Long-term Interaction with Relational SIAs

Jacqueline M. Kory-Westlund, Hae Won Park, Ishaan Grover, and Cynthia Breazeal

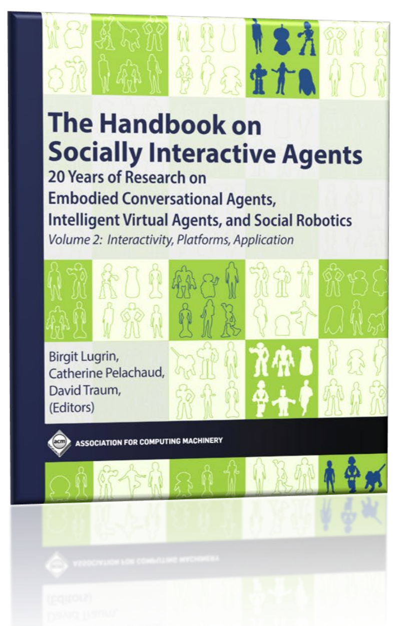

Author note:

This is a preprint. The final article is published in "The Handbook on Socially Interactive Agents" by ACM.

Citation information:

J. M. Kory-Westlund, H. W. Park, I. Grover, and C. Breazeal. (2022). Long-term Interaction with Relational SIAs. In B. Lugrin, C. Pelachaud, D. Traum (Eds.), *The Handbook on Socially Interactive Agents – 20 Years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics*, Volume 2: Interactivity, Platforms, Application (pp. 195-260). ACM.

DOI of the final chapter: 10.1145/3563659.3563667

DOI of volume 2 of the handbook: 10.1145/3563659

Correspondence concerning this article should be addressed to Jacqueline M. Kory-Westlund jacqueline.kory@gmail.com

19 Long-term Interaction with Relational SIAs

Jacqueline M. Kory-Westlund, Hae Won Park, Ishaan Grover, and Cynthia Breazeal

19.1 Motivation

In this chapter, we provide an overview of research on long-term interaction with Socially Interactive Agents (SIAs). Over the past couple decades, as SIAs have become more capable in their ability to interact and collaborate with people, they have been increasingly moving out of the laboratory and into the world with the goals of helping people and bringing benefit people's lives. Chapters in this book have surveyed applications in education, health and wellness, aging, therapy, entertainment, and more (see also Chapter 21 on "Pedagogical Agents" [Lane and Schroeder 2022], 22 on "Socially Interactive Agents as Peers" [Cassell 2022], 23 on "Socially Interactive Agents for Supporting Aging" [Ghafurian et al. 2022], 24 on "Health-Related Applications of Socially Interactive Agents" [Bickmore 2022], 25 on "Autism and Socially Interactive Agents" [Nadel et al. 2022], 26 on "Interactive Narrative and Story-telling" [Aylett 2022], 27 on "Socially Interactive Agents in Games" [Prada and Rato 2022], 28 on "Serious Games with SIAs" [Gebhard et al. 2022] of this volume of this handbook. Over time, as SIA technology has improved, their competence and effectiveness in engaging with and supporting a variety of human activities and goals is being actively investigated in schools, hospitals, assisted living facilities, museums, shopping malls, the workplace, and people's homes.

With this progress, both technical and scientific, people are starting to not just interact with SIAs, but to *live* with them and work with them as part of daily life. In a number of important application areas—e.g., helping children progress through an education curriculum, coaching people to change behavior to better manage a chronic disease, or providing therapeutic support for children with autism—lasting benefits take time to realize, often requiring repeated encounters over days, weeks, or months. In other application areas, such as aging-in-place or acting as a digital assistant, people may call upon the SIA on a daily basis for a variety of reasons, including companionship, accessing digital information or services, or help with tasks in the home or workplace context. Therefore, advancing the state-of-the-art in long-term interaction is a growing area of research interest, and our research community should address technical and user experience challenges in sustaining long-term engagement and supporting intelligent and adaptive decision-making to help people achieve long-term goals. This may include learning about people over long-periods of time to provide personalized engagement and support, as doing so can enable SIAs to better optimize the benefit provided to different people as they grow and change over time. These challenges raise important research problems in the design principles of robot platforms and interaction scenarios, human modeling and evaluation methods, and algorithmic approaches.

The technological world is changing very fast. Today, we have digital assistants in our smartphones and smart speakers, embodied conversational agents on screens that entertain and inform us, and social robots on our countertops as helpful companions—all supporting a variety of skills and services. With so many AI-enabled, socially interactive, and collaborative technologies entering everyday life, we need to deeply understand how these technologies affect us in the long term. What are the long-term consequences of having such technology in our lives—whether benefits or detriments? How can these technologies be used to promote human flourishing? How do we mitigate ethical concerns—and there are many—about the use of social technology and AI in our lives?

19.1.1 Beyond Interaction to Relationship

In this chapter, we look beyond human-SIA interaction to focus on the question of what kinds of *relationships* people are forming SIAs. Whereas rapport can be established in short-term interactions, relationships take time and change over repeated encounters. We know that there are many benefits that come along with positive relationships. For example, better student-teacher or student-tutor relationships often result in better learning outcomes [e.g., Sinha and Cassell 2015a,b, Wentzel 1997]. Better doctor-patient, therapist-patient, or patient group relationships often lead to better health outcomes [Horvath and Luborsky 1993, Mallinckrodt 1989, Wampold 2015, Yorke et al. 2008]. Better coach-client relationship lead to better engagement and behavior change outcomes [Frates et al. 2011]. In SIA research, we are witnessing that people *are* forming relationships, of some kind, with socially responsive technologies [Dautenhahn et al. 2002, Desteno et al. 2012, Kory-Westlund 2019, Kory-Westlund et al. 2018]. What benefits might we see if people form positive relationships with SIAs? What properties and characteristics are designed into SIAs that foster the ability of SIAs to form and maintain a long-term relationship—particularly one that provides benefit to people, such as improved teamwork, wellness, learning, behavior change for therapeutic improvement, and more? How do we design SIAs to build sustained, adaptive, personalized, positive relationships with people?

In the next sections, we explore what counts as a long-term interaction with an SIA. We then discuss relationship models and approaches, with an eye toward how we can design SIAs as relational partners for long-term interaction. With this background, we survey the field: what social robots, virtual agents, voice assistants, and other SIAs have been developed for long-term interaction so far? How has the field evolved? We discuss trends over time, as well as the similarities and differences between agents. Finally, we catalogue several of the current most pressing challenges in developing long-term interaction with SIAs and point out directions for future research.

19.2 What is Considered Long-Term Interaction?

Where does one draw the line between a short-term interaction study and a long-term one? Is it based on the number of sessions? The accumulated amount of time in interaction? The total elapsed time from first encounter to last? Another way of asking this, is *when does novelty wear off?* It is not disputed that novelty has some kind of an effect on initially increasing engagement, positive affect, and excitement about interacting with a SIA, and thus is a push for more longitudinal work to "get past" the novelty effect [e.g., Baxter et al. 2016].

In longitudinal SIA studies, for instance, researchers have reported that novelty has worn off after after 1–2 sessions to 1–2 weeks, as assessed by increased boredom [e.g., Salter et al. 2004] or decreases in interaction time [e.g., Gockley et al. 2005, Kanda et al. 2004]. That is, after some amount of time, the pattern of interaction changed. In the first survey paper on long-term interaction in Human-Robot Interaction (HRI), Leite et al. [2013] equated novelty with familiarization and habituation, suggesting that novelty has worn off when familiarization or habituation with the robot is stable, i.e., when a person does not react as much to it and starts preferring novel behaviors. Leite et al. [2013] suggested using gaze or looking time to determine habituation. Using gaze, as well as other behavioral measures, is a reasonable suggestion. For instance, some research suggests that children look less at a familiar peer, and look longer at an unfamiliar peer; they also may play more cooperatively with a familiar peer, and show more behaviors such as seeking attention, asking questions, and showing affection [Doyle et al. 1980, McCornack 1982].

But is gaze a reliable measure? For instance, in a 8-week child-robot study by Kory-Westlund [2019], they examined children's gaze patterns during an Anomalous Picture Task, comparing children's gaze in a between-subjects study where the anomalous pictures were introduced by a person or a social robot. If the robot was perceived as more novel than the human, and more novel at the pretest than at the posttest, one would expect to see a decrease in children's amount of looking time between the robot and the human, and from the pretest to the posttest. However, this was not what was observed. Children did spend more time looking at the robot than at the human experimenter at both times, but they also looked at the robot more at the posttest than at the pretest. This could be that the relationship between novelty and children's gaze patterns is not as simple as decreasing novelty leading to decreased gaze. It may be that children were looking more at the robot as an attention-seeking behavior, which may also be related to greater familiarization, or perhaps they were looking longer because they knew it was their last session with the robot. Either way, it seems that more work is needed to understand how children's gaze patterns relate to their perception of novelty.

Others have proposed interaction frameworks comprised of stages characterized of interaction patterns to study long-term interaction. For instance, [Sung et al. 2009b] studied how people interacted with a Roomba vacuum cleaner over a 6-month period in a home setting. Although the Roomba is not an SIA, the authors proposed a long-term framework comprised of 4 temporal phases comprised of key interaction patterns: pre-adoption, adoption, adaptation and use and retention. Kidd and Breazeal [2008] present a 6-week study of a social robot weight management coach in a home setting. They describe 3 phases of human-SIA relationship state: initial, normal or repair. The initial phase lasts the first 4 days of interaction before transitioning into the normal stage for another 4 days. After that, the robot begins to asks 2 questions from the short form of the Working Alliance Inventory each day (a measure commonly used in therapy and other helping relationships that tracks trust and belief in a common goal of helping that the therapist and patient have for one another [Horvath 1989]) to calculate a relationship score. If the relationship score falls below the normal threshold, the robot enters the repair phase where it changes its relational dialog behavior in an effort to rebuild the working alliance.

Because there are no consistent ways of measuring novelty, nor for determining whether observed interaction patterns were in fact a result of novelty wearing off (versus merely being boring after doing the activity a couple times, say), it is actually hard to know whether reported effects are due to long-term interaction or due to novelty. Thus, for the purposes of deciding what work to include in this chapter, we admittedly have to be a bit arbitrary. We include works where people interact with an SIA for at least 5 sessions over any length of time per session, and for any elapsed time from the first encounter to the fifth or more. We shall resume our discussion of novelty in the section on Future Challenges.

19.3 Relationship Models and Approaches

People form a wide variety of relationships with others. For example, people have different types of relationships with their friends, parents, children, managers, employees, colleagues, etc. People also form meaningful and beneficial relationships with non-human beings, such as companion animals. People also have different kinds of relationships with technology, though for the purposes of this chapter, we focus on people's ability to form interpersonal relationships with SIAs, and how these relationships can be maintained over time.

From numerous studies in HCI and HRI, we are seeing that people of all ages can construe SIAs as relational others if designed appropriately—with responsiveness and interactivity. That is, these *relational SIAs* are part of the broad category of things with which one can have a relationship. They are more than playful objects [Ackermann 2005] or transitional objects [Winnicott 1953] since either category would imply that they are only artifacts for projecting onto, for exploration and learning, rather than for *being with*. For example, work with social robots has shown that people perceive them to be social and relational [e.g. Darling et al. 2015, Kory-Westlund et al. 2018, Turkle et al. 2006, Zawieska et al. 2012]. Often, they are seen as having some of the properties of pets, toys, computers, artifacts, assistants, and friends, but not exactly the same properties as any of these [e.g., Bartlett et al. 2004, Kahn et al. 2002, 2012, Kory-Westlund et al. 2018, Melson et al. 2009, Weiss et al. 2009]. They are frequently ascribed social presence [e.g., Biocca et al. 2003, Leite et al. 2009] and can evoke rapport (see Chapter 12 on "Rapport Between Humans and Socially Interactive Agents" [Gratch and Lucas 2021] of volume 1 of this handbook [Lugrin et al. 2021]), attachment, trust, and emotion [e.g., Batliner et al. 2011, Bickmore et al. 2010, Desteno et al. 2012, Hancock et al. 2011, Kidd and Breazeal 2008, Turkle et al. 2006, Weiss et al. 2009]. Dynamic social interaction factors such as the contingency of the SIA's nonverbal behavior and its expressivity impact engagement, trust, learning, and judgments of the SIA's credibility [e.g., Breazeal et al. 2016, Kennedy et al. 2017a, Kory-Westlund and Breazeal 2019b, Kory-Westlund et al. 2017b, Lubold et al. 2016, 2018]. In fact, people seem to apply a range of social judgments to SIAs and are willing to treat them in a variety of roles with different levels of authority relative to themselves. For instance, in an educational context, researchers are actively exploring a range of roles for SIAs from an expert tutor, to a peer-like playmate, to a novice to be taught [e.g., Chen 2018].

The fact that children to older adults are treating SIAs as social-relational others brings new opportunities for engaging people in activities involving technology that plays an effective social-relational role. For example, numerous researchers have been exploring social robots to help children develop and practice skills that are best learned in social contexts—such as language or social and emotional skills [e.g., Bernardini et al. 2014, Clabaugh et al. 2018, Kanda et al. 2007, Kennedy et al. 2016, Kim et al. 2013, Kory and Breazeal 2014, Robins et al. 2005, Scassellati et al. 2018a,c, Vogt et al. 2019].

19.3.1 Dyadic model

In the social sciences, relationships are modeled in numerous ways. These models provide a theoretical and empirical foundation from which to design SIAs that can socially interact and build relationships with people. One common model is the social system, the simplest example of which is a dyad. In a dyad, a relationship is defined as a pattern of interaction, e.g., the interaction of two people whose behavior is interdependent [Berscheid and Reis 1998, Csikszentmihalyi and Halton 1981, Kelley et al. 1983]. Critically, this model can be applied to human-object relationships, since non-human objects can also significantly influence our patterns of interaction and behavior [Csikszentmihalyi and Halton 1981].

The long-term education studies with children and robots paint an intriguing example of dyadic, interdependent interaction. With sufficient variation in the robot's behavior, children easily engaged for many sessions, up to several months [e.g., Kanda et al. 2007, Kory and Breazeal 2014, Kory-Westlund 2019, Leite et al. 2014]. They frequently learned new words from robots. They treated the robots as social others, frequently appeared to grow more comfortable and closer to the robots over time, and often called them their friends. In one study, children engaged longer and were less likely to grow bored if they treated the robot as a peer-like friend [Kanda et al. 2007]. This suggests that children's relationships affect how interested they are in interacting and playing—as one might expect, children like playing with their friends.

19.3.2 Dimensional model

Another important relationship model is the dimensional model, in which relationships are defined in terms of various relational characteristics, including power, social distance, and trust [Berscheid and Reis 1998, Bickmore and Cassell 2001, Burgoon and Hale 1984, Cassell and Bickmore 2000, Fogg and Tseng 1999, Spencer-Oatey 1996, Trope and Liberman 2010]. The dimensional model is important because these characteristics can be manipulated by non-human objects as well to influence the relationship. For example, Desteno et al. [2012] observed human-human behavior during an economic exchange game, identifying a set of nonverbal cues that were predictive of human cooperative behavior during the game. Then, they experimentally manipulated the nonverbal cues used by a social robot that played the same economic game with a human. They showed that the robot's use of the set of nonverbal cues affected human perception of the robot's trustworthiness and cooperative behavior in the game.

19.3.3 Provisional Model

Other models include provision models, in which relationships are discussed in terms of what people provide for one another [e.g., Duck 1991], as well as economic models, such as social exchange theory, in which relationships are modeled based on perceived costs and benefits of the relationship [e.g., Brehm 1992]. Important in relation to provision models is social support theory, which describes how social relationships influence people's cognition, emotions, and behavior [Lakey and Cohen 2000]. Social support theory becomes particularly relevant if we conclude that people can have social relationships with non-human objects.

In health-related domains, therapeutic alliance (or working alliance) refers to the relationship between a healthcare professional and a client or patient. It is the means by which a therapist and a client hope to engage with each other to effect beneficial change in the client. It consists of three parts: tasks, goals and bond. Tasks are what the therapist and client agree need to be done to reach the client's goals. Goals are what the client hopes to gain from therapy. The bond forms from trust and confidence, and the belief that the tasks will help the client achieve their goals. Research on the working alliance suggests that it is a strong predictor client outcomes.

