Group Human-Robot Interaction: A Review

Alexander Wisowaty
Advisor Marynel Vázquez

April 22, 2019

Abstract

This thesis reviews past and current literature in the subfield of Human-Robot Interaction (HRI) dedicated to groups, Group Human-Robot Interaction (Group HRI). While several HRI reviews exist, none has been dedicated specifically to the research and challenges of Group HRI. First, HRI will be introduced, covering what the field is, what other disciplines contribute to it, and consideration of its interdisciplinary nature. Here, the main argument of the work is presented. Then, background material will be covered, looking at historical context, work from recent conferences and in relevant application areas. Two different drivers of research are considered–engineering solutions and establishing principles. Also in this section, necessary background research in social psychology pertaining to physical space is covered. Next, Group HRI will be reviewed in detail, dividing work into sections–abstraction from social phenomena, maintaining a presence, scaling up, group membership and responsibility, gaining acceptance, experimental methodologies, and verbal communication. Finally, implications of the reviewed work are considered. The argument is made that Group HRI faces a challenge from differences between social psychology and its own findings. Some features are pointed out that might account for this, and suggestions for future directions of HRI research are provided.

Contents

1 Introduction

1.1 What is HRI?

Human Robot Interaction (HRI) is a multidisciplinary field that seeks to better understand the nature of interaction between humans and machines. Admittedly, several parts of this definition are vague: 'multidisciplinary', 'interaction', and 'machines'. This delineation is so large as to include anything from getting frustrated with your Roomba for not vacuuming the kitchen floor in its entirety, to considering psychological implications of scenes from science fiction like *Ex Machina*, in which the (human) protagonist falls in love with a machine and is emotionally manipulated by it. The former example is especially of interest to technical members of HRI, who might want to know more about how the machine fails to map the layout of the user's kitchen. How can its sensors, movement mechanisms, and internal representations of surrounding space be improved? How can such a device be designed to take feedback? (Roomba, you missed this area. Roomba, please relearn the layout of my new kitchen now that I've moved.) The latter example is of interest to theorists and social scientists in the field who want to know what it means to bring robots into our lives. To what extent is companionship possible? Can robots be role models or encourage better behavior across a whole society? How has this relationship changed over time, and where might the trend be going?

HRI touches on many different types of questions whose answers and solutions stem from distinct, and historically unrelated, academic disciplines. Through the above examples, it also becomes clear how interaction can take a multitude of forms–master-servant dynamics, or purported equals in love. Similarly, robots come in many physical forms. The Roomba vacuum cleaning device and the near-perfect human replica both fall into the category of robot.

This is not to say that HRI is overreaching, or disorganized. Rather that it is a nascent field, built up relatively recently, and still in the process of self-organizing to meet these practical and theoretical questions with a wide enough scope to handle the important implications for our future. Any new technology that becomes commonplace in society ought to be understood. As many new technologies sprout up, hoping to become the next major solution to some issue in our lives; and as science fiction yields ever more floral predictions of brave new worlds, HRI serves both as the botanist, and the gardener: trying to classify and understand the new presences, while also tending to them, ensuring their development toward a greater good.

1.2 Who Contributes to the Field?

Human Robot Interaction exists at the overlap of several academic disciplines. Building a robot requires skills from computer science, electrical engineering, and design. Programming its behavior might additionally require an understanding of human cognition and psychology. Allowing the robot to interface with a human through language brings in linguistics and natural language processing. Situating the robot correctly in relation to its users draws on sociology and social psychology. Obtaining meaningful metrics about an interaction with a robot from the perspective of the human again draws on social sciences. Research in the field usually entails coordination between several domains. Given the hybrid nature of HRI, the field is younger than its integral parts. For a better understanding of the field today, it is worth reviewing briefly how it came into existence.

1.2.1 Collaboration Emerges

It was not until the end of the 20th century that Human Robot Interaction was formally instantiated as an academic field. In a broad 2007 review, Goodrich and Schultz cite 1992 as the year of the field's first academic meeting. It was around this time that HRI emerged as an independent field as researchers from various disciplines came to realize the importance of collaboration on this front. While scientists and scholars were eager to conduct research, there was no official conference for this purpose. That first meeting hosted by the Institute of Electrical and Electronics Engineers (IEEE) was named International Symposium on Robot & Human Interactive Communication (RoMan). This event recurs annually to this day. The IEEE also created a similar event in Japan together with the Robotics Society of Japan. Many other workshops, conferences, and symposia have been created since then. Today, the main HRI conference is put on jointly by the IEEE and Association for Computing Machinery and is known as the International Conference on Human-Robot Interaction (Goodrich et al., 2008).

1.2.2 Social Science as Starting Point

Roboticists recognize that it is inefficient and unnecessary to constantly reinvent the wheel every time when it comes to understanding the human behavior they might be trying to emulate or complement. HRI has not expanded to so many disciplines out of a philosophical imperative for inclusivity, but out of a real need for the knowledge of these fields.

Frequently in HRI the researchers include a section on the social scientific background of the subject matter, (Vázquez et al., 2017; Strohkorb et al., 2015; Admoni et al., 2013) and use this as a basis for their work. A substantial number of the papers encountered cite at least one work from psychology. One example makes this connection particularly explicit. Gonsior et al. 2012 studied "transfer of theories on prosocial behavior from Social Psychology to [HRI] in terms of helpfulness shown by humans towards a robot" (Gonsior et al., 2012, p. 101).

Similarly, much of the research reviewed below will pull knowledge from the social sciences to construct computational models of human behavior (Hüttenrauch et al., 2006; Hung et al., 2008; Shi et al., 2011; Yousuf et al., 2012; Strohkorb et al., 2015), formulate hypotheses (Chang et al., 2012; Fraune and Šabanović, 2014; Fraune et al., 2015b, 2017a; Wallisch et al., 2018), and borrow methodologies (Chang et al., 2012; Huttenrauch et al., 2006; Joosse et al., 2014). Social science provides a theoretical foundation and toolkit of methodologies.

1.3 Contributions of the Present Work

In this work, I conduct a systematic review of the Group HRI literature, dividing the research into themes that each pose unique research challenges. In discussion of the implications of this work, *I argue that the field faces a puzzling discontinuity between findings from social psychology and Group HRI, and suggest that some features unique to HRI research might account for this.* Proposals are made to avoid confounds in future work, and continue laying a solid empirical foundation of Group HRI.

2 Background

2.1 Historical Context

A brief survey of the histories of artificial intelligence and robotics will help to provide context for the current work undertaken by researchers in HRI. First, we will cover the intellectual history, looking at what AI and Robotics meant to early thinkers. Parallels are drawn to contemporary thought. Second, we will cover the technological history of these fields to provide context for their current state.

2.1.1 Brief Intellectual History of AI and Robotics

Robotics and AI both originate in mythology, as fantastical speculations. At this early stage, the robots and the programmed behavior that drove them were mostly conflated. Typically, artificial intelligence was not described as existing without embodiment. Both robotics and artificial intelligence make their first appearances in ancient Greek mythology in the form of Talos. According to the *Argonautica*, an epic poem by Apollonius of Rhodes, Talos was a giant automaton made of bronze who circled the island of Crete, protecting it from invasion (Rhodios, 2008).

But even in these early accounts, Talos could be seen either as an android (human-resembling robot), or a member of a living species from a past era. In *Works and Days*, Hesiod describes Talos as survivor from the "Age of Bronze Men" (Stallings et al., 2018). Talos' source of power and intelligent behavior was an internal system running on ichor, the blood of the gods. The sorceress Medea defeats Talos by confusing him telepathically, and causing the ichor to drain from a vulnerable point in Talos' ankle. These traits, a mind that can be clouded and a cardiovascular system, resemble a human's. But the same mythical figure is described as having super-human strength, hurling boulders at invading forces, and in close-quarter combat, embracing the enemy and heating its body to "roast them alive" (Mayor, 2018, p. 7). When the guardian is ultimately defeated, it lies on the ground as a heap of bronze. From these accounts we can see that a dichotomy arises about the nature of man and machine, sentience and soullessness early in the cultural history of robotics and AI.

Just as robots can take many forms today, Homer's *Iliad* describes a very different sort of robot. In book XVIII, Hephaestus is described making twenty tripods with golden wheels that could move on their own to serve the gods (Butler et al., 1999). Unlike Talos, these robots bore no resemblance to people, emphasizing again that robots could be conceived of as distinctly non-human, even as tools.

Referencing this, Aristotle later wrote in book one, part IV of his *Politics* that such creations would render slavery and subordination useless. "For if every instrument could accomplish its own work, obeying or anticipating the will of others, [...] chief workmen would not want servants, nor masters slaves," (Barnes et al., 1984, Book 1, Part IV). Already in the 4th century B.C. it was thought that the automation of industry could lead to a truly just society.

Similarly, extending into modern feminist scholarship, Donna Haraway wrote in 1985 that people's dependence on machines has the potential to do away with gender differences. What results is a post-gender society that, similar to Aristotle's vision, is free of injustices endemic to gender (Haraway, 2006).

Other thinkers have much darker predictions about robots and artificial intelligence. One of the earliest accounts of AI turning on its creator is from an ancient Burmese collection of tales. According to legend, King Ajatashatru, who reigned in 4th century B.C., had robots built by a craftsman who stole plans from the Romans. When the Buddha's remains were distributed throughout the kingdom, weaponized automata were built to protect the shrines. King Ajatashatru's descendant, King Asoka, was predestined to gather up all the remains and redistribute them further. When Asoka tries to do this, the automata turn against him, and he is not able to get to the remains until he recruits the help of the craftsman's son. Once they defeat the army of robots, Asoka acquires such an army himself and commands it.

Talos from Greek mythology, and the automata from this tale both depict robots as potential threats to humans, either by design or turning hostile toward their creators. More recently, Nick Bostrom picks up this line in his book *Superintellligence* in which he warns of the possibility of AI, if not restrained appropriately, taking over the world. Serious efforts are ongoing, backed by important figures in the technology industry in the United States, to prevent and prepare for these scenarios. However, many disagree about the reality of this risk (Bostrom, 2017; Friend, 2017; Urban, 2017).

For a long time, ideas about robots and AI have circulated in fascinating tales. Both in antiquity, and in modern scholarship, the boundary between man and machine is blurry. Its implementation brings promise of greater equality, at the risk of losing control of our autonomous creations. Similar theoretical questions have surrounded AI and robotics for as long as the idea has existed. These same theoretical questions surround current research into AI, and Robotics, as well as fields at their intersection, like HRI.

2.1.2 Brief Technological History of AI

We can divide the history of artificial intelligence into pre-computer and post-computer periods. The former period was sustained largely by myth and speculation, as well as cleverly devised mechanisms, tricks, and formal reasoning frameworks. The latter era, the one we find ourselves in today, consists of developing increasingly more powerful computers, and trying to mold their processing functionality into something resembling human thought–if not in process, at least in output.

Early artificial intelligence was typically embodied in the form of mechanized statues. There are many accounts of automatons from mythology that have already been discussed in the preceding section. Today it is difficult to determine what was purely mythology, mechanisms that truly existed as described, or some mixture–perhaps mechanisms with less advanced capabilities that took on mythical meanings.

One example of early analog computation is the south-pointing chariot, which existed in the third century A.D. The south-pointing chariot was a two-wheeled cart with a sculpted figure on top that always pointed to the south. At the beginning of a journey the pointer would be set to point south, and any subsequent differential movement of the wheels–if the left wheel turned more than the right–would adjust the position of the pointing figure (Qinggang, 2016).

Sometimes, though, the intelligent capabilities of these mechanisms were feigned, as with the Mechanical Turk. This machine was a chess playing automaton, seated at a boxy table that purportedly contained its mechanical cognition. As the Turk toured Europe, it played against human opponents including Napoleon Bonaparte and Benjamin Franklin (both were defeated). In reality, the Turk's source of intelligent behavior was a person who sat inside the table, granted, beside a great deal of complex machinery to direct the movement of the humanoid (Standage, 2002).

An important advance in artificial intelligence before the modern computer was formal reasoning. Both early computational devices and the computers of today work in fundamentally different ways from the human brain. In order for machines to emulate human thought, common ground needs to be found. The roots of formal reasoning again trace back to antiquity, and it was a philosophical subject in its own right. More recently, in the first half of the twentieth century, mathematicians and logicians sought to formalize mathematical thought.

An early example of formal thought, again drawing from Greek antiquity, is Aristotle's syllogism. This is an argument that follows a particular structure: two premises and a conclusion. For example: all men are mortal; no gods are mortal; no men are gods (Smith, 1989).

Nineteenth century British mathematician George Boole took up and expanded on Aristotelian logic. According to John Corcoran, who writes of the similarities between the two thinkers, Boole "found Aristotle's logic to be flawless as far as it went," (Corcoran, 2003, p. 270). But Boole went further, and re-framed the syllogism as mathematical corollary.

