steels. The artificial life roots of artificial intelligence. *Artificial Life*, 1(1/2): 75–110, 1993. doi: 10.1162/artl.1993.1.1_2.75. URL https://doi.org/10.1162/artl.1993.1.1_2.75.

Nancy L. Stein and Keith Oatley. Basic emotions: Theory and measurement. *Cognition & Emotion*, 6(3-4):161–168, 1992. doi: 10.1080/02699939208411067. URL https://doi.org/10.1080/02699939208411067.

Mariëlle Stel, Rick B. Van Baaren, and Roos Vonk. Effects of mimicking: Acting prosocially by being emotionally moved. *European Journal of Social Psychology*, 38(6):965–976, 2008. doi: 10.1002/ejsp.472. URL https://doi.org/10.1002/ejsp.472.

Sofia Strömbergsson, Anna Hjalmarsson, Jens Edlund, and David House. Timing responses to questions in dialogue. In *Interspeech*, pages 2584–2588, 2013. URL http://www.isca-speech.org/archive/archive_papers/interspeech_2013/i13_2584.pdf.

Osamu Sugiyama, Takayuki Kanda, Michita Imai, Hiroshi Ishiguro, and Norihiro Hagita. Natural deictic communication with humanoid robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1441–1448. IEEE, 2007. ISBN 978-1-4244-0911-2. doi: 10.1109/IROS.2007.4399120. URL https://doi.org/10.1109/IROS.2007.4399120.

Ja-Young Sung, Lan Guo, Rebecca E. Grinter, and Henrik I. Christensen. “my Roomba is Rambo”: Intimate home appliances. In *9th International Conference on Ubiquitous Computing*, UbiComp ’07, pages 145–162, Berlin, Heidelberg, 2007. Springer-Verlag. ISBN 978-3-540-74852-6. doi: 10.1007/978-3-540-74853-3_9. URL https://doi.org/10.1007/978-3-540-74853-3_9.

JaYoung Sung, Rebecca E. Grinter, and Henrik I. Christensen. “Pimp my Roomba”: Designing for personalization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’09, pages 193–196, New York, NY, 2009. ACM. ISBN 978-1-60558-246-7. doi: 10.1145/1518701.1518732. URL http://doi.acm.org/10.1145/1518701.1518732.

Siddharth Suri and Duncan J. Watts. Cooperation and contagion in web-based, networked public goods experiments. *PloS One*, 6(3):e16836, 2011. doi: 10.1371/journal.pone.0016836. URL https://doi.org/10.1371/journal.pone.0016836.

Daniel Szafir, Bilge Mutlu, and Terry Fong. Communicating directionality in flying robots. In *The 10th Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 19–26. ACM, 2015. ISBN 978-1-4503-2883-8. doi: 10.1145/2696454.2696475. URL https://doi.org/10.1145/2696454.2696475.

Tajika Taichi, Miyashita Takahiro, Ishiguro Hiroshi, and Hagita Norihiro. Automatic categorization of haptic interactions—what are the typical haptic interactions between a human and a robot? In *6th IEEE-RAS International Conference on Humanoid Robots*, pages 490–496. IEEE, 2006. ISBN 1-4244-0199-2. doi: https://doi.org/10.1109/ICHR.2006.321318. URL 10.1109/ICHR.2006.321318.

Leila Takayama, Doug Dooley, and Wendy Ju. Expressing thought: Improving robot readability with animation principles. In *Proceedings of the 6th International Conference on Human-Robot Interaction*, pages 69–76. ACM, 2011. ISBN 978-1-4673-4393-0. doi: 10.1145/1957656.1957674. URL https://doi.org/10.1145/1957656.1957674.

Leila A. Takayama. *Throwing voices: Investigating the psychological effects of the spatial location of projected voices*. PhD thesis, Stanford University, 2008. URL https://searchworks.stanford.edu/view/7860025.