For instance, Kidd and Breazeal [2008] performed a six-week study in which participants worked with either a co-present robotic weight loss coach, a standalone computer, or a standard paper log, to see which would most effectively help participants achieve and maintain their weight loss goals. They used a measure of working alliance to estimate the user's relationship with the robot coach or the computer, and found that taking specific actions to improve the relationship when working alliance was low led to greater engagement and improved health outcomes. Bickmore et al. [2005] found that using relational behaviors, such as social dialogue, empathy, nonverbal cues, and relationship-building actions, in a computer health interface led to increased interaction and working alliance, compared to one that did not use these behaviors.

19.3.4 Relationships with Animals

Animal-assisted therapy (AAT) recognizes the health benefits that can arise from human-animal relationships. The goal of AAT is to improve a patient's social, emotional, or cognitive functioning. The biophilia hypothesis suggests that if we see animals at rest or in a peaceful state, this may signal to us safety, security and feelings of well-being which in turn may trigger a state where personal change and healing are possible. People construe companion animals as being non-judgemental, comforting, and welcoming. Animals can also be supportive of educational and motivational objectives for people as a supportive, positive presence. For instance, canine-assisted reading programs are used to help children with special educational needs. The calm, non-judgmental, happy characteristics of dogs helps the process of reading to become more meaningful and enjoyable for children. Numerous studies have found that one-on-one or free-form interaction with the seal robot Paro in assisted living centers and nursing homes can increase positive affect and quality of life [e.g., Bemelmans et al. 2015, Lane et al. 2016, Moyle et al. 2018].

19.3.5 Attachment Theory

Finally, attachment theory concerns the relationships between people, often in the context of young children and their adult caregiver [Ainsworth 1969, Ainsworth and Bell 1970, Bowlby 1958]. Attachment means an "affectional bond" or tie between an individual and an attachment figure. Such bonds may be reciprocal between two adult. Between a child and a caregiver these bonds are based on the child's need for safety, security and protection. Attachment theory has also been discussed in relation to the formation and maintenance of relationships with both humans and objects [Bretherton 1992, Passman and Halonen 1979].

19.4 Designing Relational SIAs

Considering these various models, we can see a variety of features that tend to be associated with relationships. First, relationships tend to unfold over time and generally involve multiple interactions. This may be on short timescales, such as repeated encounters over the span of minutes or days, or it may be on longer timescales, such as months or years. Even in short timespans, people's behavior can interdependently influence each other [Davis 1982]. These repeated interactions build up shared experiences—i.e., activities done together in the past or are performing together now. Shared experiences influence later interactions, and are often referenced and remembered later on. These shared experiences can involve trying to perform tasks or achieve goals together.

There is also some amount of responsiveness and commitment. Those we form relationships with respond to us, e.g., with social cues in the moment, or social support in response to life events. Attachment and emotion often come into play; we may feel positively or negatively about interacting with certain people (see Chapter 10 on "Emotion" [Broekens 2021] of volume 1 of this handbook [Lugrin et al. 2021]). Empathy plays an important role—the ability to sense other people's emotions, coupled with the ability to imagine what someone else might be thinking or feeling. Affective empathy refers to the sensations and feelings that arise in response to others' emotions. Cognitive empathy has to do with the ability to take the other person's perspective to identify and understand their emotions. Friendship relations often involve positive feelings, trust, and attachment, such as enjoying one another's company and depending on one another. Friendship relations often involve reciprocity as well, such as exchanging favors, reciprocating contact, dialogue, and connection, and being responsive in turn.

The examples of research presented above suggests that people can construe SIAs as social agents with whom they can form friendships and relationships. And these relationships can persist over time and bring value to people. Furthermore, people also appear to understand that SIAs are not quite the same as their other human friends, nor quite like their pets, or teachers, or nurses, or mechanical toys. But when technology is designed to act as social agents, people interact with them as social agents. They share gaze, mirror emotions, show affection, help the robots or interactive characters, take turns, and disclose information—all behaviors associated with friendships and close relationships [Gleason and Hohmann 2006, Hartup et al. 1988, Newcomb and Bagwell 1995, Rubin et al. 1998].

Even with robots or devices that are arguably less social (e.g., without the capability for speech), people can still attribute intelligence and talk to them and about them as if they have social capabilities [e.g., Bemelmans et al. 2015, Chang and Šabanović 2015, Fink et al. 2012, Moyle et al. 2018, Sung et al. 2007, Wada and Shibata 2007].

These observations lead to more questions: How are people's relationships with SIAs different than their relationships with other entities? What features of SIAs impact the relationship people can develop? Can SIAs actively try to build a relationship, and if so, how would this affect people's engagement and perception of the relationship?

19.4.1 Designing SIAs as Relational Partners

If SIAs can provide similar kinds of interaction opportunities and features associated with relationships that people (or animals) can provide for each other—and it seems that they can—then they, too, can be relational. Relational is different than just being social—it is the behaviors that contribute more directly to building and maintaining an ongoing relationship. This may include numerous social behaviors, such as the use of nonverbal cues and contingency, but is a larger category that includes additional behaviors, which we discuss further below.

Much of the work in exploring relationship factors with SIAs are done over short-term encounters, often over a single session. For example, people pay attention to the verbal and nonverbal social cues of agents to build rapport and to collaborate on a variety of activities and contexts with them. People largely seem to respond to these interpersonal social cues much as if they are being exchanged by another person (e.g., Chapters 3 on "Social Reactions to Socially Interactive Agents and Their Ethical Implications" [Krämer and Manzeschke 2021], 9 on "Theory of Mind and Joint Attention" [Perez-Osorio et al. 2021], and 12 on "Rapport Between Humans and Socially Interactive Agents" [Gratch and Lucas 2021] of volume 1 of this handbook [Lugrin et al. 2021]).

Even over short encounters, we are beginning to understand the deeper implications of these construed relationships. For example, trusted and likeable SIAs can be more persuasive on human judgements and behaviors [e.g., Desteno et al. 2012]. In a learning context, we are starting to see that children will socially-emulate the behaviors and attitudes of their social robot peer-like playmates with respect to modeling curiosity, affect, growth mindset, creativity, and linguistic expression [Gordon et al. 2015, Kory and Breazeal 2014, Kory-Westlund et al. 2017b, Park et al. 2017b, 2019].

19.4.2 Long-Term Relational SIAs

SIAs can be created with long-term interaction in mind, with features such as memory and personalization that evolve over time from repeated encounters with users [e.g., Bickmore and Picard 2005, Breazeal et al. 2019, Lee et al. 2012a, Leite et al. 2013, 2017, Ostrowski et al. 2019, Singh 2018]. Based on the human social support literature and SIA research in a growing number of studies, a key aspect of why SIAs can benefit human outcomes (e.g., learning, health, wellness, etc.) is their nature as a relational technology, especially over long-term encounters.

We use term relational SIAs to refer to the broader category of relational, personified agents—i.e., all socially interactive agents that can build long-term, social-emotional relationships with users. To enable SIAs to reach their full potential as relational technologies, especially for deployment during long-term interactions in real-world contexts, they need to be autonomous. This increasingly dovetails the design of SIAs with increasingly advanced artificial intelligence. Bickmore and Picard [2005] first introduced the concept of relational agents to refer to virtual humans, primarily explored in healthcare contexts with adults. Kory-Westlund [2019] introduced the term relational AI to refer to autonomous relational technologies, recognizing the expanding range of personified technologies to include social robots, digital assistants, conversational devices, etc. that people are using on a daily basis over extended periods of time. While one could argue that Alexa, for instance, is more transactional than relational today, studies show that people would like to see conversational AIs become more relational [e.g. Lopatovska et al. 2018, López et al. 2018, Ostrowski et al. 2019, Sciuto et al. 2018, Singh 2018].

There are important human-centered design considerations for relational SIAs given that they should support familiar social and relational behaviors in order to be more understandable and relatable to humans. They should be designed to treat people in humanistic ways, and build and maintain relationships in a way that are natural, appropriate, and ethical for people (see Krämer, Chapter 3 on ”Social Reactions to Socially Interactive Agents and Their Ethical Implications” [Krämer and Manzeschke 2021] of volume 1 of this handbook [Lugrin et al. 2021]). From a computational perspective, relational SIAs need computational models, algorithms, and mechanisms to update a model of the person(s) with whom it operates in order to build and maintain a relationship over time.

Kory-Westlund [2019] identifies the following features associated with relational AIs (SIAs) to be human-centered, collaborative, interpersonal, relational, and reciprocal. Features that are necessary and sufficient to be relational per her criteria include repeated encounters, shared experiences, mutual change, responsiveness, emotion and positive affect, and reciprocity. These are all features that tend to be associated with relationships.

Repeated encounters. Relationships are longitudinal—they generally develop through time and involve multiple interactions. Relational SIAs should be designed to handle repeated interactions with users through time.

Shared experiences. Humans generally have a sense of past, present, and future, which is reflected in our relationships. We acknowledge our shared experiences through time via references to our past and present together, as well as looking forward to future activities we might do together. For example, sharing a humorous experience during an initial encounter with a stranger led to increased ratings of closeness [Fraley and Aron 2004]. Relational SIAs should create and reference a shared narrative with users. This may require an internal state that represents the user over time that can be updated during interactions.

Mutual Change. As part of creating and referencing shared experiences, relational AI should change over time. More specifically, relational AI should change as a result of the interaction with the user over time—it is not enough to follow a changing but scripted storyline [e.g., Gockley et al. 2005]. The change has to be perceived as “meaningful” in that the activities performed with the user (i.e., shared experiences over repeated encounters) must be clearly seen to affect the relational AI’s outward attitudes, emotions, or behavior. For example, people in close relationships may converge toward similar emotional reactions to events [e.g., Anderson et al. 2003] or similar choices of food [Bove et al. 2003]. Again, this may require an internal state that represents the user over time. It is not sufficient that the person changes over time in response to the SIA, but it should also change. There is a growing number of studies are examining autonomously changing/personalizing the robot’s behavior and/or the task content as a result of the child’s behavior or performance [e.g., Gordon et al. 2016, Lubold 2017, Lubold et al. 2016, 2018, Park et al. 2017a,b, 2019, Ramachandran and

Scassellati 2015, Scassellati et al. 2018a]. These studies have shown that personalization (i.e., a particular kind of change) can increase children's engagement and learning, and have opened many questions about how personalization and change might affect the child-robot relationship.

Responsiveness. Relational SIAs should ideally model a positive relationship. One element of successful, positive human relationships is rapport [Berscheid and Reis 1998], which is often indicated via behavior such as entrainment/mirroring and social reciprocity [Davis 1982, Dijksterhuis 2005, Dijksterhuis and Bargh 2001]. These behaviors are part of being responsive to users. Relational SIAs should respond and react to users, e.g., by using appropriate social cues in the moment, or personalizing its feedback, entrainment, or behavior for individual users [e.g., Cassell et al. 2007b, 2009, Sinha and Cassell 2015b].

Emotion and positive affect. As a human-centered technology, relational SIAs should respond appropriately to users' emotional states. Prior work has found that mismatches between users' emotions and the reactions of technology can negatively affect user perceptions and performance during interactions [Jonsson et al. 2005]. Promoting trust can be important for many kinds of applications. As one example relevant to education, trust can affect who children treat as credible informants [Harris 2007, 2012]. Relational AI designed to act as a friend-like agent may also need to promote positive affect or attachment as well [e.g., Leite et al. 2012a, 2014]; as discussed earlier, children's friendships often involve empathy and affection [Gleason 2002].

Reciprocity. The idea of social reciprocity relates back to responsiveness as well as shared experiences through time. As discussed earlier, relationships often involve various reciprocal behaviors, such as disclosing information, helping, conversing and engaging in activities together, and providing companionship. Relational SIAs should use these kinds of reciprocal behaviors, and attempt to recognize and be affected by the user's use of these behaviors in turn.

19.5 History of Long-Term Interaction with SIAs

We present a brief history/overview of long-term interaction with SIAs, where we shall focus on work where a person interacts with a SIA for at least 5 sessions (see our discussion earlier), in a real-world environment, and where the SIA is autonomous. We include a range of social embodiments in this survey, from social robots, to virtual agents, and voice assistants housed in smart devices. We focus on prior works that study the social or relational aspects of the SIA and its impact on human behavior, engagement, and desired outcomes. Hence, we exclude research about long-term deployments that study multiple one-time encounters in which the SIAs interact with many different people but do not form a long-term relationship with individuals (e.g., tour guide robots or information kiosk agents in public spaces). By doing so, we hope to provide a different lens on the evolution of long-term SIA research—in contrast to other chapters in this handbook that survey SIAs in specific applications, or survey the design of social interaction capabilities (often investigated over short-term or single session encounters). We first highlight long-term SIA work within each category of SIA embodiment (i.e., social robots, virtual agents, and voice assistants) as the work in each area has evolved differently. We then present larger trends across and between each category, highlighting key milestones in long-term SIA research.

19.5.1 Long-Term Interaction with Social Robots

(a) Various forms and functions of long-term social robots for educational tasks. From left to right, Tega exchanges storytelling with a child [Park et al. 2019] (image by ©2021 Hae Won Park), NAO supports second language learning [Vogt et al. 2019] (image from H2020 L2TOR project funded by the European Commission http://www.l2tor.eu/), and Keepons [Leite et al. 2015] support group learning.

(b) Various forms and functions of long-term social robots for health and wellness. From left to right, an emotional wellness companion Paro [Shibata et al. 2009] (image by ©2014 PARO Robots U.S), a daily conversational companion Jibo [Ostrowski et al. 2019] (image by ©2021 Erin Patridge), and SYMPARTNER [Gross et al. 2019] (image by ©2021 SIBIS Institute Berlin), a home robot companion for older adults in single-person households.

Figure 19.1: Long-Term Interactions with SIAs: Social Robots

Research in long-term Human-Robot Interaction (HRI) began in the early 2000's. The first paper published on long-term interaction with robots was an exploratory study with a fetch-and-carry robot called CERO in a work environment [Severinson-Eklundh et al. 2003]. One of the interesting findings was that bystanders also needed to be able to interact with the robot beyond the main user, but didn't know how. The authors raised important design issues such as personality design, natural interaction via voice, and collaboration with the main user as well as with a small group. Sustaining engagement was also an raised as an important issue.

Leite et al. [2013] presents the first survey on long-term HRI covering 24 papers that met their criteria for inclusion from 2003–2011. They used the keywords “long-term interaction”, “social robots”, and ”study”. They only included papers that presented sufficient detail on the capabilities of the robot and study details where the robots were deployed real-world environments such as offices, public places, schools, homes, and healthcare facilities. To expand this survey, we performed a search for relevant empirical work using the list of papers in Leite’s survey as a starting place, and we searched for more recent papers that referenced these earlier works. We also used various literature search tools such as Google Scholar and keywords such as “long-term interaction,” “longitudinal,” “repeated encounters,” “HRI,” “social robot,” “time,” etc.

We found 67 papers published from 2003–2020. Of these, 17 were in the domain of education; 17 were in healthcare, 12 in eldercare, 3 were classified as entertainment applications, and others for general assistance. In Figure 19.1, some examples of robots with various forms and functions are illustrated from the education and health and wellness domains.Fifty four of the 67 studies used fully autonomous robots, and the rest were teleoperated or used shared autonomy. Of the fully autonomous robots, only 25 used advanced autonomy where AI or machine learning was used by the robot to interact (e.g., perception, adaptation, dialog, decision-making, etc.). The remainder of the robots were relatively simple (e.g., reactive behaviors or hard-coded rule-based systems).

This section is not intended to be an exhaustive review, but more to highlight a few key domains and questions in which long-term HRI has been investigated with social robots. We focus on papers that used fully autonomous robots. In the remainder of this section, we focus on two major application domains of interest in long-term HRI research: 1) social robots designed to help people learn, and 2) social robots that help people stay healthy and improve emotional wellbeing.

19.5.1.1 Social robots and children

Social robots have been designed to support long-term interactions with children across a range of applications such as education e.g., [Hyun et al. 2010, Kanda et al. 2004, Lee et al. 2011, Tanaka et al. 2007], therapeutic support such as for autism or physical rehabilitation e.g., [Barakova et al. 2015, François et al. 2009, Kozima et al. 2009, Scassellati et al. 2018b], and health e.g., [Coninx et al. 2016, Kruijff-Korbayová et al. 2015, Short et al. 2014]. A wide range of robot embodiments have been explored from small humanoids (e.g., NAO, QRIO), to expressive characters that move according to principles of animation, to zoomorphic forms, mechanical forms, and more. The majority of long-term studies with children have been in the context of educational activities, with language-skill learning perhaps being the most common [Belpaeme et al. 2018, Randall 2019].

