Foundations and Trends[®] in Robotics Vol. 6, No. 4 (2016) 211–323 © 2018 E. Cha, Y. Kim, T. Fong and M. J. Mataric DOI: 10.1561/2300000057



A Survey of Nonverbal Signaling Methods for Non-Humanoid Robots

Elizabeth Cha University of Southern California echa@usc.edu

Yunkyung Kim NASA Ames Research Center yunkyung.kim@nasa.gov

Terrence Fong NASA Ames Research Center terry.fong@nasa.gov

Maja J. Matarić University of Southern California mataric@usc.edu

Contents

1	Intr	oduction	212
	1.1	Levels of Interaction	214
	1.2	Nonverbal Communication	216
	1.3	Motivation and Goals	217
	1.4	Organization	219
2	Bac	kground	221
	2.1	Attribution of Mental States	222
	2.2	Reasoning About the World	223
	2.3	Acquiring Awareness	224
	2.4	Effects of Culture, Media, and Experience	224
	2.5	Summary	226
3	Use	Cases	227
	3.1	Autonomous Vehicles	227
	3.2	Flying Robots	230
	3.3	Industrial Robots	233
	3.4	Service Robots	235
	3.5	Telepresence Robots	236
	3.6	Summary	238
4	Inte	rnal Robot States	240

	4.1	Condition	240		
	4.2	Knowledge	247		
	4.3	Activity	251		
	4.4	Summary	256		
5	Sign	al Characteristics	257		
	5.1	Criticality	257		
	5.2	Information Content	262		
	5.3	Summary	263		
6	Sign	al Generation	264		
	6.1	Gaze	264		
	6.2	Gesture	267		
	6.3	Implicit Motion	269		
	6.4	Sound	271		
	6.5	Visual Displays	273		
	6.6	Lights	275		
	6.7	Summary	277		
7	Discussion				
	7.1	Considerations for Nonverbal Signaling	278		
	7.2	Open Research Areas	288		
8	Conclusion				
Ac	Acknowledgements				
Re	References				

iii

Abstract

The goal of this survey is to inform the design and usage of nonverbal signals for human-robot interaction. With robots being increasingly utilized for tasks that require them to not only operate in close proximity to humans but to interact with them as well, there has been great interest in the communication challenges associated with the varying degrees of interaction in these environments. The success of such interactions depends on robots' ability to convey information about their knowledge, intent, and actions to co-located humans. In this work, we present a comprehensive review of literature related to the generation and usage of nonverbal signals that facilitate legibility of non-humanoid robot state and behavior. To motivate the need for these signaling behaviors, we survey literature in human communication and psychology and outline target use cases of non-humanoid robots. Specifically, we focus on works that provide insight into the cognitive processes that enable humans to recognize, interpret, and exploit nonverbal signals. From these use cases, we identify information that is potentially important for non-humanoid robots to signal and organize it into three categories of robot state. We then present a review of signal design techniques to illustrate how signals conveying this information can be generated and utilized. Finally, we discuss issues that must be considered during nonverbal signaling and open research areas, with a focus on informing the design and usage of generalizable nonverbal signaling behaviors for task-oriented non-humanoid robots.

E. Cha, Y. Kim, T. Fong and M. J. Mataric. A Survey of Nonverbal Signaling Methods

for Non-Humanoid Robots. Foundations and Trends[®] in Robotics, vol. 6, no. 4, pp. 211–323, 2016.

DOI: 10.1561/230000057.

1

Introduction

One of the primary goals of robotics is enabling robots to work alongside and with humans. This requires robots to be capable of not only navigating and manipulating in human environments but communicating and collaborating with humans as well [Khatib et al., 1999]. As autonomous robots are increasingly required to operate in concert with humans, the communication challenges associated with varying levels of proximate interaction must be explored [Arras and Cerqui, 2005, Dautenhahn et al., 2005, Forlizzi and DiSalvo, 2006].

In recent years, a significant portion of human-robot interaction (HRI) research has employed humanoid robots that possess highly anthropomorphic forms and features [Goodrich and Schultz, 2007]. However, there are also many robots which lack not only those features but humanoid form entirely. *Non-humanoid robots* typically have simpler embodiments that are targeted towards specific tasks or domains [Coeckelbergh, 2011, Terada et al., 2007]. As a result, these robots, such as those in Figure 1.1, are utilized in a wide variety of settings and applications, including healthcare, agriculture, space, industry, automotive, and service [Billingsley et al., 2008, Bualat et al., 2015, Fong et al.,



Figure 1.1: Non-humanoid robots, such as the Amazon warehouse robot, Waymo autonomous car, Savioke Relay, and NASA Astrobee (clockwise), are being regularly deployed and must use nonverbal signals to interact with co-located humans.

2013, Forlizzi and DiSalvo, 2006, Kittmann et al., 2015, Urmson et al., 2008, Wurman et al., 2008].

Historically, many commercial robots have been non-humanoid, with origins in industrial applications [Hockstein et al., 2007]. Consequently, past HRI research with these robots has largely focused on issues relating to safety, operation, and control within the context of highly regulated environments [Goodrich and Schultz, 2007, Thomaz et al., 2016]. However, as robots are deployed in other environments, their behavior must evolve to account for the needs and expectations of humans [Lee et al., 2010, Paepcke and Takayama, 2010].

Since non-humanoid robots are often more machine-like in appearance, humans are likely to have different expectations of them than of humanoid robots [Goetz et al., 2003]. Unlike other machines, however, robots can also operate as autonomous, intelligent agents that act in surprising ways. Therefore, it is vital that robots are capable of accu-

213

rately conveying their knowledge and capabilities in ways that humans can easily interpret and respond to.

A key challenge for utilizing intuitive and descriptive signals is the limited communication modalities available to most non-humanoid robots. Compared to human channels, many of these modalities are simplistic, making it challenging to generate a wide range of unique and recognizable signals [Bethel and Murphy, 2008, Harrison et al., 2012]. As a result, non-humanoid robots must also carefully consider how to optimally utilize the limited number of signals they have available.

The overarching goal of this work is to facilitate the development and usage of communicative robot signals that support varying levels of interaction while increasing transparency of a robot's internal state. This requires careful consideration of a number of factors, such as colocated humans and the robot's environment, in order for the robot to act intelligently.

Currently, more research is needed to understand these challenges in the context of non-humanoid robots. In particular, we focus on nonhumanoid robots with few to no anthropomorphic features that perform primarily functional (i.e., non-social tasks). As this encompasses a wide range of platforms utilized in research and commercial applications, we also present a selection of use cases to better illustrate factors that must be considered when designing communicative signals for such robots.

1.1 Levels of Interaction

Since humans and robots can interact in many different ways, we first propose the varying degrees of interactions that are of interest in this work [Yanco and Drury, 2004]. We focus on *proximate interactions* in which the human and robot are co-located, share the same environment, and can interact physically [Goodrich and Schultz, 2007]. Importantly, the robot's *internal state* is not readily accessible by the human. Instead, humans must utilize cues emitted by the robot to reason about the robot's internal state, similar to the cognitive processes that occur when interacting with another human [Klein et al., 2005].

1.1. Levels of Interaction

- **Coexistence** The lowest level of proximate interaction requires the robot and human to *coexist* in the same physical space. Coexistence does not require direct communication, but information is naturally exchanged through unconscious, nonverbal cues. In this lowest level of interaction, humans take on the role of bystanders or observers and do not directly interact with the robot.
- Coordination- A higher level of interaction involves the human and robot not only sharing space but coordinating their actions in time or space to deal with shared resources [Lorenz et al., 2011]. Often, the desired result of successful coordination is greater efficiency or the prevention of conflicts between agents [Fuks et al., 2007]. The term *coordination* is often used for single actions rather than long-term activities. In robotics, it is primarly used for physical tasks, such as two agents navigating the same area or jointly moving an object. To create fluid, seamless coordination, each agent must communicate enough information to create a degree of shared attention to support action prediction and planning [Bauer et al., 2008, Sebanz et al., 2006].
- Collaboration- The highest level of interaction, collaboration, requires high degrees of coordination as each party works actively towards an agreed upon, shared goal [Bauer et al., 2008]. To achieve a successful collaboration, the human and robot must communicate such that their actions are complementary and coordination between them is fluid [Bauer et al., 2008, Mutlu et al., 2013]. Human-robot collaboration (HRC) is important for HRI researchers as it leverages the strengths of both the human and robot for overall benefit [Fong et al., 2003c] and may facilitate robot deployment into environments where they are not yet capable of performing all tasks autonomously.

HRC research often focuses on methods for achieving fluid, effortless coordination [Bauer et al., 2008, Dragan et al., 2015, Hoffman and Breazeal, 2007, Strabala et al., 2013, Unhelkar et al., 2014]. Literature has emphasized the role of nonverbal communication in enabling this synchronization among multiple agents [Breazeal et al., 2005]. However, *nonverbal communication* also plays an important role in all levels of interaction, facilitating intent prediction and the generation of more accurate mental models [Breazeal et al., 2005, Dragan et al., 2015]. As a result, HRI research has started to explore nonverbal signaling for a wide range of robot states and task-related information.

1.2 Nonverbal Communication

Communication is the process of transmitting information from one agent to another [Mortensen, 2017]. The sender is the source of information and encodes the information into a form, such as a signal, that is transmitted to the other agent, the receiver, across a channel or medium [Shannon and Weaver, 1998]. The channel (i.e., signaling modality) is subject to noise or environmental disturbances (e.g., ambient sounds) which can interfere with the receiver's ability to decode the message and obtain the signal's contained information Figure 1.2. The sender can utilize the receiver's response or reaction as feedback about the message.



Figure 1.2: The robot encodes and sends information via a signal to a human receiver who decodes it and responds with feedback.

In human-human communication, the sender is a human who encodes information using available communication channels. Much of this process is done using nonverbal channels, such as facial expression,

216

body language, and gaze [Knapp et al., 2013]. Due to the prominence of nonverbal cues in human communication, humans are quite adept at consciously and subconsciously noticing, interpreting, and reacting to them [Knapp et al., 2013]. However, as these cues are somewhat abstract, they can still be misinterpreted if the receiver is unfamiliar with a specific cue, receives conflicting information, or is a noisy environment.

Nonverbal communication is a powerful tool that has been studied both in the context of human-human and human-robot interactions [Breazeal et al., 2005]. In this work, we define *nonverbal cues* as anything the receiver perceives as providing information. This includes both the *implicit*, unconscious behaviors that we can derive meaning from as well as *intentional signals* from the sender [Breazeal et al., 2005]. However, due to robots' goal-oriented nature and complexity of control, most robots do not naturally emit many of the cues humans do [Mutlu et al., 2009b]. Instead, much of their communication must take the form of planned nonverbal signals.

Many envisioned tasks require the robot to convey information about their internal state for human collaborators and bystanders [Knoblich et al., 2011, Mutlu et al., 2013, Breazeal et al., 2004]. Nonverbal cues have already been shown to improve humans' understanding of robots' behavior and improve trust [Kiesler, 2005, Knoblich et al., 2011]. While recent work in HRI represents an important first step towards enabling non-humanoid robots to utilize nonverbal signals, there are still many challenges that require further research and development.

1.3 Motivation and Goals

Human and robot behavior often differ, particularly in cases for which actions should take into account nearby humans [Alami et al., 2005, Kruse et al., 2013, Saez-Pons et al., 2014]. Although co-located humans are typically considered for safety, robots often fail to consider their presence from an interaction perspective. The resulting breakdowns have led to recent efforts to incorporate social norms into robot behavior in order to improve interactions with humans [Breazeal and Velasquez, 1999, Mead and Matarić, 2017, Montreuil et al., 2007, Nakauchi and Simmons, 2002, Shiomi et al., 2014].

Humans are highly expressive; we possess several modalities that continuously emit information, such as body posture and gaze [Mutlu et al., 2009b]. This natural expressiveness is combined with intentional, learned signals to make humans very effective communicators [Knapp et al., 2013].

To be complementary and effective interactors, robots must also exploit a broad range of signaling mechanisms to achieve a necessary degree of expressiveness [Bauer et al., 2008, Szafir et al., 2015]. While the most straightforward solution is for robots to employ similar nonverbal cues as humans (e.g., arm gesture, posture, facial expression), this is often impossible for non-humanoid robots due to physical limitations of their embodiment [Bethel and Murphy, 2008, Szafir et al., 2014].

Instead, non-humanoid robots can utilize alternative nonverbal signals on commonly available signaling modalities (e.g., light). Although more limited than human nonverbal behaviors, these signals can improve the readability, trustworthiness, and acceptance of robots and enable various levels of interaction [Takayama et al., 2011, Beer et al., 2011, Desai, 2012].

Increased expressiveness can also improve a robot's chances of success at its functional task [Fischer et al., 2014]. By utilizing nonverbal signals, the robot can elicit assistance while minimizing disruption and human annoyances, as it manages non-deterministic scenarios that arise from the dynamic and unpredictable nature of its environment and task [Cha and Matarić, 2016, Fischer et al., 2014, Saulnier et al., 2011]. In addition, research has shown the benefits of state-expressive nonverbal signals for robotic applications that typically do not involve direct human interaction, such as navigation and manipulation [Dragan et al., 2013, Kruse et al., 2012].

When deciding how to convey information about the robot's internal state, we must consider which communication modality best fits the situation [Rehbinder and Sanfridson, 2004, Weaver, 1953]. Although

218

1.4. Organization

HRI research often promotes speech as a natural and flexible form of communication, it is not well suited for many settings and applications. In loud environments, speech may be subject to environmental masking due to high ambient noise levels. Conversely, in sound-sensitive areas such as schools and hospitals, speech can be disruptive. Speech is also inefficient and overly verbose when transmitting small amounts of information (e.g., low power) [Cha et al., 2015, Kiesler, 2005].

Hence, when choosing signaling actions, the robot should utilize information about the environment and its interactors. For instance, in a dimly light environment, the robot can utilize a bright light to enhance visibility. Humans are experts at modulating their behavior to dynamically respond to the world state, such as increasing their speech volume in loud environments or utilizing several modalities for increased effectiveness.

The goal of this review is to survey past works to inform researchers in designing and utilizing nonverbal signaling behaviors for effective HRI. In addition to exploring signaling mechanisms for non-humanoid robots, this work also seeks to understand the challenges and important research areas of this field.

This survey also builds on previous surveys [Fong et al., 2003b, Yanco and Drury, 2004] to inform the young and rapidly growing field of HRI. The quickly evolving nature of HRI makes it challenging to apply past and ongoing work towards open research problems. Thus, our hope is that utilizing insights from this diverse set of past works across different fields can inspire new approaches for researchers working in this area.

1.4 Organization

In Section 2, we provide background from related fields, particularly in the cognitive processes that humans utilize during communication. The purpose of that section is to highlight past research that serves as the foundation of nonverbal communication in HRI. Although many alternative concepts exist, the ones reviewed are primarily utilized by the HRI literature surveyed in this work as well other works related to our proposed areas of signaling research.

Section 3 presents selected use cases of non-humanoid robots that are relevant both in research and commercially. For each use case, we identify several applications and discuss potential signaling needs in the context of these scenarios. This analysis enables identification of common signaling scenarios across platforms, environments, and applications, which can facilitate the development of signaling standards.

In Section 4, we identify three categories of internal robot state inspired by these use cases that are relevant to HRI. These categories enable us to present features of robot state that are often communicated together. While there are also alternative categorizations, our approach is intended to be broadly applicable to the wide range of non-humanoid robots discussed in this work.

Section 5 presents key signaling characteristics that must be considered when designing and utilizing nonverbal signals. These descriptors enable us to specify the intent of a signal in features relating to the signal's design and usage.

Section 6 surveys past work in signal generation. While nonverbal signaling is a relatively young field of research in HRI, similar signals have been utilized in other fields and applications. We draw inspiration from these signaling mechanisms, including those found in everyday life, and discuss how similar methods can be applied to signals for non-humanoid robots.

Finally, Section 7 presents considerations that must be addressed to enable successful nonverbal signal design and usage. These considerations are motivated by challenges encountered in prior work in HRI and human-computer interaction (HCI). We suggest three open areas of research that will help to address these issues and discuss approaches. We conclude by summarizing our key insights in Section 8.

2

Background

In this section, we provide background on key concepts from human communication, psychology, and cognitive science that are often utilized in HRI research, with a focus on processes relating to nonverbal communication and signaling. This list of concepts is neither exhaustive nor mutually exclusive. Instead, the purpose of this section is to provide the fundamental knowledge necessary to better understand the past works included in this survey and to provide a foundation for robot signaling research.

The presence of robots in human environments continues to grow each year. It is expected that in the future many different types of robots will operate around a diverse set of humans with vastly differing levels of robotics knowledge, experience, and training. Non-experts or those with limited experience with robots may have difficulty understanding robot control policies and recognizing or responding to robot internal states [Kiesler, 2005, Scassellati, 2002].

2.1 Attribution of Mental States

Research in psychology, cognitive science, and communication supports the notion that humans naturally attempt to ascribe mental states, such as beliefs and intents, to others [Frith and Frith, 2005, Scassellati, 2002]. This cognitive ability is called *theory of mind* and allows us to understand that other humans have knowledge, desires, and perspectives different from our own [Scassellati, 2002, Frith and Frith, 2005]. Without this metarepresentational ability, it can be difficult to understand others and their behavior, an important component of effective communication and collaboration [Scassellati, 2002, Bauer et al., 2008].

By attributing these cognitive states to others, we are able to develop *mental models* or abstract representations of the surrounding world, including other humans [Gentner and Stevens, 2014, Johnson-Laird, 1983]. These models enable humans to more accurately estimate others' knowledge and capabilities and to predict the effects of their actions on the world [Buckner and Carroll, 2007, Kiesler, 2005]. During communication, we utilize these models for establishing *common ground*, a set of mutual knowledge, beliefs, and assumptions that enables efficient communication and minimizes misunderstandings [Kiesler, 2005, Stalnaker, 2002].

The underlying processes for developing mental models of robots is a growing area of research in HRI, but the results of these works often vary due to the wide range of appearances, capabilities, and roles robots can have [Phillips et al., 2011, Powers and Kiesler, 2006, Walters et al., 2008]. Some researchers have suggested that in the absence of real-life experience humans default to models of human behavior [Beer et al., 2011, Kiesler, 2005]. Other works have shown humans are strongly influenced by cultural factors and media portrayals of robots [Bartneck, 2004, Bartneck et al., 2007, Haring et al., 2014, MacDorman et al., 2009]. These mixed results show that it is essential for researchers to better understand how robot characteristics and actions affect human perceptions in order to appropriately design robot behavior [Kiesler, 2005].

When seeking to understand these challenges, we can gain insight from the issues that arise when humans have difficulty with these processes [Scassellati, 2002]. Individuals who suffer from theory of mind impairments, such as in Autism Spectrum Disorders, have difficulty determining the beliefs, knowledge, desires, and intentions of others [Baron-Cohen et al., 1985, Mitchell, 1997]. Deficits in theory of mind have also been closely linked to the inability to engage in *perspective-taking*, the process in which an individual attempts to perceive and understand a situation from another individual's point of view [Michelon and Zacks, 2006].

2.2 Reasoning About the World

Perspective-taking is used by humans to better understand another agent's point of view. In *perceptual perspective-taking*, an agent attempts to imagine the sensory (e.g., visuo-spatial) perspective of another [Schober, 1993, Trafton et al., 2005a]. This can process be used to determine the extent of an agent's sensory range or locations of objects relative to their location. It has already been shown to improve communication by enabling references to objects in the environment in an efficient and non-ambiguous manner [Michelon and Zacks, 2006, Schober, 1993] and has been applied to HRI problems in the context of giving instructions, learning from demonstrations, and collaborating on physical tasks [Berlin et al., 2006, Breazeal et al., 2006, Trafton et al., 2005a,b].

Cognitive perspective-taking is a more complex process that involves simulating another agent's mental point-of-view by using information about the agent and the world to reason about the agent's possible cognitive states. This type of perspective-taking enables humans to understand others' thought process, anticipate their actions, and understand how our own actions will affect others' cognitive state [Batson et al., 1997, Bauer et al., 2008, Norman, 1983].

Similarly, *sensemaking* is a process in which humans attempt to understand unfamiliar or uncertain situations [Siino and Hinds, 2005]. Unlike perspective-taking, sensemaking does not focus on a particular person but instead looks at situations and events. Sensemaking often occurs when there is something out of the ordinary, such as a surprising event or discrepant cues from the environment, which causes us to actively think about the situation [Siino and Hinds, 2004]. In unfamiliar situations with robots, humans utilize cues from the robot and environment to try to comprehend the robot's behavior. This can be difficult when the robot's internal state is complex or there are inconsistencies between the robot's behaviors, embodiment, or both. The result of the sensemaking process is a mental representation or schema of the world around us [Klein et al., 2006].

2.3 Acquiring Awareness

Situation awareness (SA) can be viewed as one of the resulting knowledge states from cognitive processes that reason about the world. SA involves the ability to understand what is happening in the agent's vicinity [Endsley, 2016]. In contrast to sensemaking, building situation awareness is instantaneous, effortless, and has no active trigger. The most common model of SA utilized in HRI consists of three levels: perception of environmental elements and events, comprehension of their meaning, and projection of their status [Endsley, 1995].

SA provides a useful framework for grounding functional robot tasks since most tasks require at least the lowest level of SA, perception, to achieve success [Drury et al., 2003, Steinfeld et al., 2006]. The higher two levels of SA require the robot to take into account and comprehend environmental elements and events and thus, lend themselves to approaches in which the robot utilizes a model of the world when determining its actions (e.g., Markov Decision Process). In robotics, SA is often explored in the context of collaborative human-robot tasks as it required for high levels of coordination [ada, Drury et al., 2003, Gombolay et al., 2017, Riley et al., 2004, Scholtz et al., 2004].

2.4 Effects of Culture, Media, and Experience

One of the key processes that affects human behavior towards robots is the degree to which we anthropomorphize them [Beer et al., 2011, Duffy, 2003, Fussell et al., 2008, Phillips et al., 2011, Salem et al., 2014]. *Anthropomorphization*, the process of ascribing human characteristics to a nonhuman agent, has been of particular interest for HRI due to the wide range of human-like features robots can possess [Duffy, 2003, Epley et al., 2007]. Even simple systems, such as moving shapes, can be anthropomorphized through behavioral or visual cues [Aggarwal and McGill, 2007, Luczak et al., 2003, Mather and Murdoch, 1994]. Past research has explored how appearance, vocal signals, and other behavioral cues affect anthropomorphization as well as its effects on empathy, expectations, and other attitudes towards robots [Eyssel et al., 2012, Fussell et al., 2008, Riek et al., 2009, Sundar et al., 2016].