Boole went on to establish what is now called Boolean algebra, which is a branch of algebra for evaluating binary, true/false values (Boole, 1854). The Boolean algebra is fundamental to modern computer science because computation is performed with binary values. Many current programming languages have Boolean data types for representing 'True' and 'False'.

An important theoretical leap was made in the middle of the 20th century by Alonzo Church and Alan Turing. Their joint thesis, the Church–Turing thesis, posited that though not all mathematical thought could be represented as formal logic (Smullyan and Smullyan, 1992), all math that could be represented in this way, could also be implemented mechanically (Copeland, 1997). This theoretical proof of feasibility and power of mechanical computation sparked research that would lead to the creation of modern computers.

During World War II the first computers were built. Once these first functional models existed, and technology quickly developed, the idea of implementing human-like intelligence in such a machine spread. John McCarthy at Dartmouth College first established artificial intelligence as an academic pursuit in 1956, with the proposal for a conference. While that first conference did not immediately spark widespread research efforts into AI, it did put many of the foundational figures of the field in contact with each other (Garnham, 2017).

From this point onward, artificial intelligence went through the 'Golden Years', 'First AI Winter', the 'Boom', the 'Second AI Winter', and since 2011 large strides in research. Today's AI capability is especially attributed to big data, or the availability of vast sets of information, especially about people, as well as deep learning, which is an implementation of a neural network with more layers of nodes (wiki, 2019).[1]

The history of AI and robotics originates in antiquity and has been a source of not only technological innovation, but also theoretical speculation and fantasy. By understanding the long history of these fields, it becomes clearer why HRI research exists today. Not only is there great promise in technology that can interact with and assist humans, but we are also continuing a long-existing interest in replicating ourselves and our ability to think in machines.

2.2 Current Research

To get an understanding of where the field has gone since its early days, we turn now to a brief overview of active research areas and lines of inquiry. First, we will consider submissions from recent conferences to get a handle on what themes and topics are relevant to the academic community. The ACM/IEEE International Conference on Human-Robot Interaction conveniently puts research into four distinct categories that we will borrow: user studies, technical advances, design, and theory. Second, we will look at various application areas attracting work and research from the field. Some notable areas are service, personal, and educational robotics. This raises a discussion about two motivational forces in HRI.

2.2.1 User Studies

HRI Conference definition: "This theme targets research that provides data on and analysis of human-robot interaction, in a laboratory or in-the-wild setting."

A recent paper in this category is by Bršćić et al., titled "Escaping from Children's Abuse of Social Robots". This work exists within the broader theme of robot abuse (Brscić et al., 2015).

[1] A helpful source for the broader timeline covered in this section can be found **here**. (Accessed 4/19/2019)

The term ‘robot abuse’ was coined by Bartneck in 2005 (Bartneck et al., 2005) and touches on interesting questions of fairness and violence in human psychology. The Bršćić study replicated the Milgram Obedience experiment in which a participant is asked by a person who is perceived to be an authority figure to deliver increasingly stronger electric shocks to a confederate (who merely acted as though the shocks were being administered). The startling finding of that study was that 40% of people were willing to deliver the highest shock level, which was understood as life threatening. This behavior arose because the participants simply followed the instructions of the authority figure, usually without protest. The 2015 study replaced the human receiving shocks with a robot, and the participants were children. In this variation, all participants, not just 40% as in the original study–went all the way to the highest voltage. One interpretation is that people are more willing to do harm to a machine, perhaps because they perceive it as incapable of feeling pain.

Another paper by Moshkina et al., titled “Social engagement in public places: a tale of one robot from 2014”, explores the Computers Are Social Actors (CASA) paradigm–which originated in 1996. The study tests the paradigm in the real world by deploying humanoid robots in public venues. It was found that the use of social cues led to more engagement from human participants than a robot that exhibits no such behavior. Game playing did not increase interactions as predicted, though. This user study addresses questions about robot behavioral design for increased social engagement (Moshkina et al., 2014).

Both of these papers come from separate lines of inquiry in HRI, on the one hand trying to understand how to avoid destructive and violent behavior towards robots, and on the other, understanding how to attract interaction to begin with. A shared foundation can be found in “The Media Equation”, written by Byron Reeves in 1996, which speculated about the social perception of, and attitudes toward new forms of media (Reeves and Nass, 1996). Since HRI by definition includes humans, user studies have long been and will continue to be a central avenue of research.

2.2.2 Technical Advances

HRI Conference definition: “This theme targets research providing novel robot system designs, algorithms, interface technologies, and computational methods supporting human-robot interaction.”

A paper in this category from 2016 is by Yi et al., titled “Homotopy-Aware RRT*: Toward Human-Robot Topological Path-Planning”. The authors present an algorithm for planning optimal paths while navigating the physical world. Path-finding is a large area of research within HRI, and real-world implementation of ideas from computer science and graph theory (Yi et al., 2016). This research has an extensive history. Dijkstra’s Shortest Path First algorithm was first conceived in

1956 (Dijkstra, 1959).

Nikolaidis et al. also contributed to technical advances in 2015 with a paper titled "Efficient Model Learning from Joint-Action Demonstrations for Human-Robot Collaborative Tasks". In this work, the authors are addressing the topics of learning from humans for the purpose of collaboration on a task (Nikolaidis et al., 2015).

Not every advance in HRI necessitates a new robot. Though robots can differ in virtually every way, from form factor, to degrees of freedom, to human similarity, they share a common computational core. An algorithm devised for one platform, can generally be implemented in another platform equipped to operate in the same problem domain.

2.2.3 Design

HRI Conference definition: "This theme targets research that makes a design-centric contribution to human-robot interaction."

A 2016 submission to this category by Azenkot et al. is titled "Enabling Building Service Robots to Guide Blind People". Unlike the previous two categories, the substance of the paper is the result of design sessions, not experimental results or algorithm development. Nonetheless, these questions about service robotics are necessary to furthering these technologies (Azenkot et al., 2016).

A second contribution to HRI design, also from 2016, is by Johns et al., titled "Exploring Shared Control in Automated Driving". The authors of this paper ran an empirical study, putting participants into a virtual driving setting. They wanted to explore how a driver would best perceive the actions of an automated steering system providing haptic feedback instructions (Johns et al., 2016).

Robots are very broadly defined, and can take as many forms as people are able to dream up and build. Guiding this process of design is vital to ensuring that robots are built to adequately serve their purpose and encourage human interaction.

2.2.4 Theory and Methods

HRI Conference definition: "This theme targets research contributing to the understanding and study of fundamental HRI principles that span beyond individual interfaces or projects."

A 2016 paper by Sequeira et al. is titled "Discovering Social Interaction Strategies for Robots from Restricted-Perception Wizard-of-Oz Studies". This paper tries a variation on a foundational methodology in HRI, the so-called Wizard-of-Oz (WOZ) design. A WOZ paradigm means that the robot in a given study or interaction is controlled by a human, typically hidden from view. Here, the authors propose allowing the teleoperator to have only the sensory feedback that the robot takes in, not an additional third-person view like from a camera perched high above. This allows the designer of a robot to be aware of its constraints during interactions, for when it is eventually designed to be autonomous (Sequeira et al., 2016).

Malle et al. authored "Sacrifice One For the Good of Many?: People Apply Different Moral Norms to Human and Robot Agents". Here, the authors explore how "people apply moral norms to robot agents and make moral judgments about their behavior"(Malle et al., 2015). This taps into the broader issue of ethics and morals of HRI.

Given the large variety of robots in existence, and a large body of literature on human psychological phenomena, it is important to guide research and find methodologies that connect robots and humans in valuable ways. Like any academic field, self-reflection is important to ensure that lines of inquiry are empirically sound and in pursuit of significant findings.

2.3 Application Areas

For another angle on active areas of research in HRI, it is helpful to look at application areas. Fong et al. 2003 conducted a broad review of the many facets of HRI research, using real-world applications as an organizing principle. The main areas covered are service, personal, therapy, and education robotics. These application areas do not encompass all of robotics, only ones that put robots in direct, interactive contact with people.[2]

2.3.1 Therapeutic and Service Robotics

An area in which robots can be immensely helpful to people is service. Technology is not limited in many ways that people are–for example in the amount of time they can work uninterrupted, and emotional resilience. Of course, technology is limited in many ways that people are not, but ongoing projects seek to close this gap, and to provide services that complement what people can offer.

Perhaps the most actively researched area in service robotics are solutions for elder care. The introduction of vacuum robots was seen as both an amusement and a nuisance by residents of a care facility in Denmark (Hansen et al., 2010). As a more interpersonal intervention, PARO the zoomorphic seal robot has been studied as a way to soothe dementia patients since the early 2000s. Interaction with PARO can improve agitation, but it was found that providing a simple plush toy with no interactive functionality had similar positive effects, and was a better value for money (Mervin et al., 2018). Some as yet unaddressed concerns remain about the robot's hygiene and risk for spreading infection (Martyn et al., 2018). Due to the potential medical benefits to patients with

[2] At the time of the review by Fong et al., industrial and assembly-line robotics would have been excluded, but more recently, these robots are also being developed into more interactive collaborators, as opposed to mere tools.

dementia and the burden that these benefits might lift off service workers, research is continuing to address these concerns and improve these solutions.

2.3.2 Personal and Entertainment Robotics

This application area of robotics encompasses toys and novelties, as well as commercial solutions meant to be available to the public. One of the most successful robots in this area is Sony's AIBO, a dog-like electronic pet. Being the first robot of this kind, and accounting for development costs, AIBO was expensive. Other toy manufacturers recognized the popularity and subsequent competition drove down the price (Fong et al., 2003). Since the debut of AIBO, entertainment robotics have become ever cheaper and ubiquitous.

Another example is work by Kidd 2008 to develop a robotic nutrition coach. Several participants maintained a nutrition regimen–some using a computer or paper log, others interacting with the robotic coach. While the actual weight loss results were not significantly different between groups, the participants interacting with the robot maintained the regimen longer than those using traditional logs. Participants also experienced a closer sense of alliance with the robot (Kidd, 2008) The HRI implications of this work are important and further inform us about potential for companionship between humans and robots. The commercial implications are clear as well.

2.3.3 Educational Robotics

Robots have also been shown to have pedagogical value. Kanda et al. 2004 conducted the first field trial of robots as language tutors and companions to young school children. Many challenges were encountered relating to the field trial methodology (covered at length below), but they found that the robots indeed had a positive educational influence (Kanda et al., 2004).

More recently, Westlund et al. 2016 ran a similar two-month robot deployment study, this time focusing on the teachers' experiences and evaluations. Their main findings were that the expectations of teachers were generally misaligned with their actual experience. One main concern was that the robots would be a disruptive addition to the class environment, but this turned out not to be the case (Kory Westlund et al., 2016).

Conversely, the study of robots also has educational value. Robotics competitions for example invite creative innovation under tight technical constraints. For several years, at the Association for the Advancement of Artificial Intelligence (AAAI) Conference, teams competed to develop a robot that could serve hors d'oeuvres at the conference's reception. Many of the central challenges in HRI arise, like navigation and recognition of group interactions (Michaud and Gustafson, 2002). Thus robots also have value as educational subject matter, as well as educators.

2.4 Establishing Principles versus Engineering Solutions

Having looked at the research topics and application areas relevant to Human-Robot Interaction today, it is possible to identify two different motivational forces at work: the establishment of principles, and the engineering of solutions.

Establishing principles occurs mostly in academia. Here, researchers are incrementally expanding knowledge about how humans and robots interact. As we have seen above, these steps forward might take the form of new robots serving novel purposes, but could also be new algorithms or theoretical advances. The test faced by new work is the empirical method that the researchers use to attain significant findings, as well as the peer review process. The driving motivation is furthering the field itself. What results is a rigorously analyzed new piece of knowledge about HRI, a new principle.

Solutions are engineered through a different process and with different motivations. As we have seen robots have great potential for commercial value if a technology can be provided that balances cost efficiency and utility. Given an application area with enough demand, a robotic solution can be developed to try to meet it. The main test faced here is the free market. Motivations are not exclusively economic, though, as finding solutions–for example in elder care–have broader positive externalities.

Neither of these forces driving HRI work forward is better than the other. Important work has come from both the academic and the commercial sides of HRI, and much of it spans both. It is nonetheless important to recognize these forces as they produce results through different processes.

2.5 Social Space

One important facet of group interaction–whether between humans, robots, or some combination–is physical space. Here we briefly review work from social psychology on the use of space in social contexts that has been of particular importance to Human-Robot Interaction.

2.5.1 Physical Embodiment

An important question surrounding HRI research is what effect, the physical presence of a robot has on a person. This is important to address because much HRI research presumes that a participant or group of participants will be sharing physical space with a robot.