Kanda et al. [2004] presents one of the earliest explorations of social robots in a school setting where a small humanoid robot engaged Japanese elementary grade children over 18 days as a peer-like tutor to help children learn English words. The robot spoke the English words during playful interactions, such shaking hands, hugging, playing rock-paper-scissors, and playing a body-parts naming game. It used RFID nametags to identify individual children and included basic speech recognition and motion control. Children who remained engaged over 2 weeks learned the most English words, but children's engagement substantially waned over time, raising the challenge of sustaining long-term engagement.

A number of studies followed with several results regarding children's long-term engagement. A study by Tanaka et al. [2007] study showed that very young children (10–24 months) socialized with a robot that engaged them in social play (e.g., dancing, giggling in response to touch, sitting and standing, moving its hands). They noted that children appeared to bond with it, and the robot became part of the social ecology of the classroom. Touch-based interactions were among the most enduring, and also led to interesting teacher-child interactions, such as teachers showing the children how to treat the robot gently. This study suggested that children's relationship with the robot helped maintain their engagement.

In other work, Salter et al. [2004] found that children grew bored of a robot that was designed for physical play even within the first few sessions if the robot's behavior was too repetitive. They changed the robot's speech and behavior in later sessions and found that the increased variation improved interaction. Selecting different activities based on the child's interests is another way to increase the variation in the interaction and can also improve engagement. For example, Coninx et al. [2016] found that switching between several different activities helped engage children in diabetes education over time. Further, as different children prefer different activities, adapting to switch activities to suit individuals' preferences was also found helpful in maintaining engagement. Leite et al. [2015] highlights the importance of affective engagement, responding to users expressed emotions, and reinforcement-based adaption on long-term engagement in weekly sessions over 5-weeks.

Baxter et al. [2017] investigated the impact of personalizing a peer-like robot's social behavior on children's learning of novel or familiar subjects in a primary school classroom over a 2-week period. They found that social personalization (referring to the child by name, adapting speed of response) and adapting the speed of progression through educational material improved children's acceptance of the robot as well as children's learning of novel material. Over time, more sophisticated personalization algorithms based on reinforcement learning have been deployed in randomized controlled trials in schools over multiple months to improve children's affective engagement over time [Gordon et al. 2016], and even personalizing to improve both engagement and learning simultaneously [Park et al. 2019]. Kory-Westlund and Breazeal [2019b] discovered a positive correlation between personalization and the quality of children's reported relationship with a peer-like robot learning companion and children's resulting vocabulary learning outcomes. Namely, children who reported a closer relationship with a social robot learned more, and children who had a close relationship with a personalized social robot learned the most.

19.5.1.2 Social robots for health and wellness

The first long-term randomized control trial (RCT) study in the home was with Autom, a robot health coach that helped people to manage their weight [Kidd and Breazeal 2008]. However, the majority of papers that explore social robots for health and wellness has focused on aging. The rising number of older adults, and a growing shortage of clinical and non-clinical care providers, has motivated the development of social robots to address a wide range of issues and opportunities in the aging domain. Overall, people prefer to age-in-place in their own home for as long as they can, and there is a need and desire for affordable technological solutions that help older adults age with independence. A number of social robots have been developed to provide support through social means to address physical decline, cognitive decline, health management, or psychosocial issues such as chronic loneliness or depression. The ability of a social robot to play the role of a motivating coach that can build rapport, track progress toward goals, and provide reminders has been explored in the context of physical rehabilitation, medication adherence, or serving as a health coach. Providing an educational function is relevant for healthcare, but also is generally important to provide cognitive stimulation to help mitigate cognitive decline. To support people's emotional well being, social robots can entertain, help to socially connect people via telepresence or other means, as well as provide a sense of companionship.

These functions are relevant in assisted living contexts, too, where supporting the care staff is another important design dimension. Social robots for older adults have been designed with different physical embodiments and roles. For instance, zoomorphic robots have been studied over long-term encounters as pet therapy surrogates where affiliative touch is an positive emotion-eliciting interaction [Sung et al. 2015, Wada and Shibata 2007]. Other designs have an anthropomorphic embodiment that are capable of sharing gaze, making gestures, or using emotive expressions—sometimes with a accompanying touch screen to support a GUI integration in addition to spoken interaction. Designs can be more functional like a mobile kiosk such as SYMPARTNER [Gross et al. 2019]. Or, they can merge the qualities of a helpful ally with those of a companion, such as Jibo [Ostrowski et al. 2019]. The ability to express emotion through multiple modalities (e.g., body movement, sound, and facial expression) contributes to people's willingness to form an emotional bond with them as well as for the robot to provide emotional support. For instance, Jeong et al. [2020] reports on a 1-week study of the Jibo robot in undergraduate dormitories to serve as an emotional wellness coach to help students cope with stress and to promote emotional resilience. In general, careful consideration and alignment of how the physical appearance of the robot matches its intended function is important for user acceptance and long-term adoption.

The most extensive long-term studies with older adults have been with Paro, often with people experiencing cognitive decline or loneliness. These studies have verified Paro's ability to provide pet-like companionship and emotional support similar to the benefits of animal-assisted therapy, e.g., reducing stress and anxiety, reducing loneliness, and improving positive mood [Wada and Shibata 2007]. Mobile or legged social robots have also been explored in the context of being a component of a smart home designed for aging-in-place where acceptance by users over time has been of primary interest [Doering et al. 2016, Gross et al. 2012, 2019, Hebesberger et al. 2017, Schroeter et al. 2013, Torta et al. 2014]. These robots often have a touch screen with a graphical user interface (GUI) to supplement a virtual user interface (VUI), or the GUI is the only mode of input. Given the exploratory nature of these long-term studies (lasting a couple of days, with multiple short-term encounters), the main research question focused on acceptance of the robot and what factors could lead to longer-term adoption and successful use.

Similar to what has been learned in long-term robot-child studies, it has been found that a change in the robot's speech and behavior can help maintain user engagement and build a long-term relationship in aging, health, and wellness applications [Kidd and Breazeal 2008, Lee et al. 2012a]. Moreover, it will be important that social robots support more robust and flexible conversational abilities, they should be able to adapt and personalize to the users, and they need to support a variety of functions and tasks (both for the main users as well as caregivers). As social robot platforms have matured, akin to the smart speaker market, recent studies have deployed commercial social robots with more robust VUIs and a multitude of skills—e.g., information skills, entertainment skills, and social/persona driven skills. Ostrowski et al. [2019] reports on a 3-week study where the Jibo robot was deployed in community areas of an assisted living facility to support daily interactions with an among residents. They report that the robot served as a social catalyst to positively influence the social connectedness within the community of older adults as well as enhance senior citizens' community engagement.

19.5.1.3 Living with Consumer Robots

There are a few long-term co-habitation studies of robots sold in the consumer market for the home. The first such studies explored how people adopted and used an autonomous floor vacuum cleaner (Roomba) over several weeks [Forlizzi and DiSalvo 2006] up to 6 months [Sung et al. 2009a]. Forlizzi and DiSalvo [2006] used an ethnographic approach to investigate how people interacted with the Roomba and how the presence of the robot changed their housekeeping practices. Although the Roomba is not a social robot, nonetheless, they found that half the participants developed a relationship with it. For instance, people named their Roomba, talked to it even though it had no voice interaction capability, made attributions about how pets related to the robot, demonstrated politeness to the robot (e.g., saying "excuse me" if they bumped into the robot), and would collaborate with the robot to clean as a team (e.g., either helping the robot by picking stuff off the floor, or cleaning the room while the robot cleaned the floor). de Graaf et al. [2017] studied why participants stopped using a desktop robot (Kartoz) over a 6-month period. They report that a positive emotional experience with the product is important to capture users in the short-term (e.g., the first couple of weeks). However, relevant functionality is key to retain usage in the longer-term (e.g., the first couple of months). Over the long-term (e.g., 6-months) users would replace the robot and use a different technology for the same function if they found the robot experience to be too annoying, too cumbersome or frustrating to use, or too boring and repetitive in comparison to the alternative.

More recent cloud-connected social robot platforms (such as Jibo) incorporate the technology advances of voice assistants (e.g., wake word, far-field speech, natural language understanding, regular content or skill updates, in-depth persona design, etc.) to address many of the shortcomings of these older consumer robot technologies. Singh [2018] presents comparative results from a 1-month in-home study with the Jibo social robot where they studied different engagement patterns between generations (children, younger adults and older adults), as well as compared engagement and usage patterns to the Alexa smart speaker (see Section 19.6).

19.5.2 Long-Term Interaction with Virtual Agents

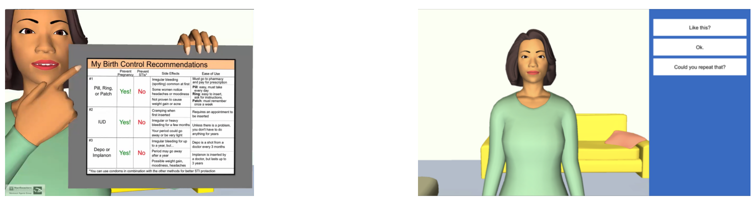

Examples of long-term embodied virtual agents for health change and wellness. From left to right, preconception care [Jack et al. 2020] and depression counseling [Ring 2017]. Images by ©2021 Timothy Bickmore.

Figure 19.2: Long-Term Interactions with SIAs: Embodied Virtual Agents.

Research in long-term interaction with virtual agents also began in the early 2000s. Bickmore [2003] first introduced the concept of *relational agents*: computational artifacts that build long-term, social-emotional relationships with users. For the purposes of this section, we surveyed papers on Google Scholar by searching on the key words "embodied agents", "conversational agent", " virtual agent", and "IVA" with the terms "long-term" or "longitudinal". Based on our interaction threshold of 5 interactions or more, we found 24 papers published from 2003-2020. Bickmore and colleagues' work has dominated research in investigating long-term interaction between users and virtual humans. Much of this work has been applied in the health and wellness domains for adults and older adults (see Figure 19.2 for examples). The approach has focused on the design and evaluation of virtual humans that provide constructive and therapeutic support to users through verbal and non-verbal behaviors that build rapport, motivate, coach, and educate. Overall, use cases of virtual relational agents have primarily focused three main areas of application: 1) health behavior change through motivation, 2) therapeutic/social companionship, and 3) health behavior change through education. We provide research highlights below, but this is not intended to be comprehensive.

19.5.2.1 Virtual agents and health change

The first application relational agents investigated the use of a virtual human as a health coach to help motivate its user to sustain an exercise program, namely to walk for at least 30 minutes a day for most days of the week [Bickmore et al. 2005]. A 1-month randomized controlled trial was carried out where users interacted with one of three interventions: a virtual human with relational behaviors, a virtual human without relational behaviors, and baseline with no agent. The agent in the relational condition used verbal rapport-building behaviors (social dialog, humor, etc) and non-verbal immediacy behaviors (close conversational distance, gaze, facial orientation, etc). All conditions included the standard behavioral interventions, self-monitoring, and educational content. They found that there was no significant difference in the amount of exercise that people performed between the virtual agent conditions, but people did exercise more compared to baseline. Furthermore, subjects in the relational condition liked the agent more, reported a closer relationship, and responded more favorably to continuing to work with the agent.

One drawback of Bickmore et al. [2005] was that the agent's dialog was deemed too repetitive, reducing motivation among many participants. A follow-on study lasting several months explicitly examined the effect of dialog repetition on behavior change [Bickmore and Schulman 2009]. Users were divided into two groups where half interacted with an agent with a variable dialog structure for several months before switching to interact with an agent with a non-variable style for about the same amount of time. In the variable dialog condition, the agent could pick one of 5 different dialog structures (e.g., "Looks like you met your exercise goal of 5,000 steps. Great job!", "Looks like you got your walking in and met your goal of 5,000 steps!", etc.) during each interaction while in the non-variable condition, the agent used the exact same dialog structure in every situation. The other group did the reverse order. Results showed that participants were significantly more likely to have and continue a conversation with the agent in the variable condition. However, participants walked a significantly greater number of steps in the non-variable condition.

A similar pattern was observed in another follow-on study to explore the effect of the virtual agent having personal back story on user's engagement [Bickmore and Schulman 2009]. The agent in the first-person condition presented back stories in the interaction as its own. In the third-person condition, the agent presented back stories about a friend. Participants in the first person condition reported significantly greater enjoyment and willingness to interact again. However, again subjects in the third-person condition walked significantly more steps. Yin and Bickmore [2018] evaluated the effects of cultural adaptation on behavior change in a 4-month study with latino adults in an under-served population. The agent itself was culturally adapted in multiple ways: it resembled a latina woman, participants could choose their preferred language between English and Spanish, and as part of social dialog, the agent demonstrated knowledge about Latino culture. Results showed that participants reported a high satisfaction with a significant increase in minutes of walking per week compared to the control arm.

In perhaps the longest RCT study with an SIA, Bickmore et al. [2013] performed a large-scale, 1-year study comparing the effectiveness of relational agents against computer-tablet based intervention with older adults. In the the control condition, participants uploaded their pedometer readings to a computer. In the virtual exercise coach condition, participants uploaded pedometer data to a computer for the first 2 months, and then interacted with the virtual coach daily for the next 10 months in a kiosk in a clinic waiting room. They found that participants in the virtual agent condition walked significantly more than those in the control condition after the first couple of months, but this trend waned over time, and there was no significant difference in the two groups at the end of 12 months. A closer look revealed that participants with adequate health literacy in the virtual coach condition significantly benefited from the interaction and walked more than the control condition at the 2-month and 12-month interview period, while those with inadequate health literacy showed little to no improvement even after interacting with the agent at either interview point.

A number of studies followed where virtual agents were studied in the context of promoting health literacy and medical adherence [Gardiner et al. 2013, Jack et al. 2015, 2020, Kimani et al. 2016]. For instance, Gardiner et al. [2017] developed a virtual agent to promote mindfulness and lifestyle education among urban women. The agent provided mindfulness exercises, positive ways of managing stress, suggestions to increase physical activity, and motivation to eat healthy. In a 1-month RCT, participants in the experimental condition interacted with the virtual agent. In the control condition, participants were provided with written educational information and meditation audio files to listen to. Results showed that women in the experimental condition significantly reduced their alcohol consumption while increasing their intake of fruits by an average of 2 servings.

19.5.2.2 Virtual agents and wellness

Emotional wellness has been another application domain for long-term studies with virtual agents with affect-aware capabilities to explore their potential benefit in therapeutic/social companionship. For instance, research has shown that older adults with strong social connections have decreased health risks and mortality, raising the question of how SIAs might be able to help alleviate problems associated with social isolation. For instance, in a 1-week study with older adults, Ring et al. [2012] developed a relational agent to help mitigate loneliness. The agent used relational behaviors, motivational dialogue, short anecdotal stories in its interactions, and sensed the users mood to provide appropriate emotional feedback. The study explored whether the virtual agent should be passive (wait for the user to initiate interactions) or proactive where the agent could initiated interactions, too. The results showed that while there weren't significant differences in the length of interactions between the two conditions, participants in the proactive group saw a greater decrease in loneliness. There was also positive correlation between comfort with the agent and time spent with the agent. Finally, the more lonely a participant was, the more they interacted with the agent, reducing their loneliness as a consequence.

In another long-term study, Ring [2017] explores the use of a virtual agent for depression counselling. Participants were divided into three conditions: control, standard, and affective. In the affective condition, the virtual agent responded to its user based on his/her emotional state (e.g., feelings of anger, shame, fear anxiety, etc), and it detected emotional discrepancies to provide appropriate feedback for emotionally-sensitive dialog. In the standard condition the agent did neither, and participants did not interact with any agent in the control condition but simply filled out weekly PHQ-8 and state anxiety questionnaires. While no differences were detected in PHQ-8 scores, the results indicated significant reduction is state anxiety scores in the affective condition compared to the standard and control conditions. No significant differences were detected between the affective and standard conditions in terms of likeability, willingness to continue, interest with all measures tending in the positive direction for both conditions. In both conditions, participants found the agents to be caring and saw them more as a friend than a stranger. These results further motivate the use of virtual agents as social companions to alleviate anxiety and provide therapeutic benefits to users.

19.5.3 Long-Term Interaction with Internet of Things (IoT) Voice Assistants

Personified voice assistants, e.g., digital assistants with a voice interface, embedded in smart devices such as smartphones, speakers, or displays are by far the most common SIA with millions of products sold. In contrast to social robots and virtual humans, they often have no visual form (or a simple, abstract visual representation) and voice is the primary way the persona is conveyed. Siri was a spin-off project originally developed at SRI International, was first released as a voice app in 2010, was quickly acquired by Apple, Inc., and then released in the iPhone 4S in 2011. The first release of a voice assistant on a smart speaker was in 2014 with Alexa on the Amazon Echo.