Imbuing robots with human-like characteristics, such as intelligence, motivations, intentions, and emotions, also creates certain expectations for a robot's behavior [Cha et al., 2015]. When inconsistencies cause expectations to go unmet, acceptance of the robot can be negatively impacted [Cha et al., 2015, Goetz et al., 2003]. Underestimation of the robot, however, can lead to inefficiency and under-utilization. Matching human expectations to the robot's actual capabilities is challenging as there are many factors, including anthropomorphization, that affect human perceptions of robots [Goetz et al., 2003]. In addition, past surveys on robot expectations have revealed a large range of opinions on robot appearance, capability, and roles that are often tied to characteristics of the human (e.g., age, gender) [Oestreicher and Eklundh, 2006].

Other important influences on human perception are culture and prior experiences with robots [Bartneck, 2004, Bartneck et al., 2007, Haring et al., 2014, MacDorman et al., 2009]. While robots have existed for many years, different countries have different levels of exposure to robots both in life and in the media. Researchers have found that cultural differences can strongly affect expectations as well as concerns about robots [Bartneck, 2004]. For instance, in the US, robots taking over the world or stealing humans' jobs are often cited as key concerns with future robot adoption [Bartneck, 2004, Bartneck et al., 2007]. In contrast, there is more concern in Japan with the social impacts of robots, particularly the emotional aspects of interaction [Bartneck, 2004, Bartneck et al., 2007]. These differences are theorized to be due to a number of factors, including cultural differences, media portrayals of robots, and the presence of robots in everyday life [Bartneck, 2004, Bartneck et al., 2007, Beer et al., 2011].

Since these concepts all contribute to the way humans perceive and behave towards robots, they are important to consider when designing algorithms for interaction. Moreover, they can also impact the acceptance and adoption of robots as exposure to robots in everyday life increases [Beer et al., 2011]. Hence, while our primary focus in this work is on exploring methods for generating nonverbal signals that facilitate successful communication and collaboration, we also strive to consider how these behaviors affect the larger scope of the robot's operations.

2.5 Summary

In this section, we reviewed key concepts from human communication, psychology, and cognitive science that are utilized in HRI. We also surveyed past work exploring the effects of experience, including robot depictions from culture and media, on humans' attitudes and behavior towards robots. These past works can inform new signal design approaches and enable anticipation of human behavior in response to a robot's designed signals. In the next section, we present a set of nonhumanoid robot use cases and discuss their signaling needs.

3

Use Cases

To ground this review in real world applications, in this section, we present a set of selected use cases of commercial and research nonhumanoid robots. We focus on robots that perform functional (i.e., nonsocial) tasks and whose appearance contains very few if any humanoid features. This is not an exhaustive list of use cases but rather it serves to illustrate the range of concerns that must be taken into account when designing signaling behaviors. We limit the use cases we discuss to those in which a robot interacts with co-located humans. We take inspiration from these use cases to identify common signaling needs for a large variety of non-humanoid robots.

3.1 Autonomous Vehicles

Autonomous or self-driving vehicles are capable of navigating the environment with no human input. This technology constitutes a rapidly growing area of robotics research, development, and commercialization. Some examples of autonomous vehicles and their interaction challenges are shown in Figure 3.1.



Figure 3.1: Autonomous vehicles, including those developed by Nissan, Uber, and Waymo (left to right), will have to interact with pedestrians, cyclists, and human drivers, creating the need for expressive signaling behaviors.

Current vehicles range in autonomous capabilities, with much of the research in this area focusing on the technical challenges associated with removing human input and improving safety, robustness, and reliability across different driving environments [Luettel et al., 2012, Rödel et al., 2014, Urmson et al., 2008]. While autonomous vehicles will become more common in the near future, they will continue to operate alongside human-driven vehicles for many years to come [Dresner and Stone, 2007, Richtel and Dougherty, 2015]. Therefore, signaling by selfdriving vehicles for the purposes of communication with human drivers, pedestrians, and bicyclists is an important use case to consider [Anaya et al., 2014, Clamann et al., 2017, Mirnig et al., 2017, Richtel and Dougherty, 2015].

Research in self-driving vehicles has explored three applications: passenger vehicles, commercial transportation, and agricultural and field vehicles. While all three types of vehicles share the general goal of safe and efficient autonomous navigation, it is also important to consider their varying usage and forms. For instance, agricultural and field vehicles typically do not need to operate within traffic rules, and commercial transport vehicles have different sizes and shapes than passenger vehicles, impacting their behavior on the road. Since field vehicles have fewer interactions with humans, we primarily focus on the communication needs of passenger and commercial vehicles.

One of the primary challenges for autonomous vehicles is coordinating actions with other drivers and pedestrians [Sadigh et al., 2016]. Failures can range in consequence from inefficiency to physical harm of the system, other vehicles, or humans. Current research utilizes models that, above all, aim to safely and reliably avoid humans and other vehicles. However, this can result in actions that are overly functional and lack expressiveness [Koo et al., 2015, Sadigh et al., 2016]. To coordinate with humans, self-driving vehicles, like any collaborative robot, must express information about their actions and intent to ensure safety and facilitate trust [Luettel et al., 2012, Mirnig et al., 2017, Sadigh et al., 2016].

Human drivers effectively glean information not only from explicit signals such as turn signals and brake lights but also from subtle cues found in the car's motion (e.g., small speed changes, slight swerving) [Broadhurst et al., 2004, Michon, 1985, Shimosaka et al., 2014]. When a vehicle's motion is erratic and difficult to predict, other drivers cannot coordinate their actions and become frustrated [Broadhurst et al., 2004, Keskinen et al., 1998]. Since the vehicle's primary functionality is motion, information about its motion intent is an important aspect of effective communication and a significant concern for real world implementation.

Signaling is also important when established rules are insufficient for resolving ambiguities, such as at intersections unregulated by stop lights [Haddington, 2010]. Since there is no set of standardized signals specifically designed for these situations, humans must find creative ways to use the signaling modalities available to them. For instance, drivers often use explicit signals such as hand gestures or flashing the car's lights to clarify their intent or signal to pedestrians. Similar explicit signals are used when vehicles have conflicting goals (e.g., lane changes into the same lane), need to go outside of the standard rules (e.g, larger vehicle must back up to make a tight turn), or operate outside of regulated areas (e.g., fields or non standard roads).

Since autonomous vehicles are more complex in operation than human-driven vehicles, they may also require a greater range of signals than standard signaling mechanisms, such as brake lights, can support. This issue is also complicated by most humans' lack of experience with autonomous vehicles which makes it challenging for them to predict how such robots will act in different scenarios. One potential solution is to gradually introduce new signals and modalities to human observers. This may be possible as each year more vehicles are incorporating autonomous features (e.g., assisted lane following, parallel parking) while still relying on human drivers.

Discrepancies in the behavior of autonomous and human-driven vehicles will likely cause humans to have different expectations of these vehicles. Since these two types of vehicles may appear identical to one another, this can cause difficulty recognizing whether the vehicle is operating autonomously or a human is providing control input. This is further complicated by the varying levels of autonomy the robot can operate under as well as the ability to return control to a human driver when necessary. As the current expectation is that vehicles are human-driven, effectively communicating the robot's control mode and capabilities is a key signaling requirement for appropriate expectation setting.

3.2 Flying Robots

Flying robots are aerial vehicles that operate without humans on board and can fly autonomously or be piloted remotely [Szafir, 2015, Valavanis and Vachtsevanos, 2014]. These robots include quadrotors, drones, and unmanned aerial vehicles (UAVs) and are being utilized for a wide variety of applications including security, surveillance, search and rescue, mapping, agriculture, entertainment, and item transportation [Szafir, 2015, Szafir et al., 2017, Valavanis and Vachtsevanos, 2014]. Some examples of flying robots and their applications are shown in Figure 3.2.

While these robots have not traditionally been flown in close proximity to people, the advancements in small, inexpensive systems is expanding free-flyer usage into human environments for both commercial and recreational uses [Szafir, 2015, Valavanis and Vachtsevanos, 2014]. This often requires robots to operate in close proximity of humans, making communication a priority for safety. Free-flying robots are being designed for a number of applications that may require direct human interaction, such as search-and-rescue and item delivery. In

3.2. Flying Robots



Figure 3.2: Flying robots can be utilized for a variety of human-oriented applications in a diverse set of environments. Operating conditions as well as noise created by the robot presents several signaling challenges. (Credit: RTS Labs, DARPA, Amazon, NASA)

these scenarios, the robot may need communication information, such as indicating it is going to hand over an item, to complete its intended task.

Mobility is the main feature and safety concern of free-flying robots; consequently, many of their signaling needs are similar to autonomous vehicles. However, several factors unique to free-flying robots affect their communication and signaling needs. Since airspace is significantly less occupied than ground surfaces where most autonomous vehicles operate, free-flying robots have fewer restrictions and conventions for their movement. Their 3D movements are less familiar to observers and often have fewer constraints (e.g., not restricted to a road network), making them more difficult to predict. Moreover, interactions with these robots are less expected as there are significantly fewer free-flying robots in use than ground vehicles and they tend to operate for shorter, irregular periods of time. Therefore, these robots may need to signal information about their location to make humans aware of their presence and able to anticipate further actions.

Currently, free-flying robot operations create significant noise and disturbance which indicate their presence in the nearby environment. These cues are limited in the information they provide about the robot's current operations and can also be masked by ambient noise. Since a goal of future free-flyer design is reduced noise for greater usability in human environments, a better long-term solution is to utilize signals designed to communicate specific information about the robot's state. Deliberate signaling can also take into account the interaction scenario, which can improve communication efficiency with humans.

Explicit task-oriented signals have already been utilized to provide information about free-flying robot motion intent [Szafir et al., 2015]. Flying robots employed for search-and-rescue and other dangerous tasks may need to indicate information about the environment for human safety. Since these robots are also envisioned to be utilized as companions or helpers, there has also been research exploring affective communication. Other recreational application, such as photography or videography, may require communication about data recording (e.g., taking a picture or collecting video).

This application highlights another major concern with flying robots: their historical use for surveillance and security [Yao et al., 2017]. Humans may have concerns with trust in the presence of such robots, especially around more private areas including the home. Since it is difficult to discern many of the robot's capabilities, transparency of these robots' operations (e.g., the data being gathered or presence of a human observer) may facilitate acceptance into everyday human environments. It can also help to set appropriate expectations since free-flying robots can operate with varying levels of autonomy.

As current free-flying robots do not carry humans on board, human input is typically provided remotely. However, in cases where the robot requires assistance, bystanders and non-expert operators are often the most conveniently located. Hence, such robots need to provide information about their error state for both safety and eliciting assistance. Unlike interaction with ground vehicles, interaction between humans and flying robots is a relatively new area of research. Consequently, there are fewer conventions, and the expectations from human interactors are less predictable. There are also significant safety risks involved with operating free-flying robots around humans or other aircraft, especially as the number of inexperienced human users, such as hobbyists, grow. Effective communication can help reduce safety concerns but the increasing presence of these robots in everyday life necessitates greater research into how free-flying robots and humans can and will interact.

3.3 Industrial Robots

Industrial robots, as shown in Figure 3.3, are mobile robots and manipulators that operate in industrial settings, such as factories and warehouses, to increase productivity and augment human labor [Cesetti et al., 2010, Heyer, 2010, Hvilshøj et al., 2009]. Many of these robots function as transports, moving materials and products, whereas others act as specialized tools [Cesetti et al., 2010, D'Andrea, 2012, Datta et al., 2008, Liang et al., 2015]. Like flying robots, these robots typically do not have any humans on board but are designed to work alongside them. Without a payload, industrial mobile robots tend to be significantly shorter in height than adult humans and have a mechanical appearance, designed with the robot's functionality in mind [Wurman et al., 2008]. Mobile manipulators combine a mobile base with one or more robotic arms for manipulating objects (e.g., moving items or assisting workers) [Datta et al., 2008, Hvilshøj et al., 2012].

Since mobile robots work in tandem with humans, they must function in both direct and indirect interactions [Heyer, 2010]. As payloads can potentially be quite heavy or large, safety is a primary concern [Hägele et al., 2002, Vasic and Billard, 2013]. Signaling needs of these robots largely relate to their movement and physical coordination with co-located humans, which requires them to express information about their presence, location, and motion intent [Chadalavada et al., 2015, Chadalavada, 2016]. Since they function in private, indus-



Figure 3.3: Industrial mobile bases and manipulators, such as the Adept Lynx, Kuka KMR iiwa, Amazon Robostow, and the ABB Yuma (clockwise), are being increasingly utilized to provide support to human workers to improve efficiency.

trial environments, these robots primarily interact with human workers that have a related goal and greater knowledge of the environment. Co-located humans may also be expected to collaborate, provide assistance, and resolve issues the robots encounter, which require information about the robot's errors, faults, needs, and tasks [Tang et al., 2014].

The environments these robots work in present unique design considerations; factory environments are typically noisy and highly dynamic, with many machines and humans moving about. Robot size and shape can also create visibility issues; short robots can be difficult to see, while large ones can obscure important visual cues. Unlike the previous use cases, however, industrial mobile robots are expected to interact with known users, so training and repeated exposure can be



Figure 3.4: Service robots, such as the Cobalt, Neato, Savioke Relay, and Lowesbot (clockwise), are being utilized for many applications in human environments, including security, cleaning, hospitality, and retail. These tasks require the robot to interact with and co-exist alongside humans.

utilized to facilitate signal design. With increasing numbers of these robots being integrated into factories and warehouses, their ability to coordinate, and hence their communication capabilities with humans, will become more important.

3.4 Service Robots

Service robots, as shown in Figure 3.4, perform service-oriented roles in non-industrial settings such as restaurants [Acosta et al., 2006], shopping malls [Datta et al., 2011, Kanda et al., 2009, Shiomi et al., 2009], hotels [Osawa et al., 2017, Zalama et al., 2014], healthcare centers [Ahn et al., 2014], and homes [Fong et al., 2003a, Forlizzi and DiSalvo, 2006, Scopelliti et al., 2004]. They are used for tasks such as greeting customers, providing information, taking orders, guiding people, cleaning,

maintaining security, and delivering items [Glas et al., 2013, Fong et al., 2003a, Forlizzi and DiSalvo, 2006].

Robots in the home are a unique case in that they typically interact with the same set of humans but may perform similar tasks to other service robots [Forlizzi and DiSalvo, 2006, Scopelliti et al., 2004]. Unlike industrial settings, service robot environments are dynamic, less structured, and often contain unfamiliar users. Many of their intended tasks also require the robots to directly interact with humans in close proximity.

Consequently, a major challenge for service robots is conforming to social conventions [Nass and Moon, 2000]. Past work has found that violating social norms creates a negative experience for human interactors and, in some cases, causes them to reject such robots [Bartneck and Forlizzi, 2004, Lee et al., 2010, Sardar et al., 2012]. Since many service robots operate in dense, crowded environments, signals that convey information about motion are important for avoiding collisions, especially for humans unfamiliar with the robots. However, these signals must also avoid social conflicts by adhering to the social expectations of a scenario, including turn-taking conventions (analogous to but often more complex than intersections) and apologies in cases of errors [Lee et al., 2010]. These signals also increase humans' perceptions of robots as polite and service-oriented, improve user experiences, and promote acceptance [Lee et al., 2010].

3.5 Telepresence Robots

Telepresence robots, as shown in Figure 3.5, combine the communication capabilities of video conferencing with the advantages of mobility to enable human operators to freely move about a remote environments [Kristoffersson et al., 2013, Tsui and Yanco, 2013]. These robots represent the remote embodiment of the human and facilitate communication with humans co-located with the robot [Tsui and Yanco, 2013]. Telepresence robots are typically human height and contain a screen that continuously displays the operator's face [Kristoffersson et al., 2013, Tsui and Yanco, 2013]. Unlike the prior use cases, these



Figure 3.5: Telepresence robots, such as the RP-Vita, Vgo, AVA 500, and Ohmni (clockwise), are enabling robot-mediated communication in remote environments for human operators in settings such as hospitals, schools, offices, and homes.

robots are mostly teleoperated, with some autonomous features [Coltin et al., 2011, Kristoffersson et al., 2013, Tsui and Yanco, 2013], and designed to be perceived and treated as the human operator [Johnson et al., 2015, Kristoffersson et al., 2013, Rae et al., 2014, Tsui and Yanco, 2013].

Since the main purpose of these systems is communication, telepresence robots have been used in a variety of settings including offices, hospitals, classrooms, and homes [Beer and Takayama, 2011, Desai et al., 2011, Kristoffersson et al., 2013, Tsui and Yanco, 2013]. The wide range of environments they operate in contain diverse sets of humans that may interact with the robots. In homes or schools, these robots typically encounter fixed sets of interactors that will become more accustomed to the presence of the robot. In offices, hospitals, and larger service settings, the robot may be utilized by several humans and may hence, interact infrequently with a more diverse set of humans in the remote environment. Potential interactors' range of experience and familiarity must be taken into account when designing signals for these robots.

Since telepresence robots move about the environment, they share many of the signaling requirements of mobile bases and service robots. However, the continuous presence of the human operator alters perceptions and expectations of the robot's behavior [Lee and Takayama, 2011, Newhart, 2014, Rae et al., 2015]. Humans in the remote environment may expect the robot to behave in accordance with social norms and be less tolerant of socially inappropriate behavior [Lee and Takayama, 2011]. As it can be difficult for the operator to understand when their remote actions violate social norms (e.g., volume in the remote environment), this can lead to decreased acceptance of telepresence robots. Therefore, greater consideration must be given to how to support the difference in human attitudes towards these robots when designing communicative signaling behaviors [Tsui and Yanco, 2013].

Although telepresence robots are designed to facilitate communication, they lack many of the nonverbal modalities humans typically use for communicating, such as head motion and gesture [Kristoffersson et al., 2013, Tsui and Yanco, 2013]. While the human face is expressive, recognizing subtle facial cues on the screen can be difficult due to visibility issues [Takayama and Harris, 2013, Tsui and Yanco, 2013]. Therefore, additional signals are typically needed to augment telepresence operator's audio and video, but they must not interfere with the operator's experience or contradict the operator's intent. Since these signals also affect humans' perceptions of and behavior towards telepresence robots, operator considerations are a key part in designing telepresence robot signaling behaviors.

3.6 Summary

In this section, we reviewed selected use cases of non-humanoid robots and discussed their signaling needs. While the robots in these scenarios have different capabilities and operate in a variety of environments, there are still common signaling needs across platforms. Hence, designing generalized signals for certain robot states can facilitate standardization and reduce the workload for robot designers. In the next section, we draw from these use cases to identify relevant robot internal states and discuss how they can be communicated to facilitate human-robot interaction.

4

Internal Robot States

In this section, we outline internal robot states relevant to the nonhumanoid robot use cases discussed in the previous section and provide support from literature. We separate these states into three categories: condition, knowledge, and action. These states represent a significant portion of the information about the robot that is relevant to co-located humans. There are many alternative categorizations of this information; our approach simply aims to group information that is related and often signaled together.

4.1 Condition

A robot's *condition* encompasses a wide set of characteristics relating to its general operation, including its health, control mode, and faults. For safety and successful coexistence, it is vital for a robot co-located with humans to effectively communicate its condition to enable both collaborators and bystanders to assess whether the robot is safely operating or whether it can benefit from humans intervention [Breazeal et al., 2005, Fischer et al., 2014, Takayama et al., 2011]. Human assistance can also be used to help robots overcome limitations in sensing



Figure 4.1: Diagram depicting the relationship between internal state features related to the robot's condition, as discussed in Section 4.1.

and functionality, enabling faster deployment into dynamic human environments [Rosenthal and Veloso, 2012]. The various features of the robot's internal state related to its condition are depicted in Figure 4.1.

4.1.1 Operational

The operational state indicates the level at which the robot is capable of perception, cognition, and action in relation to its optimal state. Without all three functionalities, the robot may lose the capability to perform its intended tasks. When the robot is turned off, it lacks these capabilities and is not a concern for co-located humans. If the robot is capable of movement (i.e., action), humans should be aware so that they can coordinate their own movements, prevent collisions, and avoid being startled if the robot suddenly starts to move [Takayama et al., 2011]. Humans may also want to know if the robot is actively perceiving the nearby environment, as this impacts their privacy [Lee et al., 2011, Lin et al., 2011, Syrdal et al., 2007].

241

Being fully operational also implies that the robot is functioning properly and does not require a human's attention or assistance. Humans utilize a variety of nonverbal cues to convey this information to one another, such as nonverbal sounds, body language, microexpressions, and eye gaze [Ekman, 2007, Mehrabian, 1972, Morency, 2010]. Subtle, non-functional motion, such as small head, arm, and eye movements, have been explored for the same purpose on humanoid robots and virtual agents [Lamarre, 2002]. Communicating a robot's operational state is more challenging for non-humanoid robots, however, as they have fewer degrees of freedom and are often unable to replicate such subtle motion cues. Moreover, utilizing available motion channels, such as base motion, for signaling can interfere with the robot's task and cause additional strain on its hardware [Zeglin et al., 2014].

An alternative solution is to utilize auditory signals similar to the cues naturally produced by the operation of machinery or biological functions of humans (e.g., breathing) [Knapp et al., 2013]. However, while these sounds indicate the robot is "alive" and operating, they do not provide information about its level of functionality. Conveying various levels of functionality using sounds requires humans to discern between auditory signals which are subtly different or may not intuitively map to the robot's state. Continuous auditory signals may also be disruptive in certain environments or masked in those with high ambient noise levels (e.g., factory or loud restaurant).

Mechanical systems and appliances often utilize continuous indicator lights to signal levels of functionality [Harrison et al., 2012]. The simplest signal is a binary light-emitting diode (LED) that indicates whether a device is turned on (e.g., computer power or power strip). Systems can utilize multiple lights or modulations in color, frequency, intensity, and timing to provide greater detail about their functionality and state [Harrison et al., 2012]. Due to their high visibility and common usage in everyday life, light indicators are often utilized on robots as well. As the number of features (e.g., lights, colors, patterns) grows, however, humans have to spend increasing amounts of time to interpret the signals. Hence, as the complexity of content increases, systems tend

242

to utilize modalities that leverage text (e.g., visual display or speech) or several modalities to increase expressiveness.

4.1.2 Power

Closely related to the operational state is the robot's *power* level. As power is necessary for the robot to function and perform tasks, the robot should warn humans when its available power will no longer be sufficient to continue operating. This is especially important when the robot is engaged in a task as it negatively impacts task efficiency and can endanger itself or co-located humans. Similar to the operational state, many systems utilize a light indicator to provide information about power levels [Harrison et al., 2012].

Devices that contain a display often utilize alerts, images, or text to convey greater detail about the power level, such as its remaining percentage or time until it runs out [Bloom et al., 2004, Rahmati et al., 2007]. In most cases, humans do not need to concern themselves with the robot's power until it is close to running out [Rahmati et al., 2007]. Since humans are not always in close proximity and visually focused on displays or indicators, devices such as cell phones and smoke detectors also utilize an auditory signal to draw attention when the battery level falls below a certain threshold [Rahmati et al., 2007, Smith, 1993]. Auditory signals increase the salience of the warning and can be used sparingly to minimize disruption and annoyance [Edworthy et al., 1991].