The empirical study of the effects of physical robot embodiment began with work by Wainer et al. They designed an experiment to tease apart the differences between a robot that was physically present, and one that was simulated on a screen. They also wanted to observe the effects of physical co-location with a robot. This initial work encountered some difficulty because of the confounding effect of novelty–people were interested in the robot they had physically in front of them. Also, the interactive task was a puzzle. Working on this task was not as social as anticipated, and made it difficult to gain insights. However, preliminary findings showed that physical embodiment might lead to a more positive interaction (Wainer et al., 2006).

Bainbridge et al. addressed these concerns in a follow up to the previous work. For their experiment, the participants engaged in a book-moving task with the robot, which was either on a live video feed, or in the same physical space. Participants both reported more positive interactions, and were more likely to complete a request by the robot when they were physically co-located with the agent (Bainbridge et al., 2011). From this work, we can infer that the physical embodiment of a robot is more likely to lead to positive interactions. This lays a good foundation for HRI research putting humans in physical contact with robots without moral qualms, or concerns about strong *a priori* bias.

2.5.2 Spatial Behavior

In social psychology, research in spatial behavior began with proxemics. The focus of proxemics is on “the perception, use, and structuring of space” (Harrigan et al., 2008). In humans, proxemics informs us about relationships between interactants, how best to approach a stranger to initiate interaction, territoriality, and what–beyond physical constraints–makes a space feel crowded. One result of this research is the establishment of culturally varied personal space zones that influence social interactions. A conversation between persons that are not close friends, for example, occurs at a greater distance than one between intimate lovers (Walters et al., 2009). The same is not true of inanimate objects, as investigated by Horowitz et al. 1964. People did not mind approaching a hatrack at close distances (Horowitz et al., 1964). Robots are somewhere in between other humans and inanimate objects, because they are recognized as distinctly not human (Mori et al., 2012), but can be seen as social agents (Reeves and Nass, 1996). Where, exactly, they lie is not so clear, and depends on factors including behavior, appearance, and many others.

2.5.3 Social Robot Navigation

These social forces have helped roboticists to address the problems faced when building a robot that can navigate a world full of people. At its beginnings in 1980, robot navigation was considered a path planning problem (Brooks, 1984; Lozano-Perez, 1990; Moravec, 1980). This approach however called for high-level operations to provide information to low-level movement control. Leaving navigational decisions to the high-level planning faculties resulted in slow reaction times. Khatib 1985 made major advances in the field by shifting the burden of planning to low-level computation by proposing a real-time obstacle avoidance algorithm that modeled the surrounding space as a field of forces (Khatib, 1985).

More recent work in this area like Luber et al. 2012 employs modern machine learning approaches, trained on human interaction datasets from surveillance cameras. The change in methodology from observing human interaction behavior in lab settings allows for data that arose spontaneously in the real world, and is in much greater quantity. Luber et al. break from traditional proxemics work to create their model, that shortened path length, and resulted in smoother paths that did not adhere to proxemics so strictly (Luber et al., 2012). The problem of social robot navigation remains an active area of research.

2.5.4 F-formations

An important conceptual framework used in HRI is the F-formation. This is a description of naturally-occurring arrangements when individuals interact with one another. Perhaps without realizing, we have all observed and participated in an F-formation. For example, during a reception or cocktail party, groups of people naturally form, persist, adjust, and ultimately dissolve. What people commonly observe is several circles of interacting people, distributed across the space. In 1990, Adam Kendon set out to define these phenomena formally in his book, "Conducting interaction: Patterns of behavior in focused encounters". Kendon posits that an "F-formation arises whenever two or more people sustain a spatial and orientational relationship in which the space between them is one to which they have equal, direct, and exclusive access," (Kendon, 1990, p. 209) The behaviors that arise naturally to sustain such a formation are called the F-formation system.

The F-formation has a few components. First, the transactional segment is that area in which a single person acts. For example, while painting, the transactional segment encompasses the easel or work surface in front of the painter. If there is a model or still life, the transactional segment also includes the space between the painter and the subject. The transactional segment can be understood as the area a passerby should not enter, as it disrupts the activity. When the activity is a conversation between two people, there is overlap between transactional segments. Regardless of whether they are facing each other, sitting side-by-side, or in an L-arrangement, there is a space considered to be shared between them. Kendon calls this overlap the o-space. The o-space is not limited to dyadic interactions, and scales to larger ones. As an example, if eight people sit equally spaced around a circular table, the o-space is roughly congruent with the surface of the table. Importantly, though, the table is not necessary to define o-space, neither does it need to be circular–though this configuration often arises because it enables participants all to see each other. Just outside of o-space is the space that the participants physically occupy. This Kendon refers to as the p-zone. Outside of this zone, behind the backs of participants is a buffer known as r-space. This area arises as a recognition of the integrity of the F-formation. As such, an outsider cannot easily join the formation–by entering the p-zone–until recognized by a participant and brought in from r-space. An easy way to conceptualize this framework is to think of the F-formation as a biological cell, where the p-zone represents the membrane and the o-space the interior.

The F-formation system was originally introduced in social psychology as a "means of defining a social encounter as a unit of analysis." With this framework, it is possible to understand chaotic interactions between many actors as discrete formations whose tendencies to form, change or disband can be quantified. Says Kendon, "any investigation into the relationship between social behavior and physical space will find it necessary to consider this organization."(Kendon, 1990, p. 210) Understanding the interaction between humans and robots is such an investigation, and this conceptual tool lends itself very well to the task.

One of the many challenges robots face when deployed into the world is how to inhabit physical space. As we have seen, these considerations of physical space are steeped in psychological forces that emerge from people's use of space. Recognizing these forces and understanding them are a crucial prerequisite for group HRI research.

2.6 Conclusion: Background

Already, it is clear that Human Robot-Interaction is a broad field whose research and applications permeate several aspects of our lives, from whimsy to medical intervention. Historically, we can trace the origins of the field back to ancient Greek and Indian mythology. Research motivations can be understood as an extension of a longstanding interest in designing ever more sophisticated and capable machines. Today, this interest is expressed both in the academy, where principles of interaction are being established, as well as in industry, where commercial solutions face the test of real world interactions and cost-benefit analysis. Just as HRI has roots in history, it also has roots in various academic fields, especially social psychology, which HRI draws from for a foundational understanding of social space. Moving forward, we turn to work specifically in Group Human-Robot Interaction.

3 Group HRI

Human social interaction is chaotic. When people join together, they communicate along a variety of channels: physical arrangement, speech, gesture, even eye gaze. As covered in section 2.5.4, in a cocktail party environment people will arrange into groups, that can be seen from above as several circles, but those arrangements constantly change. Sometimes, what is said can be taken at face value. Other times, irony or intonation reveal crucial information about true sentiment. People generally tend to take turns speaking, but sometimes the beginnings and ends of statements overlap. Sometimes, entire statements. Physical gestures complement verbal communication and are important for understanding the nature and content of an interaction. Similarly, eye gaze can transmit information between interactants as well. In order for a robot to enter into these situations effectively, it must have awareness of these modes of communication. Where and how to begin making sense of all this are important considerations.

Much of the time, when people communicate in these subtle ways, they do so subconsciously. We do not always feel like we make intentional decisions who we set our gaze on, or exactly when we begin speaking. We probably don't think all too carefully about how we orient our lower body versus our head as the focus of attention moves between members of the group. The fact that we are spared from making these decisions consciously is under-appreciated. Especially once robots are developed to interact with groups of people, and all of these subtleties become the obstacles that determine whether a particular robot succeeds or fails in its interactions with people.

3.1 Why study groups in HRI?

As we continue into Group Human-Robot Interaction specifically, let us take a moment to assess why it makes any sense to consider this subfield of HRI on its own. First, drawing a parallel to one of the foundational fields of HRI–Psychology–it becomes clear, at least theoretically, how this divide is warranted. In the study of human behavior, many effects exist that hinge upon our inclusion or exclusion from a group. Examples include feeling emotions on behalf of persons we understand as in our group (Yzerbyt et al., 2003), neurological evidence for physical pain when rejected from a group (Eisenberger and Lieberman, 2004), and higher levels of aggression between groups of people versus between individuals (Meier and Hinsz, 2004). There is, of course, an extensive body of work beyond this sample. Psychologically, there are different principles at play in group interactions versus ones between two individuals.

Since psychology is an important component of HRI, and often an intellectual guide or starting point for HRI research, it is important to consider group interactions with robots as possibly subject to similar effects. As Selma Sabanovic (Šabanović, 2010) argues in the paper that we look at in detail below, trained social scientists should be among those who evaluate robots designed to interact in groups. In other words, HRI should take Psychology and other social sciences into account as progress is made toward more seamless interactions. Psychology says dyads and groups differ, so HRI should consider studying this dichotomy as well.

Another motivation for treating group interaction as a distinct area of HRI is due to obstacles encountered in early observations of robots left to fend for themselves in real-world environments. In 2006, Sabanovic et al. completed an observational study of two robots, GRACE and Roboceptionist, that were tasked with interacting with any persons they came in contact with. One purpose of the study was to perform a field trial, sacrificing strict control of experimental conditions in order to face the robots off against the challenges that exist outside the lab. Sabanovic argues that it is important to run studies like this, taking into account corner cases that the designers of the robot might not have been able to anticipate in the isolation of the lab. Notably, they observed that people who interacted with both of these robots, often did so in groups. Neither GRACE, nor Roboceptionist were designed to handle these cases, in part because it was not expected that people who wanted to interact would make an effort to do so in groups. GRACE's sensory technology only allowed for recognition of one person at a time. The Roboceptionist on the other hand, was able to sense a single main subject of interaction as well as the surrounding observers, but its user interface (a keyboard) allowed for interaction with only one person at a time. The study concludes that GRACE and Roboceptionist would have both benefited from having a better read on the surrounding crowd dynamics and the ability to switch attention between interactors–instead of conflating two into a single person (Sabanovic et al., 2006).

In a 2004 study, Kanda et al. encountered similar difficulties with group interaction after sending robots into the wild. They took on an especially tall order by testing the efficacy of two Robovie robots as English tutors for Japanese primary schoolers. During two trials, the researchers placed two robots in the common area outside the classrooms of first and sixth grade students. These stayed for two weeks as children were free to interact with them and review English phrases. Experimenters administered a language test before the trial, after the first week, and after the second week. For some students, learning with the robot was productive, but generally only if they had some basic knowledge of the language to begin with. Once again, the issue of group interaction arose. Sixty-three percent of first graders interacted with the robot in groups of two or more, and seventy-two percent of sixth graders. The researchers acknowledged as a flaw the fact that social contextual information was not taken into account by the robot. Like the Roboceptionist, Robovie had the capability of sensing a main interactor and surrounding observers, but similarly had the limitation of interacting only one at a time (Kanda et al., 2004). Additionally, understanding and using information about the relationships between the children would have allowed for richer interactions.

Both Sabanovic's and Kanda's groups emphasized the importance of field trials for making strides in Human-Robot Interaction. They sent their robots into the real world, and observed the results. For the present review, it is noteworthy that some of the biggest limitations both projects encountered had to do with interactions involving multiple people. Even if a technology is intended for interaction with a single user–like a receptionist scenario, or an English tutor–users will come as they are, namely in groups.

Kanda and Sabanovic–among others–set the trajectory that today's research continues on. These two studies are worth pulling out of the larger body of HRI work from that time because they were both field trials. These are particularly insightful as they stress-test the technologies. Allowing issues to emerge is itself the purpose, and these new issues are added to the agenda of researchers.

3.2 Abstracting from and Modeling Social Phenomena

The first step is for a robot to make sensory input into something meaningful; to extract signal from noise. An approach taken by HRI researchers has been to leverage existing social phenomena, or patterns of human behavior that arise naturally. If a robot can identify these patterns, it can conceptualize interaction in a way that it can compute. This process of turning real-world observations into discrete computable units is called modelling. One example of this is the identification of F-formations in human-human, and human-robot interaction.

In 2005, Brdiczka and colleagues made an initial step by implementing a Hidden Markov Model (HMM) to parse interaction groups. An HMM is a statistical method for finding underlying hidden trends from observed output states. The researchers' implementation could determine–to an accuracy rate of 84.8%–whether four people in a room were speaking in groups of two and two, or one of four, and switch when that configuration changed. Again, the researchers recognized the perception of interaction groups as a crucial part of robot design. As noted, the HMM was effective, but only in the controlled environment of the lab. Given the many variations that human interaction groups might take on–far beyond groups of two or four–the particular model is not expected to perform well in field trials. Partly accounting for its limitation, this method used only speech data, captured from microphones among the four subjects, to determine the group configuration. As the number of surrounding people increases, noise does as well, and the model would likely deteriorate. The researchers admitted these limitations and recommend additional modalities as a way of making the model more robust. Nonetheless, they showed that human verbal behavior could be parsed by machine intelligence, to extract information about the interaction from raw audio (Brdiczka et al., 2005).