Fumihide Tanaka and Takeshi Kimura. Care-receiving robot as a tool of teachers in child education. *Interaction Studies*, 11(2):263–268, 2010. doi: 10.1075/is.11.2.14tan. URL https://doi.org/10.1075/is.11.2.14tan.

Fumihide Tanaka, Aaron Cicourel, and Javier R. Movellan. Socialization between toddlers and robots at an early childhood education center. *Proceedings of the National Academy of Sciences*, 104(46):17954–17958, 2007. doi: 10.1073/pnas.0707769104. URL https://doi.org/10.1073/pnas.0707769104.

Adriana Tapus, Maja J. Mataric, and Brian Scassellati. Socially assistive robotics [grand challenges of robotics]. *IEEE Robotics & Automation Magazine*, 14(1):35–42, 2007. doi: 10.1109/MRA.2007.339605. URL https://doi.org/10.1109/MRA.2007.339605.

Adriana Tapus, Andreea Peca, Amir Aly, Cristina Pop, Lavinia Jisa, Sebastian Pintea, Alina S. Rusu, and Daniel O. David. Children with autism social engagement in interaction with Nao, an imitative robot: A series of single case experiments. *Interaction Studies*, 13(3):315–347, 2012. doi: 10.1075/is.13.3.01tap. URL https://doi.org/10.1075/is.13.3.01tap.

Serge Thill, Cristina A. Pop, Tony Belpaeme, Tom Ziemke, and Bram Vanderborght. Robot-assisted therapy for autism spectrum disorders with (partially) autonomous control: Challenges and outlook. *Paladyn*, 3(4):209–217, 2012. doi: 10.2478/s13230-013-0107-7. URL https://doi.org/10.2478/s13230-013-0107-7.

Frank Thomas, Ollie Johnston, and Thomas Frank. *The illusion of life: Disney animation*. Hyperion, New York, NY, 1995. ISBN 978-0786860708. URL http://www.worldcat.org/oclc/974772586.

Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT Press, Cambridge, MA, 2005. ISBN 978-0-2622-0162-9. URL http://www.worldcat.org/oclc/705585641.

Jonas Togler, Fabian Hemmert, and Reto Wettach. Living interfaces: The thrifty faucet. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, pages 43–44. ACM, 2009. ISBN 978-1-60558-493-5. doi: 10.1145/1517664.1517680. URL https://doi.org/10.1145/1517664.1517680.

J. Gregory Trafton, Nicholas L. Cassimatis, Magdalena D. Bugajska, Derek P. Brock, Farilee E. Mintz, and Alan C. Schultz. Enabling effective human-robot interaction using perspective-taking in robots. *IEEE Trans. on Systems, Man, and Cybernetics. Part A: Systems and Humans*, 35(4):460–470, 2005. doi: 10.1109/TSMCA.2005.850592. URL https://doi.org/10.1109/TSMCA.2005.850592.

Robert Trappl, Paolo Petta, and Sabine Payr. *Emotions in humans and artifacts*. MIT Press, Cambridge, MA, 2003. ISBN 978-0262201421. URL https://mitpress.mit.edu/books/emotions-humans-and-artifacts.

Rudolph Triebel, Kai Arras, Rachid Alami, Lucas Beyer, Stefan Breuers, Raja Chatila, Mohamed Chetouani, Daniel Cremers, Vanessa Evers, Michelangelo Fiore, et al. Spencer: A socially aware service robot for passenger guidance and help in busy airports. In *Field and service robotics*, pages 607–622. Springer, 2016. ISBN 978-3-319-27700-4. doi: 10.1007/978-3-319-27702-8_40. URL https://doi.org/10.1007/978-3-319-27702-8_40.

Alan M. Turing. Computing machinery and intelligence. *Mind*, 59(236):433–460, 1950. doi: 10.1007/978-1-4020-6710-5_3. URL https://doi.org/10.1007/978-1-4020-6710-5_3.

Sherry Turkle. *Reclaiming conversation: The power of talk in a digital age*. Penguin, New York, NY, 2016. ISBN 978-0143109792. URL http://www.worldcat.org/oclc/960703115.