4.1.3 Operating Mode

Many non-humanoid robots can operate in different modes with varying levels of autonomy, human monitoring, and constrained capabilities. Autonomy can be defined as the level of human input provided to the robot; it can range from manual control, where the robot depends on a human operator for continuous input, to full autonomy, where the robot acts independently. The level of human monitoring is also relevant as many robots are designed to allow a human to monitor its sensors and intervene when necessary. Some robots also possess distinct functional
modes that constrain the robot's capabilities [Bualat et al., 2015, Stentz et al., 2015]

Transparency about the presence and role of a human-in-the-loop promotes more accurate mental models and greater trust. This is especially critical for robots expected to regularly operate with varying levels of autonomy (e.g., autonomous cars). As robots are still relatively rare in everyday life, it may also help to build trust as people may feel initially more comfortable knowing a human is monitoring the robot's behavior and can take control when necessary. Human interactors may also have different levels of trust and attitudes depending on the robot's level of autonomy.

Knowledge about the robot's operations is also a critical component of privacy as personal data may be collected or exposed by the robot, creating significant concern for sensitive or personal environments (e.g., homes, hospitals) [Lee et al., 2011, Rueben et al., 2017, Sharkey and Sharkey, 2012, Syrdal et al., 2007]. With concerns about cybersecurity and electronic surveillance, robots are another potential device where privacy may be unintentionally or purposefully violated. In privacysensitive settings, in particular, there are also ethical and legal issues with the ability of other humans to access any collected data, either in real time or at a later date.

Hence, humans may feel more at ease in the presence of a robot if they have greater understanding of the data being collected, who will have access to the data, and how the data will be used. While much of this information should be given to humans before engaging with the robot, some information such as intermittent observation of the robot's sensors by a human observer or operator can be indicated during the interaction with a clear warning signal.

4.1.4 Physical Condition

The *physical condition* of the robot relates to the state of the robot's hardware and its degree of functionality. This state includes whether motors, sensors, and other components of the robot's physical form are operational and any potential damage or issues. Information about the

244

4.1. Condition

robot's physical condition helps humans to predict if the robot is likely to malfunction and whether it is well suited towards a specific task.

Conveying information about the robot's physical condition presents several challenges. A human observer may not be able to easily perceive and distinguish among the multitude of signals needed to cover all aspects of a robot's hardware. Unfamiliar users may also have difficulty understanding how to use this information. Hence, the human's knowledge and experience should be taken into account as well when designing such signals.

This consideration also affects the modality as detailed, complex information is best conveyed using more flexible communication channels such as visual displays or speech. Humans with greater familiarity with the robot may be able to interpret simpler cues or utilize prior information to greater effect. For instance, if the human sees the robot physically collide with another object, they may anticipate that there is physical damage on the robot.

4.1.5 Faults

Faults are deviations the robot experiences relative to expectations of its functionality and can constitute an undesirable condition or error. These off-nominal conditions can come about through issues with the robot's software or hardware. Physical faults describe deviations that relate to the robot's physical embodiment or hardware (e.g., chassis, arm, wheel, connector, sensor, computer), such as a loose connector. Software or cognitive faults describe issues that arise in the robot's software such as planning failures, wireless communication connectivity problems, or computational errors. Unlike physical faults, these issues do not usually require physical actions to resolve.

Deviations do not always cause faults as those within a small threshold are considered normal due to uncertainty in the robot's hardware (e.g., sensors). Faults of a greater magnitude, especially those that prevent the robot's operations, are often called errors. The level of information provided for different types of faults varies by several factors, such as whether it poses safety risks or if assistance is required from a human, with faults that do not impede task completion or safety not always communicated. The robot can also consider any human interactor's availability or interruptibility when deciding if, when, and how to present faults.

Typically, faults that impact safety require the most salient communication signals. In dangerous situations, the signal may simply aim to elicit a fast reaction by getting nearby humans' attention and alerting them of an error without providing specific details. For less time-critical faults, the signal may convey greater levels of detail about the issue including instructions on how to resolve the fault. Verbal communication, either spoken by the robot or displayed on a screen, is often the best method for conveying large amounts of information, including providing background and context for the error. However, the optimal amount of information depends on the expertise of the human [Wogalter et al., 2002].

Error or warning signals are used across many types of machinery, appliances, and computing devices. In highly regulated applications, such as aircraft control, these signals typically have distinct, learned meetings that humans are trained to recognize and respond to [Veit-engruber, 1978]. In contrast, products found in the home (e.g., microwaves, washing machines) often present only a small set of faults directly to human users due to the assumption that they have limited expertise in resolving more complex issues.

When operating in dynamic, unstructured environments, robots are expected to encounter many types of faults and hence, need to convey a wider range of information than many other machines and devices. Humans exploit a wide range of explicit and implicit cues to express that they have encountered an unexpected situation. These cues can range in ambiguity, from affective cues (e.g., body language) to definite learned signals (e.g., open palm stop signal). Most non-humanoid robots lack the capability to replicate a large number of these highly expressive but subtle signals. In the absence of these signals, issues that can easily be overcome with human intervention will negatively impact the robot and/or co-located humans.

For example, when power levels fall below a certain percentage or set amount, a warning or error can prevent the robot's loss of power

4.2. Knowledge

in the middle of a task. If the robot is operating in a scenario where it can endanger itself or co-located humans, conveying relevant faults is especially important. Information about faults can also help co-located humans to resolve issues that are causing the robot to fail at its task, lowering the technical barriers for adoption in complex human-centric domains. These considerations make communication of faults and errors a high priority for enabling HRI in the real world.

4.1.6 Affect

Affect is the emotional state of the robot. In this work, we define three levels of affect. The highest level, true affect, involves the robot actually experiencing emotion, which is currently beyond the state of the art in robotics and artificial intelligence. The middle level uses affect to convey information about the robot's state in a human-accessible way. For instance, if a robot cannot physically complete a task, it can display sadness or frustration to indicate its failure. The lowest level uses affective signals to appropriately react and respond to a human but does not actually reflect any information about the robot's state. These lower levels utilize affect as a communication tool to make the robot a more effective interactor and are hence, not real affective states.

This review is focused on non-humanoid robots that perform functional tasks, which have had limited usage of affect to date. However, this is likely to change in the future and merits careful study but is currently outside the scope of this survey.

4.2 Knowledge

The robot's *knowledge* encompasses the information it possesses about itself, the task, and the surrounding world. Knowledge can be preprogrammed, acquired, or derived, and ranges broadly in type, from sensor data and maps of the environment to user preferences and past behavior patterns [Adams et al., 2014, Nikolaidis and Shah, 2013, Salas et al., 1995, Sadigh et al., 2016]. Knowledge estimation is a key component of constructing the representations or models of humans used during interactions and enables grounding, an essential part of successful and efficient communication [Eilan, 2005, Endsley, 1988, Kiesler, 2005, Norman, 1983]. The various features of the robot's internal state related to its knowledge are depicted in Figure 4.2.



Figure 4.2: The relationship between internal state features related to the robot's knowledge, as discussed in Section 4.2.

4.2.1 Internal

Internal knowledge is information that the robot possesses about itself, including its form, capabilities, and functions. This information enables the robot to plan its physical actions and determine whether or not it can perform a task. A common form of internal knowledge is a model of the robot's physical form, including its sensors and actuators. Typically, robots are preprogrammed with this type of information as it is required for basic functionalities, including movement and perception.

Understanding a robot's capabilities is becoming increasingly important as robots become more pervasive in everyday life [Cha et al., 2015]. While robots are a novel and diverse technology, preconceptions

248

4.2. Knowledge

about their capabilities can be highly inaccurate as they are influenced by portrayals of robots in fiction, entertainment, and the media [Bartneck, 2004, Sundar et al., 2016, Šabanović, 2014].

Since encoding robot capabilities can vary in abstraction, from complex skills to physical joint limitations, expressing this knowledge to humans is challenging, especially for capabilities that are affected by external conditions. For instance, a robot may possess the functionality to navigate around obstacles but can fail to find a successful path if the environment is too cluttered or there is interference with its sensors. Experience can facilitate interpretation of the robot's signals by providing humans with important background or repeated exposure.

4.2.2 External

External knowledge is information that the robot possesses about the world; it includes information that is preprogrammed, acquired through sensing, or derived from other information the robot possesses. Preprogrammed information often includes models of the world, objects, and people the robot may encounter and interact with. Perceived information includes any information the robot obtains through its sensors, such as video data, depth maps, or audio. The robot can integrate this information with its other knowledge to derive new information about itself and the world, such as its location in a world map. Deriving new information is essential as robots often operate in dynamic settings.

External information also enables situation awareness (SA) which is needed for the robot to interact with the world [Drury et al., 2003, Endsley, 1995]. Human's can utilize information regarding the lowest level of SA, perception, to anticipate the robot's capability and performance at tasks. As robots possess different sensory capabilities than humans, conveying information about their perceptual knowledge, as humans do with gaze, also enables humans to provide valuable feedback, including assistance [Bugnyar et al., 2004].

The next level of SA, comprehension, requires the robot to merge information it possesses to form a greater understanding of the world. Comprehension can range in complexity from inferring one's location (e.g., kitchen) using objects in the environment (e.g., food items, stove, and a refrigerator) to deriving complex relationships between objects and inhabitants in the world. If the robot is assumed to have a higher level comprehension of its environment than it possesses, humans may become frustrated when the robot makes mistakes based on its lack of knowledge. Conversely, if humans assume the robot has no higher level understanding of the environment, the robot may be underutilized. Humans typically assume that there is a base level of comprehension that other humans possess. For assessing others' knowledge of more complex relationships, humans typically rely on multiple cues from prior interactions.

The highest level of SA, projection, takes the information derived through perception and comprehension and predicts the effects of future actions on the world. Robots often undergo this process when planning actions that physically affect the environment. For instance, manipulation algorithms can utilize a physics model with object properties, such as weight and shape, to understand how a specific manipulation action will effect the object. This type of projection is subject to set physical rules and is, thus, easier for robots to perform. More complex projections including those involving humans are less deterministic due to the variability of human reactions. As with comprehension, humans typically utilize experiences gathered over time when determining this type of knowledge.

Robot capabilities vary based on many factors including the robot's price, manufacturer, environment, and application. For instance, many applications do not require the robot to have certain types of external knowledge to successfully perform its task. Robot algorithms developed by different manufacturers often have varying degrees of success affecting the robot's optimal capability. This is also affected by the cost of the robot, as more expensive platforms will have better sensors and more precise actuators.

Expectations of a robot's knowledge can also be greatly affected by its appearance and physical form; humanoid robots may be compared to humans in capability and intelligence. Non-humanoid robots are more unfamiliar and thus, may elicit a wider range of expectations and mental models. Hence, it is important for robots to convey their knowledge in order for humans to have an accurate understanding and expectations of robots behavior.

4.3 Activity

The *activity* state of the robot encompasses information relating to actions the robot is attempting to execute, typically within the scope of a task the robot is engaged in. Since the focus of this work is on functional, task-oriented robots, actions are assumed to contribute towards the completion of a goal-oriented task. Information about the robot's goal-oriented task provides context for the robot's behavior and enables humans to anticipate the robot's future actions. These are key components of collaboration and allow humans to coordinate their own actions with those of the robot in both time and space. The various features of the robot's internal state related to activity are depicted in Figure 4.3.



Figure 4.3: The relationship between internal state features related to the robot's activity, as discussed in Section 4.3.

4.3.1 Availability

The robot's *availability* relates to whether it is able to engage in a task. Understanding the robot's availability is particularly important in scenarios where humans may want to assign the robot a new task, such as in service or human-robot collaboration contexts. Availability is a complex state that can be difficult to understand. Humans are capable of doing multiple tasks simultaneously. Robots are capable of either engaging in a sequence of tasks or limited to accepting one task at a time. This distinction is important for applications in which robots have many responsibilities and humans must consider both the robots' current and future availability.

A robot may also be incapable of accepting a new task for reasons other than workload; for instance, the task can be estimated to require more time than the robot's power levels can support, meaning the robot will be unable to complete the task. Although information about the robot's condition can be directly communicated to a human (e.g., battery percentage), they may have difficulty applying these constraints to the robot's ability to complete a task (e.g., how much battery is required for a task). While the robot may be unavailable for tasks that violate a current constraint, it is still able to accept and complete other tasks.

Humans are experts at conveying and assessing availability and using a range of explicit and implicit signals. Since humans are not purely task-oriented, functional agents, this information can be conveyed using many types of nonverbal cues (e.g., looking around, yawning, slouching). For robots, signaling different types and degrees of availability is challenging as communication modalities can vary widely across platforms.

4.3.2 Task

When the robot is engaged in a *task*, it is taking actions to achieve a desired goal. We define a task as a plan, or a set of a primitive actions, with a well defined objective. This also encompasses information related to the execution of a task, primarily, the robot's actions. Information about the robot's actions enables transparency, coordination, and better anticipation of the robot's future behavior.

Goal

The robot's goal is the desired outcome of its action(s) in the context of a task. As most robots behave in a goal-oriented manner, knowledge of their goal enables prediction of their future actions [Dragan et al., 2015] and facilitates coordination and collaboration [Bauer et al., 2008]. In robotics, the terms goal and intent are often used interchangeably to denote a desired result [Dragan et al., 2013]. However, in this work we distinguish intent as the desired immediate result of a single action, while a goal relates to the overall task [Fishbein and Ajzen, 1975]. In simple tasks, goal and intent can become equivalent.

Work in HRI has explored methods for conveying a robot's goal to nearby human observers [Dragan et al., 2013] in order to provide transparency and enable humans to plan complementary actions for collaboration, coordinate their actions to improve safety, prevent collisions, or provide assistance to the robot. As many tasks require robots to engage in physical actions, motion can be an effective communication modality for conveying goals, in the form of explicit signals, such as gesture, or through implicit cues, such as legible goal-directed motion. Legible motion aims to create robot trajectories that are readable and enable observers to predict the robot's motion intent without additional signals. To convey more complex goals, gestures are often used to provide a wide range of options [Holladay et al., 2014]. While some gestures are a part of everyday vocabulary (e.g., pointing/deictic gestures), others are specially designed for specific applications and require training (e.g., aircraft control gestures).

Humans utilize indirect, unconscious signals that take on many forms and often involve a combination of multiple modalities, such as gaze, facial expression, and posture. Since non-humanoid robots often do not possess those channels, researchers have instead explored other modalities for goal expression, including visual displays, lights, and sound [Szafir et al., 2015, Szafir, 2015]. These modalities are flexible and easy to use but often lack intuitiveness.

Plan

While a goal is the desired outcome of the task, a *plan* is how that outcome is achieved. A plan is a set of primitive actions that the robot must engage in to complete a task. The scope of primitive actions can vary, but each has an intent and a subplan. Primitive actions can range from a single physical change in the robot's configuration to a complex process, such as picking up an object. The scope of a primitive action is defined by the programmer and depends on the scenario and the robot's control architecture. If the primitive action is a robot's base movement, the intent is the robot's final base configuration, and the subplan is its trajectory. Often, more than one subplan or plan can achieve the same outcome, allowing for optimizing criteria such as efficiency or goal-prediction.

A robot can take physical or non-physical actions in both the immediate and remote environments. Physical actions consist of the robot utilizing its actuators to generate motion that affects the tangible environment. Non-physical actions (e.g., speech) are typically designed not to have a noticeable effect on the physical environment but to effect the inhabitants of the robot's immediate environment. While robots have the capability to take actions (e.g., electronically transmit information) to create remote effects, the focus of this work is on interactions with co-located humans.

Research in motor cognition focuses on understanding how physical actions are perceived and understood [Haggard, 2005]. Non-physical actions often have a more straightforward goal in their execution. Since many early robots were controlled by human operators, interfaces for providing task information, such as the robot's plan, often took the form of intricate displays or text output. However, these interfaces are not suitable for co-located humans who typically possess limited knowledge of the robot's complex operations or do not provide control inputs for the robot. Interpreting these displays also requires significant time and effort. Hence, much of HRI work focuses on conveying the subset of the robot's plan that is most relevant to the human at that time. This approach may negatively affect long term efficiency as the human has limited knowledge of the robot's long term plan.

4.3. Activity

Task Progress

Task progress is the movement or development towards the task goal. In the context of this work, it is loosely defined as how far the robot has advanced in its plan. When the desired outcome has been achieved, the robot's progress is at its maximum and the task is complete. If the robot has not yet started execution of its task, its progress is at its minimum. Similarly, primitive actions once started progress towards their intent.

However, task progress is not always easy to determine. While the most straightforward method is to map progress directly to advancement in the task plan, plans are not always predetermined and can shift when conditions in the world change. Moreover, tasks can also terminate without the goal being met. The conditions for completion of a task, called postconditions, can also include faults (e.g., loss of communication) that cause the robot to stop executing, as well as partial completion of the goal.

Understanding the robot's progress can help humans to predict when and how long a robot might be performing actions in a shared environment. This is especially helpful in environments that require significant coordination and collaboration. Moreover, if a human is working with the robot, either providing tasks or contributing towards a shared objective, assessing whether the robot is making progress is critical to the human.

Success

When a robot finishes executing task, it has either succeeded and met the goal condition or has failed and halted through an alternative postcondition. This information is important to humans who depend on the outcome of the robot's task. Information about task success can also include greater detail about the termination condition, such as whether the goal was partially or completely met or faults that may have prevented further progress towards the goal. As with task progress, success also applies to primitive actions with the intent replacing the goal.

4.4 Summary

In this section, we presented an outline of internal robot states. While these states are not exhaustive, they provide a method for organizing information that is important for humans interacting with the robot in a variety of scenarios. In the next section, we present a review of prior work in signal generation and discuss how signaling methods can be utilized by non-humanoid robots to convey these internal robot states.

5

Signal Characteristics

The goal of this section is to provide roboticists with a unified vocabulary with which to describe nonverbal robot signals. We define two key signaling characteristics, criticality and information content, that describe the effects of a robot's signal on a human observer and the robot's communication intent. These characteristics represent design choices that must be addressed when developing new nonverbal robot signals and are derived primarily from past work in signal generation for HCI [McCrickard et al., 2003, Pousman and Stasko, 2006]. The key differences taken into account when defining these terms for HRI are the robot's' ability to physically interact with the environment and its potential to act and be perceived as an intelligent, embodied agent.

5.1 Criticality

Criticality is the measure of how important it is for a human to observe a robot and respond to its nonverbal signal. In HCI literature, this characteristic is often called the notification level and includes the desired attention level and response of a human user [McCrickard et al., 2003, Pousman and Stasko, 2006, Tomitsch et al., 2007]. It also corresponds to the criticality of the signal's contained information and assumes that a human will only respond to the system in the most critical cases [Pousman and Stasko, 2006]. Since robots can benefit from human intervention in non-critical scenarios as well, we describe criticality using two distinct features, *notification level* and *action response*, that better encompass robot signaling needs.

5.1.1 Notification Levels

Notification levels are defined in HCI literature as the degree at which a system alerts the user and requires a response [McCrickard et al., 2003, Pousman and Stasko, 2006]. Prior work in ambient displays defines five notifications levels: ignore, change blind, make aware, interrupt, and demand action [Matthews et al., 2003, Tomitsch et al., 2007]. Higher notification levels are meant to grab the user's attention quickly while lower notification levels should not attract conscious attention [Pousman and Stasko, 2006].

Ignore represents information that should not be displayed or signals that are completely ignored unless sought out. Change blind refers to signals that have marginal importance and are on the peripheral of user attention [Intille, 2002]. Make aware signals should be noticed and consume some of the user's attention but not all (i.e., divided attention). Interrupt and demand action represent signals that are highly attention-grabbing and are typically differentiated by the need for the user to respond to the system to cease a "demand action" signal.

To meet HRI needs, we define notification level to only designate how important it is for a human to observe the robot's signal, which relates to the signal's salience or how attention-grabbing it is. Since the highest notification level "demand action" also concerns the observer's response, we only utilize the first four signal notification levels as defined above. As the notification level specifies the amount of attention desired from co-located humans, it is particularly helpful when choosing a signaling modality, as certain communication channels may have higher salience due to human perception or environmental conditions. Notification levels can also be considered in conjunction with signaling needs when designing a robot's embodiment [McCrickard et al., 2003].

5.1. Criticality

Humans utilize a range of cues that can draw attention at varying levels [Saulnier et al., 2011]. The most overt signals include large body motions (e.g., waving), sounds (e.g., yelling, whistling), and direct verbal communication (e.g., calling someone's name). Other less conspicuous cues include proxemics, eye gaze, and subtle gestures [Admoni et al., 2011, Saulnier et al., 2011, Takayama and Pantofaru, 2009]. These cues vary in salience and duration but can also be modulated (e.g., sound intensity) to increase awareness.

Robot signal design can draw inspiration not only from human communication behavior but also from alarms and alerts utilized by other machines and devices [Cha and Matarić, 2016, Ulfvengren, 2003]. Past research has explored the principles of alarm design, which often draw on psychophysical perceptions of urgency [Edworthy et al., 1991, Guillaume et al., 2002]. Popular modalities for alarm signals include auditory signals, lights, and visual displays [McCrickard et al., 2003]. The use and modulation of certain signal modalities (e.g., acoustic properties) have been well studied in order to understand their effects on human perception [Ulfvengren, 2003]. As these signals are regularly encountered in everyday life, utilizing similar mechanisms for robot signal design allows humans to draw from their common experiences, making signals more intuitive and easy to interpret.

A key challenge for robot signaling is taking into account the world state, including the human and robot, on the design and usage of signals. Past work in alerting has already shown that the human's current state, such as cognitive load, and past experiences are important factors to consider when deciding how and when to signal a human [Fischer, 2001, Horvitz et al., 1999]. Moreover, non-humanoid robots can vary widely in their physical appearance, functionalities, and autonomy, which can limit their signaling capability [Bartneck et al., 2009, Bethel and Murphy, 2008, Cha et al., 2015, Hinds et al., 2004, Steinfeld et al., 2006, Walters et al., 2009]. In scenarios where there are several relevant pieces of information, the notification level can help to prioritize signals.

5.1.2 Action Response

Action response is the attribute of desiring the human observer to perform some action in response to the signal. This action acts as a type of feedback to the robot's communication Figure 1.2. We can categorize the desired human response by two features, *necessity* and *urgency*, which describe if and when the human should respond to the robot's signal.

While prior work typically assumes the human's response directly involves the system, we assume that these actions can be independent of the robot [Pousman and Stasko, 2006]. For example, if the robot is moving at a fast pace towards the human, the human can move out of the robot's path to prevent a collision.

As robots encounter many unpredictable and non-deterministic scenarios, human actions involving the robot or its environment can be beneficial, even in cases where the action is not strictly necessary or has low urgency [Rosenthal and Veloso, 2012]. For instance, if the robot is taking a long time to find a particular object it requires, a human can identify the item to help the robot complete the task more quickly. Different action response levels also provide greater flexibility for the human by enabling them to respond at a later time if the action is not urgent.