Other companies have followed with their versions of voice assistants. For instance, the Google Assistant was unveiled in 2016 in the Google Home, and many other types of smart devices with voice assistants have since entered the IoT consumer market. These smart devices are "always listening" for their wake word and support far-field speech. These innovations have proven to lower the barrier to access making it easy for all sorts of people, from young children to older adults, to quickly and easily launch skills (akin to mobile apps) through simple voice commands. Developer ecosystems have proven successful in populating these smart devices with literally thousands of digital skills and services—from playing music, setting timers, controlling other IoT devices, ordering pizza, getting news, Q&A, and much more.

Significant effort has gone into the persona design of these agents, and people enjoy having digital assistants tell jokes and engage in small talk or chit-chat with them. Consumers have shown broad willingness and pleasure in asking these agents to offer their own “opinions” on a wide range of topics, to express their own likes or dislikes, and to express their associated “emotions”). It has been noted that users readily personify voice assistants, ascribe a gender to them, and refer to them in human-like terms such as “friend” and “someone to talk to” [Pradhan et al. 2018, Turk 2016]. They are cloud connected and support over-the-air (OTA) updates, so such devices are constantly updated with new content, skills, and personality quips.

Despite that millions of these devices are in people’s homes, the long-term use of these devices has not been extensively studied or understood by the HCI research community. There are a limited number but growing number of long-term study papers. For instance, [Bentley et al. 2018] analyzed voice history logs of over 65 thousand interactions with Google Home devices in 88 early-adopter homes over about 3-months of use. They identified particular patterns of activities by four distinct user groups based on the type of skills and the time of day they used the smart speaker. They also found that users settled pretty quickly into these patterns of use for which commands they tended to favor, and these did not change much over time past the first 3-weeks of use. They also identified that different demographics by age tended to use the Google Home differently. Younger adults were the most active users (age 18–44) favoring music, home automation, chit-chat, Q&A, setting timers/alarms and getting the weather were pretty consistent over time. The device was not nearly so engaging over the long-term for older adults (age 45–64), however music, chit-chat, setting timers/alarms and weather were among the most used features. The authors note that there are opportunities in the design of such devices to help introduce users to new skills in new domains, to support multi-modal interactions (e.g., spoken interface with a screen), to anticipate user’s patterns of use to proactively offer information, and to support deeper agent-based interactions through richer conversation-based interaction.

Users may project a relationship onto voice assistants and form emotional attachments, as was seen for some users in a recent 1-month home study with the Alexa agent [Singh 2018]. Pradhan et al. [2019] studied how older adults perceive personified smart speakers as social agents or companions verses objects such as appliances. They deployed Amazon Echo Dot devices into the homes of adults over the age of 65 for a period of 3 weeks. They found that these personified devices were often treated as social agents, revealed through use of pronouns, or polite behaviors (such as saying “please” and “thank you” to the agent). Interestingly, older adults would fluidly talk about the agent as both being human-like and object-like, depending on the specific encounter. For instance, greeting the user by name encouraged personification whereas interactions that conveyed a lack of “personal touch” encouraged objectifying the agent. Having the voice agent engage in small talk and greetings were more important to users who desired companionship (e.g., were more prone to feelings of loneliness).

Today's commercial voice assistants generally do not attempt to build or maintain a social-emotional relationship, yet. Bentley et al. [2018] notes that the current style of interaction is far more transactional, where users task the agents using voice commands rather than engaging in collaborative conversation. However, research studies with virtual humans and social robots reveal users' desire for relationship and companionship with these personified technologies, especially in areas where deeper engagement is needed—i.e., to promote long-term quality of life and learning outcomes.

19.6 Similarities and Differences in Social Robots, Virtual Agents, Voice Assistants, and other SIAs

When we compare the long-term work using different relational SIAs—such as physical robots, virtual agents, voice agents, and more—we can examine a number of dimensions, including:

- The domain—e.g., healthcare, education, therapy, entertainment;
- The population—e.g., children, adults, the elderly;
- The agent's embodiment — e.g, physical or virtual, humanoid or non-humanoid, etc.
- What was the impact of the SIA on human outcomes?
- What were the broader implications of the work, e.g., for ethics or design?

Both social robots and virtual agents have been developed for a wide range of application domains and for use by a wide range of user populations. The embodiment of a SIA affects the kinds of behaviors it is capable of using, the tasks it can be used for, and the kinds of relationships people may form with it. In terms of long-term studies, a wider range of morphologies have been explored in social robots so far, including zoomorphic—e.g., Paro (seal-like), Aibo (dog-like), and Pleo (dinosaur-like). There are also robots with functional embodiment (e.g., Roomba or Cero), as well as anthropomorphic characters that express according to principles of animation, e.g., Tega (a squash-and-stretch robot) or Jibo (that can strike expressive postures through its line-of-action). Finally, there are humanoid or even android forms—e.g., the humanoid Nao, and the Geminoid that emulates the human form with skin, teeth, and hair. The physicality of robots also affords touch-based interactions, which are particularly observed when robots take a pet-like or character-like embodiment. The wider variety of robot morphologies may lead to a wider range of possible human-agent relationships—pet-like companions, vacuums that people collaborate with to clean a room, learning companions for children, health coaches, and more. Virtual agents, on the other hand, are most often portrayed as virtual humans in long-term studies, though nothing precludes them from taking other forms. The majority of long-term studies with virtual humans has been in health and wellness domains where the agent often serves as a health coach and utilizes relational features to build rapport to sustain engagement [e.g. Bickmore et al. 2018, Bickmore and Picard 2005, Ring et al. 2015, Sidner et al. 2018, Vardoulakis et al. 2012]. There are also a growing number of long-term studies with smart devices with digital assistant personas—e.g., Siri with smartphones, the Google Assistant with smart speakers, Amazon Alexa with smart displays, etc.

We recently surveyed 79 peer-reviewed publications (from 2003–2020) comprising 87 unique comparative studies in human-robot interaction that compared virtual agents to co-present and telepresent robots and smart speakers [Kory-Westlund et al. in review]. The vast majority of these are in the context of short-term encounters. We also performed a meta-analysis of 59 of these studies. We categorized the SIAs studied based on their embodiment (physical, virtual, or a mixed agent with some physical and some digital components) and their perceived physical presence (co-present or distant). Multiple studies also included other kinds of agents, such as humans, tablets, or laptops without an SIA, or a voice-only agent embedded in a static device. A total of 54 studies compared two agents (62.1%); 28 compared three (32.2%); and 6 compared four or more (6.9%). We also categorized whether the types of tasks performed with the agent included physical components (e.g., a Towers of Hanoi or block-stacking physical puzzle), social components (e.g., conversation, storytelling, judging emotions), or digital components (e.g., tasks shown on a screen, digital puzzles).

Overall, examining at the results from both the survey and meta-analysis favoring each agent type, we see a trend that physically embodied co-present robots affect humans more deeply and strongly than virtual agents, telepresent robots, or smart speakers. However, this can be modulated somewhat by the type of tasks performed with the agent. Humans and physically present robots were favored most often, and most often during socially interactive tasks and physical tasks. During digital tasks that focus on information, the results generally did not favor one agent over another. In addition, humans and co-present physical robots led to stronger increases in important social metrics including attention, attraction/liking, empathy, persuasion, and trust. The results of our survey suggest that the robot's co-presence and embodiment—but presence more so than embodiment—seem especially important for interpersonal, social tasks.

To date, there are only a handful of long-term comparative studies that examine different SIA embodiments. While not all measured emotional engagement or relationships—overall, the physical presence of social robots seems to support deeper emotional engagement and often stronger relationship scores.

For example, Kidd and Breazeal [2008] compared a robotic weight loss coach, to a non-embodied computer coach, to self-report using a standard paper log. The goal was to determining which embodiment would be more effective at sustaining engagement. The exact same dialog model and screen interface ran on the robot weight-management coach as on the computer. It was hypothesized that all interventions would help people lose weight, but the real challenge is helping people to keep the weight off. Hence, long-term engagement was the main outcome studied. They performed a six-week in-home study with 45 adults. The results showed that participants were significantly more likely to adhere to the weight loss program with the social robot. Participants also reported a stronger emotional bond with the robot as well as higher ratings for the robot of working alliance, trust, credibility, and engagement.

In a health and wellness application, Sidner et al. [2018] developed a SIA system for home use comparing a social robot (Reeti) to a virtual human (a female avatar) with older adults. The study looked at health, wellness, and social engagement outcomes as measured through conversation, on-screen games, and filling out forms. They tested the system with 26 older adults, each of whom interacted daily with either the robot or the virtual avatar for 30 days. Sidner et al. [2018] asked participants about their satisfaction with the agents, measured overall system usage and time with agent, and asked about a variety of other attributes of the agent (e.g., likeability, trustworthiness). They found that participants generally displayed reasonably positive attitudes toward both agents, and there was a trend toward participants finding the robot more trustworthy and wanting to have more conversations with the robot than with the virtual agent.

In the domain of education, Vogt et al. [2019] developed a humanoid language learning tutor to teach English words to Dutch children, using a tablet-based game. In a 7-session study, 194 children played the tablet game one-on-one with the robot or with the tablet alone. Language skills tests showed no differences in learning outcomes. This may have been because the tablet game played a crucial role in presenting the learning content; the robot's presence may not have been as important. This study did not measure children's relationship or social engagement with the robot or tablet.

Singh [2018] reports a comparative study on how families and individuals interacted with either a social robot (Jibo) or a smart speaker (Amazon Echo) for one month in the home. This study examined differences in how different generations interact with VUI agents over time: children (under 18 years old), adults (between 18–65 years old), and older adults (aged 65 and older). Both agents could perform a variety of tasks from entertainment (music, jokes, etc.), information (weather, news, Q&A, etc.) and social (sharing opinions, likes/dislikes, etc.). Overall, the Amazon Echo could perform many more skills than Jibo (Alexa had thousands of skills while Jibo had around 20). However Jibo was a far more emotively expressive and companion-like agent. For instance, Jibo had a persistent life-like presence, responded to being petted, could dance, express emotions through body posture, turn to look at people, proactively greet users in the morning through face recognition, and could inquire and remember about how family members slept to provide personalized, contextually-relevant responses. Participants in the study could interact with the agent when they liked; they could use any or all of the agent's functionality (classified as social, entertainment, or functional). Singh [2018] found that children and older adults tended to use all three categories of skills on Jibo. Children primarily only used the entertainment skills on the Amazon Echo.

With Jibo, children's reactions and preferences showed that they were drawn to the robot's ability to be a social, other and they treated Jibo more like a companion. Interestingly, overall engagement was sustained better for children and older adults with Jibo where the companion-like social-relational capabilities seemed to be an important driver for this trend. Usage of the Amazon Echo by older adults dropped over time. In contrast, younger adults tended to favor the utility of the Amazon Echo and practical skills tended to drive their usage over time. Overall, the social robot was often classified as a friend or member of the family, as opposed to being classified as an assistant. The robot was also seen as more open, agreeable, and extroverted than Alexa.

In sum, there are few long-term comparative studies, so we cannot draw any strong generalizations at this point. However, these prior works (both over short and long term studies) suggest that the physical co-presence and expressive behaviors of social robots enhances their ability to emotionally engage people, and this can be an important factor in sustaining engagement over time (more so than for virtual agents or smart devices). For tasks that focus on information, or where relational capabilities are less important, we do not see a significant difference in human behavior across embodied agents. However, we could anticipate that for tasks where establishing a long-term relationship is important for outcomes, social robots could lead to better results. Further studies are needed to verify whether or not this is the case.

19.7 Trends in Long-term SIA Research Over the Past 20 Years

From this brief survey on long-term interaction with social robots, virtual humans, and voice assistants we can identify a number of trends in SIA research. First, early studies in long-term interaction were exploratory in nature, focused on basic questions such as user adoption, identifying sustained patterns of use, and reasons that underlie engagement or abandonment. As these factors have become better understood, we see more effort on running long-term randomized controlled trial (RCT) studies to understand how to design SIAs to bring about intended desired outcomes for people. Currently, there are still few studies that last beyond a couple of months even though this may not be long enough to get past the novelty period. We discuss the novelty issue in more detail in Section 19.9.1. Those lasting 6 months or more are very rare. Second, we see the SIAs increasing in their AI and algorithmic sophistication over time. Earlier systems were fairly scripted in their verbal behavior or reactive in their physical behavior, and people would quickly lose interest, especially if there wasn't sufficient practical functionality to entice usage. Recent papers are starting to investigate algorithmic innovations in long-term contexts such as data-driven personalization or context-aware adaptation. Third, the design of SIAs early on tended either focus on information/decision support or emotional support—SIAs were either tutors/coaches or they were pet-like companions. More recently, SIAs are more often dovetailing cognitive, social and emotional support—and are increasingly combining affective computing methods and reinforcement learning techniques with more advanced conversational AI. Finally, the technological platforms are evolving to become far more capable and robust. While earlier systems in the 2000s were pre-coded with fixed behaviors, modern platforms are cloud connected with over-the-air updates, and SDKs to support developer eco-systems. It is now much easier to deploy social robots, virtual agents or voice assistants in the field and capture much finer grained, continuous user data (under Institutional Review Board ethics protocols), and update algorithms or skills on-the-fly. This trend is making it easier run larger and longer-term studies—deploying multiple SIAs over longer-periods of time with more participants (although this work is by no means easy). This platform trend is also making it possible to develop more sophisticated AI methods to understand long-term usage across users (as well as with a specific user) to drive adaptation, personalization, and continuous expansion (e.g., persona backstory, natural language models, knowledge about its users, etc). This mitigates the interaction from getting "stale" as well as providing a repertoire of functional skills. Both help to address key factors that have diminished long-term engagement in the past. This also opens to door to much more sophisticated relational AI research (see Section 19.9.2). It also raises critical issues around their ethical and responsible design (see Section 19.9.3).

19.8 Current Challenges

Developing SIAs for long-term interaction presents a wide range of research questions, design challenges, and shall require algorithmic and technical innovations. In this section, we focus on four key challenges for the research community to address that were raised in the previous section. For each, the challenge is to achieve these beyond a couple of months (with proof-points to date) to 6-months or more (which has been more elusive).

- Sustaining long-term engagement,
- Supporting flexible, engaging conversation,
- Adapting and personalizing to people effectively,
- Achieving long-term beneficial outcomes using relational properties.

19.8.1 Long-Term Engagement

Maintaining engagement over long-term interactions of weeks, months, or years can be critical to the success of the SIAs, especially in domains such as education and healthcare. Numerous factors can help foster long-term engagement.

With respect to SIAs, four of the most important factors are (1) change over time, (2) shared experience, (3) backstory, and (4) design as a social agent. Change can be "scripted"—such as variation in how the agent speaks or acts, activities performed, or backstory revealed over time—as well as "unscripted", such as personalizing different aspects of the interaction in response to the user's behavior. Change over time has been shown to increase engagement, and help maintain and build relationships [e.g. Bickmore et al. 2010, Gordon et al. 2016, Kidd and Breazeal 2008, Kory-Westlund and Breazeal 2019b, Lee et al. 2012b].

Shared experience can be considered part of change and personalization. It can contribute to the sense that the agent "knows you" and help build a relationship. For example, prior work on long-term child-robot interactions has found that children responded positively to the robot referencing shared experience—e.g., using their name, talking about activities performed together, mentioning facts learned about the child such as their favorite color [Kory-Westlund 2019]. Other work has found that including a memory system that can track and reference prior interactions with the user can be beneficial for engagement and positive affect [e.g., Kasap and Magnenat-Thalmann 2010, Leite et al. 2017].

How a SIA is introduced, the stories told about it, and the story told by it all influence human perception of the SIA and their behavior with it [e.g. Darling et al. 2015, Klapper et al. 2014, Kory-Westlund et al. 2016, Stenzel et al. 2012]. Backstory can also be used to add interesting variation to dialogue to help maintain interest and engagement over time, e.g., as was done with the robot receptionist [Gockley et al. 2005]. The agent's story can also be used to help shape users' expectations about the agent through sharing the agent's history, capabilities, and limitations. The story can be used to establish the agent's character, in the same way we learn about other people through conversation and disclosure. This story can be told by people who lead interactions with the agent, such as experimenters, as well as by the agent itself during conversation.