In 2006 Hüttenrauch and colleagues further investigated spatial relationships in HRI, and introduced an important conceptual tool. They ran a Wizard of Oz study to see whether participants spontaneously created F-formations with the robot. If this phenomenon arises between a human and a robot, it would be valuable information a robot could leverage to computationally model the interaction. As mentioned in section 2.5.4, the F-formation framework serves to delineate units of interaction. In the study, a “home tour” design was used where a human participant led a robot around a room and explained what various objects are. It was found that as people interacted with the robot, F-formations did indeed occur (Huttenrauch et al., 2006). This is valuable to HRI researchers, as it opens up the F-formation framework to be used as a way of modeling interactions. Granted, this study only provided evidence for F-formations in a dyadic interaction.

Taking up the same challenge as Brdiczka and colleagues, in 2017 Vázquez applied the idea of the F-formation to group HRI. She concluded that F-formations not only arise dyads, but also in situations with multiple humans and one robot. This makes the F-formation available as a valuable source of information to robots in group settings. A more nuanced understanding of the social surroundings can be gained by applying the F-formation system to situations encountered in the wild.

In a subsequent analysis of their own 2006 study Hüttenrauch et al. noted another phenomenon that a robot could leverage to better understand an interaction. They found that the full interaction could be split up into another form of unit, interaction episodes. The episodes seen in their study were ‘follow’, in which the robot was guided; ‘show’, in which the user taught the robot information; and ‘validate’, in which the robot’s knowledge was tested. The researchers also noted the transitions between these episodes were often signaled by the participant’s body language. For example, when transitioning from show to follow, the human would turn away from the robot. In the language of F-formations, this is when the o-space dissolves. Parsing a social interaction into unitary episodes and knowing when (and with what probability) transitions are being made is another way in which robots could tap into social phenomena (Hüttenrauch et al., 2006).

Shi and colleagues similarly drew on F-formations to create a model for the initiation of conversation in their 2011 research. To begin, they ran scenarios between two people–a shopkeeper approaching a customer, and a host greeting a visitor. The researchers reviewed the videos of ten pairs of undergraduates, acting out the scenarios several times each, modeling the natural behavior. They extracted three factors that influenced the initiation: next plan, visibility, and initiation position. The next plan referred to what the host wanted to do upon initiating conversation. If the intent was simply to greet, there is no next plan, but if the intent was to bring the visitor to another location, there is a next plan. Visibility is whether the host has been noticed by the visitor before the initiation of conversation. Initiation position could be either the closest position to the visitor, or the position from which the next target (the plan) was visible. Thus, the host can approach the visitor with the intent to show an item or without, the visitor might have already seen and noticed the approaching host or not, and finally the host might position himself in the location of closest convenience or such that the visitor reorients towards the host and the item.

Drawing on the data they had collected from the scenarios, the researchers then implemented this model of human behavior in a robot, and ran trials between the robot and a human participant. According to the evaluations given by participants after the robot interaction, the behavior of the new model of interaction was considered better and more appropriate than two baseline models. This shows the promising possibility of drawing out interaction episodes from natural scenarios, and then using it as a basis for appropriate robot behavior (Shi et al., 2011).

In 2012, Yousuf et al. built a museum tour guide robot that explicitly sought to establish an F-formation before explaining an exhibit. The starting point was again to look at video footage of natural behavior, in this case in a museum during a tour. Researchers picked out the instantiation of F-formations, and the guide's tendency to wait for attention and restart the explanation when needed. When this model for guiding behavior was implemented in a robot, and trials were run in a lab setting, participants rated the robot using the proposed behavioral model higher than a conventional approach that did not account for F-formations (Yousuf et al., 2012).

Taken together, the work reviewed here indicates that the use of F-formations is a significant aid to robots in dyadic as well as in group interactions. Psychological phenomena such as interaction episodes, transitions, initiation position and timing have similarly been parsed out and factored into a more complex model of interactive behavior. Work in this area shows promising progress toward solving the problem of obtaining meaningful information about human social interactions from sensory input.

3.3 Maintaining a Presence

Separate and apart from the detection of conversation groups is the question of how to behave within one. Even if a robot has perfect information about who is party to an interaction, the difficult task of being a productive member remains. What the most appropriate behavior is depends largely on the particulars of the situation, but one constant is the shared focus of the group. To maintain this, certain universal behaviors exist. Here we look at how body orientation, eye gaze, and gesture each play a role. We will also consider how cultural differences change what physical behavior is appropriate in groups.

Vroon and colleagues found in 2015 that the body orientation rated most comfortable by participants was when the robot faced the center of the group. The robot used was a Giraff telepresence robot, and did not distinguish between upper and lower body–the entire machine could only face in one direction. This finding has limitations, however, as the type of robot used might have an effect. A telepresence robot is an alias for a human whose face can be seen on the screen. This partial inclusion of a human might be a confound. From this it is difficult to draw definitive conclusions about ideal robot body orientation (Vroon et al., 2015). Study into this topic is ongoing.

Vázquez and colleagues used reinforcement learning in a simulated environment to tune the simulated robot's policy for body orientation. This was a proof-of-concept for the viability of putting reinforcement learning on an autonomous platform, but also an opportunity to use the resulting policies from RL trials for future implementations. The advantage of reinforcement learning in this application is that the policy does not need to be fine-tuned by hand, and scales more easily (Vázquez et al., 2016).

In 2017 Vázquez et al. proposed another approach. The robot used in this study similarly had a fixed head, but also had eyes that could change the direction of their gaze. They found that the robot's body orientation affected participants' perception of gaze, and gaze affected perception of body orientation. These two behaviors work in tandem, and increase the participants' sense of inclusion in the group. The gaze and orientation behaviors that maximized this were not middle orientation and random gaze, but so-called attentive orientation and attentive gaze. This incorporated not just the center of the group's o-space, but applied a bias toward the focus of attention–for example, the participant currently speaking. This provides evidence that mere center orientation is less optimal behavior, compared to a more complex policy that accounts for more social nuance (Vázquez et al., 2017).

Kirchner et al. examined in 2011 the possibility of using gaze to communicate to one individual in a group. The scenario used is one that we might be familiar with from human interaction: passing an object while trying to communicate nonverbally with the recipient. Often among humans, it can be ambiguous who the intended recipient is, if not called out by name. The researchers designed a robot that handed out freshly-made cotton candy to a randomly selected onlooker among the crowd that gathered. The robot would either set its gaze on the intended recipient, or cast its gaze downward. It was found that more successful hand-overs occurred when the robot's gaze aided communication. This is encouraging evidence for the possibility of nonverbal gaze-based communication between robots and humans (Kirchner et al., 2011).

In work from 2009, Mutlu et al. found that gaze behavior was also able to communicate to participants what role they had in the conversation–addressee, bystander, or overhearer. Using a Robovie humanoid robot with eye gaze capabilities, a three sided interaction was set up: the robot and two human participants seated in a triangular arrangement. Participants were able to identify correctly 99% of the time when the robot signaled their speaking turn. Eye gaze, as we have seen previously, was effective in communicating this. Further, the disproportionate attention given to participants was perceived as well. Bystanders, who were at least acknowledged by the robot had higher ratings than overhearers, who were essentially invisible in the situation. Addresses had a clear understanding of their role, too. This provides evidence that a robot party to a group conversation has the ability to communicate participant's status in the conversation nonverbally (Mutlu et al., 2009). Below, we will discuss implications of this finding relating to responsibility to the group, and moral implications.

Studying cultural differences in human-robot interaction, Joosse and colleagues found that when a robot approaches a group of people, different behaviors are preferred between China, the U.S., and Argentina. The results showed that in China, closer approaches were considered to be more appropriate, whereas in the U.S. and Argentina, more space is preferred. This cross-cultural survey shows that just as human behavior should vary to adapt to local cultural norms, robots' behavior should as well (Joosse et al., 2014). This adds another layer of nuance to the optimal behavior of a social robot.

Research in this area shows how many variables factor into the maintenance of physical presence in a group. When it comes to body orientation, it is not just a matter of orienting toward the center of the group, but a bias should be calculated in towards the focal point of attention. If a robot has different upper and lower body components, the orientation of both of these need to be taken into account. Eye gaze has also been found to communicate effectively in groups, and physical movements should be coordinated with it. This research on the one hand clarifies, because it allows researchers to understand all the factors that play a role in group interactions. But on the other hand, it heaps many engineering challenges onto the task of creating a robot that communicates effectively in groups.

3.4 Scaling Up

Another important aspect of human-robot interaction in groups is the composition of those groups: the number and proportion of humans and robots. Just as an interaction between one robot and one person is different from one robot and four people, an interaction with one robot and a thousand people is another thing altogether. Or, for that matter, one person and a thousand robots. Here we look at research with a focus on the effects of group size and proportion.

In 2015 Leite et al. explored the effect of group size on children's memory of an emotional story told by two robots. In one condition, children were alone with the two robots, and in the other they were in a group of three. They found that individual participants came away from the interaction with a better grasp of the plot and semantic details of the story. When it came to recall of the emotional content of the story, there was no difference between individuals and groups (Leite et al., 2015a).

Performing further analysis on their study, Leite et al. looked at the effect of using a computational model trained in groups in an individual interaction, and vice versa. The model in question predicts disengagement in the participating child or children. They found that if trained in a group environment, it could perform better with an individual than a model trained with an individual in a group environment. This is in line with earlier discussion of work by Kanda and Sabanovic: individual HRI solutions do not effectively scale to groups. It was noted, though, that the ideal implementation would be to have separate predictors for individuals and groups (Leite et al., 2015b).

Chang et al. 2012 ran a between subjects study, focusing on the effect of group size and proportion on cooperative behavior. Extrapolating from the interindividual-intergroup discontinuity effect in psychology–which predicts interactions between groups will be more competitive than between individuals–they wanted to see if the same might occur with robots. The conditions pitted one or two humans against one or two robots in a prisoner's dilemma-like game that allows for co-operation or defecting. The researchers found that the number of people in the interaction had an effect on competitive behavior that agreed with the interindividual-intergroup discontinuity effect. This is to say, a group of people is more competitive against one or two robots than a single person is. However, the number of robots did not have an effect (Chang et al., 2012). This particular psychological phenomenon does not map over to HRI.

In a 2015 study, Fraune et al. measured the effect of varying group size and type of robot. The types studies were anthropomorphic, zoomorphic, and mechanomorphic. They expected number and type to have a strong effect on attitudes towards robots, but found that the type of robot and number interacted to influence attitudes. For example the mechanomorphic iCreate robots led to more negative attitudes in groups, whereas anthropomorphic NAO robots were more positively viewed in groups. Positive and negative attitudes were not affected by the presence of one versus multiple zoomorphic Pleo robots, but in groups were considered less likely to be coworkers or companions than when encountered alone (Fraune et al., 2015b). This study shows that attitudes are affected by a multitude of factors, not just number. More work has been undertaken by Fraune to find these.

In another study from the same year, Fraune et al. did a field trial of trash can robots that roamed a university cafeteria either alone or as a group of three. They sought to study the potential effect of number on attitudes towards the robot, as well as the effect of country (Japan versus the United States), and the effect of behavior (functional versus social). Results indicated that groups were interacted with more frequently, but liked the same as individual robots. Social robots were rated as more friendly than functional robots. Interestingly, single social robots were more positively rated than group social robots, but group functional robots more positively than single functional robots. Finally, it was observed that in general robots were more liked in Japan than in the United States (Fraune et al., 2015a). Again, there are several factors in addition to number that influence the perception of robots in these interactive settings.

In 2017 Fraune et al. studied the alternating effects of group entitativity of robots on people's attitudes. They found that robot groups who act as an entity–behaving the same way, and moving as a group–are viewed more negatively than a diverse group of robots. The diverse group was also viewed as having more mind. One way to account for this is that humans might find an entity of robots intimidating or threatening (Fraune et al., 2017b). Further implications of the work are covered below.

The literature on the effects of group size is not so clear cut. In some cases, people feel more positively about a group of robots than an individual, and in other cases, the opposite is true. The valuable advances in this work look into more specific questions in more targeted domains, like social trash robots or the effect of robot group entitativity. There is much more work to be done in this field since the effects found are complex, and not exact correlates to social psychology.

3.5 Group Membership and Interpersonal Dynamics

If robots are able to robustly perceive groups of people interacting, orient toward the focus of attention, and maintain a seamless physical presence, a new set of questions arises: how should the robot understand its place in the group, and what responsibilities does it have within it? Again, the right answers to these questions depend largely on context, like the group's shared purpose. But developing ways of addressing the question in the first place–by modeling social dominance hierarchies, understanding ingroup versus outgroup dynamics, and leveraging features of effective leadership–is an active area of research in HRI.

Kuzouka and colleagues ran a study in 2010 to find out how a robot's body orientation affected the positioning of people around it. They found that when a robot led someone through a makeshift exhibition, the person would position themself to maintain an F-formation with the robot. If the robot's lower body rotated, it would be taken as a cue to reconfigure the F-formation. Though this was run as a dyadic interaction, the conclusion is drawn that a robot can influence arrangement of an F-formation by adjusting its lower body orientation (Kuzuoka et al., 2010). Robots can influence how a human participant occupies physical space.