Since the urgency and necessity of the action response can vary significantly, we describe four types of action response: informative, beneficial, necessary, and urgent. In the first type, *informative* (Figure 5.1 top, left), no action is needed by the human observer, and the signal is primarily designed to provide information. Several signals that aim to provide transparency into the robot's behavior and condition fall into this category (e.g., an indicator that the robot is fully operational). They also enable more accurate mental models and can reinforce the robot's default state (e.g., sufficient power).

Beneficial indicates that the robot can benefit from a human observer taking action, but the action is not necessary for the robot's current operations (Figure 5.1 middle, left). For example, if the robot is having difficulty navigating, a human can remove obstacles or manually move the robot to help it advance more quickly. These actions

260



Figure 5.1: The action response is the desired response from a human observer to a robot's signal; it is characterized by the necessity and urgency of the human's action.

may occur if the robot and human are collaborating on a task, or if the robot is impeding the human's own task (e.g., blocking a path).

For scenarios in which a response action is not required, we assume that the urgency of the response is low (Figure 5.1 top, right) as there are limited cases where an action response is not required but still urgent.

The last categories of action response, *necessary* and *urgent*, both require action to be taken by the observer but differ in the urgency of the action(Figure 5.1 bottom). Necessary indicates that the robot

will eventually require a human to perform the response action, but not immediately (Figure 5.1 bottom, left). For instance, a robot that is running low on power may need to be charged but can sustain itself for some period of time before turning off.

Urgent, on the other hand, requires that the human observer take action immediately (Figure 5.1 bottom, right). For instance, if a robot is physically stuck during a time-sensitive task, they require immediate assistance from a human. Similarly, if the robot is about to collide with an object or human in the environment, an imminent action is necessary to ensure the robot's and/or human's safety.

It is important to note that the type of action response the robot desires may escalate over time if the urgency or necessity of a response increases. In the low power example, as the robot's power level becomes closer to insufficient operating levels, the action response escalates in type from necessary to urgent.

5.2 Information Content

The other signal characteristic we consider is *information content*, the amount and type of information contained in a signal. Information content helps to determine appropriate communication modalities for a nonverbal signal. In cases where greater flexibility or depth are needed, text, either through speech or a visual display, is often the fastest, and most efficient method for transmitting information [Pousman and Stasko, 2006, Tomitsch et al., 2007]. If an error requires the human to take very specific actions to resolve it, it may be easier to illustrate the steps with an image. In such scenarios, brief, nonverbal signals are poorly suited for conveying the type and amount of information the human requires.

Conversely, there are many pieces of information that can be easily conveyed by a nonverbal signal. Certain information, such as power levels, may need to be shown continuously, but repeated speech would become highly annoying while a textual display would not be salient. Moreover, the use of speech facilitates anthropomorphization and can create expectations of intelligence or capability that the robot does not

5.3. Summary

meet. Hence, considering the best form for conveying information is a vital part of creating robot communication behaviors.

This issue is also closely related to the concept of *representational fidelity*, which describes a system's display components and how data from the world is encoded into patterns, pictures, words, or sounds on the display [Pousman and Stasko, 2006]. While some systems represent information in very precise ways, such as a textual display, others are more abstract in their representation (e.g., blinking light). Humans also have different capabilities when it comes to recognizing and differentiating signals based on the signaling modality. The psychophysical differences are often hard to compare but are important to consider as they are also affected by the environment, and hence, the context of the scenario.

Similarly, we must also consider how many distinct signals a robot should utilize across all of its signaling modalities. As humans are able to learn and intuitively react to a limited number of signals, the robot's signaling vocabulary must be constrained to match human capabilities and preferences. Although training can help to surpass this limitation, in many of the HRI scenarios presented in Section 3, we cannot assume that human interactors have training or prior experience. An additional constraint is efficiency, as it is undesirable for humans to spend significant time or effort identifying and interpreting a signal. Due to the numerous features in a robot's state, certain types of information will need to be prioritized to meet the robot's goals.

5.3 Summary

In this section, we defined characteristics which describe a signal's content and desired feedback in the form of a human observer's action. These features help to describe a signal and can affect important design decisions, such as the signaling modality, and how a signal should be utilized (e.g., timing). In the next section, we review signal generation techniques and discuss how they can be applied to creating nonverbal signals with varying levels of these characteristics.

6

Signal Generation

We presented a selection of non-humanoid robot use cases (Section 3), an overview of the internal robot states relevant for HRI (Section 4), and a set of key signal design characteristics (Section 5). While Section 3 and Section 4 primarily served to illustrate what information is important for a robot to signal, Section 5 was a first step towards understanding how to generate and utilize nonverbal signals. The goal of this section is to further inform signal design by discussing specific methods for generating nonverbal signals. Towards this goal, we survey past works in signal generation, categorized by signaling modality, and discuss their application to signal design for non-humanoid robots.

6.1 Gaze

Although humans possess several senses, much of their perception relies on vision. Consequently, human visual behavior or gaze is a powerful nonverbal communication tool that can be used to reveal information about a human's mental state, such as their goals, knowledge, and intent [Admoni, 2016, Argyle and Cook, 1976, Kleinke, 1986, Srinivasan and Murphy, 2011]. By focusing on a human's gaze, we are often able to

6.1. Gaze

discern not only whether they have perceived an object in the environment but also the focus of their attention. Gaze can also be used to draw others' attention to an object or event in the environment [Mukawa et al., 2001, Yoshikawa et al., 2006]. This is a key component of achieving *joint attention*, or shared attentional focus on a common object, which enables physical coordination [Admoni et al., 2011].

Gaze also plays a strong role in social interaction by conveying important social and emotional information, such as affect, engagement, and interpersonal attitude, and providing conversational cues [Admoni, 2016, Andrist et al., 2014, Argyle and Dean, 1965, Argyle, 1972, Argyle and Cook, 1976, Kleinke, 1986, Rich et al., 2010]. Accordingly, gaze behavior has been extensively studied to provide insight into the complex dynamics of nonverbal human communication [Admoni, 2016, Argyle and Dean, 1965, Argyle, 1972, Kleinke, 1986, Mehrabian, 1972]. Social sciences literature has largely looked at measurement and interpretation of human behavior, but other areas of research have explored the modeling and generation of gaze behavior for animation, virtual agents, and robots [Admoni, 2016, Andrist et al., 2012, Thiebaux et al., 2009].

Recent research in robotics and virtual agents seeks to replicate human gaze behaviors on different characters and embodiments [Admoni, 2016, Andrist et al., 2012, Garau et al., 2003, Mutlu et al., 2006]. Much of this work has sought to create models for generating expressive eye gaze behavior. Unlike animated characters and virtual agents, however, robot communication capabilities are limited by physical embodiment. As robots utilize cameras to capture visual data, "eyes" and head movement may not be indicative of the robot's actual focus or visual field. Hence, recent work in HRI has generally looked at two problems: 1) how can gaze behaviors be utilized by current robots, and 2) how does robot gaze behavior affect interactions with humans [Admoni, 2016, Andrist et al., 2012, Ham et al., 2011, Huang and Mutlu, 2013b, Mutlu et al., 2006, Thiebaux et al., 2009].

Past works in HRI have studied several types of gaze behavior. In addition to joint attention, there is mutual gaze, referential or deictic gaze, and gaze aversion [Admoni, 2016, Andrist et al., 2014, Choi et al., 2013, Mutlu et al., 2006, Okumura et al., 2013]. *Mutual gaze* is

often described as "eye contact" or when two agents have reciprocal eye gaze directed towards one another [Argyle and Cook, 1976]. *Referential* or *deictic gaze* refers to gaze directed at an object or location in space [Senju et al., 2008]. *Gaze aversions* are shifts of gaze away from the main direction of gaze and can be intentional (i.e., gaze avoidance) or unintentional (i.e., gaze omission) [Kleinke, 1986]. These different types of gaze have been shown to demonstrate engagement [Rich et al., 2010], control the flow of conversation [Mutlu et al., 2009a], and refer to the environment [Okumura et al., 2013] during human-robot interactions.

Some works have also found shortcomings in researchers' attempts to replicate human gaze behavior. In comparison to human gaze, robot gaze behavior can lack saliency, and thus, be overshadowed by other aspects of the robot's behavior, such as its movement [Admoni et al., 2011, 2014, Okumura et al., 2013]. This can be compounded by the robot's appearance, as many platforms do not have eyes that are human-like in appearance or capability. However, even when varying the level of anthropomorphism, robot gaze behavior can fail to elicit the same response as a human's, making it challenging to utilize for many types of interaction [Admoni et al., 2011].

Although few works explore the use of gaze-inspired behavior on non-humanoid robots, there have been some promising results on applying gaze-inspired behaviors to other signaling modalities [Szafir et al., 2015]. Inspiration can also be drawn from fictional non-humanoid robots such as *Star Wars'* R2D2 and BB-8 or *Meet the Robinson's* Doris. These robots have more mechanical appearances and lack humanoid shape entirely, yet are shown to successfully utilize gazeinspired nonverbal behaviors across a range of scenarios. The channels utilized by these fictional robots and in prior HRI work include head motion, body movement, and LED "eyes."

Gaze behavior can also be designed using psychophysical principles, as well as drawing on common user experiences from everyday life. For instance, to increase the saliency of a non-humanoid robot's gaze, abrupt luminance, color, or shape changes can be utilized to create a significant pop-out effect (e.g., red light on unlit black back-

266

6.2. Gesture

ground) [Theeuwes, 1995]. We can also take advantage of common signal conventions, such as light color, or other perceptual methods, such as directional cues created by changing light color or shifting location [Dobkins and Albright, 1993].

Prior work in HRI suggests that the use of exaggeration, either in motion amplitude or alterations in temporal properties, can enhance the recognizability of and response to such signals [Hill and Pollick, 2000]. Exaggeration and timing are both key principles of animation that have been applied to robot behavior to enhance noticeability, readability, and effectiveness [Ribeiro and Paiva, 2012, Takayama et al., 2011, Thomas et al., 1995, Van Breemen, 2004, Zhou et al., 2017].

One of the key challenges for utilizing gaze-inspired behavior is incorporating an appropriate signaling channel into a non-humanoid robot. Instead of true human-like eyes, many robots utilize simple, low fidelity signaling mechanisms to mimic gaze. This often makes human gaze models ineffective or difficult to apply on such platforms. Hence, animation principles and knowledge from psychophysics can inform the design of gaze-inspired behavior and compensate for the shortcomings of such platforms.

6.2 Gesture

Gestures are a popular form of nonverbal communication that include visible movements of the hands, arms, face, and other parts of the body to express an idea or meaning [Knapp et al., 2013, McNeill, 2008, Mehrabian, 1972]. They are used across a wide range of interactions such as conversation, narration, instruction, and formal speeches [Argyle, 1972, McNeill, 2008]. Researchers have identified broad categories of gestures including emblems, beat gestures, iconic gestures, deictics, metaphoric gestures, regulators, and affect displays [Allan, 1988, Argyle, 1972, Beattie, 2004].

Emblems, iconic, and *metaphoric gestures* have strong semantic ties and are also referred to as representative gestures as they tend to depict concrete ideas or meanings [McNeill, 1992, 2008]. *Deictic gestures* are used to point to objects or areas in the environment and can be used to disambiguate, show interest, or create emphasis [Brooks and Breazeal, 2006, McNeill, 1992, Sauppé and Mutlu, 2014]. Similarly, *beat gestures* can function to emphasize different points of speech or regulate its rhythm [Argyle, 1972, McNeill, 1992, Salem et al., 2012, So et al., 2012]. Unlike representative and deictic gestures, beats are almost always used in combination with speech.

Gestures often evolved differently across cultures, meaning that a particular gesture that is common in one culture can be unrecognizable or considered offensive in another [Archer, 1997, Efron, 1972, Mc-Neill, 2000]. Hence, gestures often vary not only in form but usage and social significance. Although gestures can communicate semantic information, regulate speech, illustrate concepts, and enhance conversation, the value and role of gestures differ across interactions [Argyle, 1972, McNeill, 1992, Mehrabian, 1972]. For instance, when teaching or providing instruction, gestures help listeners to better comprehend the information, which can affect people's attitudes and perceptions of the speaker [Huang and Mutlu, 2013a]. In collaborative scenarios, deictic or referential gestures can disambiguate references and enable higher levels of coordination [Sauppé and Mutlu, 2014].

Similar to gaze, researchers have primarily focused on how robots can replicate human gesture to improve interactions with humans [Huang and Mutlu, 2013a, Salem et al., 2012]. Past work has explored gestures in the context of specific application such as teaching, direction giving, and coordinating joint activities. [Chidambaram et al., 2012, Okuno et al., 2009]. Other works have focused more on specific interaction outcomes such as task performance, robot persuasiveness, conversational fluidity, and perceptions of the robot [Huang and Mutlu, 2013a].

While there have been many advances in creating effective gestures for HRI, the gestures utilized in these studies are often realized through movement of a robot's arms or head, which non-humanoid robots either do not possess or have a limited degree of expressiveness. Moreover, these movements can interfere with a robot's primary goal-oriented behavior. Fewer works in HRI have looked at methods for replicating these gestures on more abstract platforms. As there is a large body of

268

6.3. Implicit Motion

work investigating how gestures affect people's interactions and perceptions of different robots [Huang and Mutlu, 2013a, Salem et al., 2012, Sauppé and Mutlu, 2014], it may be worthwhile to explore the creation of new gestures on alternative communication modalities to achieve similar outcomes on non-humanoid platforms.

One relevant body of work is animation, as characters often take a wide range of shapes and forms and have varying levels of expressiveness. The technique of *motion retargetting* explores different methods for adapting motion from one form to another. This can be done through mapping a few key poses between a human and the character and interpolating between the poses or building a statistical model for the mapping [Yamane et al., 2010]. Alternatively, important features of the original motion can be used as constraints while generating new motion on the character [Gleicher, 1998]. These methods are often computationally difficult and take significant amounts of time, requiring preprocessing that is not always possible for robots. For many nonhumanoid robots, retargeted gestures rely on whole body movements which create additional complexity and safety concerns. As the primary purpose of a gesture is to convey a specific meaning, it may be easier to create new signals that convey the same information as the gesture without being constrained to a specific motion shape or communication channel [Szafir et al., 2015].

6.3 Implicit Motion

A recent area of HRI explores algorithms for generating implicit motion cues for the robot [Bodden et al., 2016, Dragan, 2015]. Motion can be a highly expressive form of communication and studies have shown that even babies attribute complex properties such as intentionality to motion [Buren et al., 2016]. In this section, we focus on implicit motion cues contained in the robot's goal-directed motion as explicit motion cues are discussed in the context of gaze and gesture.

Since traditional path planning algorithms primarily focus on quickly generating a collision-free trajectory from the start to the goal, the resulting motion can be unpredictable, complicated, and unnatural [Dragan, 2015, Dragan et al., 2015]. Several recent works in HRI have explored approaches to making robot motion more human-like, predictable, intent-expressive, or legible [Kruse et al., 2012, Dragan et al., 2013, Szafir et al., 2014] for use in more human-oriented applications. Past studies found that this more expressive motion promotes safety, enhances trust, and enables coordination making it a valuable tool for HRI [Dragan et al., 2015, Lichtenthäler and Kirsch, 2016].

A common approach optimizes the robot's trajectory according to desired motion properties while preserving the start and end positions and/or configurations. The quality of the motion is highly dependent on the cost or reward function, which can be generated in several ways. One method is to use human movement data to extract a cost function directly using a technique such as inverse optimal control [Mombaur et al., 2010]. This also enables the robot to incorporate different levels of social convention in its behavior [Kruse et al., 2012].

However, human-inspired motion does not always contain the desired properties for a human-robot interaction. For instance, human arm trajectories for goal-directed motion, such as grasping, often prioritize efficiency over intent expression. In such cases, it may be better to directly create motion that emphasizes the desired characteristic rather than attempting to replicate human motion behavior. In the previous example, legible or intent-expressive motion can be generated by optimizing the probability of a motion-goal based on its trajectory from an observer's view [Dragan et al., 2013, Dragan and Srinivasa, 2013].

Other methods rely on a greater amount of human input, such as choreographing the robot's motions or directly controlling the robot's trajectories [Knight and Simmons, 2014]. Although there are numerous benefits to having a human-in-the-loop, many applications do not have an available human. Hence, recent work has emphasized the need for more generalized methods that achieve the desired interaction outcomes while moving away from the need for human input.

As with using gesture, there are also additional challenges associated with using robot motion for non-humanoid robots. Many of the non-humanoid robots presented in the target use cases possess only a few degrees-of-freedom which can lack the capability to generate highly expressive motion. Moreover, altering the robot's goal-directed motion may negatively affect its task performance (e.g., increased task time). Hence, utilizing these techniques to mimic motion through lights or a visual display may offer more flexibility and alleviate some of these concerns. As there has been little work in this area, more research is necessary to understand how other communication modalities can achieve similar effects.

6.4 Sound

There are many ways to leverage sound as a mechanism for nonverbal communication. Many robots naturally generate noise while operating, which expresses information about their internal state (e.g. motor sounds when moving). Robots can also use explicit auditory cues, ranging from simplistic nonverbal utterances (e.g., beep or chirp) to speech [Cha and Matarić, 2016, Fischer et al., 2014, Martelaro et al., 2016, Read and Belpaeme, 2014b]. Although speech is a natural mechanism for human communication, it can be inefficient for information easily conveyed by concise nonverbal signals. The usage of speech may also raise expectations of the robot's capabilities [Cha et al., 2015] and hearing a robot continuously vocalize its state information may become tiresome or disruptive to nearby humans [Kiesler, 2005].

Humans typically generate two types of auditory cues defined by their source and method of production. *Internal auditory cues*, such as snoring, sighs, or breathing, are generated entirely by the human body. While certain sounds are involuntary and correspond to a biological function in the body, others are a more conscious form of communication (e.g., groan). *External auditory cues* are produced by the physical interaction of a human with elements in the environment, such as footsteps on the floor. These noises often provide functional, task-oriented information as they are generally a result of our goal-driven physical actions. Internally produced sounds, on the other hand, tend to convey information about a human's mental state, such as their mood or emotional state. A subset of these internal sounds are referred to as *nonverbal utterances* (NLUs), sounds produced by the vocal track that do not form words or language. Some examples include moans, sighs, grunts, and laughs. These sounds occur naturally in human communication and often rely on variations in human vocal dynamics. For functional nonhumanoid robots, especially those with mechanical appearances, these cues may appear inconsistent with their embodiment and behavior.

Instead, non-humanoid robots can utilize synthetic NLUs, such as beeps, whistles, or chirps [Read and Belpaeme, 2014b]. These types of sounds are often utilized by fictional non-humanoid robots, such as *Star Wars'* R2D2 or *WALL-E's* M-O, instead of speech. Typically, NLUs are individually generated by a human designer, making it difficult to create an extensive vocabulary of signals to fit a larger range of situations.

Recent work has attempted to overcome this challenge by exploring how specific acoustic properties of NLUs affect humans' attribution of meaning to these cues [Read and Belpaeme, 2012, 2014a]. Preliminary results seem to indicate that interpretations tend to be inconsistent and coarse in the absence of situational context [Read and Belpaeme, 2012, 2014a]. This is unsurprising as these auditory cues are abstract and mostly utilized in visual media which provides context and reinforcement through other characters' reactions, music, and visual effects.

Many deployed non-humanoid robots instead utilize auditory signals similar to those employed by machinery, appliances, and computing devices. These sounds often take the form of alerts or alarms and usually signal predefined events. This approach works well for robots that operate in a limited number of scenarios and only require a few auditory cues. However, for more complex use cases, the number of abstract sounds needed becomes greater than humans are capable of easily learning and distinguishing. Additionally, many robots are expected to interact with naive humans who require intuitive signals to appropriately react and respond to.

An alternative approach is to create signals that mimic natural sounds produced by the robot's operations (i.e., auditory icons). Many sounds produced by humans, such as breathing or walking, are eas-

272

ily mapped to information about their internal state and are hence, intuitive to listeners. Depending on the specific platform and environment, the robot may already produce some of these cues. Since these unintentional cues can be inconsistent, even from the same robot, the regular use of auditory icons for signaling can produce a more reliable experience.

Consistency in the use of auditory signals enables other cognitive processes such as auditory localization. These signals can also be modulated to create a larger set of signals to match the world state (e.g., the robot, human, and environment). For example, the volume or intensity of a cue can be altered depending on the robot's proximity to humans. This lowers the amount of unnecessary noise, minimizes disruption, and creates more intelligent, adaptive behaviors for the robot.

Other areas of work that provide insight into auditory signal design include warning design, human factors, and psychophysics. Combining these signals with other cues such as light has shown positive results in overcoming the abstract nature of sound to create more understandable and readable signals [Cha and Matarić, 2016]. As the focus of prior work in auditory cues for robots has primarily focused on speech, more research is needed in generating nonverbal auditory signals that support interaction for a larger range of platforms and scenarios.

6.5 Visual Displays

Visual displays have become increasingly popular for many robotic applications with the rise of touchscreen platforms, such as smartphones and tablets. Visual displays are able to utilize a wide range of signals with varying representational fidelities [Pousman and Stasko, 2006]. This versatility can be extremely helpful when trying to express large amounts of information as these displays can utilize text and high resolution images. Visual displays can also replicate simpler signals, such as point light sources, making them extremely flexible. However, these displays are still costly and have significant power and weight requirements. Moreover, many of the signals presented on visual displays require significant time and effort to interpret and require the human to be in close proximity.

Visual displays range in the amount of attention and interaction they are intended to have with humans. Displays, such as computer screens, are designed such that humans will directly interact with and focus the majority of their attention on the display. Some devices utilize their display at different attention levels; most smartphones have ambient notification systems which provide alerts when the user is not directly interacting with the system.

Robots utilizing visual displays, such as the LoweBot or Savioke Relay, enable humans to interact with and provide commands to the robot via the display, reducing the chance of error from using speech or gesture. Since many non-humanoid robots are designed to perform functional tasks, the inclusion of visual display for providing and receiving task-relevant information can be advantageous. However, such displays are not as well suited for providing information to bystanders or users while the robot is in operation or mobile.

Ambient displays are on the periphery of a user's attention and do not involve direct interaction between the human user and system. Instead, they aim to present information without distracting the user from their primary task. Since a significant advantage of nonverbal communication is the efficiency with which it conveys information, prior work in ambient displays can offer many relevant design insights for robot signal generation. By utilizing images, color, shape, and other characteristics, ambient displays have been able to convey a wide range of information including weather forecasts, physical activity, and proximity [Fortmann et al., 2013, Hegarty et al., 2010, Ishii and Ullmer, 1997, Vogel and Balakrishnan, 2004]. Research on visual displays has also explored the effects of particular signal characteristics on the interaction outcomes of human users and the system.