Sherry Turkle. *Alone together: Why we expect more from technology and less from each other*. Basic Books, New York, NY, 2017. ISBN 9780465031467. URL https://www.basicbooks.com/titles/sherry-turkle/alone-together/9780465093663/.

George E. Vaillan. *Triumphs of experience: The men of the Harvard Grant Study*. Belknap Press, Cambridge, MA, 2015. ISBN 978-0674503816. URL http://www.worldcat.org/oclc/910969527.

Albert van Breemen, Xue Yan, and Bernt Meerbeek. iCat: An animated user-interface robot with personality. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 143–144. ACM, 2005. ISBN 1-59593-093-0. doi: 10.1145/1082473.1082823. URL https://doi.org/10.1145/1082473.1082823.

Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv*, 2016. URL http://arxiv.org/abs/1609.03499.

Jan BF Van Erp and Alexander Toet. How to touch humans: Guidelines for social agents and robots that can touch. In *Humaine Association Conference on Affective Computing and Intelligent Interaction*, pages 780–785. IEEE, 2013. ISBN 978-0-7695-5048-0. doi: 10.1109/ACII.2013.145. URL https://doi.org/10.1109/ACII.2013.145.

Kurt VanLehn. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4):197–221, 2011. doi: 10.1080/00461520.2011.611369. URL https://doi.org/10.1080/00461520.2011.611369.

Gentiane Venture, Hideki Kadone, Tianxiang Zhang, Julie Grèzes, Alain Berthoz, and Halim Hicheur. Recognizing emotions conveyed by human gait. *International Journal of Social Robotics*, 6(4):621–632, 2014. doi: 10.1007/s12369-014-0243-1. URL https://doi.org/10.1007/s12369-014-0243-1.

Janet Vertesi. *Seeing like a rover: How robots, teams, and images craft knowledge of Mars*. University of Chicago Press, Chicago, IL, 2015. ISBN 978-0226155968. URL http://www.worldcat.org/oclc/904790036.

Gianmarco Veruggio, Fiorella Operto, and George Bekey. Roboethics: Social and ethical implications. In Bruno Siciliano and Oussama Khatib, editors, *Springer handbook of robotics*, pages 2135–2160. Springer, 2016. ISBN 978-3-319-32550-7. doi: 10.1007/978-3-319-32552-1. URL https://doi.org/10.1007/978-3-319-32552-1.

Walter G. Vincenti. *What engineers know and how they know it: Analytical studies from aeronautical history.* Johns Hopkins studies in the history of technology. Johns Hopkins University Press, Baltimore, MD, 1990. ISBN 0801839742. URL `http://www.worldcat.org/oclc/877307767`.

Anna-Lisa Vollmer, Robin Read, Dries Trippas, and Tony Belpaeme. Children conform, adults resist: A robot group induced peer pressure on normative social conformity. *Science Robotics*, 3(21):eaat7111, 2018. doi: 10.1126/scirobotics.aat7111. URL `https://doi.org/10.1126/scirobotics.aat7111`.

Karel Vredenburg, Ji-Ye Mao, Paul W Smith, and Tom Carey. A survey of user-centered design practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 471–478. ACM, 2002. ISBN 1-58113-453-3. doi: 10.1145/503376.503460. URL `https://doi.org/10.1145/503376.503460`.

Kazuyoshi Wada and Takanori Shibata. Living with seal robots—its sociopsychological and physiological influences on the elderly at a care house. *IEEE Transactions on Robotics*, 23(5):972–980, 2007. doi: 10.1109/TRO.2007.906261. URL `https://doi.org/10.1109/TRO.2007.906261`.

Jeffrey J. Walczyk, Karen S. Roper, Eric Seemann, and Angela M. Humphrey. Cognitive mechanisms underlying lying to questions: Response time as a cue to deception. *Applied Cognitive Psychology*, 17(7):755–774, 2003. doi: 10.1002/acp.914. URL `https://doi.org/10.1002/acp.914`.