Finally, designing SIAs from the ground up for social interaction with humans will go a long way toward maintaining engagement over time. Social design includes all aspects of the robot's social behavior and communicative abilities—including whether and how it speaks, how it moves, its nonverbal behavior, and its social contingency. Designing a SIA from the ground up with social interaction in mind means considering how to make the agent's facial expressions, movement, gaze, dialogue, and other behaviors understandable to humans. It enables the SIA to be responsive, expressive, and social. All these social behaviors contribute to people's engagement with the SIA as a social other, as well as their trust, relationship, and engagement with it.

The design of social robots as social agents may be more important for agents that interact with children than those that interact with adults. For example, in a recent study of in-home use of a social home robot and a voice-only home assistant, Singh [2018] found that children were more drawn to the entertainment and social capabilities of the agents, while adults were more interested in the agents' functionality and usefulness.

19.8.2 From Voice Interfaces to Engaging Conversation

Most of the SIAs mentioned in this chapter utilize rule-based forms of multi-modal dialog flow (also see Pieraccini, Chapter 5 on "Natural Language Understanding in Socially Interactive Agents" [Pieraccini 2021] of volume 1 of this handbook [Lugrin et al. 2021]; Traum, Chapter

15 on ”Socially Interactive Agent Dialogue” [Traum 2022] of this volume of this handbook). In many of the long-term health coach SIAs mentioned, the user chooses their response from a set of pre-defined options, and the agent decides its response based on rules set by the designer of the agent. This provides for a rather rigid social experience with the agent. More modern NLU approaches use a flow editor tool to design dialog flows, and may use machine learning methods to train models to recognize different intents from user utterances and automatically generate responses (see Chapters 5 on ”Natural Language Understanding in Socially Interactive Agents” [Pieraccini 2021], 6 on ”Building and Designing Expressive Speech Synthesis [Aylett et al. 2021], 7 on ”Gesture Generation” [Saund and Marsella 2021], and 8 on ”Multimodal Behavior Modeling for Socially Interactive Agents” [Pelachaud et al. 2021] of volume 1 of this handbook [Lugrin et al. 2021]).

19.8.2.1 Flexible Dialog

For truly social dialog, SIAs must be able to move away from brittle rule-based systems to more robust and flexible approaches that allow for a more free form conversation. There is a significant body of research in the natural language understanding and generation community that focuses on statistical machine learning approaches to generate next utterance from a given context and history of conversations by relying on a large corpora of conversations. However, these approaches fail when there exists no corpora for the context of the conversation. This is often the case for SIA research deployments. In such cases, a mix of rule-based and machine learning approaches have been proposed. Researchers have also employed approaches where responses are crowd-sourced, for example by having people perform collaborative tasks in multi-player game scenarios, and then applying machine learning or reasoning methods to generate agent responses and plan networks based on human utterances and behavior [Breazeal et al. 2013, Orkin and Roy 2008, 2009].

For example, Kennedy et al. [2017b] developed an embodied agent for social chit-chat that could self-author responses and grow its knowledge and ability to respond overtime. It has been deployed on a robotic platform and as a virtual agent on a phone. They integrate ideas from rule-based, machine learning and crowd-sourcing approaches for automatic dialog generation. In their approach, every conversation between an agent and user translates into a dialog tree that becomes a part of the agents graph database. When an utterance is received, the agent finds the nearest node in the graph (using cosine distance between average word2vec embeddings of words in the sentence) and randomly chooses a child node (i.e., the response to the utterance). When the agent fails to find a nearest neighbor (constrained by a threshold), it ends the conversation and attempts to grow its graph. Apart from the natural growth of the graph from having different conversations, the agent also crowd-sources responses, so that it doesn’t fail the next time it is in the same state (i.e., at the same node). The model was evaluated in a 12-day study where users were encouraged to chat with the system multiple times a day. Results show that with time the number of conversations increased significantly with number of failures significantly lower on the last day compared to the first day. There was also a trend of increased conversation length with time. These results are encouraging, and much more research in conversational AI with SIAs is needed for them to support robust and flexible social conversation that satisfies users expectations and desires.

19.8.2.2 Non-Verbal Cues

Use of appropriate nonverbal cues such as gesture, gaze, displays of affect, behavior mimicry, and turn-taking can drastically improve how engaging, understandable, interactive, and believable conversation with SIAs can be. For example, use of nonverbal mirroring and behavioral mimicry increased a virtual agent's likability and persuasiveness [Bailenson et al. 2005], and use of nonverbal mirroring and affective support decreased frustration and increase flow [Burleson and Picard 2007]. With robots, use of appropriate social cues, social contingency, nonverbal immediacy, vocal entrainment, and expressivity have led to increased learning and trust in the robot as an informant [e.g., Breazeal et al. 2016, Kennedy et al. 2017a, Kory-Westlund et al. 2017a,b, Lubold 2017, Lubold et al. 2018].

Nonverbal cues are important in establishing joint attention, trust, and rapport in both human-human relationships [e.g., Chartrand and van Baaren 2009, Dijksterhuis 2005, Dijksterhuis and Bargh 2001, Harris 2007, 2012, Lakin et al. 2003, Rotenberg et al. 2003, Semin and Cacioppo 2008, Tickle-Degnen and Rosenthal 1990, Valdesolo and DeSteno 2011, Wiltermuth and Heath 2009] (also see Chapter 8 on ”Multimodal Behavior Modeling for Socially Interactive Agents” [Pelachaud et al. 2021] of volume 1 of this handbook [Lugrin et al. 2021]) and in human-agent relationships [e.g., Bell et al. 2003, Breazeal 2002, Breazeal et al. 2016, Gordon et al. 2016, Levitan et al. 2016, Suzuki and Katagiri 2007].

Many SIAs currently use nonverbal cues that are, at least in part, scripted—e.g., rule-based systems that direct the agent to look at the user when speaking or look down at a shared work surface during pauses, display certain facial expressions upon detecting particular user facial expressions, use beat gestures when certain lines of dialogue are played back. More recent approaches use machine learning to automatically generate non-verbal cues (see Chapter 7 on ”Gesture Generation” [Saund and Marsella 2021] of volume 1 of this handbook [Lugrin et al. 2021]) and even to learn personalized policies for non-verbal cue generation. An increasing number of comparative studies examine different nonverbal behavior generation systems and different methods of directing gaze or gesturing in an effort to determine more effective and natural ways for SIAs to communicate. Like with conversation and dialogue, to achieve truly social nonverbal behavior, we need to move from brittle rule-based systems to approaches that allow for more robust, free-form interaction.

One promising direction has been taken by Justine Cassell and colleagues. They have performed several long-term human-human studies in which they have collected data on how humans coordinate behavior during conversation and use nonverbal cues to, e.g., establish rapport and a positive relationship [Cassell et al. 2007a, Sinha and Cassell 2015a,b]. They have then used this data to build models for SIAs in human-agent interaction [Zhao et al. 2014, 2016], thus enabling the SIA to be more reactive and interactive in human-like ways.

19.8.3 Long-Term Adaptation and Personalization

One important aspect of several long-term studies so far is personalization. For instance, tailoring educational content to individuals can lead to greater engagement and improved learning outcomes. This has been seen in HRI with children [Gordon et al. 2016, Kory and Breazeal 2014, Leite et al. 2012b, Palestra et al. 2016, Park et al. 2019, Scassellati et al. 2018a] as well as in other learning contexts, e.g., with virtual agents or with older children and adults [D'Mello et al. 2012, Gordon and Breazeal 2015, Kasap and Magnenat-Thalmann 2012, Leyzberg et al. 2014, Ramachandran and Scassellati 2015, Thrun et al. 1999] (also see Chapter 18 on "Adaptive Artificial Personalities" [Janowski et al. 2022] of this volume of this handbook). So far, personalization has been studied far more often in longitudinal studies than in one-session studies. This is likely because nearly all personalization studies so far have focused on providing personalized educational content or feedback, using the results of the previous sessions to plan out the content or feedback types for the next session. For example, in Leyzberg et al. [2014], two different models of personalization were used to determine which lessons individuals received about how to solve logic puzzles over the course of 4 sessions. One model tallied positive and negative demonstrations of a relevant skill; the other model used Bayesian updates to model the probability of mastering a relevant skill. They found that receiving personalized lessons significantly improved participants' performance in the puzzle-solving task.

In the AutoTutor intelligent tutoring system, the system monitored students' affective and cognitive states and selected actions to increase learning and help students regulate negative emotional states [D'Mello et al. 2012]. It modeled human tutor dialogue styles and used semantic matching algorithms and conversation rules to pick next dialogue moves in the curriculum script. It detected learning-centered emotions, including engagement, boredom, confusion, and frustration, using facial feature tracking, body posture measurements, and contextual cues. It provided feedback via the virtual tutor's affective facial expressions and verbal responses. They found that the supportive tutor increased students' deep learning, but primarily for low-domain knowledge students, and only the first session—i.e., after there was sufficient context to know the student had problems and actually needed support.

Leite et al. [2009, 2012b, 2014] studied how enabling a robot to express empathy and support during a chess-playing activity might increase children's engagement over time. The robot detected children's affect and made assessments about the child's emotional state using facial expressions and the chess game's state. It used this information to select appropriate supportive behaviors, such as providing advice or guidance, reinforcing the child's sense of competence, or showing expressions of caring and empathy. It also stored information about prior interactions with the child and used reinforcement learning to learn what support strategies worked best with each child. In addition, in the earlier work [Leite et al. 2009], a human instructor chose level-appropriate chess exercises for each child. This work showed that personalizing the robot's supportive behaviors to individual children increased children's engagement and their ratings of the robot's social presence and helpfulness.

Multiple studies were preliminary, in that they presented personalization strategies but did not test them in full experimental studies or did not report all results as yet. For example, Serholt and Barendregt [2016] used information about children's affective states to determine the pedagogical strategy. Although that paper did not report learning results, they found that children expressed significant social engagement, and the robot's personalization appeared to increase engagement. In a preliminary case study, Palestra et al. [2016] scaled up the difficulty of several social skills games played by three children with autism (e.g., about eye contact, joint attention, and body mimicry), and stopped leveling up when children were unable to complete a task. Two of the children appeared to benefit from the leveling.

More recently, Park et al. [2019] conducted a 8-session study where 44 children aged 4–7 interacted one-on-one with a fully autonomous robot. The robot told stories and children were asked to retell the stories, thereby practicing language skills and learning new vocabulary. This work found that personalizing the robot's story curriculum improved children's engagement in the interaction that also led to higher vocabulary learning. Children were given language assessments prior to the study, which were used to select the first stories children heard and to select curricula for children in the non-personalized condition. In the personalized condition, children's story retells, task behavior (e.g., answering dialogic questions during the robot's narration), and affective arousal were used as input for a Q-learning algorithm, which selected personalized storybooks for each child at an appropriate syntactic and lexical level, while maximizing for engagement and learning. In this study, it was observed that children who reported of having closer relationship with the robot also achieved higher learning gains and vice versa, and this trend was more significant when the robot interaction was personalized to the child [Kory-Westlund et al. 2018]. This work showed how closer relationship between SIAs and their users can achieve higher long-term beneficial outcomes, and further motivates why relational properties need to be considered in the interaction design of SIAs.

All of the work so far on long-term interactions and personalization provides evidence for several takeaways. First, personalization of curricula, support, and feedback can improve students' learning, engagement, and positive emotions. Agents that provide support and feedback may be seen as having greater social presence and as being more helpful. The relationships students developed with the agents appeared to influence their engagement and interest in further interaction. Including change and variation in the agents' behavior and the learning content over time can also increase engagement and social interaction.

19.8.4 Achieving Long-Term Beneficial Outcomes from Relational Properties

One difficulty in using relational properties in SIAs is determining which properties to use in a particular SIA, and how to design and use said properties effectively to achieve beneficial outcomes. Although research so far suggests that relational properties can indeed lead to, e.g., increased user engagement, learning, or adherence to health-related programs, it is unclear which relational properties may be most helpful in promoting particular desired outcomes—or which may not contribute positively at all. While many studies found social contingency and social interaction being positively associated with increased trust and learning [e.g., Breazeal et al. 2016, Kennedy et al. 2017a, Kory-Westlund and Breazeal 2019b, Lubold et al. 2018], Kennedy et al. [2015] observed that an excessive amount of social behavior by a robot may detract from children's learning during a math-learning activity.

In another example, [Kory-Westlund 2019] found that using a variety of relational properties in an educational social robot lead to increased engagement and learning, compared to a robot that did not use any relational properties. However, the benefits appeared to be moderated by children's gender and overall affiliation with the robot—i.e., children who formed a stronger relationship with the robot tended to engage more and learn more, regardless of whether the robot they played with was using relational properties.

SIAs may need to be designed to use different behaviors—e.g., dialogue, emotional reactions, ways of expressing information—in order to connect with different users and meet their needs in ways that work best for them. More research is needed to disentangle how different relational properties contribute to the achievement of different outcomes. For example, acknowledging shared experiences and showing mutual change/personalization may be more effective at creating a sense that the SIA "knows" you than responsiveness and use of appropriate emotion—or it may be that responsiveness and building rapport may contribute just as much, for different people, or in different situations.

Another challenge in using relational properties in SIAs is measuring them, and tracking the changes over time. So far, relationship measures in SIA research were heavily reliant on participants' self reported surveys or experimenter conducted interviews. The Godspeed Questionnaire Series (GQS) is one of the most frequently used questionnaires in HRI [Bartneck et al. 2009, Weiss and Bartneck 2015]. The GQS consists of five scales that are relevant to evaluating the perception of the interaction with the robot; anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety. Some studies have borrowed measures from social psychology to assess response to SIAs, such as the Working Alliance Inventory [Horvath 1989] or the Dyadic Interaction System [Broadbent et al. 2007, Ickes et al. 1990]. However, these popular scales are focused on evaluating SIAs as useful tools, and they lack in measuring and tracking the change in the relationship and bonding between users and SIAs. Few works have designed new metrics to specifically measure human-SIA relationship, such as measuring young children's relational perception toward a peer-like robot learning companion [Kory-Westlund 2019, Kory-Westlund et al. 2018]. Kory-Westlund et al. [2018] proposed an interview protocol to measure children's perception of robot as social-relational other and developed questionnaires to measure perceived closeness to the robot (inclusion of other in self), a narrative description task comparing existing relationships with others, pets, and toys to relationship with the robot, and a self-disclosure task that measures if children would share something they are not good at to a robot that discloses its incompetence. However, most of these interviews and surveys were conducted by the experimenter pre or post to the study, and there still is a huge lack of automatic behavioral measures for relational properties. Zhao et al. [2014, 2016] showed some potential by building computational models of humans building, maintaining, and destroying rapport through the use of conversational strategies with verbal and nonverbal behaviors. Being able to measure relational cues autonomously during interactions will open up for much advanced SIA personalization and adaption for long-term beneficial outcomes.

19.9 Future Directions

In this section we offer three future directions that will be important for SIA research to address in the coming years. The first is understanding and measuring when SIAs are beyond the novelty period. This is a critical question for long-term SIA research: indeed, when is the SIA truly engaged in long-term interaction? Presently, this is poorly understood, standard metrics don't exist, and long-term studies don't quantify it. The second concerns advancing AI methods for SIAs to realize more socially and emotionally intelligent relational agents with more sophisticated, flexible and robust behavior. This also includes developing more tools to measure and assess the quality of relationship and its impact on human engagement and behavior. The third topic concerns the responsible and ethical design of relational SIA technologies. This is particularly important as many applications and intended benefits are to support vulnerable populations. And given the popularity of voice assistants, with millions of units already sold, this creates a glide path to relational agents in the near future. There is still much to be understood in terms of how our evolving relationships with SIAs will shape our behavior, attitudes, policies, communities, and more.

19.9.1 Understanding Novelty

Currently, there is no consensus on what amount of exposure is needed to say that novelty has worn off, nor consensus on exactly what novelty *is* or how to measure it. Sung et al. [2009a] studied the adoption and sustain use of the Roomba in the home and found that stable usage patters developed after the 2 month mark. de Graaf et al. [2017] studied the long-term use and reasons for abandonment of a simple desktop robot-information device over 6-months in the home. They found that participants had different reasons for ceasing use, and caution against defining the novelty period based on a fixed amount of time. Rather, they recognize that the novelty period will be different for different devices and different use cases. According to the psychology of behavior change literature, it takes about 2 months on average to establish a new behavior—however, the amount of time ranged from 3-weeks to a year depending on the new type of behavior being established [Lally et al. 2010].