Visual displays can be combined with other communication modalities to provide greater detail or flexibility in certain scenarios. For instance, if a complicated error occurs, a robot can utilize sound or motion to draw humans' attention to its display. The display can then provide information that is too detailed or difficult to convey with more limited nonverbal modalities. This enables the robot to utilize a more diverse set of communication behaviors, similar to humans, while still providing higher levels of detail when necessary.

6.6 Lights

Lights are a promising visual mechanism that has been utilized on a wide range of mechanical devices such as appliances, computers, cars, airplanes, and mobile phones [Harrison et al., 2012]. On such devices, lights are primarily used as simple indicators. Mobile phones indicate new notifications through the use of a single LED while aircraft control boards signify the state of different controls and equipment through color and whether the LED is turned on.

Although light is simpler and carries less information content than other visual indicators, such as text, it is easily discernible and can vary in noticeability [Hansson and Ljungstrand, 2000, Matviienko et al., 2016]. This enables light to be used as a signaling mechanism across different notification and criticality levels. Several lights can also be combined to increase information capacity and create different configurations that fit a range of robot forms and applications.

Prior work in visual perception and psychophysics has shown that dynamic visual cues can convey complex properties even in simple animations or light motions [Dittrich and Lea, 1994]. Moreover, people attribute not only animacy, but intent to these simple signals. Other work in lighting displays also shows that light can evoke complex social responses, including emotion [Dittrich et al., 1996]. Light colors already have strong connotations and are utilized in everyday applications such as traffic control.

Since non-humanoid robots are often more machine-like in appearance, LEDs are a natural choice of signaling mechanism. Current nonhumanoid robots, such as the Sphero or iRobot Ava (Figure 6.1), often utilize LEDs of varying colors to indicate when the robot is turned on or in different states (e.g., error). Autonomous vehicles are already equipped with lights (i.e., brake lights, head lights, tail lights) which act as indicators of the robot's state.



Figure 6.1: Lights have been utilized on fictional and commercial robots, including (clockwise) Eve from Pixar's *WALL-E*, the Zeroids from *Terrahawks*, the Nissan Esporo, and the iRobot Ava, to increase expressiveness and provide state information.

Light signal design in the context of non-humanoid robots also requires consideration of the robot's variability in form, animacy, and application. Roboticists have explored utilizing light to convey motion intent, task state, and affect [Baraka et al., 2015, Szafir et al., 2015]. Many of these works draw design inspiration from common user experiences with light signals (e.g., lighthouse, jet engine flames) [Baraka et al., 2015, Szafir et al., 2015]. Other works have utilized light to reinforce another communication mechanism, such as speech or facial expressions [Funakoshi et al., 2008, Kobayashi et al., 2011].

Another source to draw inspiration from are fictional robots, as shown in the top of Figure 6.1. Some common examples of light signal usage by robots in popular media include Eve from Pixar's WALL-E, Kitt from the television show *Knight Rider*, the Zeriods from *Terrahawks*, and the mechanical Cylons from *Battlestar Galactica*. Some of these robots utilize continuous variations in lights to give the appearance of animacy or life; the Cylons have a red light that pulses in a

6.7. Summary

cyclic pattern to mimic movement from one end of the robot's head to the other while the Zeriods have "eyes" with 3x3 grids of LEDs that turn on and off.

Despite the large number of commercial and research robots that employ LED lights, few works have actually researched design methods for light signals. The approach of designing for individual products limits the generalizability of the developed signals and lowers user friendliness of the platform. This problem can be even worse for non-humanoid robots as their usage and embodiment vary significantly compared to other product (e.g., smartphones). The size of the design space of light also limits testing to a small set of signals in constrained scenarios. This indicates the need for greater research into general design principles that can be applied across robot platforms and scenarios to enable generation of light signals with little to no human input.

6.7 Summary

In this section, we discussed how to generate signals for non-humanoids and presented an overview of past work in signal generation, categorized by signaling modality. Although there are many areas of work to draw inspiration from, there are still many challenges for leveraging this past work in generating non-humanoid robot signaling behaviors. In the next section, we discuss considerations for nonverbal signaling for nonhumanoid robots and open research areas.

7

Discussion

The works included in this survey suggest that nonverbal signaling for non-humanoid robots is a growing research effort in the HRI community. This topic is becoming more important as increasing numbers of non-humanoid robots are deployed into common human environments. For many humans, these robots will act as an introduction to future robot technologies and may significantly influence their perceptions. Therefore, it is important that robots are able to effectively communicate their internal state in order to set appropriate expectations and support interaction. In this section, we outline challenges and open research areas, as shown in Figure 7.1, to aid researchers in advancing the field of HRI for non-humanoid robots.

7.1 Considerations for Nonverbal Signaling

In this section, we identify and discuss issues that must be taken into account when developing nonverbal signals and signaling algorithms for non-humanoid robots. Since the research in this area is still young and emerging, many of these considerations are open-ended. We expect that



Figure 7.1: From the works surveyed, we identify eight considerations for nonverbal signaling and three open research areas.

as the field grows and gains structure, these considerations will evolve but continue to pose ongoing challenges for HRI research.

7.1.1 Embodiment

The design of signals for non-humanoid robots is primarily constrained by the robot's embodiment (i.e., its physical form). As the use cases in Section 3 illustrate, non-humanoid robots vary widely in capability and usage, which also affects their size and shape. This variance can also alter the dynamics of interaction as humans often perceive and interact with robots differently based on their appearance and intended task. For instance, humans are likely to have preconceptions about autonomous vehicles as they largely resemble current human-driven ones.

Non-humanoid robots typically have physical forms targeting specific tasks in order to make them more reliable, robust, cost-effective, and manageable. Although we have identified several common signal-
ing scenarios, differences in embodiment make it challenging to create a standard set of signals utilizing only the available hardware.

These limitations in signaling capability also prevent most nonhumanoid robots from duplicating human nonverbal behaviors, because they either lack the physical hardware entirely or cannot achieve a reasonable degree of fidelity. One common example is human gaze behavior; many robots utilize gaze via an actuated camera, sensor array, or head-like construct. However, it can be difficult for this imitated "gaze behavior" to achieve the same communication outcomes or the high degree of precision and saliency as human gaze [Admoni et al., 2011, Al Moubayed and Skantze, 2012, Imai et al., 2002].

Since non-humanoid robots typically have fewer degrees of freedom, signals can rely on components that are typically utilized while performing functional tasks. Using these signals can interfere with the robot's primary capabilities and negatively impact its task efficiency. For example, a robot that uses base or arm motion for both signaling and moving objects may be unable to do both simultaneously. In these cases, the task is typically given higher priority. As communication is a vital part of working with and around humans, more research is needed to explore other communication channels that can easily be utilized by many different non-humanoid robots without affecting their primary functionality.

7.1.2 Intuitiveness

The prevalence of nonverbal cues in everyday communication has made humans experts at noticing and interpreting such cues. Often signals utilized by non-humanoid robots are either poor imitations of human nonverbal behavior (e.g., gaze) or simplistic indicators, such as those used by electronic devices (e.g., LEDs). This has resulted in signals that are not particularly intuitive or iconic. Instead, humans have to learn by rote what each signal means on each robot.

Training is typically utilized to help humans learn large sets of signals is settings where the same humans and machines are expected to interact (e.g., industrial settings, aircraft control). However, there are many use cases in which we cannot expect humans to spend the effort

280

and time to learn a robot's signals prior to interacting. Inexperienced humans who are exposed to the robot for the first time must still be able to sufficiently understand the robot's internal state for safety and interaction.

Making signals intuitive and easy to decode is also necessary for designing a larger set of signals that can be applied more broadly to different non-humanoid robots and applications. Key to this goal is exploiting common user experiences that can inspire and inform signal design. Many abstract signals, such as police sirens, are so ubiquitous that humans react and respond automatically, without significant thought or effort. These signals may be learned through instruction (e.g., turn signals) or from regular exposure (e.g., fire alarms). Leveraging similar signaling mechanisms can help robot nonverbal signals to be learned and adopted more rapidly.

A problem with this approach is that it still requires a large number of distinct signals to indicate all of the features of a robot's internal state. Signal modulation can help to overcome this challenge by altering the properties of a signal to slightly alter its meaning. For instance, error signals are often modulated in intensity or tempo to convey varying amounts of urgency. This can reduce the total number of signals to learn and can increase human efficiency in reacting to those signals.

We can also aim to minimize the number of signals humans are exposed to, especially in cases where they are unfamiliar with the robot or are unlikely to regularly interact. For instance, a robot that is deployed in the home may utilize only the most important signals to prevent humans from becoming overwhelmed. As the human becomes more efficient at decoding the robot's signals, the robot can increase its signaling vocabulary. Alternatively, if a service robot employed in a public setting (e.g., restaurant) is experiencing a fault, such as low battery failure, the robot can delay signaling this information until a human worker is available to provide assistance.

7.1.3 Complexity

As human-robot interactions become more frequent and complex, signals will have to convey increasing amounts of information without becoming significantly more difficult to understand and process. Abstract communication modalities employed by non-humanoid robots must often balance limitations in their information capacity with signal efficiency. For instance, a common signal utilized in automation and robotics is a binary light indicator that can be turned on or off. In order to convey more information about the robot, multiple lights, each representing a component of its internal state, can be utilized. However, this is a suboptimal method for fast and fluid interaction as humans will likely require conscious thought, effort, and time to perceive and interpret multiple indicators.

Designing signals that are rich and can convey more complex information is challenging for a number reasons. Many available signaling modalities are coarse and lack distinctiveness. In cases where there is a significant amount of information to communicate, such as the robot's detailed plans, it may be more effective to utilize text through visual displays or speech. However, in other cases, nonverbal signals can relay information significantly faster (e.g., low battery alert). Text also requires certain perceptual conditions (e.g., distance from source, ambient noise) which can be detrimental in many scenarios.

One approach to increasing signal expressiveness is to combine signals across several modalities. Multimodal signals are often utilized in human and machine communication and can provide redundancy (e.g., traffic signal changing light position and color). However, these signals require additional processing, planning, and effort to coordinate and can be difficult to utilize for robots that possess only a few communication channels.

Alternatively, the robot can utilize a set of signals to convey information in a sequential manner. If the individual signals employed are concise and easy to recognize, this may actually result in greater efficiency over one signal that contains the entirety of the communication. For instance, humans often use this approach to initially draw attention to themselves (e.g., arm waving, gaze) before communicating their message. Without the first signal, the receiver may not actually notice the sender's second communication. Robots can also utilize this approach to deal with humans of varying experience or knowledge. For instance, the robot may prefer to communicate simpler information to start (e.g., going to start moving) if it is unsure of the human's ability to decode more complex messages (e.g., moving to the left).

7.1.4 Generalizability

Another common problem with signal design is generalizability as signals are often tied to a robot's specific form and capabilities. Designing specialized signals is not only inefficient but can prevent common standards from emerging or being adhered to. Although interactions involving different non-humanoid robots vary greatly, adopting common design standards can facilitate more rapid integration of robots into regular human environments.

Signals that follow set standards also facilitate initial interactions. Since the presence of non-humanoid robots in the world is still limited, humans are likely to rely on prior experiences with non-humanoid robots when encountering new platforms. For many other applications, signaling standards have enabled humans to easily interact with a range of machine types (e.g., different car models) without trouble.

In some cases, standards emerge due to their frequent usage in everyday products (e.g., touchscreen motions). In others, organizations and companies will set out to create these rules (e.g., ISO standards from International Organization for Standardization, Apple UI guidelines, and Google Material Design Guidelines) to facilitate more rapid standardization. This leads to an important question for HRI of how such standards should be developed and agreed upon.

Generalizability also requires the ability to replicate signals across robot platforms. This can either be accomplished by employing similar signaling modalities across robots or finding methods to generate analogous signals on different modalities. The ability to translate signals across modalities can facilitate the design of large sets of signals with limited human input. Common design standards can aid in this goal by narrowing the design space and providing insights that will push researchers towards a more cohesive signaling vocabulary.

7.1.5 External Factors

When deciding how and when to communicate, a robot should take into account information about the world state, including the environment and human interactors. This context is important for creating efficient and intelligent communication and requires that the robot have some awareness of the world around it. Using these observations, the robot can determine its actions, such as what information is important to signal to humans, what types of signals will be optimal for the desired feedback from humans, and when to signal.

Information about the environment is especially important for signaling due to the subjective nature of human perception of visual and auditory stimuli. For certain signaling modalities, such as light and sound, environmental conditions, both current and past, play a huge role in how a human perceives the stimuli. The field of psychophysics has characterized many of these relationships, but HRI often fails to leverage these phenomena for better communication.

For instance, the robot should consider ambient noise levels when deciding whether to utilize an auditory signal. If the environment is very loud, the robot's normal signal may be masked. The robot can attempt to modulate the signal's intensity or pitch, using psychophysics principles, to make the sound discernible over the ambient noise. However, this solution introduces new issues. The modulated signal may cause humans annoyance or surpass reasonable noise limits. Alternatively, the robot can utilize a different modality or wait until the sound levels have lowered.

Choosing the optimal action for a particular scenario may require the robot to not only model the environment but humans as well. Modeling humans enables the robot to better consider their needs when signaling. In the previous example, the robot considers the human's comfort when choosing whether to utilize a modulated signal. The robot can also decide if the human is busy, to delay its signal until a more appropriate time rather than interrupting them and causing annoyance.

Although this work has primarily discussed signaling in the context of one human and one robot, we can also consider when there are multiple humans in the environment. The robot can choose one human to signal to based on some criteria (e.g., distance or most helpful) or can broadcast its signal more widely. Signaling multiple humans can be helpful in cases where the robot requires a response immediately or the information it is signaling pertains to more than one person.

The ability to take in information about the world and balance the human's and robot's needs when signaling is an important component of intelligent communication. We expect that as robots increase in their roles and responsibilities around humans, these skills will become increasingly critical to the success of HRI.

7.1.6 Perceptions

One of the primary research questions of HRI concerns how humans perceive robots and their behavior. In the previous section, we discussed perceptions relating to a human's ability to identify and interpret information from their physical senses (e.g., visual system). In this section, we concern ourselves with mental perceptions or our impressions of other agents, including humans and robots.

These mental perceptions affect human behavior and outcomes when interacting with a robot [Bartneck et al., 2009]. Understanding the factors affecting how humans perceive and react to robots enables expectation setting, the development of more accurate mental models of the robot, and the robot to act more effectively [Goodrich and Olsen, 2003, Paepcke and Takayama, 2010, Powers and Kiesler, 2006]. Past work in HRI has explored how attributes of the robot (e.g., appearance) and the robot's behavior (e.g., apologizing) affect human perceptions and attitudes towards the robot [Eyssel et al., 2012, Goetz et al., 2003, Lee et al., 2010, 2014].

Since robots with more anthropomorphic features are often viewed and treated quite differently than machine-like platforms, this is an important area of research for non-humanoid robots. If the robot fails to match human expectations, users can become disappointed and unwilling to interact with a robot [Beer et al., 2011, Cha et al., 2015, Goetz et al., 2003]. Researchers have advocated for matching the robot's appearance, capabilities, and behavior to avoid this issue [Goetz et al., 2003].

When creating an agent or character, designers often think about a wide range of features. These features include the character's appearance, how they speak, and their body language, behavior, and backstory. This promotes consistency in the character's personality, appearance, and actions. However, task capability is often the greater concern when designing and building robots. A systematic approach that incorporates information about the robot's interactions can enable consistency in humans' perceptions and mental models of the robot, which can aid acceptance [Beer et al., 2011].

7.1.7 Exposure

As robots become an everyday product in human lives, we must also consider how interactions between humans and robots will evolve over time. Prior work has already shown that human behavior and perceptions shift as humans continue to interact with a robot [Huttenrauch and Eklundh, Leite et al., 2013]. We can envision changes that arise from increased duration and frequency of interactions.

Humans that are exposed to a robot infrequently will have more difficulty learning a large set of nonverbal signals. Conversely, humans that interact with the robot more often can learn a larger set of signals and may be able to better utilize them. If interactions are infrequent, however, these effects may be lessened.

Ideally, increased exposure to the robot results in more accurate mental models of the robot, however, this may not always be the case [Dragan and Srinivasa, 2014, Koenig et al., 2010]. Once a human becomes familiarized and/or habituated with the robot, they may become overconfident in the robot's capabilities or forget that the robot can act in unpredictable ways.

These conditions also affect how humans perceive and react to a robot's nonverbal signals. Habituation to nonverbal signals may cause signal fatigue, reducing human responsiveness to the robot's communication. This phenomenon has already been shown to be an issue with

286

other machines that emit nonverbal signals or alarms, such as medical systems Meredith and Edworthy [1995].

To prevent these issues, robots should utilize signals intelligently. Using important signals sparingly can help to prevent humans ignoring the signal over time. Moreover, signals can also be designed to better fit the situation, enabling the robot to have a diverse and rich vocabulary. The signaling criteria mentioned in Section 5.1 can be used to define different levels of human awareness to a signal and to identify the importance of the robot's encoded communication. Modulation of a basic signal can also convey varying levels of certain properties, such as urgency, and can be a powerful tool for robot communication. Altering the robot's signals in subtle but noticeable ways may also create similar versions of the same signal which can prevent repetitiveness in the robot's actions.

Information about the environment and the observer are also critical to creating intelligent robot behavior. For instance, to prevent overexposure, a nonverbal signal with a set duration can be ended early if a human observer starts responding to the robot. Consideration of ambient conditions also enable the robot to utilize less salient signals.

7.1.8 Evaluation

Evaluating how robot signaling behaviors affect human perceptions and consequently, interactions with the robot is an essential step in designing signals and signaling algorithms in order to verify their effectiveness and efficiency. Since many nonverbal behaviors are utilized quickly and unconsciously, it can be difficult to create a scenario where the use of such behaviors is natural and their effects can be easily perceived. This problem extends beyond signaling behaviors to all of HRI as designing user studies to evaluate HRI in realistic settings has been a challenge for many years.

As robots are complex machines still being developed, researchers often have trouble creating uniform experiences for user study participants. Thus, they often rely on Wizard-of-Oz studies that utilize a human to perform some portion of the perception, planning, and control. While this minimizes the chance of error or inconsistencies, these studies are often criticized for their lack of realism. This issue is also exacerbated by the scenarios employed in most studies, since they are often simplified and involve a small number of short interactions.

Lastly, tools and metrics for measuring the success of these interactions are limited. In cases where the robot and human have to coordinate their actions, overall time and fluidity can be measured, but many other scenarios depend on surveys or other measures of self reporting. These methods are imperfect and often have difficulty capturing the complexities of interaction. Without accurate tools to measure and evaluate interactions, it is difficult to gather and validate feedback on new signaling behaviors.

7.2 Open Research Areas

Our review has shown that while the field continues to grow, ongoing and past research often focuses on specific scenarios (i.e., platform, application, use case). However, as Section 3 and Section 4 also showed, many non-humanoid robots and applications share common signaling needs. Many of the limitations discussed in the previous section also suggest that a more unified approach is needed to advance the field.

To aid in this goal, we present three open research areas in this field. We also discuss potential approaches that take into account the considerations from the previous section. As shown by the diversity of literature surveyed in this work, many different research fields can provide significant insights. Thus, we also aim to inspire future research to leverage these results to further advance signaling by non-humanoid robots.

7.2.1 Non-Humanoid Robot Design

Since the variability of non-humanoid robots makes it difficult to create signals that apply to many different robots, we can also consider signaling needs when designing new robot platforms. One approach is to utilize modalities which can be form-agnostic. This means that modalities should be able to be added to platforms easily without significant modification or interference to the robot's main capabilities. Some communication channels that have been used in past works and can potentially fit this description include light, sound, visual displays, projectors, non-manipulating arms, and haptics [Baraka et al., 2015, Cha and Matarić, 2016, Szafir et al., 2015].

Modular versions of these signaling channels can also be designed to be easily added to and removed from the robot in order to make the robot more versatile. Modular versions will need to have low power and weight requirements to be utilized across many platforms. Current platforms often contain certain modalities, such as light, sound, of visual displays, which can benefit from modularity in order to minimize the effects of platform variations since placement and/or shape can have large effects on signal design.

Certain robots contain high degree-of-freedom manipulators or other expressive modalities that may be difficult to add or replicate on other platforms. While non-manipulating, expressive arms have been suggested for robots that involve frequent communication, such as telepresence platforms, the design and creation of such mechanisms is more costly than many other signaling channels due to movement capabilities. Moreover, these arms have few degrees-of-freedom, which limit the number of gestures that can be utilized. They also require greater power and modification to the robot than many of the other modalities listed above. For some non-humanoid robots, such as autonomous vehicles, proper placement of arms may be difficult due to visibility, safety, and regulatory concerns.

Modalities already integrated and utilized by robots such as sound, light, and visual displays have fewer barriers to usage and have significant related work to provide insight. Sound requires little additional hardware and can have high noticeability. It has also been well studied and utilized across many applications (e.g., cinema, alarms, appliances, computing notifications). However, it is often coarse in interpretation, making it difficult to signal precise information without having users memorize precise sounds. Sound is also easily blocked out or distorted by environmental noise and obstacles. Thus, manipulating sound to be more perceivable or combining sound with other modalities (e.g. fire alarm) can help to overcome these obstacles. Lights and visual displays are often included on robots as there is less chance of perceptual errors. Visual displays are a dynamic and easy way of displaying large amounts of information through text or images. Displays are also sometimes utilized to act as the "face" of the robot. The popularity and ready availability of smartphones and tablets has increased their inclusion on robots as they offer cost-effective computing with a display. Lights are simpler and typically convey less information but are inexpensive and have been utilized frequently enough to create some common user perceptions. Lights can also mimic many of the properties of motion; combining multiple lights, colors, and levels of brightness can produce many combinations of light behaviors. Projectors are similar to visual displays but enable a larger viewing area and potentially greater interaction with the environment. Typically, projectors require more hardware than a visual display, and are limited in use to certain light conditions, making them less versatile.

Currently, few standards exist for non-humanoid robot design of functional platforms. We envision that as research in nonverbal signal design advances, these insights will be integrated into future iterations of robot platforms. This knowledge can also facilitate the inclusion of standardized features or signaling hardware across different types of non-humanoid robots, which may provide other benefits, such as meeting regulatory or safety requirements.

7.2.2 Nonverbal Signal Design

Much of the work discussed in this review centered on the challenges associated with designing nonverbal signals. The "semantic" problem of communication, or how precisely the robot's signals convey the desired meaning, has been the focus of the majority of HRI research in nonverbal signaling [Weaver, 1953]. Designing a rich vocabulary of signals which are intuitive, efficient, and applicable to non-humanoid robots presents many challenges, as discussed in the prior section.