Justin Walden, Eun Hwa Jung, S. Shyam Sundar, and Ariel Celeste Johnson. Mental models of robots among senior citizens: An interview study of interaction expectations and design implications. *Interaction Studies*, 16(1):68–88, 2015. doi: 10.1075/is.16.1.04wal. URL `https://doi.org/10.1075/is.16.1.04wal`.

Michael L. Walters, Kerstin Dautenhahn, René Te Boekhorst, Kheng Lee Koay, Christina Kaouri, Sarah Woods, Chrystopher Nehaniv, David Lee, and Iain Werry. The influence of subjects' personality traits on personal spatial zones in a human-robot interaction experiment. In *IEEE International Workshop on Robot and Human Interactive Communication*, pages 347–352. IEEE, 2005. ISBN 0-7803-9274-4. doi: 10.1109/ROMAN.2005.1513803. URL `https://doi.org/10.1109/ROMAN.2005.1513803`.

Michael L. Walters, Dag Sverre Syrdal, Kheng Lee Koay, Kerstin Dautenhahn, and R. Te Boekhorst. Human approach distances to a mechanical-looking robot with different robot voice styles. In *Robot and human interactive communication (RO-MAN)*, pages 707–712. IEEE, 2008. doi: 10.1109/ROMAN.2008.4600750. URL `https://doi.org/10.1109/ROMAN.2008.4600750`.

Michael L. Walters, Kerstin Dautenhahn, René Te Boekhorst, Kheng Lee Koay, Dag Sverre Syrdal, and Chrystopher L. Nehaniv. An empirical framework for human-robot proxemics. *Proceedings of New Frontiers in Human-Robot Interaction*, 2009. URL `http://hdl.handle.net/2299/9670`.

Lin Wang, Pei-Luen Patrick Rau, Vanessa Evers, Benjamin Krisper Robinson, and Pamela Hinds. When in Rome: The role of culture & context in adherence to robot recommendations. In *5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 359–366, Piscataway, NJ, USA, 2010. IEEE. ISBN 978-1-4244-4893-7. doi: 10.1109/HRI.2010.5453165. URL `https://doi.org/10.1109/HRI.2010.5453165`.

Rebecca M. Warner, Daniel Malloy, Kathy Schneider, Russell Knoth, and Bruce Wilder. Rhythmic organization of social interaction and observer ratings of positive affect and involvement. *Journal of Nonverbal Behavior*, 11(2):57–74, 1987. doi: 10.1007/BF00990958. URL `https://doi.org/10.1007/BF00990958`.

Miki Watanabe, Kohei Ogawa, and Hiroshi Ishiguro. Can androids be salespeople in the real world? In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, pages 781–788, New York, NY, 2015. ACM. ISBN 978-1-4503-3146-3. doi: 10.1145/2702613.2702967. URL `https://doi.org/10.1145/2702613.2702967`.

Adam Waytz, John Cacioppo, and Nicholas Epley. Who sees human? The stability and importance of individual differences in anthropomorphism. *Perspectives on Psychological Science*, 5(3):219–232, 2010. doi: 10.1177/1745691610369336. URL `https://doi.org/10.1177/1745691610369336`.

Blay Whitby. Sometimes it's hard to be a robot: A call for action on the ethics of abusing artificial agents. *Interacting with Computers*, 20(3):326–333, 2008. doi: 10.1016/j.intcom.2008.02.002. URL `https://doi.org/10.1016/j.intcom.2008.02.002`.

Andrew Whiten, Jane Goodall, William C. McGrew, Toshisada Nishida, Vernon Reynolds, Yukimaru Sugiyama, Caroline E. G. Tutin, Richard W. Wrangham, and Christophe Boesch. Cultures in chimpanzees. *Nature*, 399(6737):682–685, 1999. doi: 10.1038/21415. URL `https://doi.org/10.1038/21415`.