Work by Hung and colleagues in 2008 produced a method for determining dominance in a group conversation using only microphone input. The system does this without previous knowledge of the members of the conversation, and works real-time as the conversation goes on. Some limitations, though, are that the predictions of the system correlate strongly with speaking length–the more you speak, the more dominant you are in a group–but this is not necessarily true. It also does not account for the content of what is said. The addition of video input is acknowledged as a potential enhancement. With video input it would also be possible to take nonverbal communications like gestures into account (Hung et al., 2008). The importance of these communication channels has been acknowledged in section 3.2.

In 2009, Jayagopi and colleagues worked on precisely this issue. They tested the audio input from a conversation, video input, and both in combination. It was found that audio cues were more accurate in classifying dominance than visual cues. At best, the latter could perform on par with audio cues in some instances. Similarly, the fusion of audiovisual cues was no better than audio alone. The researchers point out that this method is able to classify the most or least dominant persons in a group, but cannot determine whether two or more persons are equally dominant as a subset of the group (Jayagopi et al., 2009).

Jie and Peng sought to improve upon the previous work in dominance detection by emphasizing nontraditional features in speech. They were able to attain similar dominance classification accuracy using total speaking length and total speaking energy as cues. With a more pointed focus on these cues, a less time-complex algorithm could be developed that yields the same results (Jie and Peng, 2010).

Strohkorb and colleagues took this research further in 2015 by running a study with children successfully using verbal and nonverbal features in conjunction as an indicator for social dominance. Some of the verbal cues used were the same as previous work, like total talking time, but here types of utterance and utterance addressee were factored in as well. The nonverbal behaviors taken into account were gestures, physical coercion, and gaze. Combined verbal-nonverbal behaviors were looking while listening and talking to colleagues and robots. The overall accuracy of the approach was 89% with the support-vector machine and 81% with logistic regression, in the ballpark of previous research. Interestingly, though, the misclassifications of both models were accounted for by a handful of individuals. That is, most participants were classified with no miscalculations, but three were misclassified 100% of the time. The researchers speculate this might be due to subtleties of the preexisting relationships between some individuals or being in the presence of an exceptionally dominant or submissive participant (Strohkorb et al., 2015). It suggests that there is more about modeling social dominance to be understood.

Thus far, visual input has not proven to be much more effective than audio input. Because we know that substantial communication can be nonverbal, we can infer that the methods for interpreting video input are not effective enough yet. That is, if visual information was understood as a human would, a robot using a combination of audio and visual input would outperform just the one or the other. Research by Strohkorb shows that researchers are getting closer to achieving this, but exposes how preexisting social relationships, and likely other factors as well, can be confounds.

Moving from dominance hierarchies within a group to dynamics between groups, Fraune et al. 2017 compared attitudes toward ingroup robots versus outgroup humans. This particular question is of interest because previous work has found people to be less aggressive toward members of their ingroup, and more aggressive towards robots than fellow humans. In a price guessing game, it resulted that people valued team membership more than humanness. This is promising because it suggests that humans and robots can form bonds in competitive game settings. The authors point out that in some contexts, like under scarce resources, favoring an ingroup robot might come at a cost to outgroup humans (Fraune et al., 2017a). Whether this is morally right remains an open question.

In their 2017 study, Short and colleagues tested a robot's ability to serve as moderator for a group collaborative task. The moderator acted in one of two ways, performance reinforcing (encouraging the highest total score in the group) or performance equalizing (encouraging equal point contributions from all participants). Results indicated that performance equalizing reduced group cohesion measures but improved task performance, and the reverse was true for performance reinforcing. This was the opposite of what the researchers expected to find. However, there was evidence that a robot moderator can improve group social dynamics, though understanding exactly how remains an open question (Short et al., 2016).

Work by Strohkorb et al. investigated the influence a robot can have on two children collaborating on a task. It was found that by posing questions, the overall performance on the task could be improved, as well as the participants' perception of their performance, but the robot was unable to improve cohesion between the collaborators (Strohkorb et al., 2016).

The Bršĉić et al. study from section 2.2.1 shows that people can influence each other to harm robots in a Milgram experiment paradigm. Another important psychological work on behavioral conformity is Asch's work from 1951. In this experiment, a single participant sat at a table with several other people, all of whom were confederates. A researcher stood before them with a visual cue on a board: one vertical line on the left side, three on the right. Importantly, the line on the left matched one of the three on the right in length. The other two were noticeably too long or too short. The researcher asked each of the seated individuals (all confederates, save one) to say which line on the right (labelled A, B, C) was the same length as the line on the left. The correct answer to this question is easy to identify, as the differences between the lines is clear to see. But one by one, the confederates would give the same wrong answer. Finally, when it was the participant's turn, it was found that they conform to the group's incorrect consensus one third of the time. This is an especially striking finding compared to the control condition where no pressure was put on the participant to conform. In that condition, the correct answer was given 99% of the time.

Recent work by Salomons et al. explored whether the same effect can occur in HRI. In a variation of the Asch study, the confederates were replaced with two robots. The task in this study did not have a correct or incorrect answer, rather the players chose a drawing that best matched an abstract word (e.g. 'irony'). It was found that people conform at an average rate of 29%, which is significant, compared to the control condition of 6%. While several studies preceding this one that more strictly followed the structure of the original Asch experiment were unable to induce conformity (Beckner et al., 2016; Brandstetter et al., 2014), this variation was able to do so (Salomons et al., 2018).

HRI is mostly concerned with how people interact with robots, but another consideration made in Group HRI is how people interact with each other around robots. Tan et el. studied whether a human participant would intervene on behalf of a robot when a confederate abused it. They found that most participants did indeed come to the defense of the robot. It was found that the most effective way to induce bystander intervention was by making the robot briefly shut down in response to the abuse. Participant responses suggested this was because it caused emotional concern (Tan et al., 2018).

The research reviewed in the area of group membership shows several areas where more research is needed. While it is possible to parse social dynamics from audio and visual input, in adults as well as children, more work needs to be done to merge the two effectively. Evidence that an ingroup robot can be closer than an outgroup human is both promising for human-robot companionship, and raises some moral concerns about priority. Finally, evidence shows that a robot can moderate social dynamics in a group much like they are able to influence physical arrangements, however how these effects work is not fully understood. Continuing to address these areas of research is especially important because of the moral implications they have. We will further consider these moral implications in section 4 below.

3.6 Gaining Acceptance

Though the presence of robots in most people's daily lives is not yet common, as we have seen, work is being done to make it a reality across a variety of domains. The following studies look at how people feel about robots, and the prospect of welcoming them into daily life.

In 2010 Sabanovic wrote that the people developing new robotic technologies intended for interaction with humans must allow for and consider feedback from society. Quoting the 1933 Chicago World Fair, Sabanovic illustrates how new technologies typically enter our daily lives: "Science finds, industry applies, man conforms". However, for human-robot interaction to thrive, Sabanovic calls for a "mutual shaping and co-production framework" (Šabanović, 2010). Developing robots in the isolation of the lab leads to technologies that are deemed important by developers, and that the rest of society have to conform to and integrate into their lives. Instead, according to Sabanovic, the users who are ultimately the target audience should have their preferences and concerns considered, with high priority.

Kim et al. 2010 found that in-group/out-group effects from social psychology apply to human-robot interactions as well. They presented participants with photographs of humans, animals, products, and robots, each of which is either an in-group item or an out-group item. The in-group animal, for example was a Jindo dog, common in Japan where the study was run. The out-group counterpart was a type of hound more commonly found in the West. It was found that the familiarity, reliability, and preference scores for humans had the greatest in-group out-group difference. This was followed by robots, then animals, and finally products. This is evidence that in-group/out-group psychological effects apply to robots, so some extent (Kim et al., 2010).

Wallisch and colleagues ran a 2018 study to determine whether contact with a robot through interaction would affect the emotional evaluation of a robot, or willingness to interact. It was found that after a five to ten minute interaction, participants who had contact with the robot had slightly higher willingness to interact with a robot, but emotions toward robots were the same as the control condition. The lack of effect on the emotional measure might be a result of the robot's minimal social behavior. Participants also commented after trials that the robot did not seem intelligent as a result of doing a simple task. The conclusion is that contact can improve attitudes toward robots, however the strength of the effect might depend on what the robot is capable of doing. Simple contact is not enough to improve emotions toward robots (Wallisch et al., 2018).

Fraune and colleagues in 2014 explore the question of what effect indirect exposure to robots might have on attitudes toward them. Participants sat in a room, completing a task on a computer as three iCreate robots roamed the room for the purported pretense of collecting data. It was found that after a thirty-minute exposure of this kind, there was no effect on the attitudes toward robots, even with different communication styles between them. In the analysis of the results, the researchers call the 'negative attitude toward robots scale' (NARS) into question (Fraune and Šabanović, 2014). Perhaps it is too simple a measure and ought to be complemented by others.

Short et al. 2017 studied interactions between a robot and a group comprised of one older adult, one adult, and one child. The researchers found that the majority of participants, in all three generational categories had a positive interaction with the robot. The sample size of the study was small (18 participants), but adequate for the researchers to determine whether there is interest for bringing robots into family homes. Participants mostly came away with a positive experience, saying they could conceive of having a robot at home (Short et al., 2017).

Recent research into attitudes of integrating robots into group settings points to a promising future. Direct contact with robots has the ability to change people's attitudes, but indirect contact does not. Families might be willing to bring robots into their homes because they have something to offer to three different age groups. The user's perception ought to be one of the primary metrics used to evaluate interactive robotics. Because of this, scales like NARS must provide accurate measurements. Some concerns have come up in recent research. This will be addressed further in the following section, and in section 4.

3.7 Methodologies

HRI research faces the challenge of using and developing adequate empirical methodologies. One of the simplest, and yet most revealing methodologies in HRI is the field trial. As we have seen in the 2004 and 2005 work by Kanda and Sabanovic, putting robots out in the world and observing the interactions closely can yield important preliminary findings. They both found that group interactions are an inevitable part of HRI, even if the robot is incapable of such interactions.

Kanda observed that school children interacted with the Robovie tutors significantly less in the second week of the study. This might be due to the robot's novelty effect, the authors suggested. Initially, it attracts attention, but later interest wanes and fewer interactions happen in the long term. This was an important finding because some aspects of HRI research are concerned with long term interactions, and must find ways to overcome post-novelty effect disinterest.

Leite et al. 2013 conducted a survey of long-term HRI findings in social robotics. They, too, observed the novelty effect across a variety of studies. In fact, they consider an interaction to be long term not by the measured length of the interaction, but if the interaction continues after the initial novelty effect wears off. At this point, familiarity with the robot becomes stable. They make a connection to habituation from developmental psychology, or the point where a novel stimulus ceases to draw the observer's attention simply for being there. Leite et al. make suggestions for future technologies across a range of traits from appearance to memory and adaptation. From the many broad field trials surveyed, the authors could distill several actionable recommendations (Leite et al., 2013).

Sabanovic's field trial also revealed the importance of timing and rhythmic behavior. One of the observed obstacles to the interactions was delay on the part of the Roboceptionist. Often the robot's attention would be drawn to a passer by, but would react only once they were out of range of the interaction. This highlights the importance of timely responses, as well as effective initiation of interaction. As we have seen in section 3.2, this challenge was addressed by Shi and colleagues in 2011.

Field trials have been very useful to HRI, because of their ability to bring in relevant observations about a wide range of topics. Exploratory studies like these show researchers what obstacles remain, and allow for new lines of inquiry to address them.

The aforementioned work by Shi et al. 2011 is an example of another experimental method in HRI. In this study, researchers first ran the scenarios with human participants who, beyond simple instructions to act out a particular role, acted naturally. From these observed scenarios, the researchers coded interaction behaviors, and used them to implement a machine that could more effectively start an interaction than robot with baseline behavior.

This idea of borrowing not just from the study of human behavior, but directly from observed human behavior is a promising one. The study by Luber et al. mentioned in section 2.5.3 is an example of this. Here the researchers used footage of pedestrians crossing a public square to train a model to generate optimal paths. The paths made by the model that learned from real-world data were more efficient, and smoother than what proxemics predict, and closely resembled the paths actually taken. Given the efficacy of modern machine learning approaches, and the ever-growing availability of data on human behavior, this methodology of using real world observations as a starting point or training set shows promise.

Breazeal et al. used a similar approach by first running a collaborative task as an online game, then using the interaction data found to create a model of behavior-generation. That model was implemented in a robot and tested with participants in vivo. The autonomous behavior from this model matched human-directed behavior in a WOZ condition (Breazeal et al., 2013). This shows that models for behavior can be built from data that is crowd sourced.