It is also worth questioning whether novelty ever does completely “wear off.” In human-human relationships, there seems to always be the potential for some novelty in some level. While interactions may reach some kind of steady state with less novelty, there is potential, e.g., for one’s spouse of thirty years to still cause surprises (e.g., “keeping things fresh”). Furthermore, published works in human-SIA interaction rarely measure when repeated encounters move past a novelty period. Given this is the case, how do we know that research about long-term effects is reporting accurate results?

For instance, is novelty different than unfamiliarity, and if so, how? What are we contrasting novelty with—familiarization? Habituation? Boredom? How does novelty relate to engagement? For example, we could define the novelty effect in SIAs as engagement that is due to newness rather than due to intrinsic qualities of a thing (e.g., a virtual agent, a robot, a talking speaker) being engaging or fun. Then the question is, when does boredom or engagement overtake engagement-from-novelty? However, novelty is not intrinsically associated with either positive or negative valence and could lead an individual into a curious/interested state, or a state of threat/risk [Gillebaart 2012]. There are likely many individual differences in preference for seeking novelty, e.g., children may prefer novel toys and pictures over familiar ones, which may promote development and acquisition of new concepts [Gillebaart 2012]. Boredom-prone people may be more focused on novel experiences and may find them more interesting. Given this, defining novelty in terms of engagement may not make sense.

One useful framework for examining novelty may be Novelty Categorization Theory (NCT). NCT suggests that appraisal of events as novel relates to categorization, in that an event is novel if it does not fit in any existing categories one has [Förster 2009, Förster et al. 2010]. Novel events are processed in a more global processing style that uses broader, more inclusive mental categories in order to assimilate the novel information and integrate it into existing mental categories or knowledge structures. Assimilated knowledge becomes more familiar and likeable. We could use this framework to evaluate people’s categorizations of robots. For example, there are multiple tasks that measure global versus local processing on which people have performed differently when the task is framed as novel versus as familiar, such as the Gestalt Completion Task [discussed in Gillebaart 2012]. Perhaps one could administer these kinds of tasks either during or following a SIA interaction to learn whether people are using a more global or local processing style, and thus, whether they are perceiving the activity with the robot as more or less novel.

Clearly this is a very important topic for long-term interaction research with SIAs. Advancing how we define and assess novelty (or when an interaction is past the novelty phase) is important work that is relatively unexplored in long-term human-SIA work. In the future, any paper claiming long-term effects should empirically determine that the findings are not due to novelty.

19.9.2 Relational AI for SIAs

A great deal of work remains be done to advance relational AI to improve the socio-emotional intelligence and response of relational SIAs over longitudinal encounters to help people collaboratively achieve long-term goals. There is active research in developing and studying key real-time behaviors such as gaze, reciprocity, entrainment, dialog patterns, affective response, and the like that contributes to building rapport, trust, and working alliance. Algorithmic methods need to be advanced to make these behaviors more flexible, robust, adaptable and socially appropriate in terms of style and timing. Relational behaviors, for instance, can be divided up in many different ways: by timescale (e.g., behaviors developing in the present on shorter timescales in matters of seconds or minutes, versus behaviors that develop over longer times, such as days, weeks, or years), or by modality (e.g., verbal vs. nonverbal cues, linguistic vs. non-linguistic). Investigating how these contribute to building an engaging and appropriate relationship over time—beyond days and weeks to months and even years—is needed. Adaptation to context also matters, whether that is about the task, the larger social context such as the setting, or specifics about the individual?

Much work remains to be done on computational methods for personalization to an individual's needs and differences, as well as adaptation to changing contexts. What can and should a relational SIA learn and remember about you (or, forget)? Perhaps the user should have control over this explicitly. The memory capability of SIAs raises both technical and ethical issues. What kinds of memory are needed—for example, memory of facts, events, emotions, and personal preferences or other details? At what point in time should the relational SIA probe a person again to gain more information or reassess what it knows about the user? If people grow and change, the relationship will need to as well. Part and parcel of learning and adaptation is data capture, privacy, security, and data ownership issues.

Also the automatic perception of relational properties in the interaction and tracking changes over time is a remaining challenge. In human to human interactions, use of verbal and nonverbal expressions are direct observable cues that reflect the relationship between the involved parties [Canary and Stafford 1994]. When relationship changes over time, so do the use and style of language, prosodic cues, and facial and body expressions. While much prior work relied on interviews and self-reported surveys for SIA relationship measures, being able to detect and track the behavioral relational properties in real-time during interactions through SIAs will impact the way agents can personalize and adapt its behavior policies for better user engagement and greater beneficial interaction outcomes in the long term.

These kinds of question show that we need better ways of measuring and assessing the state and quality of the human-SIA relationship. Such measures need to also be designed to be appropriate for the human counterpart (e.g., considering age, task, setting, personality, and other contextual factors, etc.). At what point do relationships move from novelty to familiarity and habituation, and how does novelty continue to play a role in relationship continuation (e.g., to keep the relationship interesting?).

19.9.3 Ethical Issues and Design Practices

The promise of relational SIAs that can help and support people in humanistic, high impact ways is a commonly held goal by many researchers in the community. Because of the distinct human-engagement of relational SIAs, they have the potential to engage diverse people in innovative ways across many domains: education, therapy, healthcare, etc. This also comes with ethical concerns about their appropriate and responsible use as well as potential for (unintended) misuse. Relational SIAs such as social robots, embodied agents, and personified smart devices raise many different ethical concerns—most of which are also encountered in other technologies and domains—all at once. Many of the ethical concerns are most contentious with vulnerable populations, such as children and older adults, who potentially also have the most to benefit from relational SIAs. All of these concerns are pressing given that social and relational technology is swiftly entering the market. We highlight a few of these ethical concerns below as well as some ethical design practices to consider, especially when working with vulnerable stakeholders.

19.9.3.1 Social Bonds and Authenticity

One concern that has been expressed about relational technology is that it will replace the social bonds people have (or would have had) with other people [Turkle 2007, 2017]. One part of this concern pertains to deception—i.e., whether relational SIAs are deceptive in their display of relationship, emotions, and empathy, causing people to think, act, and believe that they have emotional and relational capabilities that they do not "really" have [Coeckelbergh Fourth 2012, Picard and Klein 2002, Turkle 2007]. Questions about deception and authenticity are, at the heart, about the effects of deception on people—i.e., deception is a problem because it causes harm to people. One possible harmful effect relates to human attachment to and reliance on relational technology. Will we come to depend on it too much, when we should not, to our social detriment [e.g., Turkle 2007]?

Coeckelbergh [Fourth 2012] argues that what robot ethicists really mean when arguing about emotional deception is either (1) that the robots intend to deceive; (2) that the emotions robots have are not real; or (3) that the robots pretend to be a kind of entity they are not. In the first case, he argues that it is not the robot that intends to deceive but the robot designer, and that designers have a long tradition across many disciplines (literature, video games, movies, etc.) to create *believable characters*. No one is fooled that these characters are "real", though, or if they are, this is generally considered an acceptable kind of deception based on widespread prior art. In the case of relational SIAs, the question is whether people are fooled—and then, whether this is necessarily a bad thing. Research has reported numerous benefits from the relational properties of SIAs from boosting children's learning outcomes, improving engagement in health protocols, serving as a catalyst to promote human-human connection in assisted living facilities, and more. People have deep conversations with chatbots and virtual therapists [Bickmore et al. 2005, Bobicz and Richard 2003, Pontier and Siddiqui 2008]; often, people consider these agents less judgmental than humans [Bickmore et al. 2005, Gratch et al. 2007, Lucas et al. 2014, Utami et al. 2017].

When researchers do probe whether these users actually believe that relational SIA's have human-equivalent emotions we see nuance in their answers, even from young children [Kory-Westlund and Breazeal 2019a, Kory-Westlund et al. 2018]. Thus far, that data suggests that people do not see SIAs as a relationship human-equivalent, but more of a "sort-of" comparison. Young children, for instance, treat robots as social others, apply social judgments to robots, and respond to their social cues in ways similar to how they respond to people. Nonetheless, when explicitly assessed, children seem to place robots in a different "in-between" ontological category than either living or non-living things [Gaudiello et al. 2015, Kahn et al. 2011, Severson and Carlson 2010]. They have shown a moral objection to the object-like treatment of robots, such as putting a robot away in a closet, because of the perception they have formed about the robot as a social other [Kahn et al. 2012]; however, they may also say that like other objects, a person made the robot, people can own robots, and that robots can break [Kory and Breazeal 2014, Kory-Westlund et al. 2016]. In several studies in the early 2000's, children categorized the robot dog Aibo as not a dog and not a robot, but as a "robotic dog"—a dog with robotic attributes [Bartlett et al. 2004, Kahn et al. 2002, Melson et al. 2009, Weiss et al. 2009].

What does it mean to have "real" or "authentic" emotions in the first place? Sherry Turkle, for example, has argued that social robots are inauthentic: they may provoke emotional attachment, trust, caring, and empathy that is not deserved because the relationship and the feelings are not reciprocal [Turkle 2007]. Must a relationship be reciprocal in a human or equal way? Reciprocality in equal measure is not a requirement even of human relationships. People are capable of having many different kinds of relationships, simultaneously: with peers, our children, our parents, our pets, etc. Human-SIA relationships may simply be one more different kind of relationship that we are still figuring out. Given that SIAs are becoming increasingly mainstream (e.g., Alexa, Siri, the Google Assistant, etc.) an important area of ongoing research is to turn these into empirical questions to understand what effects these relationships with relational AI actually have. These are far from simple topics where data captured about human behavior with relational SIAs often reveals greater nuance and complexity than expected. Speculation or theorizing about the goods and bads of authenticity, attachment, and deception are no longer enough.

19.9.3.2 Persuasion and Social Manipulation

Another ethical concern pertains to social manipulation and persuasion. Technologies often mediate and implicitly shape human interaction with and perception of the world, by encouraging or inviting some forms of actions while discouraging or inhibiting others [Verbeek 2006]. As we have discussed, some research has focused on creating SIAs for behavior change in health contexts—e.g., to help someone with particular weight loss goals to stay on track or engage isolated older adults, among others [Bickmore et al. 2018, Kidd and Breazeal 2008, Sidner et al. 2018]. Social robot learning companions that exhibit curiosity [Gordon et al. 2015], creativity [Ali et al. 2019] or a growth mindset [Park et al. 2017b] have been shown to promote the same behaviors and attitudes in children. This kind of change is usually considered acceptable: it is a "positive" change, with the goal of helping people achieve what they want to achieve. When used for "good," then persuasion, in a SIA or in a human, is often seen as a positive attribute—it gets us to the end we want. If used for "bad," it is another story altogether. For example, robots that are used to provide the elderly with shopping assistance may be seen as beneficial [Iwamura et al. 2011], but those that target potential customers may raise some eyebrows [Kanda et al. 2008]. When SIAs enter people's homes from corporations who want to nudge human behaviors to increase profits certainly raises concern, in addition to concerns around data privacy, security, and use of personal data. The IEEE guidelines for Ethically Aligned Design address some of these issues [1].

People are socially manipulative and persuasive all the time with each other—this is part of how social interaction works. However, while it may be acceptable for the car salesman to use conversation and rapport tactics to up-sell expensive features for a new car, an SIA that does the same thing could be considered alarming. The question here is whether being social manipulative or being persuasive is acceptable for technology, and if so, to what degree? Verbeek [2006] argued that persuasion is not an intrinsic property of any technology, but comes from both the designer and the user. He argued that we should assess potentially persuasive technologies on three fronts: (1) whether the intended persuasions are morally justifiable, e.g., that they do not cause harm, and promote beneficence or justice; (2) that the methods of persuasion used are morally acceptable, e.g., that they respect human autonomy; and (3) that the outcomes or consequences of persuasion are morally justifiable. The biggest challenge, here, however, is that people are often not going to agree on what is considered morally acceptable or morally justifiable.

However, we can look to other domains for inspiration on how to handle these ethical quandaries. Marketing and advertising are two domains that frequently raise similar questions about social manipulation using human-made artifacts and face similar challenges regarding lack of consensus about what is ethical behavior [Drumwright and Murphy 2009]. Some marketing agencies have adopted codes of ethics promoting transparency, honesty of relation-

[1] https://standards.ieee.org/industry-connections/ec/autonomous-systems.html, retrieved February 7, 2019.

ships, opinions, and identity—i.e., promoting the idea that they should make sure consumers know when they are being advertised to. Relational SIAs could follow this example of promoting transparency and honesty, e.g., using backstory or dialogue to explain to users what it is capable of and what its goals are for others' behavior. How the SIA talks about itself can continue to remind and reinforce this transparency and appropriate relationship, such as not answering certain questions or doing certain tasks on the grounds of those being as "only appropriate for humans."—an approach adopted by the design team of Jibo, a social robot for the home. Informing and reminding users through the interaction design of relational SIAs can be helpful, but there are additional questions we can raise about how much users can really trust the designers of the SIA technology. Who is held responsible for the behaviors of users, how we can be sure that a technology will not have undesired or unintended effects, and whether being persuasive is in itself an ethical thing for an SIA to do. The development of design guidelines, best practices, and policies are all important areas of ongoing work as SIAs move into the human environment for longer and longer periods of time.

19.9.3.3 Privacy and Security

SIAs for long-term interaction will need to collect data about users—but what data, how much, and how will it be stored and protected? E.g., we raised the question earlier of SIAs needing memory of facts, events, emotions, personal preferences, or other details in order to perform their tasks, learn and adapt, build relationships, and maintain relationships over time. For a SIA, memory is data. The capability of SIAs to monitor and surveil beyond the capacities of human sensing (e.g., through the use of infrared or ultrasonic sensors, or at ranges or distances unavailable to humans on their own) is concerning. Ryan Calo describes three areas of privacy that we should be concerned about: direct surveillance—i.e., SIAs that magnify the human capacity to observe; increased access—e.g., new access to historically protected spaces, like inside homes; and social meaning—e.g., people may act differently as a result of feeling observed and evaluated [Calo 2010].

These issues are not unique to SIAs—they arise with many current technologies, such as laptops and devices in the internet of things [Arnold 2010, Goldman 2015]. Beyond issues of privacy, we also need to be aware of security in how data is collected, transmitted, and stored: e.g., data breaches are increasingly common. Finding satisfactory solutions for issues in data capture, security, privacy, and data ownership will likely require joint action from governments and regulatory bodies regarding, e.g., what surveillance is acceptable in different circumstances (e.g., by SIAs in public spaces), accountability for anyone dealing with protected data (e.g., legal consequences for negligence in protecting user data, similar to HIPAA-protected data), and imposing standards regarding security, encryption, data forensics, and so forth. We may also need designers of SIAs to adopt an ethical code similar to codes that professionals in other fields follow that emphasizes privacy, accuracy, intellectual property, and access [Calo 2010, Riek and Don Howard 2014].

19.9.3.4 Ethical Design Practices

These ethical design questions are difficult to answer. There is compelling opportunity to include philosophers and ethicists more directly in the design of future relational SIAs. We also need designers of SIA technologies to be aware of the ethical and moral issues involved in the things they are creating, and to attempt whenever possible to create technology that supports and affirms people in becoming who they want to be—that supports human flourishing. Towards this goal, we offer a few ethical design principles of relational SIAs:

- Design responsibly. Involve philosophers and ethicists who have specific training in relevant ethical and moral frameworks and applications, in the design of new technology. Also include domain experts in the application area and user demographic to be engaged. Design with empathy and human flourishing in mind—companies are often criticized for designing for "addiction" or only to maximize profit. Rather than creating technology that serves as a "crutch" that people may become over-reliant upon, consider how to design technology that empowers and respects human agency and dignity.
- Be informed by data as well as theory. An increasing number of research studies are exploring questions highly relevant to the ethical design of relational AI, such as questions about engagement, trust, and attachment. We need to use the data from both human-human studies and human-agent studies to learn how people actually form relationships, develop trust, and interact with relational agents, and use these data to inform future design.
- Co-design and involve all stakeholders. There are emerging disciplines in the area of ethical design where stakeholders have an important role and voice at the design table. Work in the area of participatory design and design justice are becoming increasingly relevant and important as SIAs move from research labs into human environments for longitudinal time frames.
- Be transparent and honest. Inform users about what a technology can do and what it will do. Use the technology's packaging, introduction, framing, and backstory to share information and set user expectations appropriately about the technology, its capabilities, and its limitations. Verify that users actually understood the technology's capabilities and limitations.
- Implement security and privacy by design as well as safety by design. Collect only data that are needed for agent to fulfill its tasks, only data that can be sufficiently protected, and only data that are acceptable to users. Be transparent about what set of data are collected, how data are stored and transmitted, and how data are used.