Christian J. A. M. Willemse, Gijs Huisman, Merel M. Jung, Jan B. F. van Erp, and Dirk K. J. Heylen. Observing touch from video: The influence of social cues on pleasantness perceptions. In *International Conference on Human Haptic Sensing and Touch Enabled Computer Applications*, pages 196–205. Springer, 2016. ISBN 978-3-319-42323-4. doi: 10.1007/978-3-319-42324-1_20. URL `https://doi.org/10.1007/978-3-319-42324-1_20`.

Kipling D. Williams. Ostracism. *Annual Review of Psychology*, 58:425–452, 2007. doi: 10.1146/annurev.psych.58.110405.085641. URL `https://doi.org/10.1146/annurev.psych.58.110405.085641`.

Lawrence E. Williams and John A. Bargh. Keeping one's distance: The influence of spatial distance cues on affect and evaluation. *Psychological Science*, 19(3): 302–308, 2008. doi: 10.1111/j.1467-9280.2008.02084.x. URL `https://doi.org/10.1111/j.1467-9280.2008.02084.x`.

Tom Williams, Daria Thames, Julia Novakoff, and Matthias Scheutz. Thank you for sharing that interesting fact: Effects of capability and context on indirect speech act use in task-based human-robot dialogue. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, pages 298–306. ACM, 2018. ISBN 978-1-4503-4953-6. doi: 10.1145/3171221.3171246. URL `https://doi.org/10.1145/3171221.3171246`.

Paul Wills, Paul Baxter, James Kennedy, Emmanuel Senft, and Tony Belpaeme. Socially contingent humanoid robot head behaviour results in increased charity donations. In *The 11th ACM/IEEE International Conference on Human-Robot Interaction*, pages 533–534. IEEE, 2016. ISBN 978-1-4673-8370-7. doi: 10.1109/HRI.2016.7451842. URL `https://doi.org/10.1109/HRI.2016.7451842`.

Daniel H. Wilson. *How to survive a robot uprising: Tips on defending yourself against the coming rebellion*. Bloomsbury, New York, NY, 2005. ISBN 9781582345925. URL `http://www.worldcat.org/oclc/1029483559`.

Katie Winkle, Praminda Caleb-Solly, Ailie Turton, and Paul Bremner. Social robots for engagement in rehabilitative therapies: Design implications from a study with therapists. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, pages 289–297. ACM, 2018. ISBN 978-1-4503-4953-6. doi: 10.1145/3171221.3171273. URL `https://doi.org/10.1145/3171221.3171273`.

Ryan Wistort and Cynthia Breazeal. Tofu: A socially expressive robot character for child interaction. In *8th International Conference on Interaction Design and Children*, pages 292–293. ACM, 2009. ISBN 978-1-60558-395-2. doi: 10.1145/1551788.1551862. URL `https://doi.org/10.1145/1551788.1551862`.

Ricarda Wullenkord. *Messung und Veränderung von Einstellungen gegenüber Robotern-Untersuchung des Einflusses von imaginiertem Kontakt auf implizite* *und explizite Maße*. PhD thesis, University of Bielefeld, 2017. URL https://pub.uni-bielefeld.de/publication/2913679.

Ricarda Wullenkord, Marlena R. Fraune, Friederike Eyssel, and Selma Šabanović. Getting in touch: How imagined, actual, and physical contact affect evaluations of robots. In *25th IEEE International Symposium on Robot and Human Interactive Communication*, pages 980–985. IEEE, 2016. ISBN 978-1-5090-3930-2. doi: 10.1109/ROMAN.2016.7745228. URL https://doi.org/10.1109/ROMAN.2016.7745228.

Junchao Xu, Joost Broekens, Koen Hindriks, and Mark A. Neerincx. Robot mood is contagious: Effects of robot body language in the imitation game. In *International Conference on Autonomous Agents and Multi-Agent Systems*, pages 973–980. International Foundation for Autonomous Agents and Multiagent Systems, 2014. ISBN 978-1-4503-2738-1. URL https://dl.acm.org/citation.cfm?id=2617401.