Previously, we touched on the approach of taking inspiration from common user experiences with signals to overcome the lack of intuitiveness of non-anthropomorphic signaling modalities. Many everyday machines already employ a limited number of signals (e.g., light signals

290

employed by cars, auditory alerts from computers) that can be used to inform the design of robot nonverbal signals. However, robots likely require a significantly larger number of signals than appliances (e.g., toaster) or other simple machinery, necessitating greater exploration.

More complex devices, such as computers and cell phones, typically rely on a visual display that can show pictures, text, or other more descriptive elements. Such displays are present on certain types of non-humanoid robot (e.g., telepresence or service robots) but are often limited in size (due to cost) and require a human to be in close proximity to view the signal. In addition, text, diagrams, or other complex signals can take significant time and effort to interpret. Functional robots are designed to be in motion and hence, require more efficient and rapid signals.

As the complexity of robots' internal state grows, it is also becoming less feasible for a designer to create the entirety of signals a robot will need to utilize. A possible approach for incorporating design expertise while creating a wider range of signals is to have a designer create a basic set of signals and modulate them to fit different scenarios.

Signal attributes can also be tied to the signal characteristics outlined in Section 5 for more thoughtful and intuitive modulation. For instance, light color or signal pattern frequency (e.g., beeping rate) can be modulated to indicate different levels of criticality or urgency [Cha and Matarić, 2016]. This enables the robot to utilize a larger vocabulary of signals without requiring humans to individually learn each new signal. Instead, humans can learn the meaning of the modulation (e.g., increasing intensity), which can be more intuitive. Signal modulation can also be utilized to account for different and/or changing scenarios and environments, an approach we will further discuss in the next section.

By combining modalities, signals can increase their effectiveness. To create effective multimodal signals, each modality must be manipulated in a manner that maintains the coherence of the signal. One approach is coupling attributes of each modality (e.g., light color and auditory pitch). Alternatively, signals can be translated between signaling modalities. Music visualization systems typically use Fourier transforms of the sound source to generate imagery from the music. Multimodal signals can also be utilized sequentially to provide greater amounts of information.

Nonverbal signal design for non-humanoid robots still requires significant research. Previously, we laid out important issues that affect nonverbal signaling with many pertaining to signal design. As new applications and platforms emerge, these issues will become more relevant. Future research may find useful insights and references in this review, which can facilitate work towards the development of signaling standards that can help to overcome these issues.

7.2.3 Signaling Algorithms

As the environments and interactions of non-humanoid robots become more complex, there is a need for more intelligent, adaptive robot behaviors. To meet this need, robots will be required to model the world around them and reason about how to signal in order to be effective and efficient communicators. The development of signaling algorithms, which can take the state of the world, such as environmental conditions (i.e., perceptual noise) and information about robot's task and co-located humans as input and compute signaling behaviors or policies is a major goal for HRI.

Going beyond simply sensing and reacting to items and people in the environment is challenging. One of the driving motivations for doing so is safety. Humans are dynamic and can quickly change the state of the environment. Anticipating their future actions prevents collisions and improves coordination [Jarrassé et al., 2008]. The field of human activity recognition utilizes a variety of techniques to predict human motion, actions, and interactions. For a more thorough review of prediction techniques, readers can refer to Aggarwal and Ryoo [2011].

Other works specifically explore modeling a human's state to help a robot to plan its own actions based on how the human will interpret them. During this process, a robot can also integrate information about perceptual conditions [Trafton et al., 2005b]. This has been done with visual perspective-taking for obstructions in the field of view but can also be expanded to account for other perceptual noise such as poor

7.2. Open Research Areas

lighting conditions, ambient noise, and other visual distractors [Martinson and Brock, 2007, Takayama and Harris, 2013].

Integrating many types of information about the environment, including individual predictions over the human's actions, can be used to generate more optimal signals. The most basic approach should simply account for human perceptual limitations and noise. For instance, in a dark environment a brighter light should be used or a softer sound when the robot is close to the human. This can be implemented as either loose design guidelines or within a more complex model for determining robot behavior. A simpler approach is to generate a set of basic signals that can be modulated for use in a wider variety of situations.

As non-humanoid robots operate in a diverse set of environments, signal modulation can help account for the many factors that affect signal perception. An important approach we mentioned in the previous section is signal modulation or manipulation. For instance, an auditory signal can be modulated to amplify the volume or alter its pitch to stand out in a noisy environment (e.g., factory or restaurant). Moreover if the signaling robot is far away from a human observer, it can choose to alter a signal (e.g., faster blinking light) to increase its salience. Similarly, a robot can account for other environmental factors such as visibility, directionality, number of humans and can even alter the signal dynamically to account for changes in the world state.

As more information about the world is incorporated, robot models for behavior can become more complex and less transparent. If inferences about the human's behavior and actions (e.g., motion prediction) are incorrect, the robot's signaling behaviors may cause confusion or other problems. Hence, without a good idea of how humans will perceive the robot's actions it is difficult to generate more reactive, dynamic behaviors.

Modeling human internal state is a growing area of research in HRI. These models can be difficult to infer so recent research has explored active information gathering to clarify human internal state [Dragan, 2017]. Other methods seek to discover underlying reward functions that drive human behavior. Strong assumptions, such as the human and robot optimizing the same reward function during collaboration, can facilitate estimation of portions of the human's internal state but are often unrealistic of actual humans and their behavior.

Another limitation is that several works assume a specific "type" of human (e.g., novice user) and don't necessarily account for differences in experience, training, or capability. Factory workers who encounter the same robots everyday are likely to have a richer mental model of the robot than a one-time interactor. Accounting for those differences when the robot communicates is an important component of effective communication, as it enables the robot to generate messages with the correct level of detail and duration [Kiesler, 2005]. Moreover, since nonverbal signals often require certain perceptual capabilities, not considering those differences can cause the robot to fail at communicating (e.g., light signal to a blind human).

Integrating computational models of cognition into robot communication planning can create more rich and dynamic interactions with humans. However, such methods also add complexity and are time consuming and difficult to solve. As interactions become more rich and diverse, generating real-time robot actions will only become more difficult.

This research area can also benefit from applying algorithms or methods from related fields, such as artificial intelligence, machine learning, game theory, and human-computer interaction. Researchers in these areas have successfully explored similar problems within other applications or contexts. Leveraging these results can help to advance computational HRI in this area more rapidly.

One area in particular that we can draw from is attention-sensitive systems in HCI. The main goal of a signaling algorithm is effectively mediating the flow of communication from the robot to the human. Hence, if the human is engaged in a high cognitive task or does not want to be interrupted, the robot should consider this when deciding when and how to signal. Attention-sensitive alerting systems attempt to provide information that has the greatest utility to the human user; this approach also assumes that the human's attention is a valuable but limited resource [Horvitz et al., 1999]. This results in alerts that

7.2. Open Research Areas

are not transmitted due to the high cost of the signal (on the human's attentional focus) and the low utility of the information.

Despite taking in information about co-located humans, current approaches in HRI tend to use static policies for transmitting information. They also tend to assume that there is no cost to signaling to the human or that the utility of a communication is always the same. Since robot internal state contains many features, deciding which information to convey and when can be done using a similar approach as attention-sensitive alerting. One of the main challenges to applying these algorithms is the uncertainty of the inferred human state, since robots do not have direct insights into the human's activities (e.g., keypresses, open programs).

There have also been several recent works that formulate interaction as a game in order to plan robot actions with models of cognition [Dragan, 2017]. This approach enables different approximations for human behavior that have been shown to have various advantages. Using game theoretic methods can also facilitate the inclusion of multiple interactors and different types of relationships between the interactors (e.g., cooperative vs. adversarial).

While the majority of this review has been discussed in the context of one robot and one human communicating, this assumption is not realistic of the actual scenarios robots will operate in. In environments with many humans, the robot will have to consider which humans to signal to and how to select its signal to account for multiple observers in the environment. In cases where the robot needs action from any human, this can also be advantageous for obtaining the best human response.

Other more design-based methods can also be applied to areas of HRI, such as signaling [Johnson, 2014]. However, these methods may be difficult to translate to implementation as they often function as guidelines. Hence, more work is needed to understand how to incorporate the results of existing HRI works into robot planning algorithms.

8

Conclusion

In this work, we presented a survey of literature in nonverbal signaling for non-humanoid robots with the goal of identifying current open problems. Since there are a growing number of robots designed to work with and alongside humans, it is necessary to consider the problem of how these robots can effectively communicate information about their internal state to increase legibility and facilitate communication, coordination, and collaboration.

There are two fundamental challenges for nonverbal signaling that we attempt to address in this work. The first is creating intuitive and effective signals that humans quickly and effectively react and respond to [Breazeal et al., 2005]. As non-humanoid robots span a large range of physical forms, environments, and applications, it is important to consider how nonverbal signals can be generated in a manner that allows for broad usage while maintaining intuitiveness. The current approach of utilizing human designers to create specialized signals for every platform requires extensive time and effort and is not a feasible solution as the number of robots continues to grow.

We also need to consider how nonverbal signals can be utilized in an intelligent and effective manner. Many current platforms utilize simplistic algorithms for signaling that do not take into account information about the environment or potential human observers when communicating. These signaling policies can result in the annoyance of human observers, signal fatigue, and limitations in the robot's usability. Thus, another fundamental research challenge is developing signaling algorithms which enable the robot to act as intelligent and thoughtful agents during interaction.

Towards this goal, we identified some common signaling characteristics to specify the desired human reaction and response to a robot communication signal. We also presented several use cases of current non-humanoid robots to provide the necessary background and context for non-humanoid robots signaling needs. With this information, we identified internal robot states that are important for greater transparency and reviewed literature on signal generation techniques. Lastly, we identified major challenges in the development of this research area and provided some design guidelines based on the prior sections and previous research.

Research in human interactions with non-humanoid robots is growing, but there is still a lack of standards and methods for designing robot communication behaviors. This is compounded by the fact that non-humanoid robots continue to advance in their development to new forms and applications. This work aims to provide a review of relevant knowledge and structure to help inform researchers in their future efforts in designing nonverbal signals for non-humanoid robots.

Acknowledgements

This work was supported by a grant from the NASA Space Technology Research Fellowship (NSTRF) program.

- Human-Robot Interaction Design: Understanding User Needs and Requirements, author=Adams, Julie A, booktitle=The Human Factors and Ergonomics Society Annual Meeting, volume=49, number=3, pages=447-451, year=2005, organization=SAGE Publications.
- Leopoldo Acosta, Evelio González, José Natán Rodríguez, Alberto F Hamilton, et al. Design and implementation of a service robot for a restaurant. *International Journal of Robotics & Automation*, 21(4):273, 2006.
- Julie A Adams, Scott A DeLoach, and Matthias Scheutz. Shared Mental Models for Human-Robot Teams. In AAAI Fall Symposium, 2014.
- Henny Admoni. Nonverbal Communication in Socially Assistive Human-Robot Interaction. AI Matters, 2(4):9–10, 2016.
- Henny Admoni, Caroline Bank, Joshua Tan, Mariya Toneva, and Brian Scassellati. Robot gaze does not reflexively cue human attention. In *CogSci*, pages 1983–1988, 2011.
- Henny Admoni, Anca Dragan, Siddhartha S Srinivasa, and Brian Scassellati. Deliberate Delays During Robot-to-Human Handovers Improve Compliance With Gaze Communication. In ACM/IEEE International Conference on Human-Robot Interaction, pages 49–56. ACM, 2014.
- Jake K Aggarwal and Michael S Ryoo. Human Activity Analysis: A Review. ACM Computing Surveys, 43(3):16, 2011.
- Pankaj Aggarwal and Ann L McGill. Is That Car Smiling at Me? Schema Congruity as a Basis for Evaluating Anthropomorphized Products. *Journal* of Consumer Research, 34(4):468–479, 2007.

- Ho Seok Ahn, I-Han Kuo, Chandan Datta, Rebecca Stafford, Ngaire Kerse, Kathy Peri, Elizabeth Broadbent, and Bruce A MacDonald. Design of a Kiosk Type Healthcare Robot System for Older People in Private and Public Places. In *International Conference on Simulation, Modeling, and Programming for Autonomous Robots*, pages 578–589. Springer, 2014.
- Samer Al Moubayed and Gabriel Skantze. Perception of Gaze Direction for Situated Interaction. In Workshop on Eye Gaze in Intelligent Human Machine Interaction, page 3. ACM, 2012.
- Rachid Alami, Aurélie Clodic, Vincent Montreuil, Emrah Akin Sisbot, and Raja Chatila. Task planning for Human-Robot Interaction. In *Joint Conference on Smart Objects and Ambient Intelligence*, pages 81–85. ACM, 2005.
- Pease Allan. Body Language. How to Read Others' Thoughts by Their Gestures. London: Sheldon Press, 1988.
- José Javier Anaya, Pierre Merdrignac, Oyunchimeg Shagdar, Fawzi Nashashibi, and José E Naranjo. Vehicle to Pedestrian Communications for Protection of Vulnerable Road Users. In *IEEE Intelligent Vehicles Sympo*sium, pages 1037–1042. IEEE, 2014.
- Sean Andrist, Tomislav Pejsa, Bilge Mutlu, and Michael Gleicher. Designing Effective Gaze Mechanisms for Virtual Agents. In ACM Conference on Human Factors in Computing Systems, pages 705–714. ACM, 2012.
- 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.
- Dane Archer. Unspoken Diversity: Cultural Differences in Gestures. Qualitative Sociology, 20(1):79–105, 1997.
- Michael Argyle. Non-verbal communication in human social interaction, chapter 9. Cambridge U. Press, 1972.
- Michael Argyle and Mark Cook. *Gaze and Mutual Gaze*. Cambridge U Press, 1976.
- Michael Argyle and Janet Dean. Eye-Contact, Distance and Affiliation. Sociometry, 28(3):289–304, 1965.
- Kai Oliver Arras and Daniela Cerqui. Do we want to share our lives and bodies with robots? a 2000 people survey. Technical Report 0605-001, Ecole Polytechnique Federale de Lausanne (EPFL), 2005.
- Kim Baraka, Ana Paiva, and Manuela Veloso. Expressive Lights for Revealing Mobile Service Robot State. In *Robot 2015: Second Iberian Robotics Conference*, pages 107–119. Springer, 2015.

- Simon Baron-Cohen, Alan M Leslie, and Uta Frith. Does the autistic child have a "theory of mind"? *Cognition*, 21(1):37–46, 1985.
- Christoph Bartneck. From Fiction to Science–A cultural reflection of social robots. In ACM Conference on Human Factors in Computing Systems: Workshop on Shaping Human-Robot Interaction, pages 1–4, 2004.
- Christoph Bartneck and Jodi Forlizzi. A Design-Centred Framework for Social Human-Robot Interaction. In *IEEE International Symposium on Robot and Human Interactive Communication*, pages 591–594. IEEE, 2004.
- Christoph Bartneck, Tomohiro Suzuki, Takayuki Kanda, and Tatsuya Nomura. The influence of peopleâĂŹs culture and prior experiences with Aibo on their attitude towards robots. Ai & Society, 21(1-2):217–230, 2007.
- Christoph Bartneck, Dana Kulić, Elizabeth Croft, 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.
- C Daniel Batson, Shannon Early, and Giovanni Salvarani. Perspective Taking: Imagining How Another Feels Versus Imaging How You Would Feel. *Personality and Social Psychology Bulletin*, 23(7):751–758, 1997.
- Andrea Bauer, Dirk Wollherr, and Martin Buss. Human-Robot Collaboration: A Survey. International Journal of Humanoid Robotics, 5(01):47–66, 2008.
- Geoffrey Beattie. Visible Thought: The New Psychology of Body Language. Psychology Press, 2004.
- Jenay M Beer and Leila Takayama. Mobile Remote Presence Systems for Older Adults: Acceptance, Benefits, and Concerns. In ACM/IEEE International Conference on Human-Robot Interaction, pages 19–26. ACM, 2011.
- Jenay M Beer, Akanksha Prakash, Tracy L Mitzner, and Wendy A Rogers. Understanding Robot Acceptance. Technical Report HFA-TR-1103, Georgia Institute of Technology, 2011.
- Matt Berlin, Jesse Gray, Andrea Lockerd Thomaz, and Cynthia Breazeal. Perspective Taking: An Organizing Principle for Learning in Human-Robot Interaction. In AAAI, volume 2, pages 1444–1450, 2006.
- 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, 38(1):83–92, 2008.
- John Billingsley, Arto Visala, and Mark Dunn. Robotics in agriculture and forestry. In *Springer Handbook of Robotics*, pages 1065–1077. Springer, 2008.

- Lance Bloom, Rachel Eardley, Erik Geelhoed, Meera Manahan, and Parthasarathy Ranganathan. Investigating the Relationship Between Battery Life and User Acceptance of Dynamic, Energy-Aware Interfaces on Handhelds. In *Mobile Human-Computer Interaction*, pages 43–79. Springer, 2004.
- Christopher Bodden, Daniel Rakita, Bilge Mutlu, and Michael Gleicher. Evaluating Intent-Expressive Robot Arm Motion. In *IEEE International Symposium on Robot and Human Interactive Communication*, pages 658–663. IEEE, 2016.
- Cynthia Breazeal and J Velasquez. Robot in Society: Friend or Appliance? In Autonomous Agents Workshop on Emotion-Based Agent Architectures, pages 18–26, 1999.
- Cynthia Breazeal, Andrew Brooks, David Chilongo, Jesse Gray, Guy Hoffman, Cory Kidd, Hans Lee, Jeff Lieberman, and Andrea Lockerd. Working collaboratively with humanoid robots. In *IEEE-RAS International Conference on Humanoid Robots*, volume 1, pages 253–272. IEEE, 2004.
- 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*, pages 708–713. IEEE, 2005.
- Cynthia Breazeal, Matt Berlin, Andrew Brooks, Jesse Gray, and Andrea L Thomaz. Using perspective taking to learn from ambiguous demonstrations. *Robotics and Autonomous Systems*, 54(5):385–393, 2006.
- Adrian Broadhurst, Simon Baker, and Takeo Kanade. A prediction and planning framework for road safety analysis, obstacle avoidance and driver information. Technical Report CMU-RI-TR-04-11, The Robotics Institute, Carnegie Mellon University, 02 2004.
- Andrew G Brooks and Cynthia Breazeal. Working with Robots and Objects: Revisiting Deictic Reference for Achieving Spatial Common Ground. In ACM/IEEE International Conference on Human-Robot Interaction, pages 297–304. ACM, 2006.
- Maria Bualat, Jonathan Barlow, Terrence Fong, Christopher Provencher, Trey Smith, and Allison Zuniga. Astrobee: Developing a free-flying robot for the International Space Station. In *AIAA SPACE Conference and Exposition*, 2015.
- Randy L Buckner and Daniel C Carroll. Self-projection and the brain. Trends in Cognitive Sciences, 11(2):49–57, 2007.

- Thomas Bugnyar, Mareike Stöwe, and Bernd Heinrich. Ravens, corvus corax, follow gaze direction of humans around obstacles. *Proceedings of the Royal Society of London B: Biological Sciences*, 271(1546):1331–1336, 2004.
- Benjamin Buren, Stefan Uddenberg, and Brian J Scholl. The automaticity of perceiving animacy: Goal-directed motion in simple shapes influences visuomotor behavior even when task-irrelevant. *Psychonomic Bulletin & Review*, 23(3):797–802, 2016.
- A Cesetti, CP Scotti, G Di Buo, and S Longhi. A Service Oriented Architecture Supporting an Autonomous Mobile Robot for Industrial Applications. In *Mediterranean Conference on Control & Automation*, pages 604–609. IEEE, 2010.
- Elizabeth Cha and Maja Matarić. Using Nonverbal Signals to Request Help During Human-Robot Collaboration. In *IEEE/RSJ International Confer*ence on Intelligent Robots and Systems, pages 5070–5076. IEEE, 2016.
- Elizabeth Cha, Anca D Dragan, and Siddhartha S Srinivasa. Perceived Robot Capability. In *IEEE International Symposium on Robot and Human Interactive Communication*, pages 541–548. IEEE, 2015.
- Ravi Teja Chadalavada. Human robot interaction for autonomous systems in industrial environments. Master's thesis, Chalmers University of Technology, 2016.
- Ravi Teja Chadalavada, Henrik Andreasson, Robert Krug, and Achim J Lilienthal. That's on my Mind! Robot to Human Intention Communication through on-board Projection on Shared Floor Space. In European Conference on Mobile Robots, pages 1–6. IEEE, 2015.
- Vijay Chidambaram, Yueh-Hsuan Chiang, and Bilge Mutlu. Designing Persuasive Robots: How Robots Might Persuade People Using Vocal and Nonverbal Cues. In ACM/IEEE International Conference on Human-Robot Interaction, pages 293–300. ACM, 2012.
- Jung Ju Choi, Yunkyung Kim, and Sonya S Kwak. The Less You Look the More You Get: Designing Gaze Avoidance in Human-Robot Interaction. Archives of Design Research, 26(2):71–83, 2013.
- Michael Clamann, Miles Aubert, and Mary L Cummings. Evaluation of Vehicle-to-Pedestrian Communication Displays for Autonomous Vehicles. In *Transportation Research Board Annual Meeting*. Transportation Research Board, 2017.
- Mark Coeckelbergh. Humans, Animals, and Robots: A Phenomenological Approach to Human-Robot Relations. *International Journal of Social Robotics*, 3(2):197–204, 2011.

- Brian Coltin, Joydeep Biswas, Dean Pomerleau, and Manuela Veloso. Effective Semi-autonomous Telepresence. In *Robot Soccer World Cup*, pages 365–376. Springer, 2011.
- Raffaello D'Andrea. A revolution in the warehouse: A retrospective on kiva systems and the grand challenges ahead. *IEEE Transactions on Automation Science and Engineering*, 9(4):638–639, 2012.
- Chandan Datta, Anuj Kapuria, and Ritukar Vijay. A Pilot Study to Understand Requirements of a Shopping Mall Robot. In ACM/IEEE International Conference on Human-Robot Interaction, pages 127–128. IEEE, 2011.
- S Datta, R Ray, and D Banerji. Development of autonomous mobile robot with manipulator for manufacturing environment. *The International Journal of Advanced Manufacturing Technology*, 38(5-6):536–542, 2008.
- Kerstin Dautenhahn, Sarah Woods, Christina Kaouri, Michael L Walters, Kheng Lee Koay, and Iain Werry. What is a Robot Companion âĂŞ Friend, Assistant or Butler? In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1192–1197. IEEE, 2005.
- Munjal Desai. *Modeling trust to improve human-robot interaction*. PhD thesis, University of Massachusetts Lowell, 2012.
- Munjal Desai, Katherine M Tsui, Holly A Yanco, and Chris Uhlik. Essential Features of Telepresence Robots. In *IEEE Conference on Technologies for Practical Robot Applications*, pages 15–20. IEEE, 2011.
- Winand H Dittrich and Stephen EG Lea. Visual perception of intentional motion. Perception, 23:253–253, 1994.
- Winand H Dittrich, Tom Troscianko, Stephen EG Lea, and Dawn Morgan. Perception of emotion from dynamic point-light displays represented in dance. *Perception*, 25(6):727–738, 1996.
- Karen R Dobkins and Thomas D Albright. What Happens if it Changes Color when it Moves?: Psychophysical Experiments on the Nature of Chromatic Input to Motion Detectors. *Vision Research*, 33(8):1019–1036, 1993.
- Anca Dragan and Siddhartha Srinivasa. Generating Legible Motion. In *Robotics: Science and Systems*, 2013.
- Anca Dragan and Siddhartha Srinivasa. Familiarization to Robot Motion. In ACM/IEEE International Conference on Human-Robot Interaction, pages 366–373. ACM, 2014.
- Anca D Dragan. Legible Robot Motion Planning. PhD thesis, The Robotics Institute, Carnegie Mellon University, 2015.