One form of research that has been instrumental to research in HRI has been the so-called "Wizard of Oz" design. This paradigm puts the participant in contact with a robot, but instead of that robot acting autonomously, in accordance with its programmed behavior, it is aided or controlled entirely by a human who is hidden from view. What results is an experimental environment in which the participant is made to believe she is interacting with the autonomous robot. Since, from the perspective of the participant the robot is autonomous, measures of the participant's perceptions of the robot's capabilities are valid.

A review of WOZ studies by Laural Riek covered 54 papers spanning from 2003 to 2011. This survey shows that the three most commonly human-controlled behaviors are natural language processing, non-verbal behavior, and navigation. Reporting guidelines for future publications that employ a WOZ paradigm are recommended so that consistency is maintained across studies. The work also addresses the existence of the "Oz of Wizard" model, put forward by Steinfield et al (Riek, 2012). This inverse of the WOZ study simulates the human participant instead of the robot. The argument for this is that running a study with human participants has many logistical obstacles. People need to be recruited and often brought into the lab, and ethical concerns about human exposure to risk have to be cleared by an ethics committee (Steinfeld et al., 2009). Work by Vázquez et al. discussed in section 3.3 is an example of this model. Here, the robot was a machine learning model, optimizing its attentive orientation behavior, and the humans were simulated (Vázquez et al., 2016).

Seeing how robot behavior compares to human behavior in similar situations is one way to evaluate the quality of interactions. However, in HRI the robots have to interact with humans on humans' terms. For a robot to be considered capable, it must be able to act in ways that do not conflict long standing human social tendencies. In other words, what matters is how people feel about the robots. To measure this, Nomura et al. developed the Negative Attitudes towards Robots Scale (NARS) in 2003, and published in 2006. It has a five-point answer range (from strongly disagree to strongly agree), and a total of 32 items. Since the scale was developed by Japanese researchers, the items were originally in Japanese, later translated to English (Nomura et al., 2006). This scale is seen frequently in HRI research, especially between the years 2000 and 2010.

More recently, Bartneck et al. developed the Godspeed scale which intends to address general aspects of HRI: anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety. Keeping the questionnaire at a high level, and not delving into subscales for each category, allows for comparison between the results of different studies (Bartneck et al., 2009). This scale has become increasingly common in HRI research.

In 2017, Carpinella et al. develop another scale specifically to measure the social perception of robots. The Robotic Social Attributes Scale (RoSAS) is based on the Godspeed questionnaire as well as research on social perception. The 18-item assessment tests the three dimensions of warmth, competence, and discomfort. The authors explicitly state that this scale is not intended to replace general scales like Godspeed and NARS, but complements those to provide information about how robots' social capabilities are seen. Both types of measurements are very important to research because the standard for effective interaction is largely decided by ability to seamlessly coexist with humans (Carpinella et al., 2017). All of these methods have been important to making advances in

HRI. Below in Future Directions, we consider in more detail possible shortcomings, and suggestions for future studies.

3.8 Verbal Communication

Speech is the primary channel of communication between people interacting in a group. While much can be gleaned from non-verbal behavior, the content of an interaction is transferred verbally. A robot that participates effectively in a group of people must have the ability to understand and transmit verbal communications.

Oliveira et al. 2018 studied rates of gaze and verbal support behavior in a card game between two teams of two agents (human or robot). It was found that verbal communications were most frequent to human partners, followed by human opponents. Measured by dividing the number of behaviors by the length of the session, robotic partners were addressed at far lower rates (0.4 vs. 0.1), and robotic opponents less frequently still (0.02 vs. 0.19) as compared to humans. This suggests that people were overall more likely to address another person over a robot, regardless of the person's status as teammate or opponent. The authors posit that this preference for human-human interaction could be accounted for by an in-group preference effect (Oliveira et al., 2018).

Work addressed in section 3.2 by Brdiczka showed that machine learning approaches to verbal communication can parse conversation groups in a limited setting. Hung et al. 2008 expanded this to approach to also parse out dominance hierarchies from group conversations.

Lee et al. conducted a review of study to observe how people interacted with two types of robots, one in a receptionist role, and another in an information kiosk role. They found that if interactions began with a verbal greeting, people would have a more substantive interaction. Specifically, they were more likely to listen to the robot's story, tell the robot more about themselves, and be less rude. The authors argue that this initial greeting will lead to the perception of the robot as a social actor, thus evoking social behavior in return (Lee et al., 2010).

Xu et al. identified that a central challenge in verbal communication during human-robot interaction is ascertaining the user's engagement state. The authors used machine learning to develop two computational models to gauge a human's engagement intent. The models used visual data to check several features and calculate an engagement rating from them: direction of attention, distance, upper body pose, change of distance, and others. A robot was then given this model of verbal communication perception and evaluated. The behaviors designed used ratings of engagement to determine who has the conversational floor, and whether the participants had intention to engage or disengage. Xu et al. found that this model of engagement behavior improved the robot's perceived behavior over a baseline (Xu et al., 2013).

This work is progress toward improving verbal communication, but the focus is on group metadata: who is in it, in what social arrangement, and how can social behavior be prompted. This is all important information to a robot. When it comes to machines understanding human speech, and formulating responses, these are issues that fall predominantly into the field of Natural Language Processing (NLP). This is not to say that HRI is not itself concerned with verbal semantics, but the field more directly pursuing research in this space is likelier to make the important advances to improve verbal communications between humans and robots.

3.9 Conclusion: Group Human-Robot Interaction Research

In this section, work in Group Human-Robot Interaction was reviewed. Several categories of research were covered showing challenges and solutions unique to them. The use of preexisting social phenomena as a guide for computational models has been effective and shows promise for future advances. Research into the maintenance of a social presence in a group exposed nuances in how best to orient the robot's physical parts. The efficacy of eye gaze as a form of communication was also uncovered. Studying the effects of group size has yielded many additional questions about discontinuities between Human-Robot Interaction and Human-Human Interaction (HHI). Answering these questions will likely result in a deeper understanding of humans' perception of robots. It has been found that robots have the ability to influence behaviors and attitudes in groups, so questions about moral responsibility begin to arise. Research into people's willingness to accept robots reveals how ubiquitous they might be in the near future. Methodologies are an important and difficult part of HRI research. Verbal communication exposes people's preference for HHI, perhaps in the face of inadequately complex HRI communication. In the coming section, we will look into the future implications of work from these areas, and propose guidelines for robust research.

4 Future Directions

While research in group-centric human-robot interaction has come a very long way, in this section we will consider questions that arise from the work reviewed, gaps in active research areas, as well as guidelines for producing robust work in the future.

4.1 Further Research into the Effect of Group Size and Composition

In the work reviewed above, considerable efforts have been dedicated to better understanding the effects of group size and ratio of humans to robots (section 3.4). In some cases, psychological phenomena did not translate to HRI as researchers predicted, and this has exposed a gap in knowledge about the effects of group size and composition. Ongoing work by Fraune and colleagues, is illuminating this area by studying the effect of group size with a narrower scope, and replicating past work with fewer confounds.

The research by Chang et al. 2012 revealed that a group of two robots is not viewed as more competitive than a single opposing robot, though a group of humans was more competitive than a single one (Chang et al., 2012). But for clarity it is important to look closer into why competitive behavior arises between groups in social psychology.

Since the robot team was designed to play a tit-for-tat strategy, any increase in competitive behavior depends on the human participant or participants. If the human team chooses not to cooperate, that decision is mirrored by the robot team. Any change in competitive behavior can only be determined by the humans. They are making the decision for both teams, albeit unwittingly. Whether a robot is alone or in a group cannot influence the robot team's competitive behavior per se. Thus, it might be that no effect was found simply because there are fewer participants from whom competitive behavior could have originated. The measure of group effect on robots is really just a measurement of the effect of the group on human perception of the robots, not an actual change in robot behavior.

A fallacy here is to design the experiment to measure competitive behavior as though it could arise in a natural way, as between two human teams. This is not possible because a robot's behavior depends on what it is programmed to do. In some cases, treating the robot as a stand-in for a human to observe what differences might arise makes no sense. By removing a human, the decision-making processes of a human are also removed. Future research should design experiments that do not treat robots and humans as equally likely to be psychologically influenced.

Recognizing this, Fraune et al. 2019 ran a similar study, but instead of playing tit-for-tat, the robots had a hard-coded strategy. In both practice rounds, the robot or robot team would collaborate, and in the live trials, collaborate four times and defect twice. With the robots' competitive behavior held constant across trials, the effect of group size could be measured more reliably. They again found that number of humans had an effect on competitive behavior, but number of robots did not. A novel discovery, though, was that when humans faced off against an equal number of robots, one or three, "perceived threat, anxious emotion, greed motivation, and competition increased," (Fraune et al., 2019, p. 111). This reaffirms the findings of the Chang et al. work, and adds to it.

Beyond how group size effects competition specifically, much remains unknown. In the 2015 work by Fraune et al. they showed that group size and robot type are variables that interact with each other, and do not reveal clear trends. Each robot type was viewed most positively in varying quantity–mechanomorphic ones as singletons, anthropomorphic ones as groups, and zoomorphic ones evoked no strong effect. From this, it becomes clear that we cannot draw a simple conclusion about people's preferences or comfort with groups because there are many factors at play (Fraune et al., 2015b).

Looking for a deeper underlying effect, Fraune et al. continued by studying group entitativity. Since an entity of robots was viewed more negatively, this might have been a confound in the study of type-number interaction. But this leaves unexplained why the non-diverse anthropomorphic robots were viewed positively as a group than as singletons (Fraune et al., 2017b).

Ultimately in her doctoral dissertation, Fraune urges that context ought to be further explored because even the effect of entitativity is influenced by other factors. What these factors are and how they influence psychological effects that pertain to groups remains an open question. Fraune calls for more research, adjusting several variables like environment, duration of interaction, and type of robot (Fraune, 2018).

Much has been uncovered with respect to group size and composition, and several psychological principles governing groups have been found to extend to HRI, more work is needed to address unknown or ambiguous effects: how exactly do psychological effects in HRI change as the group size or proportion changes? Apart from just the number of humans and robots, what other factors exert an influence? Fraune's work is particularly directed at these questions.

4.2 Replication in Groups

Given the differences between group and individual psychology, and between group and individual HRI, it would be a worthwhile project to test how effects found in dyadic interactions transfer to group interactions.

For example, the studies by Hüttenrauch et al. was conducted in a dyadic situation. Hüttenrauch's work was followed up Yousuf et al. who expanded the tour scenario to groups of participants. It was revealed that robots could govern F-formations in groups just as Hüttenrauch had found in dyads.

Similar replications should be of other studies that focus on dyadic interaction. Any outcome to such a replication would be a valuable insight. If a particular effect transfers from individual to group HRI, then additional knowledge is gained about group interactions and can be implemented in future projects. If on the other hand, an effect does not transfer to group HRI, then it exposes a discontinuity (like the group size discontinuity discussed in the previous section) and shows where research efforts should be directed.

For example, the work by Shi et al. on the initiation of interaction could be attempted with additional human participants. The same scenarios could be run with multiple people (e.g. two customers, one shopkeeper), and parsed for interaction episodes.

A robot's interaction initiation behavior could be implemented, based on this observed data. Researchers could then compare the initial behavior from the dyadic study in scenarios with multiple customers and the new behavior extracted from the multi-party scenarios. If there is an improvement in the quality of interaction, it would serve as evidence that a robot should have initiation behavior specific to groups and to individuals. If there is no improvement, it serves as evidence that this distinction need not be made in future robot designs.

4.3 Evaluation of Methodologies

There are two challenges to designing experiments in HRI. On the one hand, researchers try to find answers to empirical questions about the nature of interactions between humans and robots. This comes with all the challenges that many academic fields face: isolating and uncovering confound-free, significant results. On the other hand, though, HRI researchers often face the additional challenge of implementing new technological developments in these studies, for example new robots or algorithms.

This is not an insignificant trait of HRI research. By comparison, in psychological research, experimenters can depend on fMRI to be reliable because it is a well-established method. Two studies along the same line of inquiry that use fMRI will be using the same technology. The ability to hold this part of the methodology constant allows researchers to adjust other variables as they look for significant effects, and minimizes the probability of confounds. By comparison, one experiment with a Robovie replicated by another lab with a Pepper adds variance to the experiment that we do not yet know how to account for.

Syrdal et al. 2009 found that NARS is vulnerable to confounds from cultural differences. Once three items were removed from the scale, the internal consistency of the scale was improved, and the authors suggest that the cultural difference could be accounted for by this. They also found that the subscale that evaluates 'Actual Interactions' did not predict participants' behavior with a robot. They suggest that this could be a more targeted 'Robot Anxiety' effect at work that the original developers of the scale have previously acknowledged (Syrdal et al., 2009).

Weiss and Bartneck 2015 surveyed the HRI literature for work that used the Godspeed Questionnaire Series (GQS). In a subset of the work included they analyzed studies that used the NAO anthropomorphic robot. Some contradictions were found, though, relating to perceived intelligence. Some studies indicated that interaction lowered perceived intelligence, while others increased it (Weiss and Bartneck, 2015). This suggests that the GQS should be further studied, perhaps with a similar systematic analysis of studies sharing a common technology.