Yuto Yamaji, Taisuke Miyake, Yuta Yoshiike, P. Ravindra S. De Silva, and Michio Okada. STB: Human-dependent sociable trash box. In *5th ACM/IEEE International Conference on Human-Robot Interaction*, pages 197–198. IEEE, 2010. ISBN 978-1-4244-4892-0. doi: 10.1109/HRI.2010.5453196. URL https://doi.org/10.1109/HRI.2010.5453196.

Fumitaka Yamaoka, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. "Life-like" behavior of communication robots based on developmental psychology findings. In *5th IEEE-RAS International Conference on Humanoid Robots*, pages 406–411, 2005. ISBN 0-7803-9320-1. doi: 10.1109/ICHR.2005.1573601. URL https://doi.org/10.1109/ICHR.2005.1573601.

Fumitaka Yamaoka, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. A model of proximity control for information-presenting robots. *IEEE Transactions on Robotics*, 26(1):187–195, 2010. doi: 10.1109/TRO.2009.2035747. URL https://doi.org/10.1109/TRO.2009.2035747.

Steve Yohanan and Karon E. MacLean. The role of affective touch in human-robot interaction: Human intent and expectations in touching the haptic creature. *International Journal of Social Robotics*, 4(2):163–180, 2012. doi: 10.1007/s12369-011-0126-7. URL https://doi.org/10.1007/s12369-011-0126-7.

James E. Young, JaYoung Sung, Amy Voida, Ehud Sharlin, Takeo Igarashi, Henrik I. Christensen, and Rebecca E. Grinter. Evaluating human-robot interaction. *International Journal of Social Robotics*, 3(1):53–67, 2011. doi: 10.1007/s12369-010-0081-8. URL https://doi.org/10.1007/s12369-010-0081-8.

Chen Yu and Linda B. Smith. Joint attention without gaze following: Human infants and their parents coordinate visual attention to objects through eye-hand coordination. *PloS One*, 8(11):e79659, 2013. doi: 10.1371/journal.pone.0079659. URL https://doi.org/10.1371/journal.pone.0079659.

Robert B. Zajonc. Attitudinal effects of mere exposure. *Journal of Personality and Social Psychology*, 9(2p2):1–27, 1968. doi: 10.1037/h0025848. URL https://doi.org/10.1037/h0025848.

Heiga Zen, Keiichi Tokuda, and Alan W. Black. Statistical parametric speech synthesis. *Speech Communication*, 51(11):1039–1064, 2009. doi: 10.1016/j.specom.2009.04.004. URL https://doi.org/10.1016/j.specom.2009.04.004.

Zhihong Zeng, Maja Pantic, Glenn I. Roisman, and Thomas S. Huang. A survey of affect recognition methods: Audio, visual, and spontaneous expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(1):39–58, 2009. doi: 10.1109/TPAMI.2008.52. URL https://doi.org/10.1109/TPAMI.2008.52.

Jakub Zlotowski, Diane Proudfoot, Kumar Yogeeswaran, and Christoph Bartneck. Anthropomorphism: Opportunities and challenges in human-robot interaction. *International Journal of Social Robotics*, 7(3):347–360, 2015. doi: 10.1007/s12369-014-0267-6. URL https://doi.org/10.1007/s12369-014-0267-6.

Jakub Zlotowski, Kumar Yogeeswaran, and Christoph Bartneck. Can we control it? Autonomous robots are perceived as threatening. *International Journal of Human-Computer Studies*, 100(April 2017):48–54, 2017. doi: 10.1016/j.ijhcs.2016.12.008. URL `https://doi.org/10.1016/j.ijhcs.2016.12.008`.

Index

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

This material has been published by Cambridge University Press as Human Robot Interaction by
Christoph Bartneck, Tony Belpaeime, Friederike Eyssel, Takayuki Kanda, Merel Keijsers, and Selma Sabanovic.
ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

ISBN: 9781108735407 (http://www.cambridge.org/9781108735407).

Notes