19.10 Conclusion

Relational SIAs have great potential to support people of all ages in areas that can profoundly contribute to quality of life and opportunity—from education and life long learning, to health and wellness, and more. These are all long-term endeavors for people, and each can benefit from having a strong relationships with a supportive ally. Relational AIs have the potential to offer personalized, high-touch support that is far more scalable, accessible, and affordable than hiring human professionals. However, this also raises critical ethical and societal questions as to the responsible and appropriate use of relational AIs—to ensure that this technology empowers people to achieve important personal goals *and* supports our human networks and stakeholders in the process. Our goal should be to use relational SIAs to help all people flourish, to augment and support human relationships, and to enable people to be happier, healthier, more educated, and more able to lead the lives they want to live. Much work remains to be done, and many questions and issues need to be understood and addressed. Long-term SIA research continues to make exciting progress with the potential for positive, transformative impact for quality of life for many.

19.11 Acknowledgements

We want to thank the members of the Personal Robots Group at the MIT Media Lab who helped to survey papers on long-term interaction with social robots. Part of this research was supported by an MIT Media Lab Learning Innovation Fellowship and the National Science Foundation (NSF) under Grants IIS-1122886, IIS-1122845, IIS-1123085, IIS-1523118, CCF-1138986, and a Graduate Research Fellowship under Grant No. 1122374.

Bibliography

E. K. Ackermann. 2005. Playthings That Do Things: A Young Kid's "Incredibles"! In *Proceedings of the 2005 Conference on Interaction Design and Children*, IDC '05, pp. 1–8. ACM, New York, NY, USA. ISBN 978-1-59593-096-5. DOI: 10.1145/1109540.1109541.

M. D. S. Ainsworth. 1969. Object relations, dependency, and attachment: A theoretical review of the infant-mother relationship. *Child Development*, 40(4): 969–1025. ISSN 0009-3920. DOI: 10.2307/1127008.

M. D. S. Ainsworth and S. M. Bell. 1970. Attachment, exploration, and separation: Illustrated by the behavior of one-year-olds in a strange situation. *Child Development*, 41(1): 49–67. ISSN 0009-3920. DOI: 10.2307/1127388.

S. Ali, T. Moroso, and C. Breazeal. Jun 2019. Can children learn creativity from a social robot? In *Proceedings of the 2019 on Creativity and Cognition*, C&C '19, p. 359–368. Association for Computing Machinery. ISBN 978-1-4503-5917-7. https://doi.org/10.1145/3325480.3325499. DOI: 10.1145/3325480.3325499.

C. Anderson, D. Keltner, and O. P. John. 2003. Emotional convergence between people over time. *Journal of Personality and Social Psychology*, 84(5): 1054–1068. ISSN 1939-1315 0022-3514. DOI: 10.1037/0022-3514.84.5.1054.

B. Arnold. 2010. Digital Handcuffs or Electronic Nannies: Children, privacy and emerging surveillance technologies.

M. P. Aylett, L. Clark, B. R. Cowan, and I. Torre. 2021. *Building and Designing Expressive Speech Synthesis*, pp. 173–211. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM, 1. DOI: http://dx.doi.org/10.1145/3477322.3477329.

R. Aylett. 2022. *Interactive Narrative and Story-telling*, pp. 463–491. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563674.

J. Bailenson, K. Swinth, C. Hoyt, S. Persky, A. Dimov, and J. Blascovich. 2005. The Independent and Interactive Effects of Embodied-Agent Appearance and Behavior on Self-Report, Cognitive, and Behavioral Markers of Copresence in Immersive Virtual Environments. *Presence*, 14(4): 379–393. ISSN 1054-7460. DOI: 10.1162/105474605774785235.

E. I. Barakova, P. Bajracharya, M. Willemsen, T. Lourens, and B. Huskens. Dec. 2015. Long-term LEGO therapy with humanoid robot for children with ASD. *Expert Systems*, 32(6): 698–709. ISSN 1468-0394. DOI: 10.1111/exsy.12098.

B. Bartlett, V. Estivill-Castro, and S. Seymon. 2004. Dogs or Robots: Why Do Children See Them As Robotic Pets Rather Than Canine Machines? In *Proceedings of the Fifth Conference on Australasian*

User Interface - Volume 28, AUIC '04, pp. 7–14. Australian Computer Society, Inc., Darlinghurst, Australia, Australia.

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

A. Batliner, S. Steidl, and E. Noth. 2011. Associating children's non-verbal and verbal behaviour: Body movements, emotions, and laughter in a human-robot interaction. pp. 5828–5831. IEEE. ISBN 978-1-4577-0538-0. DOI: 10.1109/ICASSP.2011.5947686.

P. Baxter, J. Kennedy, E. Senft, S. Lemaignan, and T. Belpaeme. 2016. From Characterising Three Years of HRI to Methodology and Reporting Recommendations. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, HRI '16, pp. 391–398. IEEE Press, Piscataway, NJ, USA. ISBN 978-1-4673-8370-7.

P. Baxter, E. Ashurst, R. Read, J. Kennedy, and T. Belpaeme. May 2017. Robot education peers in a situated primary school study: Personalisation promotes child learning. *Plos One*, 12(5): e0178126–e0178126. ISSN 1932-6203. DOI: 10.1371/journal.pone.0178126.

L. Bell, J. Gustafson, and M. Heldner. 2003. Prosodic adaptation in human-computer interaction. In *Proceedings of ICPHS*, volume 3, pp. 833–836.

T. Belpaeme, J. Kennedy, A. Ramachandran, B. Scassellati, and F. Tanaka. Aug. 2018. Social robots for education: A review. *Science Robotics*, 3(21): eaat5954. ISSN 2470-9476. DOI: 10.1126/scirobotics.aat5954.

R. Bemelmans, G. J. Gelderblom, P. Jonker, and L. de Witte. Nov. 2015. Effectiveness of Robot Paro in Intramural Psychogeriatric Care: A Multicenter Quasi-Experimental Study. *Journal of the American Medical Directors Association*, 16(11): 946–950. ISSN 1525-8610. DOI: 10.1016/j.jamda.2015.05.007.

F. Bentley, C. Luvogt, M. Silverman, R. Wirasinghe, B. White, and D. Lottridge. Sept. 2018. Understanding the long-term use of smart speaker assistants. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2(3). https://doi.org/10.1145/3264901. DOI: 10.1145/3264901.

S. Bernardini, K. Porayska-Pomsta, and T. J. Smith. Apr. 2014. ECHOES: An intelligent serious game for fostering social communication in children with autism. *Information Sciences*, 264: 41–60. ISSN 0020-0255. DOI: 10.1016/j.ins.2013.10.027.

E. Berscheid and H. T. Reis. 1998. Attraction and close relationships. In D. T. Gilbert, S. T. Fiske, and G. Lindzey, eds., *The Handbook of Social Psychology, Vols. 1-2, 4th Ed*, pp. 193–281. McGraw-Hill, New York, NY, US. ISBN 978-0-19-521376-8.

T. Bickmore. 2022. *Health-Related Applications of Socially Interactive Agents*, pp. 403–435. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563672.

T. Bickmore and J. Cassell. 2001. Relational Agents: A Model and Implementation of Building User Trust. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '01, pp. 396–403. ACM, New York, NY, USA. ISBN 978-1-58113-327-1. DOI: 10.1145/365024.365304.

T. Bickmore and D. Schulman. 2009. A virtual laboratory for studying long-term relationships between humans and virtual agents. In *Proceedings of The 8th International Conference on Autonomous*

Agents and Multiagent Systems-Volume 1, p. 297–304.

T. Bickmore, A. Gruber, and R. Picard. 2005. Establishing the computer–patient working alliance in automated health behavior change interventions. *Patient Education and Counseling*, 59(1): 21–30. ISSN 0738-3991. DOI: 10.1016/j.pec.2004.09.008.

T. Bickmore, D. Schulman, and L. Yin. 2010. Maintaining engagement in long-term interventions with relational agents. *Applied Artificial Intelligence*, 24(6): 648–666.

T. Bickmore, H. Trinh, R. Asadi, and S. Olafsson. 2018. Safety First: Conversational Agents for Health Care. In R. J. Moore, M. H. Szymanski, R. Arar, and G.-J. Ren, eds., *Studies in Conversational UX Design*, Human–Computer Interaction Series, pp. 33–57. Springer International Publishing, Cham. ISBN 978-3-319-95579-7. DOI: 10.1007/978-3-319-95579-7_3.

T. W. Bickmore. 2003. *Relational Agents : Effecting Change through Human-Computer Relationships*. Thesis, Massachusetts Institute of Technology.

T. W. Bickmore and R. W. Picard. June 2005. Establishing and Maintaining Long-term Human-computer Relationships. *ACM Trans. Comput.-Hum. Interact.*, 12(2): 293–327. ISSN 1073-0516. DOI: 10.1145/1067860.1067867.

T. W. Bickmore, R. A. Silliman, K. Nelson, D. M. Cheng, M. Winter, L. Henault, and M. K. Paasche-Orlow. 2013. A randomized controlled trial of an automated exercise coach for older adults. *Journal of the American Geriatrics Society*, 61(10): 1676–1683.

F. Biocca, C. Harms, and J. Burgoon. 2003. Toward a More Robust Theory and Measure of Social Presence: Review and Suggested Criteria. *Presence*, 12(5): 456–480. ISSN 1054-7460. DOI: 10.1162/105474603322761270.

K. P. Bobicz and D. C. S. Richard. 2003. The virtual therapist: Behavior therapy in a digital age. *Behavior Therapist*, 26(3): 265–269.

C. F. Bove, J. Sobal, and B. S. Rauschenbach. Feb. 2003. Food choices among newly married couples: Convergence, conflict, individualism, and projects. *Appetite*, 40(1): 25–41. ISSN 0195-6663. DOI: 10.1016/S0195-6663(02)00147-2.

J. Bowlby. 1958. The nature of the child's tie to his mother. *Int. J. Psycho-Anal.*, 39: 350–373.

C. Breazeal. Jan. 2002. Regulation and Entrainment in Human-Robot Interaction. *The International Journal of Robotics Research*, 21(10-11): 883–902. ISSN 0278-3649, 1741-3176. DOI: 10.1177/0278364902021010096.

C. Breazeal, N. DePalma, J. Orkin, S. Chernova, and M. Jung. Feb 2013. Crowdsourcing human-robot interaction: new methods and system evaluation in a public environment. *Journal of Human-Robot Interaction*, 2(1): 82–111. DOI: 10.5898/JHRI.2.1.Breazeal.

C. Breazeal, P. L. Harris, D. DeSteno, J. M. Kory-Westlund, L. Dickens, and S. Jeong. 2016. Young Children Treat Robots as Informants. *Topics in Cognitive Science*, pp. 1–11. ISSN 1756-8765. DOI: 10.1111/tops.12192.

C. Breazeal, A. K. Ostrowski, N. Singh, and H. Park. 2019. Designing social robots for older adults. *URL: https://www. nae. edu/Publications/Bridge/205212/208332. aspx. Accessed*, 16: 2019.

S. S. Brehm. 1992. *Intimate Relationships, 2nd Ed*, volume xviii of *The McGraw-Hill series in social psychology.* Mcgraw-Hill Book Company, New York, NY, England. ISBN 978-0-07-007443-9.

I. Bretherton. 1992. The origins of attachment theory: John Bowlby and Mary Ainsworth. *Developmental Psychology*, 28(5): 759–775. ISSN 1939-0599(Electronic);0012-1649(Print). DOI: 10.1037/0012-1649.28.5.759.

E. Broadbent, B. MacDonald, L. Jago, M. Juergens, and O. Mazharullah. 2007. Human reactions to good and bad robots. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3703–3708. IEEE.

J. Broekens. 2021. *Emotion*, pp. 349–384. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM, 1. DOI: http://dx.doi.org/10.1145/3477322.3477333.

J. K. Burgoon and J. L. Hale. 1984. The fundamental topoi of relational communication. *Communication Monographs*, 51(3): 193–214. ISSN 1479-5787 0363-7751. DOI: 10.1080/03637758409390195.

W. Burleson and R. W. Picard. 2007. Gender-specific approaches to developing emotionally intelligent learning companions. *Intelligent Systems, IEEE*, 22(4): 62–69.

R. Calo. 2010. Robots and privacy. In P. Lin, G. Bekey, and K. Abney, eds., *Robot Ethics: The Ethical and Social Implications of Robotics*, pp. 187–201. MIT Press, Cambridge, MA.

D. J. Canary and L. Stafford. 1994. Maintaining relationships through strategic and routine interaction.

J. Cassell. 2022. *Socially Interactive Agents as Peers*, pp. 331–365. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563670.

J. Cassell and T. Bickmore. Dec. 2000. External Manifestations of Trustworthiness in the Interface. *Commun. ACM*, 43(12): 50–56. ISSN 0001-0782. DOI: 10.1145/355112.355123.

J. Cassell, A. Gill, and P. Tepper. 2007a. Coordination in conversation and rapport. In *Proceedings of the workshop on Embodied Language Processing*, pp. 41–50.

J. Cassell, A. Tartaro, Y. Rankin, V. Oza, C. Tse, and F. Searle. 2007b. Virtual peers for literacy learning. *EDUCATIONAL TECHNOLOGY-SADDLE BROOK THEN ENGLEWOOD CLIFFS NJ-*, 47(1): 39.

J. Cassell, K. Geraghty, B. Gonzalez, and J. Borland. 2009. Modeling Culturally Authentic Style Shifting with Virtual Peers. In *Proceedings of the 2009 International Conference on Multimodal Interfaces*, ICMI-MLMI '09, pp. 135–142. ACM, New York, NY, USA. ISBN 978-1-60558-772-1. DOI: 10.1145/1647314.1647338.

W.-L. Chang and S. Šabanović. 2015. Interaction Expands Function: Social Shaping of the Therapeutic Robot PARO in a Nursing Home. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '15, pp. 343–350. ACM, New York, NY, USA. ISBN 978-1-4503-2883-8. DOI: 10.1145/2696454.2696472.

T. L. Chartrand and R. van Baaren. 2009. Human mimicry. *Advances in experimental social psychology*, 41: 219–274.

H. Chen. 2018. *Adaptive Role Switching in Socially Interactive Agents for Children's Language Learning*. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.

C. Clabaugh, D. Becerra, E. Deng, G. Ragusa, and M. Matarić. 2018. Month-long, In-home Case Study of a Socially Assistive Robot for Children with Autism Spectrum Disorder. In *Companion of the*

2018 ACM/IEEE International Conference on Human-Robot Interaction, HRI '18, pp. 87–88. ACM, New York, NY, USA. ISBN 978-1-4503-5615-2. DOI: 10.1145/3173386.3177018.

M. Coeckelbergh. Fourth 2012. Are Emotional Robots Deceptive? *IEEE Transactions on Affective Computing*, 3(4): 388–393. ISSN 1949-3045. DOI: 10.1109/T-AFFC.2011.29.

A. Coninx, P. Baxter, E. Oleari, S. Bellini, B. Bierman, O. B. Henkemans, L. Canamero, P. Cosi, V. Enescu, R. R. Espinoza, A. Hiolle, R. Humbert, B. Kiefer, I. Kruijff-korbayova, R. Looije, M. Mosconi, M. Neerincx, G. Paci, G. Patsis, C. Pozzi, F. Sacchitelli, H. Sahli, A. Sanna, G. Sommavilla, F. Tesser, Y. Demiris, and T. Belpaeme. 2016. Towards long-term social child-robot interaction: Using multi-activity switching to engage young users. *Journal of Human-Robot Interaction*, 5(1): 32–67. ISSN 2163-0364.

M. Csikszentmihalyi and E. Halton. Oct. 1981. *The Meaning of Things: Domestic Symbols and the Self*. Cambridge University Press. ISBN 978-0-521-28774-6.

K. Darling, P. Nandy, and C. Breazeal. Aug. 2015. Empathic concern and the effect of stories in human-robot interaction. In *2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 770–775. DOI: 10.1109/ROMAN.2015.7333675.

K. Dautenhahn, A. H. Bond, L. Canamero, and B. Edmonds. 2002. *Socially intelligent agents: Creating relationships with computers and robots*, volume 3. Springer Science & Business Media.