Jung and Hinds call for a return to the field trial in HRI research. By studying interactions in the real world, outside the lab, robots will face the challenges of complex social scenarios. Watching how a robot succeeds or fails in this environment is valuable knowledge for the furtherance of HRI. They also encourage the use of comparative studies, to observe how reactions to humans and robots differ in similar scenarios. A call is also made for the inclusion of researchers from many disciplines to develop "more robust, nuanced, and comprehensive theories of HRI," (Jung and Hinds, 2018, p. 4). This is much in agreement with Sabanovic's suggestions from 2010 (Šabanović, 2010).

4.4 Challenges in Verbal Communication

Oliveira found that people were more likely to engage with other human participants in the card game than robots, regardless of their role as teammates or opponents. A possible explanation is that, since no verbal communications during play concerned game strategy itself, there was little reason to interact with the robot. It is not made clear what the robots' capabilities are, and whether conversation with them could result in any further interaction. If, for example, one asks a robot teammate whether they are any good at the game before getting started, does the robot respond appropriately? If not, participants might have the impression that what they say to the robots is of no consequence. Therefore, why bother?

Alternatively, the discrepancy could be explained by the richness of human communication compared to communication with the robot. The competitive and cooperative robots were programmed with 419 utterances each. Humans, by contrast, communicate with infinitely generative language. Again, if the participants glean that the robot has a simple capacity for communication, the motivation to interact verbally might drop. Ultimately, the challenge of sufficiently complex verbal communication to encourage ongoing interaction is primarily the project of NLP research.

4.5 Further Research into Influence on Social Dynamics

The assertion made by research is that robots can change the way people act, both in physical space (Yousuf et al., 2012), and how they act toward other people (Salomons et al., 2018). Given this influence, and the possibility that further research might uncover other ways in which people's behavior can be changed by robots, an important moral question arises. Should robots that influence human behavior be guided by a moral code?

For example, work by Tan et al. showed that a robot can induce bystander intervention on its behalf to stave off abuse from another human. Conversely, how should a robot act if the victim of abuse, and bystander effect, is a human? The effect causes inaction in part because responsibility is diffused among all bystanders. No one person feels the need to act because there are so many other people who could do so instead. In this way, a non-intervening person is guilty of diffusing responsibility merely by being there and contributing to the effect. Arguably, this is a moral failure. If a robot is among the witnesses to an injustice, and among the bystanders, would it also be guilty of a moral failure, or does the fact of being a robot shirk responsibility? Should a robot itself intervene? If a particular robot is not capable of intervening due to physical or behavioral constraints, should it encourage others? Based on what heuristic should responsibility be redistributed?

Another moral dilemma arose in work by Fraune et al. that focused on competitive behavior between humans and robots. They found that it was possible for a human to favor an ingroup robot over an outgroup human, in the context of their study (Fraune et al., 2017a). Is it wrong to value a robot over a human? When posed this way, the question begs the intuitive answer that yes, it is indeed wrong. But is it not already the case that human life is valued less than non-living things? People gather wealth in great quantities and expend it on goods that are not essential over expending it on the well being or survival of other humans. This is not to say that practice is wrong outright, rather that the idea of valuing human life below any number of things is not a new idea and is common today.

Work by Hung et al., Jayagopi et al., Jie and Peng, and Strohkorb et al. showed that it is possible to model social dominance hierarchies based on interaction data. If a robot party to a group interaction has this knowledge of power dynamics, does it have any moral compulsion to act? For example, if bias towards or against any member arises based on gender, race, socioeconomic status or some other feature, will a robot blindly perpetuate the bias based on others' behaviors? Similar to the bystander intervention dilemma, should the robot explicitly change the existing dynamics of the group?

All of these are important considerations that need to be addressed intentionally, and not decided as a byproduct of other development priorities. Similarly, just as Sabanovic and Jung call for the involvement of multiple disciplines in HRI development, here the need for the involvement of moral philosophers becomes clear. If the goal of HRI really is to create agents that interact with humans as realistically as possible, then the moral questions that apply to human interaction apply to HRI as well.

5 Conclusion

Groups of people are complex in a number of ways. Beyond just multiplying the individual psychological agents present in a scenario, psychological effects emerge when an interaction scales. A group takes on a psychology of its own. When sharing space, people are under the influence of social forces that steer them along–towards some things, away from others. The way groups arrange themselves to converse is also predicted by frameworks from psychology. Competitive behavior similarly adjusts as group size changes. Much of this behavior happens without people's explicit knowledge.

When integrating robots into group interactions, researchers must be explicitly aware of all these factors. To overcome computational challenges, robots have to build models that accurately reflect social dynamics. To optimize the behavior of an autonomous embodied agent, one must first know how humans behave. An understanding of group size is needed to produce appropriate behavior as numbers fluctuate, and it is important to determine what sort of behavior intimidates. If robots take on roles well enough to exert influence, we must be careful that it is a positive influence, not perpetuation of the biases they encounter in the world. Finally, for any of this to succeed, robots have to be designed to be broadly accepted by society.

Looking at efforts to tackle all these issues, promising future directions can be identified. In order to best continue research along these lines of inquiry, further research into group size and composition is recommended. HRI would benefit from more clarity as to what effects of group size are unique to human-robot as opposed to human-human interaction. More research on the moderation of social dynamics is also needed. At present it remains ambiguous the extent to which humans are influenced by robots, and in what patterns. Ongoing work is beginning to address this. Given initial results that there is some influence, there is a moral imperative to consider how best to do so. Finally, replication would be a beneficial project for research in the field. This would serve to solidify the basic assumptions of research, and allow different projects with varying methodologies and technologies to rely on each other. Re-evaluating findings from group HRI in dyads, and vice versa, will help bridge the conceptual gap between the two.

6 Author Contributions

Marynel Vázquez, the advisor for this project, was extremely helpful in guiding me along. As an expert in HRI she identified the need for a review of the Group HRI literature to begin with, and encouraged me to try filling this gap with the present work. If ever I felt at a dead end in the research, she was able to provide names of researchers or specific studies to continue my investigation. Similarly, if my inquiry deviated too far into tangential work, she would help maintain the focus. When it came to formulating my own arguments about the literature, I proposed ideas and we discussed in order to refine them.

Bibliography

Admoni, H., Hayes, B., Feil-Seifer, D., Ullman, D., and Scassellati, B. (2013). Are you looking at me?: perception of robot attention is mediated by gaze type and group size. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction*, pages 389–396. IEEE Press.

Azenkot, S., Feng, C., and Cakmak, M. (2016). Enabling building service robots to guide blind people a participatory design approach. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 3–10. IEEE.

Bainbridge, W. A., Hart, J. W., Kim, E. S., and Scassellati, B. (2011). The benefits of interactions with physically present robots over video-displayed agents. *International Journal of Social Robotics*, 3(1):41–52.

Barnes, J., Jowett, B., et al. (1984). Politics. the complete works of aristotle. *Trans. Benjamin Jowett (1885). Princeton: Princeton University Press.*

Bartneck, C., Kulić, D., Croft, E., and Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International journal of social robotics*, 1(1):71–81.

Bartneck, C., Rosalia, C., Menges, R., and Deckers, I. (2005). Robot abuse–a limitation of the media equation. In *Proceedings of the interact 2005 workshop on agent abuse, Rome.*

Beckner, C., Rácz, P., Hay, J., Brandstetter, J., and Bartneck, C. (2016). Participants conform to humans but not to humanoid robots in an english past tense formation task. *Journal of Language and Social Psychology*, 35(2):158–179.

Boole, G. (1854). *An investigation of the laws of thought: on which are founded the mathematical theories of logic and probabilities.* Dover Publications.

Bostrom, N. (2017). *Superintelligence.* Dunod.

Brandstetter, J., Rácz, P., Beckner, C., Sandoval, E. B., Hay, J., and Bartneck, C. (2014). A peer pressure experiment: Recreation of the asch conformity experiment with robots. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1335–1340. IEEE.

Brdiczka, O., Maisonnasse, J., and Reignier, P. (2005). Automatic detection of interaction groups. In *Proceedings of the 7th international conference on Multimodal interfaces*, pages 32–36. ACM.

Breazeal, C., DePalma, N., Orkin, J., Chernova, S., and Jung, M. (2013). Crowdsourcing human-robot interaction: New methods and system evaluation in a public environment. *Journal of Human-Robot Interaction*, 2(1):82–111.

Brooks, R. A. (1984). Aspects of mobile robot visual map making. In *Proc. 2nd Int. Symp. on Robotics Research, Cambridge, MA, USA, 1984*, pages 325–331.

Brscić, D., Kidokoro, H., Suehiro, Y., and Kanda, T. (2015). Escaping from children's abuse of social robots. In *Proceedings of the tenth annual acm/ieee international conference on human-robot interaction*, pages 59–66. ACM.

Butler, S. et al. (1999). *The Iliad*. Courier Corporation.

Carpinella, C. M., Wyman, A. B., Perez, M. A., and Stroessner, S. J. (2017). The robotic social attributes scale (rosas): Development and validation. In *Proceedings of the 2017 ACM/IEEE International Conference on human-robot interaction*, pages 254–262. ACM.

Chang, W.-L., White, J. P., Park, J., Holm, A., and Šabanović, S. (2012). The effect of group size on people's attitudes and cooperative behaviors toward robots in interactive gameplay. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, pages 845–850. IEEE.

Copeland, B. J. (1997). The church-turing thesis. *Stanford Encyclopedia of Philosophy*.

Corcoran, J. (2003). Aristotle's prior analytics and boole's laws of thought. *History and Philosophy of Logic*, 24(4):261–288.

Dijkstra, E. (1959). Dijkstra's algorithm. *Dutch scientist Dr. Edsger Dijkstra network algorithm: http://en. wikipedia. org/wiki/Dijkstra's_algorithm*.

Eisenberger, N. I. and Lieberman, M. D. (2004). Why rejection hurts: a common neural alarm system for physical and social pain. *Trends in cognitive sciences*, 8(7):294–300.

Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and autonomous systems*, 42(3-4):143–166.

Fraune, M. R. (2018). *Examining Effects of Groups and Intergroup Contexts on Human-Robot Interaction*. PhD thesis, Indiana University.

Fraune, M. R., Kawakami, S., Sabanovic, S., De Silva, P. R. S., and Okada, M. (2015a). Three's company, or a crowd?: The effects of robot number and behavior on hri in japan and the usa. In *Robotics: Science and Systems*.

Fraune, M. R. and Šabanović, S. (2014). Negative attitudes toward minimalistic robots with intragroup communication styles. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pages 1116–1121. IEEE.

Fraune, M. R., Šabanović, S., and Smith, E. R. (2017a). Teammates first: Favoring ingroup robots over outgroup humans. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 1432–1437. IEEE.

Fraune, M. R., Šabanović, S., Smith, E. R., Nishiwaki, Y., and Okada, M. (2017b). threatening flocks and mindful snowflakes: How group entitativity affects perceptions of robots. In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI*, pages 205–213. IEEE.

Fraune, M. R., Sherrin, S., Sabanović, S., and Smith, E. R. (2015b). Rabble of robots effects: Number and type of robots modulates attitudes, emotions, and stereotypes. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 109–116. ACM.

Fraune, M. R., Sherrin, S., Šabanović, S., and Smith, E. R. (2019). Is human-robot interaction more competitive between groups than between individuals? In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 104–113. IEEE.

Friend, T. (2017). Silicon valley's quest to live forever. *The New Yorker*, 3.

Garnham, A. (2017). *Artificial intelligence: An introduction.* Routledge.

Gonsior, B., Buß, M., Sosnowski, S., Wollherr, D., Kühnlenz, K., and Buss, M. (2012). Towards transferability of theories on prosocial behavior from social psychology to hri. In *2012 IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*, pages 101–103. IEEE.

Goodrich, M. A., Schultz, A. C., et al. (2008). Human–robot interaction: a survey. *Foundations and Trends® in Human–Computer Interaction*, 1(3):203–275.

Hansen, S. T., Andersen, H. J., and Bak, T. (2010). Practical evaluation of robots for elderly in denmark: An overview. In *Proceedings of the 5th ACM/IEEE international conference on Human-robot interaction*, pages 149–150. IEEE Press.

Haraway, D. (2006). A cyborg manifesto: Science, technology, and socialist-feminism in the late 20th century. In *The international handbook of virtual learning environments*, pages 117–158. Springer.

Harrigan, J., Rosenthal, R., Scherer, K. R., and Scherer, K. (2008). *New handbook of methods in nonverbal behavior research.* Oxford University Press.

Horowitz, M. J., Duff, D. F., and Stratton, L. O. (1964). Body-buffer zone: exploration of personal space. *Archives of general psychiatry*, 11(6):651–656.