- Anca D. Dragan. Robot Planning with Mathematical Models of Human State and Action. CoRR, abs/1705.04226, 2017.
- Anca D Dragan, Kenton CT Lee, and Siddhartha S Srinivasa. Legibility and Predictability of Robot Motion. In ACM/IEEE International Conference on Human-Robot Interaction, pages 301–308. IEEE, 2013.
- Anca D Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S Srinivasa. Effects of Robot Motion on Human-Robot Collaboration. In ACM/IEEE International Conference on Human-Robot Interaction, pages 51–58. ACM, 2015.
- Kurt Dresner and Peter Stone. Sharing the Road: Autonomous Vehicles Meet Human Drivers. In International Joint Conference on Artificial Intelligence, pages 1263–1268, 2007.
- Jill L Drury, Jean Scholtz, and Holly A Yanco. Awareness in Human-Robot Interactions. In *IEEE International Conference on Systems, Man and Cy*bernetics, volume 1, pages 912–918. IEEE, 2003.
- Brian R Duffy. Anthropomorphism and the social robot. Robotics and Autonomous Systems, 42(3):177–190, 2003.
- Judy Edworthy, Sarah Loxley, and Ian Dennis. Improving Auditory Warning Design: Relationship between Warning Sound Parameters and Perceived Urgency. *Human factors*, 33(2):205–231, 1991.
- David Efron. Gesture, race and culture. Mouton de Gruyter, 1972.
- Naomi Eilan. Joint attention: Communication and other minds: Issues in philosophy and psychology. Oxford University Press, 2005.
- Paul Ekman. Emotions revealed: Recognizing faces and feelings to improve communication and emotional life. Macmillan, 2007.
- Mica R Endsley. Design and evaluation for situation awareness enhancement. In *The Human Factors Society Annual Meeting*, pages 97–101. SAGE Publications, 1988.
- Mica R Endsley. Toward a Theory of Situation Awareness in Dynamic Systems. Human Factors, 37(1):32–64, 1995.
- Mica R Endsley. Designing for Situation Awareness: An Approach to User-Centered Design. CRC press, 2016.
- Nicholas Epley, Adam Waytz, and John T Cacioppo. On Seeing Human: A Three-Factor Theory of Anthropomorphism. *Psychological Review*, 114(4): 864, 2007.

- 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 ACM/IEEE International Conference on Human-Robot Interaction Late Breaking Report, pages 125–126. ACM, 2012.
- Gerhard Fischer. User modeling in human-computer interaction. User modeling and user-adapted interaction, 11(1):65–86, 2001.
- Kerstin Fischer, B Soto, C Pantofaru, and Leila Takayama. Initiating Interactions in Order to Get Help: Effects of Social Framing on People's Responses to Robots' Requests for Assistance. In *IEEE International Symposium* on Robot and Human Interactive Communication, pages 999–1005. IEEE, 2014.
- Martin E. Fishbein and Icek Ajzen. Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research, 1975. http://people.umass.edu/aizen/f&a1975.html.
- Terrence Fong, Illah Nourbakhsh, and Kerstin Dautenhahn. A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42(3):143–166, 2003a.
- Terrence Fong, Charles Thorpe, and Charles Baur. Collaboration, Dialogue, Human-Robot Interaction. In International Symposium on Robotics Research, pages 255–266. Springer, 2003b.
- Terrence Fong, Charles Thorpe, and Charles Baur. Robot, asker of questions. Robotics and Autonomous systems, 42(3):235–243, 2003c.
- Terrence Fong, Jennifer Rochlis Zumbado, Nancy Currie, Andrew Mishkin, and David L Akin. Space Telerobotics Unique Challenges to Human–Robot Collaboration in Space. *Reviews of Human Factors and Ergonomics*, 9(1): 6–56, 2013.
- Jodi Forlizzi and Carl DiSalvo. Service Robots in the Domestic Environment: A Study of the Roomba Vacuum in the Home. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 258–265. ACM, 2006.
- Jutta Fortmann, Tim Claudius Stratmann, Susanne Boll, Benjamin Poppinga, and Wilko Heuten. Make Me Move at Work! An Ambient Light Display to Increase Physical Activity. In *International Conference on Pervasive Computing Technologies for Healthcare*, pages 274–277. Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, 2013.
- Chris Frith and Uta Frith. Theory of mind. *Current Biology*, 15(17):644–645, 2005.

- Hugo Fuks, Alberto Raposo, Marco A Gerosa, Mariano Pimentel, and Carlos JP Lucena. The 3C Collaboration Model. The Encyclopedia of E-Collaboration, pages 637–644, 2007.
- Kotaro Funakoshi, Kazuki Kobayashi, Mikio Nakano, Seiji Yamada, Yasuhiko Kitamura, and Hiroshi Tsujino. Smoothing Human-robot Speech Interactions by Using a Blinking-Light as Subtle Expression. In *International Conference on Multimodal Interfaces*, pages 293–296. ACM, 2008.
- Susan R Fussell, Sara Kiesler, Leslie D Setlock, and Victoria Yew. How People Anthropomorphize Robots. In ACM/IEEE International Conference on Human-Robot Interaction, pages 145–152. IEEE, 2008.
- Maia Garau, Mel Slater, Vinoba Vinayagamoorthy, Andrea Brogni, Anthony Steed, and M Angela Sasse. The Impact of Avatar Realism and Eye Gaze Control on Perceived Quality of Communication in a Shared Immersive Virtual Environment. In ACM Conference on Human Factors in Computing Systems, pages 529–536. ACM, 2003.
- Dedre Gentner and Albert L Stevens. Mental Models. Psychology Press, 2014.
- Dylan F Glas, Kanae Wada, Masahiro Shiomi, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. Personal Service: A Robot that Greets People Individually Based on Observed Behavior Patterns. In ACM/IEEE International Conference on Human-Robot Interaction, pages 129–130. IEEE, 2013.
- Michael Gleicher. Retargetting Motion to New Characters. In Conference on Computer Graphics and Interactive Techniques, pages 33–42. ACM, 1998.
- Jennifer Goetz, Sara Kiesler, and Aaron Powers. Matching Robot Appearance and Behavior to Tasks to Improve Human-Robot Cooperation. In *IEEE In*ternational Symposium on Robot and Human Interactive Communication, pages 55–60. IEEE, 2003.
- Matthew Gombolay, Anna Bair, Cindy Huang, and Julie Shah. Computational design of mixed-initiative human–robot teaming that considers human factors: situational awareness, workload, and workflow preferences. *International Journal of Robotics Research*, pages 597–617, 2017.
- Michael A Goodrich and Dan R Olsen. Seven principles of efficient human robot interaction. In Systems, Man and Cybernetics, 2003. IEEE International Conference on, volume 4, pages 3942–3948. IEEE, 2003.
- Michael A Goodrich and Alan C Schultz. Human-Robot Interaction: A Survey. Foundations and Trends in Human-Computer Interaction, 1(3):203–275, 2007.

- A Guillaume, C Drake, M Rivenez, L Pellieux, and V Chastres. Perception of urgency and alarm design. In *International Conference on Auditory Display*. Elsevier, 2002.
- Pentti Haddington. Turn-Taking for Turntaking: Mobility, Time, and Action in the Sequential Organization of Junction Negotiations in Cars. *Research* on Language and Social Interaction, 43(4):372–400, 2010.
- Martin Hägele, Walter Schaaf, and Evert Helms. Robot Assistants at Manual Workplaces: Effective Co-operation and Safety Aspects. In *International* Symposium on Robotics, pages 7–11, 2002.
- Patrick Haggard. Conscious intention and motor cognition. Trends in Cognitive Sciences, 9(6):290–295, 2005.
- Jaap Ham, René Bokhorst, Raymond Cuijpers, David van der Pol, and John-John Cabibihan. Making robots persuasive: the influence of combining persuasive strategies (gazing and gestures) by a storytelling robot on its persuasive power. In *International Conference on Social Robotics*, pages 71–83. Springer, 2011.
- Rebecca Hansson and Peter Ljungstrand. The Reminder Bracelet: Subtle Notification Cues for Mobile Devices. In ACM Conference on Human Factors in Computing Systems Extended Abstracts, pages 323–324. ACM, 2000.
- Kerstin Sophie Haring, Céline Mougenot, Fuminori Ono, and Katsumi Watanabe. Cultural Differences in Perception and Attitude towards Robots. International Journal of Affective Engineering, 13(3):149–157, 2014.
- Chris Harrison, John Horstman, Gary Hsieh, and Scott Hudson. Unlocking the Expressivity of Point Lights. In ACM Conference on Human Factors in Computing Systems, pages 1683–1692. ACM, 2012.
- Mary Hegarty, Matt S Canham, and Sara I Fabrikant. Thinking About the Weather: How Display Salience and Knowledge Affect Performance in a Graphic Inference Task. Journal of Experimental Psychology: Learning, Memory, and Cognition, 36(1):37, 2010.
- Clint Heyer. Human-Robot Interaction and Future Industrial Robotics Applications. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4749–4754. IEEE, 2010.
- Harold Hill and Frank E Pollick. Exaggerating temporal differences enhances recognition of individuals from point light displays. *Psychological Science*, 11(3):223–228, 2000.
- Pamela J Hinds, Teresa L Roberts, and Hank Jones. Whose Job Is It Anyway? A Study of Human-Robot Interaction in a Collaborative Task. *Human-Computer Interaction*, 19(1):151–181, 2004.

- Neil G Hockstein, CG Gourin, RA Faust, and David J Terris. A history of robots: from science fiction to surgical robotics. *Journal of robotic surgery*, 1(2):113–118, 2007.
- Guy Hoffman and Cynthia Breazeal. Effects of Anticipatory Action on Human-Robot Teamwork Efficiency, Fluency, and Perception of Team. In *ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 2007.
- Rachel M Holladay, Anca D Dragan, and Siddhartha S Srinivasa. Legible Robot Pointing. In *IEEE International Symposium on Robot and Human Interactive Communication*, pages 217–223. IEEE, 2014.
- Eric Horvitz, Andy Jacobs, and David Hovel. Attention-Sensitive Alerting. In Conference on Uncertainty in Artificial Intelligence, pages 305–313. Morgan Kaufmann Publishers Inc., 1999.
- Chien-Ming Huang and Bilge Mutlu. Modeling and Evaluating Narrative Gestures for Humanlike Robots. In *Robotics: Science and Systems*, pages 57–64, 2013a.
- Chien-Ming Huang and Bilge Mutlu. The Repertoire of Robot Behavior: Enabling Robots to Achieve Interaction Goals through Social Behavior. Journal of Human-Robot Interaction, 2(2):80–102, 2013b.
- Helge Huttenrauch and K Severinson Eklundh. Fetch-and-carry with CERO: observations from a long-term user study with a service robot. In *IEEE International Symposium on Robot and Human Interactive Communication*.
- Mads Hvilshøj, Simon Bøgh, Ole Madsen, and Morten Kristiansen. The Mobile Robot "Little Helper": Concepts, ideas and working principles. In *IEEE Conference on Emerging Technologies & Factory Automation*, pages 1–4. IEEE, 2009.
- Mads Hvilshøj, Simon Bøgh, Oluf Skov Nielsen, and Ole Madsen. Autonomous industrial mobile manipulation (AIMM): past, present and future. *Industrial Robot: An International Journal*, 39(2):120–135, 2012.
- Michita Imai, Takayuki Kanda, Testuo Ono, Hiroshi Ishiguro, and Kenji Mase. Robot mediated round table: Analysis of the effect of robot's gaze. In *IEEE International Symposium on Robot and Human Interactive Communica*tion, pages 411–416. IEEE, 2002.
- Stephen Intille. Change Blind Information Display for Ubiquitous Computing Environments. In *Ubiquitous Computing*, pages 193–222. Springer, 2002.
- Hiroshi Ishii and Brygg Ullmer. Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms. In ACM Conference on Human Factors in Computing Systems, pages 234–241. ACM, 1997.

- Nathanaël Jarrassé, Jamie Paik, Viviane Pasqui, and Guillaume Morel. How can human motion prediction increase transparency? In *IEEE International Conference on Robotics and Automation*, pages 2134–2139. IEEE, 2008.
- Matt J Johnson. Coactive Design: Designing Support for Interdependence in Human-Robot Teamwork. PhD thesis, TU Delft, 2014.
- Steven Johnson, Irene Rae, Bilge Mutlu, and Leila Takayama. Can You See Me Now? How Field of View Affects Collaboration in Robotic Telepresence. In ACM Conference on Human Factors in Computing Systems, pages 2397– 2406. ACM, 2015.
- Philip N Johnson-Laird. Mental models: Towards a cognitive science of language, inference, and consciousness. Harvard University Press, 1983.
- Takayuki Kanda, Masahiro Shiomi, Zenta Miyashita, Hiroshi Ishiguro, and Norihiro Hagita. An Affective Guide Robot in a Shopping Mall. In ACM/IEEE International Conference on Human-Robot Interaction, pages 173–180. IEEE, 2009.
- Esko Keskinen, Hiro Ota, and Ari Katila. Older Drivers Fail in Intersections: Speed Discrepancies Between Older and Younger Male Drivers. Accident Analysis & Prevention, 30(3):323–330, 1998.
- Oussama Khatib, Kazu Yokoi, Oliver Brock, K Chang, and Arancha Casal. Robots in Human Environments: Basic Autonomous Capabilities. *Interna*tional Journal of Robotics Research, 18(7):684–696, 1999.
- Sara Kiesler. Fostering Common Ground in Human-Robot Interaction. In IEEE International Symposium on Robot and Human Interactive Communication, pages 729–734. IEEE, 2005.
- Ralf Kittmann, Tim Fröhlich, Johannes Schäfer, Ulrich Reiser, Florian Weißhardt, and Andreas Haug. Let me introduce myself: I am care-o-bot 4, a gentleman robot. In *Mensch Und Computer*, pages 223–232. De Gruyter Oldenbourg, 2015.
- Gary Klein, Paul J Feltovich, Jeffrey M Bradshaw, and David D Woods. Common ground and coordination in joint activity. Organizational simulation, 53:139–184, 2005.
- Gary Klein, Brian Moon, and Robert R Hoffman. Making Sense of Sensemaking 1: Alternative Perspectives. *IEEE Intelligent Systems*, 21(4):70–73, 2006.
- Chris L Kleinke. Gaze and Eye Contact: a Research Review. Psychological Bulletin, 100(1):78, 1986.
- Mark L Knapp, Judith A Hall, and Terrence G Horgan. Nonverbal Communication in Human Interaction. Cengage Learning, 2013.

- Heather Knight and Reid Simmons. Expressive Motion with X, Y and Theta: Laban Effort Features for Mobile Robots. In *IEEE International Symposium on Robot and Human Interactive Communication*, pages 267–273. IEEE, 2014.
- Günther Knoblich, Stephen Butterfill, and Natalie Sebanz. Psychological research on joint action: theory and data. *Psychology of Learning and Motivation-Advances in Research and Theory*, 54:59, 2011.
- Kazuki Kobayashi, Kotaro Funakoshi, Seiji Yamada, Mikio Nakano, Takanori Komatsu, and Yasunori Saito. Blinking Light Patterns as Artificial Subtle Expressions in Human-Robot Speech Interaction. In *IEEE International* Symposium on Robot and Human Interactive Communication, pages 181– 186. IEEE, 2011.
- Nathan Koenig, Leila Takayama, and Maja Matarić. Communication and knowledge sharing in human–robot interaction and learning from demonstration. *Neural Networks*, 23(8):1104–1112, 2010.
- 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*, 9(4):269– 275, 2015.
- Annica Kristoffersson, Silvia Coradeschi, and Amy Loutfi. A Review of Mobile Robotic Telepresence. Advances in Human-Computer Interaction, 2013:3, 2013.
- Thibault Kruse, Patrizia Basili, Stefan Glasauer, and Alexandra Kirsch. Legible Robot Navigation in the Proximity of Moving Humans. In *IEEE Workshop on Advanced Robotics and its Social Impacts*, pages 83–88. IEEE, 2012.
- 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.
- Thomas Lamarre. From animation to anime: drawing movements and moving drawings. In *Japan Forum*, volume 14, pages 329–367. Taylor & Francis, 2002.
- Hyewon Lee, Jung Ju Choi, and Sonya S Kwak. Will You Follow the Robot's Advice? The Impact of Robot Types and Task Types on People's Perception of a Robot. In *International Conference on Human-Agent Interaction*, pages 137–140. ACM, 2014.

- Min Kyung Lee and Leila Takayama. Now, I Have a Body: Uses and Social Norms for MobileRemote Presence in the Workplace. In *ACM Conference on Human Factors in Computing Systems*, pages 33–42. ACM, 2011.
- Min Kyung Lee, Sara Kielser, Jodi Forlizzi, Siddhartha Srinivasa, and Paul Rybski. Gracefully Mitigating Breakdowns in Robotic Services. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 203–210. IEEE, 2010.
- Min Kyung Lee, Karen P Tang, Jodi Forlizzi, and Sara Kiesler. Understanding Users! Perception of Privacy in Human-Robot Interaction. In ACM/IEEE International Conference on Human-Robot Interaction Late Breaking Report, pages 181–182. ACM, 2011.
- 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.
- C Liang, KJ Chee, Y Zou, H Zhu, A Causo, S Vidas, T Teng, IM Chen, KH Low, and CC Cheah. Automated Robot Picking System for E-Commerce Fulfillment Warehouse Application. In *IFToMM World Congress*, 2015.
- Christina Lichtenthäler and Alexandra Kirsch. Legibility of Robot Behavior: A Literature Review. 2016.
- Patrick Lin, Keith Abney, and George A Bekey. *Robot ethics: the ethical and social implications of robotics.* MIT press, 2011.
- Tamara Lorenz, Alexander Mörtl, Björn Vlaskamp, Anna Schubö, and Sandra Hirche. Synchronization in a goal-directed task: human movement coordination with each other and robotic partners. In *IEEE International Symposium on Robot and Human Interactive Communication*, pages 198–203. IEEE, 2011.
- Holger Luczak, Matthias Roetting, and Ludger Schmidt. Let's talk: anthropomorphization as means to cope with stress of interacting with technical devices. *Ergonomics*, 46(13-14):1361–1374, 2003.
- Thorsten Luettel, Michael Himmelsbach, and Hans-Joachim Wuensche. Autonomous Ground Vehicles– Concepts and a Path to the Future. *Proceedings of the IEEE*, 100:1831–1839, 2012.
- 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, 2009.

- Nikolas Martelaro, Victoria C Nneji, Wendy Ju, and Pamela Hinds. Tell Me More Designing HRI to Encourage More Trust, Disclosure, and Companionship. In ACM/IEEE International Conference on Human-Robot Interaction, pages 181–188. IEEE, 2016.
- Eric Martinson and Derek Brock. Improving Human-Robot Interaction through Adaptation to the Auditory Scene. In ACM/IEEE International Conference on Human-Robot Interaction, pages 113–120. ACM, 2007.
- George Mather and Linda Murdoch. Gender discrimination in biological motion displays based on dynamic cues. Proceedings of the Royal Society of London B: Biological Sciences, 258(1353):273–279, 1994.
- Tara Matthews, Tye Rattenbury, Scott Carter, Anind Dey, and Jennifer Mankoff. A Peripheral Display Toolkit. Technical Report UCB/CSD-03-1258, EECS Department, University of California, Berkeley, 2003.
- Andrii Matviienko, Andreas Löcken, Abdallah El Ali, Wilko Heuten, and Susanne Boll. NaviLight: investigating ambient light displays for turn-by-turn navigation in cars. In *International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 283–294. ACM, 2016.
- D Scott McCrickard, Christa M Chewar, Jacob P Somervell, and Ali Ndiwalana. A Model for Notification Systems Evaluation-Assessing User Goals for Multitasking Activity. ACM Transactions on Computer-Human Interaction, 10(4):312–338, 2003.
- David McNeill. Hand and mind: What gestures reveal about thought. University of Chicago press, 1992.
- David McNeill. *Language and gesture*, volume 2. Cambridge University Press, 2000.
- David McNeill. Gesture and thought. University of Chicago press, 2008.
- Ross Mead and Maja J Matarić. Autonomous human-robot proxemics: socially aware navigation based on interaction potential. *Autonomous Robots*, 41(5):1189–1201, 2017.
- Albert Mehrabian. Nonverbal communication. Transaction Publishers, 1972.
- Christina Meredith and Judy Edworthy. Are there too many alarms in the intensive care unit? an overview of the problems. *Journal of advanced nursing*, 21(1):15–20, 1995.
- Pascale Michelon and Jeffrey M Zacks. Two kinds of visual perspective taking. Attention, Perception, & Psychophysics, 68(2):327–337, 2006.
- John A Michon. A critical view of driver behavior models: What do we know, what should we do. *Human behavior and traffic safety*, pages 485–520, 1985.

- Nicole Mirnig, Nicole Perterer, Gerald Stollnberger, and Manfred Tscheligi. Three Strategies for Autonomous Car-to-Pedestrian Communication: A Survival Guide. In ACM/IEEE International Conference on Human-Robot Interaction Late Breaking Report, pages 209–210. ACM, 2017.
- Peter Mitchell. Introduction to theory of mind: Children, autism and apes. Edward Arnold Publishers, 1997.
- Katja Mombaur, Anh Truong, and Jean-Paul Laumond. From human to humanoid locomotion-an inverse optimal control approach. Autonomous Robots, 28(3):369–383, 2010.
- Vincent Montreuil, Aurélie Clodic, Maxime Ransan, and Rachid Alami. Planning Human Centered Robot Activities. In *IEEE International Conference* on Systems, Man and Cybernetics, pages 2618–2623. IEEE, 2007.
- Louis-Philippe Morency. Modeling Human Communication Dynamics. *IEEE Signal Processing Magazine*, 27(5):112–116, 2010.
- C David Mortensen. Communication Theory. Routledge, 2017.
- Naoki Mukawa, Atsushi Fukayama, Takehiko Ohno, Minako Sawaki, and Norihiro Hagita. Gaze Communication between Human and Anthropomorphic Agent-Its concept and examples. In *IEEE International Symposium on Robot and Human Interactive Communication*, pages 366–370. IEEE, 2001.
- Bilge Mutlu, Jodi Forlizzi, and Jessica Hodgins. A Storytelling Robot: Modeling and Evaluation of Human-like Gaze Behavior. In *IEEE-RAS International Conference on Humanoid Robots*, pages 518–523. IEEE, 2006.
- 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 ACM/IEEE International Conference on Human-Robot Interaction, pages 61–68. ACM, 2009a.
- Bilge Mutlu, Fumitaka Yamaoka, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. Nonverbal Leakage in Robots: Communication of Intentions through Seemingly Unintentional Behavior. In ACM/IEEE International Conference on Human-Robot Interaction, pages 69–76. ACM, 2009b.
- Bilge Mutlu, Allison Terrell, and Chien-Ming Huang. Coordination Mechanisms in Human-Robot Collaboration. In ACM/IEEE International Conference on Human-Robot Interaction: Workshop on Collaborative Manipulation, 2013.
- Yasushi Nakauchi and Reid Simmons. A Social Robot that Stands in Line. Autonomous Robots, 12(3):313–324, 2002.
- Clifford Nass and Youngme Moon. Machines and mindlessness: Social responses to computers. *Journal of Social Issues*, 56(1):81–103, 2000.