Hung, H., Jayagopi, D. B., Ba, S., Odobez, J.-M., and Gatica-Perez, D. (2008). Investigating automatic dominance estimation in groups from visual attention and speaking activity. In *Proceedings of the 10th international conference on Multimodal interfaces*, pages 233–236. ACM.

Hüttenrauch, H., Eklundh, K. S., Green, A., and Topp, E. A. (2006). Investigating spatial relationships in human-robot interaction. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5052–5059. IEEE.

Huttenrauch, H., Eklundh, K. S., Green, A., Topp, E. A., and Christensen, H. I. (2006). What's in the gap? interaction transitions that make hri work. In *ROMAN 2006-The 15th IEEE International Symposium on Robot and Human Interactive Communication*, pages 123–128. IEEE.

Jayagopi, D. B., Hung, H., Yeo, C., and Gatica-Perez, D. (2009). Modeling dominance in group conversations using nonverbal activity cues. *IEEE Transactions on Audio, Speech, and Language Processing*, 17(3):501–513.

Jie, C. and Peng, P. (2010). Recognize the most dominant person in multi-party meetings using nontraditional features. In *2010 IEEE International Conference on Intelligent Computing and Intelligent Systems*, volume 1, pages 312–316. IEEE.

Johns, M., Mok, B., Sirkin, D., Gowda, N., Smith, C., Talamonti, W., and Ju, W. (2016). Exploring shared control in automated driving. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 91–98. IEEE.

Joosse, M. P., Poppe, R. W., Lohse, M., and Evers, V. (2014). Cultural differences in how an engagement-seeking robot should approach a group of people. In *Proceedings of the 5th ACM international conference on Collaboration across boundaries: culture, distance & technology*, pages 121–130. ACM.

Jung, M. and Hinds, P. (2018). Robots in the wild: A time for more robust theories of human-robot interaction. *ACM Transactions on Human-Robot Interaction (THRI)*, 7(1):2.

Kanda, T., Hirano, T., Eaton, D., and Ishiguro, H. (2004). Interactive robots as social partners and peer tutors for children: A field trial. *Human–Computer Interaction*, 19(1-2):61–84.

Kendon, A. (1990). *Conducting interaction: Patterns of behavior in focused encounters*, volume 7. CUP Archive.

Khatib, O. (1985). Real-time obstacle avoidance for manipulators and mobile robots. In *Proceedings. 1985 IEEE International Conference on Robotics and Automation*, volume 2, pages 500–505. IEEE.

Kidd, C. D. (2008). Designing for long-term human-robot interaction and application to weight loss. *n/a*.

Kim, Y., Kwak, S. S., and Kim, M.-s. (2010). Effects of intergroup relations on people's acceptance of robots. In *Proceedings of the 5th ACM/IEEE international conference on Human-robot interaction*, pages 107–108. IEEE Press.

Kirchner, N., Alempijevic, A., and Dissanayake, G. (2011). Nonverbal robot-group interaction using an imitated gaze cue. In *Proceedings of the 6th international conference on Human-robot interaction*, pages 497–504. ACM.

Kory Westlund, J., Gordon, G., Spaulding, S., Lee, J. J., Plummer, L., Martinez, M., Das, M., and Breazeal, C. (2016). Lessons from teachers on performing hri studies with young children in schools. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, pages 383–390. IEEE Press.

Kuzuoka, H., Suzuki, Y., Yamashita, J., and Yamazaki, K. (2010). Reconfiguring spatial formation arrangement by robot body orientation. In *Proceedings of the 5th ACM/IEEE international conference on Human-robot interaction*, pages 285–292. IEEE Press.

Lee, M. K., Kiesler, S., and Forlizzi, J. (2010). Receptionist or information kiosk: how do people talk with a robot? In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, pages 31–40. ACM.

Leite, I., Martinho, C., and Paiva, A. (2013). Social robots for long-term interaction: a survey. *International Journal of Social Robotics*, 5(2):291–308.

Leite, I., McCoy, M., Lohani, M., Ullman, D., Salomons, N., Stokes, C., Rivers, S., and Scassellati, B. (2015a). Emotional storytelling in the classroom: Individual versus group interaction between children and robots. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 75–82. ACM.

Leite, I., McCoy, M., Ullman, D., Salomons, N., and Scassellati, B. (2015b). Comparing models of disengagement in individual and group interactions. In *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 99–105. IEEE.

Lozano-Perez, T. (1990). Spatial planning: A configuration space approach. In *Autonomous robot vehicles*, pages 259–271. Springer.

Luber, M., Spinello, L., Silva, J., and Arras, K. O. (2012). Socially-aware robot navigation: A learning approach. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 902–907. IEEE.

Malle, B. F., Scheutz, M., Arnold, T., Voiklis, J., and Cusimano, C. (2015). Sacrifice one for the good of many?: People apply different moral norms to human and robot agents. In *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*, pages 117–124. ACM.

Martyn, K., Dodds, P., and Brown, M. (2018). The paro seal: weighing up. *Journal of Dementia Care*, 26(3):32–33.

Meier, B. P. and Hinsz, V. B. (2004). A comparison of human aggression committed by groups and individuals: An interindividual–intergroup discontinuity. *Journal of Experimental Social Psychology*, 40(4):551–559.

Mervin, M. C., Moyle, W., Jones, C., Murfield, J., Draper, B., Beattie, E., Shum, D. H., O'Dwyer, S., and Thalib, L. (2018). The cost-effectiveness of using paro, a therapeutic robotic seal, to reduce agitation and medication use in dementia: Findings from a cluster–randomized controlled trial. *Journal of the American Medical Directors Association*, 19(7):619–622.

Michaud, F. and Gustafson, D. A. (2002). The hors d'oeuvres event at the aaai-2001 mobile robot competition. *AI Magazine*, 23(1):31.

Moravec, H. P. (1980). Obstacle avoidance and navigation in the real world by a seeing robot rover. Technical report, STANFORD UNIV CA DEPT OF COMPUTER SCIENCE.

Mori, M., MacDorman, K. F., and Kageki, N. (2012). The uncanny valley [from the field]. *IEEE Robotics & Automation Magazine*, 19(2):98–100.

Moshkina, L., Trickett, S., and Trafton, J. G. (2014). Social engagement in public places: a tale of one robot. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, pages 382–389. ACM.

Mutlu, B., Shiwa, T., Kanda, T., Ishiguro, H., and Hagita, N. (2009). Footing in human-robot conversations: how robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 61–68. ACM.

Nikolaidis, S., Ramakrishnan, R., Gu, K., and Shah, J. (2015). Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*, pages 189–196. ACM.

Nomura, T., Suzuki, T., Kanda, T., and Kato, K. (2006). Measurement of negative attitudes toward robots. *Interaction Studies*, 7(3):437–454.

Oliveira, R., Arriaga, P., Alves-Oliveira, P., Correia, F., Petisca, S., and Paiva, A. (2018). Friends or foes?: Socioemotional support and gaze behaviors in mixed groups of humans and robots.

In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 279–288. ACM.

Qinggang, H. (2016). A history of chinese science and technology, volume 1.

Reeves, B. and Nass, C. I. (1996). *The media equation: How people treat computers, television, and new media like real people and places.* Cambridge university press.

Rhodios, A. (2008). *The Argonautika*, volume 25. Univ of California Press.

Riek, L. D. (2012). Wizard of oz studies in hri: a systematic review and new reporting guidelines. *Journal of Human-Robot Interaction*, 1(1):119–136.

Šabanović, S. (2010). Robots in society, society in robots. *International Journal of Social Robotics*, 2(4):439–450.

Sabanovic, S., Michalowski, M. P., and Simmons, R. (2006). Robots in the wild: Observing human-robot social interaction outside the lab. In *9th IEEE International Workshop on Advanced Motion Control, 2006.*, pages 596–601. IEEE.

Salomons, N., van der Linden, M., Strohkorb Sebo, S., and Scassellati, B. (2018). Humans conform to robots: Disambiguating trust, truth, and conformity. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 187–195. ACM.

Sequeira, P., Alves-Oliveira, P., Ribeiro, T., Di Tullio, E., Petisca, S., Melo, F. S., Castellano, G., and Paiva, A. (2016). Discovering social interaction strategies for robots from restricted-perception wizard-of-oz studies. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, pages 197–204. IEEE Press.

Shi, C., Shimada, M., Kanda, T., Ishiguro, H., and Hagita, N. (2011). Spatial formation model for initiating conversation. *Proceedings of robotics: Science and systems VII*, pages 305–313.

Short, E., Sittig-Boyd, K., and Mataric, M. J. (2016). Modeling moderation for multi-party socially assistive robotics. In *IEEE Int. Symp. Robot Hum. Interact. Commun.(RO-MAN 2016). New York, NY: IEEE.*

Short, E. S., Swift-Spong, K., Shim, H., Wisniewski, K. M., Zak, D. K., Wu, S., Zelinski, E., and Matarić, M. J. (2017). Understanding social interactions with socially assistive robotics in intergenerational family groups. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 236–241. IEEE.

Smith, R. (1989). *Prior analytics.* Hackett Publishing.

Smullyan, R. M. and Smullyan, R. (1992). *Gödel's incompleteness theorems*. Oxford University Press on Demand.

Stallings, A. et al. (2018). *Works and Days*. Penguin UK.

Standage, T. (2002). *The Turk: The life and times of the famous eighteenth-century chess-playing machine*. Walker & Company.

Steinfeld, A., Jenkins, O. C., and Scassellati, B. (2009). The oz of wizard: simulating the human for interaction research. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 101–108. ACM.

Strohkorb, S., Fukuto, E., Warren, N., Taylor, C., Berry, B., and Scassellati, B. (2016). Improving human-human collaboration between children with a social robot. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 551–556. IEEE.

Strohkorb, S., Leite, I., Warren, N., and Scassellati, B. (2015). Classification of children's social dominance in group interactions with robots. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, pages 227–234. ACM.

Syrdal, D. S., Dautenhahn, K., Koay, K. L., and Walters, M. L. (2009). The negative attitudes towards robots scale and reactions to robot behaviour in a live human-robot interaction study. *Adaptive and Emergent Behaviour and Complex Systems*.

Tan, X. Z., Vázquez, M., Carter, E. J., Morales, C. G., and Steinfeld, A. (2018). Inducing bystander interventions during robot abuse with social mechanisms. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 169–177. ACM.

Urban, T. (2017). The ai revolution: Our immortality or extinction. *Wait But Why*.

Vázquez, M., Carter, E. J., McDorman, B., Forlizzi, J., Steinfeld, A., and Hudson, S. E. (2017). Towards robot autonomy in group conversations: Understanding the effects of body orientation and gaze. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pages 42–52. ACM.

Vázquez, M., Steinfeld, A., and Hudson, S. E. (2016). Maintaining awareness of the focus of attention of a conversation: A robot-centric reinforcement learning approach. In *2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 36–43. IEEE.

Vroon, J., Joosse, M., Lohse, M., Kolkmeier, J., Kim, J., Truong, K., Englebienne, G., Heylen, D., and Evers, V. (2015). Dynamics of social positioning patterns in group-robot interactions.

In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 394–399. IEEE.

Wainer, J., Feil-Seifer, D. J., Shell, D. A., and Mataric, M. J. (2006). The role of physical embodiment in human-robot interaction. In *ROMAN 2006-The 15th IEEE International Symposium on Robot and Human Interactive Communication*, pages 117–122. IEEE.

Wallisch, K., Fraune, M., Sabanović, S., Sherrin, S., and Smith, E. (2018). Getting to know you: Relationship between intergroup contact and willingness to interact. In *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 275–276. ACM.

Walters, M. L., Dautenhahn, K., Te Boekhorst, R., Koay, K. L., Syrdal, D. S., and Nehaniv, C. L. (2009). An empirical framework for human-robot proxemics. *Procs of new frontiers in human-robot interaction*.

Weiss, A. and Bartneck, C. (2015). Meta analysis of the usage of the godspeed questionnaire series. In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 381–388. IEEE.

wiki (2019). History of artificial intelligence. *Wikipedia*.

Xu, Q., Li, L., and Wang, G. (2013). Designing engagement-aware agents for multiparty conversations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2233–2242. ACM.

Yi, D., Goodrich, M. A., and Seppi, K. D. (2016). Homotopy-aware rrt*: Toward human-robot topological path-planning. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, pages 279–286. IEEE Press.

Yousuf, M. A., Kobayashi, Y., Kuno, Y., Yamazaki, A., and Yamazaki, K. (2012). Development of a mobile museum guide robot that can configure spatial formation with visitors. In *International Conference on Intelligent Computing*, pages 423–432. Springer.

Yzerbyt, V., Dumont, M., Wigboldus, D., and Gordijn, E. (2003). I feel for us: The impact of categorization and identification on emotions and action tendencies. *British Journal of Social Psychology*, 42(4):533–549.