- Veronica Ahumada Newhart. Virtual Inclusion via Telepresence Robots in the Classroom. In ACM Conference on Human Factors in Computing Systems Extended Abstracts, pages 951–956. ACM, 2014.
- Stefanos Nikolaidis and Julie Shah. Human-Robot Cross-Training: Computational Formulation, Modeling and Evaluation of a Human Team Training Strategy. In Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction, pages 33–40. IEEE Press, 2013.
- Donald A Norman. Some Observations on Mental Models, chapter 1, pages 7–14. Psychology Press, 1983.
- Lars Oestreicher and Kerstin Severinson Eklundh. User Expectations on Human-Robot Co-operation. In *IEEE International Symposium on Robot* and Human Interactive Communication, pages 91–96. IEEE, 2006.
- Yuko Okumura, Yasuhiro Kanakogi, Takayuki Kanda, Hiroshi Ishiguro, and Shoji Itakura. Infants understand the referential nature of human gaze but not robot gaze. *Journal of Experimental Child Psychology*, 116(1):86–95, 2013.
- Yusuke Okuno, Takayuki Kanda, Michita Imai, Hiroshi Ishiguro, and Norihiro Hagita. Providing Route Directions: Design of Robot's Utterance, Gesture, and Timing. In ACM/IEEE International Conference on Human-Robot Interaction, pages 53–60. IEEE, 2009.
- Hirotaka Osawa, Arisa Ema, Hiromitsu Hattori, Naonori Akiya, Nobutsugu Kanzaki, Akinori Kubo, Tora Koyama, and Ryutaro Ichise. What is Real Risk and Benefit on Work with Robots?: From the Analysis of a Robot Hotel. In ACM/IEEE International Conference on Human-Robot Interaction Late Breaking Report, pages 241–242. ACM, 2017.
- Steffi Paepcke and Leila Takayama. Judging a Bot By Its Cover: An Experiment on Expectation Setting for Personal Robots. In ACM/IEEE International Conference on Human-Robot Interaction, pages 45–52. IEEE, 2010.
- Elizabeth Phillips, Scott Ososky, Janna Grove, and Florian Jentsch. From Tools to Teammates: Toward the Development of Appropriate Mental Models for Intelligent Robots. In *Human Factors and Ergonomics Society Annual Meeting*, volume 55, pages 1491–1495. SAGE Publications, 2011.
- Zachary Pousman and John Stasko. A Taxonomy of Ambient Information Systems: Four Patterns of Design. In Working Conference on Advanced Visual Interfaces, pages 67–74. ACM, 2006.
- Aaron Powers and Sara Kiesler. The Advisor Robot: Tracing People's Mental Model from a Robot's Physical Attributes. In ACM/IEEE International Conference on Human-Robot Interaction, pages 218–225. ACM, 2006.
- Irene Rae, Bilge Mutlu, and Leila Takayama. Bodies in Motion: Mobility, Presence, and Task Awareness in Telepresence. In ACM Conference on Human Factors in Computing Systems, pages 2153–2162. ACM, 2014.
- Irene Rae, Bilge Mutlu, Gary M Olson, Judith S Olson, Leila Takayama, and Gina Venolia. Everyday Telepresence: Emerging Practices and Future Research Directions. In ACM Conference on Human Factors in Computing Systems Extended Abstracts, pages 2409–2412. ACM, 2015.
- Ahmad Rahmati, Angela Qian, and Lin Zhong. Understanding Human-Battery Interaction on Mobile Phones. In International Conference on Human Computer Interaction with Mobile Devices and Services, pages 265– 272. ACM, 2007.
- Robin Read and Tony Belpaeme. How to Use Non-Linguistic Utterances to Convey Emotion in Child-Robot Interaction. In ACM/IEEE International Conference on Human-Robot Interaction, pages 219–220. ACM, 2012.
- Robin Read and Tony Belpaeme. Non-Linguistic Utterances Should be Used Alongside Language, Rather than on their Own or as a Replacement. In ACM/IEEE International Conference on Human-Robot Interaction Late Breaking Report, pages 276–277. ACM, 2014a.
- Robin Read and Tony Belpaeme. Situational Context Directs How People Affectively Interpret Robotic Non-Linguistic Utterances. In *ACM/IEEE International Conference on Human-Robot Interaction*, pages 41–48. ACM, 2014b.
- Henrik Rehbinder and Martin Sanfridson. Scheduling of a limited communication channel for optimal control. *Automatica*, 40(3):491–500, 2004.
- Tiago Ribeiro and Ana Paiva. The Illusion of Robotic Life: Principles and Practices of Animation for Robots. In ACM/IEEE International Conference on Human-Robot Interaction, pages 383–390. IEEE, 2012.
- Charles Rich, Brett Ponsler, Aaron Holroyd, and Candace L Sidner. Recognizing Engagement in Human-Robot Interaction. In ACM/IEEE International Conference on Human-Robot Interaction, pages 375–382. IEEE, 2010.
- Matt Richtel and Conor Dougherty. Google's driverless cars run into problem: Cars with drivers. *New York Times*, Sep 2015.

- Laurel D Riek, Tal-Chen Rabinowitch, Bhismadev Chakrabarti, and Peter Robinson. How Anthropomorphism Affects Empathy Toward Robots. In ACM/IEEE International Conference on Human-Robot Interaction Late Breaking Report, pages 245–246. ACM, 2009.
- Jennifer M Riley, David B Kaber, and John V Draper. Situation Awareness and Attention Allocation Measures for Quantifying Telepresence Experiences in Teleoperation. Human Factors and Ergonomics in Manufacturing & Service Industries, 14(1):51–67, 2004.
- Christina Rödel, Susanne Stadler, Alexander Meschtscherjakov, and Manfred Tscheligi. Towards Autonomous Cars: The Effect of Autonomy Levels on Acceptance and User Experience. In International Conference on Automotive User Interfaces and Interactive Vehicular Applications, pages 1–8. ACM, 2014.
- Stephanie Rosenthal and Manuela M Veloso. Mobile Robot Planning to Seek Help with Spatially-Situated Tasks. In AAAI Conference on Artificial Intelligence, page 1, 2012.
- Matthew Rueben, Cindy M Grimm, Frank J Bernieri, and William D Smart. A Taxonomy of Privacy Constructs for Privacy-Sensitive Robotics. arXiv:1701.00841, 2017.
- Selma Šabanović. Inventing Japan's 'robotics culture': The repeated assembly of science, technology, and culture in social robotics. Social Studies of Science, 44(3):342–367, 2014.
- Dorsa Sadigh, Shankar Sastry, Sanjit A Seshia, and Anca D Dragan. Planning for Autonomous Cars that Leverage Effects on Human Actions. In *Robotics: Science and Systems*, 2016.
- Joan Saez-Pons, Hagen Lehmann, Dag Sverre Syrdal, and Kerstin Dautenhahn. Development of the Sociability of Non-Anthropomorphic Robot Home Companions. In *Joint IEEE International Conference on Devel*opment and Learning and Epigenetic Robotics, pages 111–116. IEEE, 2014.
- Eduardo Salas, Carolyn Prince, David P Baker, and Lisa Shrestha. Situation Awareness in Team Performance: Implications for Measurement and Training. *Human Factors*, 37(1):123–136, 1995.
- Maha Salem, Stefan Kopp, Ipke Wachsmuth, Katharina Rohlfing, and Frank Joublin. Generation and Evaluation of Communicative Robot Gesture. *International Journal of Social Robotics*, 4(2):201–217, 2012.

- Maha Salem, Micheline Ziadee, and Majd Sakr. Marhaba, how may I help you? Effects of Politeness and Culture on Robot Acceptance and Anthropomorphization. In ACM/IEEE International Conference on Human-Robot Interaction, pages 74–81. ACM, 2014.
- Aziez Sardar, Michiel Joosse, Astrid Weiss, and Vanessa Evers. Don't Stand so Close to Me: Users' Attitudinal and Behavioral Responses to Personal Space Invasion by Robots. In ACM/IEEE International Conference on Human-Robot Interaction Late Breaking Report, pages 229–230. ACM, 2012.
- Paul Saulnier, Ehud Sharlin, and Saul Greenberg. Exploring Minimal Nonverbal Interruption in HRI. In *IEEE International Symposium on Robot* and Human Interactive Communication, pages 79–86. IEEE, 2011.
- Allison Sauppé and Bilge Mutlu. Robot Deictics: How Gesture and Context Shape Referential Communication. In ACM/IEEE International Conference on Human-Robot Interaction, pages 342–349. ACM, 2014.
- Brian Scassellati. Theory of Mind for a Humanoid Robot. *Autonomous Robots*, 12(1):13–24, 2002.
- Michael F Schober. Spatial perspective-taking in conversation. *Cognition*, 47 (1):1–24, 1993.
- Jean Scholtz, Jeff Young, Jill L Drury, and Holly A Yanco. Evaluation of Human-Robot Interaction Awareness in Search and Rescue. In *IEEE In*ternational Conference on Robotics and Automation, volume 3, pages 2327– 2332. IEEE, 2004.
- M Scopelliti, MV Giuliani, AM D'amico, and F Fornara. If I had a robot at home... Peoples' representation of domestic robots, chapter 26, pages 257–266. Springer, 2004.
- Natalie Sebanz, Harold Bekkering, and Günther Knoblich. Joint action: bodies and minds moving together. *Trends in Cognitive Sciences*, 10(2):70–76, 2006.
- Atsushi Senju, Gergely Csibra, and Mark H Johnson. Understanding the referential nature of looking: Infants' preference for object-directed gaze. *Cognition*, 108(2):303–319, 2008.
- Claude E Shannon and Warren Weaver. The mathematical theory of communication. University of Illinois press, 1998.
- Amanda Sharkey and Noel Sharkey. Granny and the robots: ethical issues in robot care for the elderly. *Ethics and Information Technology*, 14(1):27–40, 2012.

- Masamichi Shimosaka, Takuhiro Kaneko, and Kentaro Nishi. Modeling Risk Anticipation and Defensive Driving on Residential Roads with Inverse Reinforcement Learning. In *IEEE International Conference on Intelligent Transportation Systems*, pages 1694–1700. IEEE, 2014.
- Masahiro Shiomi, Takayuki Kanda, Dylan F Glas, Satoru Satake, Hiroshi Ishiguro, and Norihiro Hagita. Field Trial of Networked Social Robots in a Shopping Mall. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2846–2853. IEEE, 2009.
- 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.
- Rosanne M Siino and Pamela J Hinds. Making sense of new technology as a lead-in to structuring: The case of an autonomous mobile robot. In *Academy of Management Proceedings*, volume 2004, pages 1–6. Academy of Management, 2004.
- Rosanne M Siino and Pamela J Hinds. Robots, Gender & Sensemaking: Sex Segregation's Impact On Workers Making Sense Of a Mobile Autonomous Robot. In *IEEE International Conference on Robotics and Automation*, pages 2773–2778. IEEE, 2005.
- Charles L Smith. Smoke detector operability survey: report on findings. US Consumer Product Safety Commission Bethesda, MD, 1993.
- Wing Chee So, Colin Sim Chen-Hui, and Julie Low Wei-Shan. Mnemonic effect of iconic gesture and beat gesture in adults and children: Is meaning in gesture important for memory recall? *Language and Cognitive Processes*, 27(5):665–681, 2012.
- Vasant Srinivasan and Robin R Murphy. A Survey of Social Gaze. In ACM/IEEE International Conference on Human-Robot Interaction Late Breaking Report, pages 253–254. IEEE, 2011.
- Robert Stalnaker. Common Ground. *Linguistics and philosophy*, 25(5):701–721, 2002.
- Aaron Steinfeld, Terrence Fong, David Kaber, Michael Lewis, Jean Scholtz, Alan Schultz, and Michael Goodrich. Common Metrics for Human-Robot Interaction. In ACM/IEEE International Conference on Human-Robot Interaction, pages 33–40. ACM, 2006.

- Anthony Stentz, Herman Herman, Alonzo Kelly, Eric Meyhofer, G Clark Haynes, David Stager, Brian Zajac, J Andrew Bagnell, Jordan Brindza, Christopher Dellin, et al. CHIMP, the CMU Highly Intelligent Mobile Platform. *Journal of Field Robotics*, 32(2):209–228, 2015.
- Kyle W Strabala, Min Kyung Lee, Anca D Dragan, Jodi L Forlizzi, Siddhartha Srinivasa, Maya Cakmak, and Vincenzo Micelli. Toward Seamless Human-Robot Handovers. *Journal of Human-Robot Interaction*, 2(1):112– 132, 2013.
- S Shyam Sundar, T Franklin Waddell, and Eun Hwa Jung. The Hollywood Robot Syndrome Media Effects on Older Adults' Attitudes toward Robots and Adoption Intentions. In ACM/IEEE International Conference on Human-Robot Interaction, pages 343–350. IEEE, 2016.
- Dag Sverre Syrdal, Michael L Walters, Nuno Otero, Kheng Lee Koay, and Kerstin Dautenhahn. He Knows When You Are Sleeping-Privacy and the Personal Robot Companion. In AAAI Conference on Artificial Intelligence: Workshop on Human Implications of Human-Robot Interaction, pages 28– 33, 2007.
- Daniel Szafir, Bilge Mutlu, and Terrence Fong. Communication of Intent in Assistive Free Flyers. In ACM/IEEE International Conference on Human-Robot Interaction, pages 358–365. ACM, 2014.
- Daniel Szafir, Bilge Mutlu, and Terry Fong. Communicating Directionality in Flying Robots. In ACM/IEEE International Conference on Human-Robot Interaction, pages 19–26. ACM, 2015.
- Daniel Szafir, Bilge Mutlu, and Terrence Fong. Designing planning and control interfaces to support user collaboration with flying robots. *International Journal of Robotics Research*, pages 1–29, 2017.
- Daniel J Szafir. *Human Interaction with Assistive Free-Flying Robots*. PhD thesis, The University of Wisconsin-Madison, 2015.
- Leila Takayama and Helen Harris. Presentation of (Telepresent) Self: On the Double-Edged Effects of Mirrors. In ACM/IEEE International Conference on Human-Robot Interaction, pages 381–388. IEEE, 2013.
- Leila Takayama and Caroline Pantofaru. Influences on Proxemic Behaviors in Human-Robot Interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5495–5502. IEEE, 2009.
- Leila Takayama, Doug Dooley, and Wendy Ju. Expressing Thought: Improving Robot Readability with Animation Principles. In ACM/IEEE International Conference on Human-Robot Interaction, pages 69–76. ACM, 2011.

- Gilbert Tang, George Charalambous, Philip Webb, and Sarah R Fletcher. Users' understanding of industrial robot gesture motions and effects on trust. *Contemporary Ergonomics and Human Factors*, pages 116–123, 2014.
- Kazunori Terada, Takashi Shamoto, Akira Ito, and Haiying Mei. Reactive Movements of Non-humanoid Robots Cause Intention Attribution in Humans. In *IEEE/RSJ International Conference on Intelligent Robots and* Systems, pages 3715–3720. IEEE, 2007.
- Jan Theeuwes. Abrupt luminance change pops out; abrupt color change does not. Attention, Perception, & Psychophysics, 57(5):637–644, 1995.
- Marcus Thiebaux, Brent Lance, and Stacy Marsella. Real-Time Expressive Gaze Animation for Virtual Humans. In *International Conference on Au*tonomous Agents and Multiagent Systems, pages 321–328. International Foundation for Autonomous Agents and Multiagent Systems, 2009.
- Frank Thomas, Ollie Johnston, and Frank. Thomas. The illusion of life: Disney animation. Hyperion New York, 1995.
- Andrea Thomaz, Guy Hoffman, Maya Cakmak, et al. Computational Human-Robot Interaction. Foundations and Trends in Robotics, 4(2-3):105–223, 2016.
- Martin Tomitsch, Karin Kappel, Andreas Lehner, and Thomas Grechenig. Towards a Taxonomy for Ambient Information Systems. *Ambient Information* Systems, 2007.
- 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 Transactions* on Systems, Man, and Cybernetics, 35(4):460–470, 2005a.
- J Gregory Trafton, Alan C Schultz, Magdalena Bugajska, and Farilee Mintz. Perspective-taking with Robots: Experiments and models. In *IEEE International Symposium on Robot and Human Interactive Communication*, pages 580–584. IEEE, 2005b.
- Katherine M Tsui and Holly A Yanco. Design Challenges and Guidelines for Social Interaction Using Mobile Telepresence Robots. *Reviews of Human Factors and Ergonomics*, 9(1):227–301, 2013.
- Pernilla Ulfvengren. Design of Natural Warning Sounds in Human-Machine Systems. PhD thesis, KTH Royal Institute of Technology in Stockholm, 2003.

- Vaibhav V Unhelkar, Ho Chit Siu, and Julie A Shah. Comparative Performance of Human and Mobile Robotic Assistants in Collaborative Fetchand-Deliver Tasks. In ACM/IEEE International Conference on Human-Robot Interaction, pages 82–89. ACM, 2014.
- Chris Urmson, Joshua Anhalt, Drew Bagnell, Christopher Baker, Robert Bittner, MN Clark, John Dolan, Dave Duggins, Tugrul Galatali, Chris Geyer, Michele Gittleman, Sam Harbaugh, Martial Hebert, Thomas M. Howard, Sascha Kolski, Alonzo Kelly, Maxim Likhachev, Matt McNaughton, Nick Miller, Kevin Peterson, Brian Pilnick, Raj Rajkumar, Paul Rybski, Bryan Salesky, Young-Woo Seo, Sanjiv Singh, Jarrod Snider, Anthony Stentz, William (Red) L. Whittaker, Ziv Wolkowicki, and Jason Ziglar. Autonomous Driving in Urban Environments: Boss and the Urban Challenge. Journal of Field Robotics, 25(8):425–466, June 2008.
- Kimon P Valavanis and George J Vachtsevanos. *Handbook of unmanned aerial* vehicles. Springer Publishing Company, Incorporated, 2014.
- A J N Van Breemen. Bringing Robots To Life: Applying Principles Of Animation To Robots. In ACM Conference on Human Factors in Computing Systems Extended Abstracts, pages 143–144, 2004.
- Milos Vasic and Aude Billard. Safety Issues in Human-Robot Interactions. In *IEEE International Conference on Robotics and Automation*, pages 197–204. IEEE, 2013.
- J E Veitengruber. Design Criteria for Aircraft Warning, Caution, and Advisory Alerting System. Journal of Aircraft, 15(9):574–581, 1978.
- Daniel Vogel and Ravin Balakrishnan. Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users. In ACM Symposium on User Interface Software and Technology, pages 137–146. ACM, 2004.
- Michael L Walters, Dag S Syrdal, Kerstin Dautenhahn, René Te Boekhorst, and Kheng Lee Koay. Avoiding the Uncanny Valley - Robot Appearance, Personality and Consistency of Behavior in an Attention-Seeking Home Scenario for a Robot Companion. *Autonomous Robots*, 24(2):159–178, 2008.
- Michael L Walters, Kheng Lee Koay, Dag Sverre Syrdal, Kerstin Dautenhahn, and René Te Boekhorst. Preferences and Perceptions of Robot Appearance and Embodiment in Human-Robot Interaction Trials. In *New Frontiers in Human-Robot Interaction*, 2009.
- Warren Weaver. Recent Contributions to The Mathematical Theory of Communication. ETC: A Review of General Semantics, 10(4):261–281, 1953.

- Michael S Wogalter, Vincent C Conzola, and Tonya L Smith-Jackson. Research-based guidelines for warning design and evaluation. Applied Ergonomics, 33(3):219–230, 2002.
- Peter R Wurman, Raffaello D'Andrea, and Mick Mountz. Coordinating Hundreds of Cooperative, Autonomous Vehicles in Warehouses. *AI magazine*, 29(1):9, 2008.
- Katsu Yamane, Yuka Ariki, and Jessica Hodgins. Animating non-humanoid characters with human motion data. In Proceedings of the 2010 ACM SIG-GRAPH/Eurographics Symposium on Computer Animation, pages 169– 178. Eurographics Association, 2010.
- Holly A Yanco and Jill L Drury. Classifying Human-Robot Interaction:An Updated Taxonomy. In *IEEE Transactions on Systems, Man, and Cybernetics*, pages 2841–2846, 2004.
- Yaxing Yao, Huichuan Xia, Yun Huang, and Yang Wang. Privacy Mechanisms for Drones: Perceptions of Drone Controllers and Bystanders. In ACM Conference on Human Factors in Computing Systems, pages 6777–6788. ACM, 2017.
- Yuichiro Yoshikawa, Kazuhiko Shinozawa, Hiroshi Ishiguro, Norihiro Hagita, and Takanori Miyamoto. Responsive Robot Gaze to Interaction Partner. In *Robotics: Science and Systems*, 2006.
- Eduardo Zalama, Jaime Gómez García-Bermejo, Samuel Marcos, Salvador Domínguez, Raúl Feliz, Roberto Pinillos, and Joaquín López. Sacarino, a Service Robot in a Hotel Environment. In *ROBOT2013: First Iberian Robotics Conference*, pages 3–14. Springer, 2014.
- Garth Zeglin, Aaron Walsman, Laura Herlant, Zhaodong Zheng, Yuyang Guo, Michael C Koval, Kevin Lenzo, Hui Jun Tay, Prasanna Velagapudi, Katie Correll, et al. HERB's Sure Thing: A rapid drama system for rehearsing and performing live robot theater. In *IEEE Workshop on Advanced Robotics* and its Social Impacts, pages 129–136. IEEE, 2014.
- Allan Zhou, Dylan Hadfield-Menell, Anusha Nagabandi, and Anca D Dragan. Expressive Robot Motion Timing. In ACM/IEEE International Conference on Human-Robot Interaction, pages 22–31. ACM, 2017.