M. Davis, ed. 1982. *Interaction Rhythms: Periodicity in Communicative Behavior*. Human Sciences Press, New York, NY.

M. de Graaf, S. Ben Allouch, and J. van Dijk. 2017. Why Do They Refuse to Use My Robot?: Reasons for Non-Use Derived from a Long-Term Home Study. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '17, pp. 224–233. ACM, New York, NY, USA. ISBN 978-1-4503-4336-7. DOI: 10.1145/2909824.3020236.

D. Desteno, C. Breazeal, R. H. Frank, D. Pizarro, J. Baumann, L. Dickens, and J. J. Lee. Dec. 2012. Detecting the trustworthiness of novel partners in economic exchange. *Psychological Science*, 23(12): 1549–1556. ISSN 1467-9280. DOI: 10.1177/0956797612448793.

A. Dijksterhuis. 2005. Why we are social animals: The high road to imitation as social glue. *Perspectives on imitation: From neuroscience to social science*, 2: 207–220.

A. Dijksterhuis and J. A. Bargh. 2001. The perception-behavior expressway: Automatic effects of social perception on social behavior. *Advances in experimental social psychology*, 33: 1–40.

S. D'Mello, A. Olney, C. Williams, and P. Hays. 2012. Gaze tutor: A gaze-reactive intelligent tutoring system. *International Journal of Human-Computer Studies*, 70(5): 377–398.

N. Doering, K. Richter, H.-M. Gross, C. Schroeter, S. Mueller, M. Volkhardt, A. Scheidig, and K. Debes. 2016. Robotic companions for older people: A case study in the wild. *Annual Review of Cybertherapy and Telemedicine 2015: Virtual Reality in Healthcare: Medical Simulation and Experiential Interface*, 219: 147–52.

A.-B. Doyle, J. Connolly, and L.-P. Rivest. 1980. The effect of playmate familiarity on the social interactions of young children. *Child Development*, pp. 217–223.

M. E. Drumwright and P. E. Murphy. 2009. The current state of advertising ethics: Industry and academic perspectives. *Journal of Advertising*, 38(1): 83–108.

S. Duck. 1991. *Understanding Relationships*, volume vii. Guilford Press, New York, NY, US. ISBN 978-0-89862-758-9 978-0-89862-470-0.

J. Fink, O. Mubin, F. A. A. Kaplan, and P. Dillenbourg. 2012. Anthropomorphic language in online forums about Roomba, AIBO and the iPad. In *Advanced Robotics and Its Social Impacts (ARSO), 2012 IEEE Workshop On*, pp. 54–59. IEEE.

B. J. Fogg and H. Tseng. 1999. The Elements of Computer Credibility. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '99, pp. 80–87. ACM, New York, NY, USA. ISBN 978-0-201-48559-2. DOI: 10.1145/302979.303001.

J. Forlizzi and C. DiSalvo. Mar 2006. Service robots in the domestic environment: a study of the roomba vacuum in the home. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, HRI '06, p. 258–265. Association for Computing Machinery. ISBN 978-1-59593-294-5. https://doi.org/10.1145/1121241.1121286. DOI: 10.1145/1121241.1121286.

J. Förster. 2009. Cognitive consequences of novelty and familiarity: How mere exposure influences level of construal. *Journal of Experimental Social Psychology*, 45(2): 444–447. ISSN 0022-1031. DOI: 10.1016/j.jesp.2008.10.011.

J. Förster, J. Marguc, and M. Gillebaart. 2010. Novelty Categorization Theory. *Social and Personality Psychology Compass*, 4(9): 736–755. ISSN 1751-9004. DOI: 10.1111/j.1751-9004.2010.00289.x.

B. Fraley and A. Aron. Mar. 2004. The effect of a shared humorous experience on closeness in initial encounters. *Personal Relationships*, 11(1): 61–78. ISSN 1475-6811. DOI: 10.1111/j.1475-6811.2004.00071.x.

D. François, S. Powell, and K. Dautenhahn. Jan 2009. A long-term study of children with autism playing with a robotic pet: Taking inspirations from non-directive play therapy to encourage children's proactivity and initiative-taking. *Interaction Studies*, 10(3): 324–373. ISSN 1572-0373, 1572-0381. DOI: 10.1075/is.10.3.04fra.

E. P. Frates, M. A. Moore, C. N. Lopez, and G. T. McMahon. 2011. Coaching for behavior change in physiatry. *American Journal of Physical Medicine & Rehabilitation*, 90(12): 1074–1082.

P. Gardiner, M. B. Hempstead, L. Ring, T. Bickmore, L. Yinusa-Nyahkoon, H. Tran, M. Paasche-Orlow, K. Damus, and B. Jack. 2013. Reaching women through health information technology: the gabby preconception care system. *American Journal of Health Promotion*, 27(3_suppl): eS11–eS20.

P. M. Gardiner, K. D. McCue, L. M. Negash, T. Cheng, L. F. White, L. Yinusa-Nyahkoon, B. W. Jack, and T. W. Bickmore. 2017. Engaging women with an embodied conversational agent to deliver mindfulness and lifestyle recommendations: A feasibility randomized control trial. *Patient education and counseling*, 100(9): 1720–1729.

I. Gaudiello, S. Lefort, and E. Zibetti. 2015. The ontological and functional status of robots: How firm our representations are? *Computers in Human Behavior*, 50: 259–273. ISSN 0747-5632. DOI: 10.1016/j.chb.2015.03.060.

P. Gebhard, D. Tsovaltzi, T. Schneeberger, and F. Nunnari. 2022. *Serious Games with SIAs*, pp. 527–559. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563676.

M. Ghafurian, J. E. Muñoz, J. Boger, J. Hoey, and K. Dautenhahn. 2022. *Socially Interactive Agents for Supporting Aging*, pp. 367–402. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563671.

M. Gillebaart. 2012. *Something Old, Something New: When People Favor Novelty over Familiarity and How Novelty Affects Creative Processes*. ISBN 978-94-6191-330-2.

T. R. Gleason. Nov. 2002. Social provisions of real and imaginary relationships in early childhood. *Developmental Psychology*, 38(6): 979–992. ISSN 0012-1649. DOI: http://dx.doi.org.libproxy.mit.edu/10.1037/0012-1649.38.6.979.

T. R. Gleason and L. M. Hohmann. Feb. 2006. Concepts of Real and Imaginary Friendships in Early Childhood. *Social Development*, 15(1): 128. ISSN 0961205X.

R. Gockley, A. Bruce, J. Forlizzi, M. Michalowski, A. Mundell, S. Rosenthal, B. Sellner, R. Simmons, K. Snipes, A. Schultz, and J. Wang. 2005. Designing robots for long-term social interaction. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2005. (IROS 2005)*, pp. 1338–1343. DOI: 10.1109/IROS.2005.1545303.

D. Goldman, Feb. 2015. Your Samsung TV is eavesdropping on your private conversations. http://money.cnn.com/2015/02/09/technology/security/samsung-smart-tv-privacy/index.html.

G. Gordon and C. Breazeal. 2015. Bayesian Active Learning-based Robot Tutor for Children's Word-Reading Skills. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*.

G. Gordon, C. Breazeal, and S. Engel. 2015. Can Children Catch Curiosity from a Social Robot? In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '15, pp. 91–98. ACM, New York, NY, USA. ISBN 978-1-4503-2883-8. DOI: 10.1145/2696454.2696469.

G. Gordon, S. Spaulding, J. Kory-Westlund, J. J. Lee, L. Plummer, M. Martinez, M. Das, and C. Breazeal. 2016. Affective personalization of a social robot tutor for children's second language skill. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*. Palo Alto, CA.

J. Gratch and G. Lucas. 2021. *Rapport Between Humans and Socially Interactive Agents*, pp. 433–462. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM, 1. DOI: http://dx.doi.org/10.1145/3477322.3477335.

J. Gratch, N. Wang, J. Gerten, E. Fast, and R. Duffy. Sept. 2007. Creating Rapport with Virtual Agents. In *Intelligent Virtual Agents*, pp. 125–138. Springer, Berlin, Heidelberg. DOI: 10.1007/978-3-540-74997-4_12.

H.-M. Gross, C. Schroeter, S. Müller, M. Volkhardt, E. Einhorn, A. Bley, T. Langner, M. Merten, C. Huijnen, H. van den Heuvel, et al. 2012. Further progress towards a home robot companion for people with mild cognitive impairment. In *2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 637–644. IEEE.

H.-M. Gross, A. Scheidig, S. Müller, B. Schütz, C. Fricke, and S. Meyer. 2019. Living with a mobile companion robot in your own apartment-final implementation and results of a 20-weeks field study with 20 seniors. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 2253–2259. IEEE.

P. A. Hancock, D. R. Billings, K. E. Schaefer, J. Y. C. Chen, E. J. de Visser, and R. Parasuraman. Oct. 2011. A Meta-Analysis of Factors Affecting Trust in Human-Robot Interaction. *Human Factors:*

The Journal of the Human Factors and Ergonomics Society, 53(5): 517–527. ISSN 0018-7208, 1547-8181. DOI: 10.1177/0018720811417254.

P. L. Harris. 2007. Trust. *Developmental science*, 10: 135–138.

P. L. Harris. 2012. *Trusting What You're Told: How Children Learn from Others*. Harvard University Press.

W. W. Hartup, B. Laursen, M. I. Stewart, and A. Eastenson. 1988. Conflict and the Friendship Relations of Young Children. *Child Development*, 59(6): 1590–1600. ISSN 0009-3920. DOI: 10.2307/1130673.

D. Hebesberger, T. Koertner, C. Gisinger, and J. Pripfl. Jan. 2017. A Long-Term Autonomous Robot at a Care Hospital: A Mixed Methods Study on Social Acceptance and Experiences of Staff and Older Adults. *International Journal of Social Robotics*, pp. 1–13. ISSN 1875-4791, 1875-4805. DOI: 10.1007/s12369-016-0391-6.

A. O. Horvath and L. Luborsky. 1993. The role of the therapeutic alliance in psychotherapy. *Journal of consulting and clinical psychology*, 61(4): 561.

A. O. G. Horvath. 1989. Development and Validation of the Working Alliance Inventory. *Journal of Counseling Psychology*, 36(2): 223–33.

E. Hyun, H. Yoon, and S. Son. 2010. Relationships Between User Experiences and Children's Perceptions of the Education Robot. In *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction*, HRI '10, pp. 199–200. IEEE Press, Piscataway, NJ, USA. ISBN 978-1-4244-4893-7.

W. Ickes, V. Bissonnette, S. Garcia, and L. L. Stinson. 1990. Implementing and using the dyadic interaction paradigm.

Y. Iwamura, M. Shiomi, T. Kanda, H. Ishiguro, and N. Hagita. 2011. Do elderly people prefer a conversational humanoid as a shopping assistant partner in supermarkets? In *Human-Robot Interaction (HRI), 2011 6th ACM/IEEE International Conference On*, pp. 449–457. IEEE.

B. Jack, T. Bickmore, M. Hempstead, L. Yinusa-Nyahkoon, E. Sadikova, S. Mitchell, P. Gardiner, F. Adigun, B. Penti, D. Schulman, et al. 2015. Reducing preconception risks among african american women with conversational agent technology. *The Journal of the American Board of Family Medicine*, 28(4): 441–451.

B. W. Jack, T. Bickmore, L. Yinusa-Nyahkoon, M. Reichert, C. Julce, N. Sidduri, J. Martin-Howard, Z. Zhang, E. Woodhams, J. Fernandez, M. Loafman, and H. J. Cabral. 2020. Improving the health of young african american women in the preconception period using health information technology: a randomised controlled trial. *The Lancet Digital Health*, 2(9): e475–e485. ISSN 2589-7500. https://www.sciencedirect.com/science/article/pii/S2589750020301898. DOI: https://doi.org/10.1016/S2589-7500(20)30189-8.

K. Janowski, H. Ritschel, and E. André. 2022. *Adaptive Artificial Personalities*, pp. 155–193. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: http://dx.doi.org/ 10.1145/3563659.3563666.

S. Jeong, S. Alghowinem, L. Aymerich-Franch, K. Arias, A. Lapedriza, R. Picard, H. W. Park, and C. Breazeal. 2020. A robotic positive psychology coach to improve college students' wellbeing.

In *Robot and Human Interactive Communication, 2020. ROMAN 2020. 29th IEEE International Conference On*. IEEE.

I.-M. Jonsson, C. Nass, H. Harris, and L. Takayama. 2005. Matching In-Car Voice with Driver State: Impact on Attitude and Driving Performance. In *Driving Assessment 2005: 3rd International Driving Symposium on Human Factors in Driver Assessment, Training, and Vehicle Design*. ISBN 978-0-87414-151-1.

P. H. Kahn, B. Friedman, and J. Hagman. 2002. I care about him as a pal: Conceptions of robotic pets in online aibo discussion forums. In *CHI'02 Extended Abstracts on Human Factors in Computing Systems*, pp. 632–633. ACM.

P. H. Kahn, A. L. Reichert, H. E. Gary, T. Kanda, H. Ishiguro, S. Shen, J. H. Ruckert, and B. Gill. 2011. The new ontological category hypothesis in human-robot interaction. In *Proceedings of the 6th International Conference on Human-Robot Interaction*, pp. 159–160. ACM.

P. H. Kahn, T. Kanda, H. Ishiguro, N. G. Freier, R. L. Severson, B. T. Gill, J. H. Ruckert, and S. Shen. 2012. "Robovie, you'll have to go into the closet now": Children's social and moral relationships with a humanoid robot. *Developmental psychology*, 48(2): 303.

T. Kanda, T. Hirano, D. Eaton, and H. Ishiguro. 2004. Interactive robots as social partners and peer tutors for children: A field trial. *Hum.-Comput. Interact.*, 19(1): 61–84. ISSN 0737-0024. DOI: 10.1207/s15327051hci1901&2_4.

T. Kanda, R. Sato, N. Saiwaki, and H. Ishiguro. Oct. 2007. A Two-Month Field Trial in an Elementary School for Long-Term Human #x2013;Robot Interaction. *IEEE Transactions on Robotics*, 23(5): 962–971. ISSN 1552-3098. DOI: 10.1109/TRO.2007.904904.

T. Kanda, D. F. Glas, M. Shiomi, H. Ishiguro, and N. Hagita. 2008. Who will be the customer?: A social robot that anticipates people's behavior from their trajectories. In *Proceedings of the 10th International Conference on Ubiquitous Computing*, pp. 380–389. ACM.

Z. Kasap and N. Magnenat-Thalmann. 2010. Towards episodic memory-based long-term affective interaction with a human-like robot. In *RO-MAN, 2010 IEEE*, pp. 452–457. IEEE.

Z. Kasap and N. Magnenat-Thalmann. 2012. Building long-term relationships with virtual and robotic characters: The role of remembering. *The Visual Computer*, 28(1): 87–97.

H. H. Kelley, E. Berscheid, A. Christensen, J. Harvey, T. Huston, G. Levinger, E. McClintock, L. Peplau, and D. Peterson. 1983. *Close Relationships*. Freeman, New York. ISBN 978-0-7167-1443-9. OCLC: 902175384.

J. Kennedy, P. Baxter, and T. Belpaeme. 2015. The robot who tried too hard: Social behaviour of a robot tutor can negatively affect child learning. In *Proc. HRI*, volume 15.

J. Kennedy, P. Baxter, E. Senft, and T. Belpaeme. 2016. Social robot tutoring for child second language learning. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 231–238. IEEE.

J. Kennedy, P. Baxter, and T. Belpaeme. Jan. 2017a. Nonverbal Immediacy as a Characterisation of Social Behaviour for Human–Robot Interaction. *International Journal of Social Robotics*, 9(1): 109–128. ISSN 1875-4791, 1875-4805. DOI: 10.1007/s12369-016-0378-3.

J. Kennedy, I. Leite, A. Pereira, M. Sun, B. Li, R. Jain, R. Cheng, E. Pincus, E. J. Carter, and J. F. Lehman. 2017b. Learning and reusing dialog for repeated interactions with a situated social agent. In *International Conference on Intelligent Virtual Agents*, p. 192–204. Springer.

C. D. Kidd and C. Breazeal. 2008. Robots at home: Understanding long-term human-robot interaction. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference On*, pp. 3230–3235. IEEE.

E. S. Kim, L. D. Berkovits, E. P. Bernier, D. Leyzberg, F. Shic, R. Paul, and B. Scassellati. May 2013. Social Robots as Embedded Reinforcers of Social Behavior in Children with Autism. *Journal of Autism and Developmental Disorders*, 43(5): 1038–1049. ISSN 0162-3257, 1573-3432. DOI: 10.1007/s10803-012-1645-2.

E. Kimani, T. Bickmore, H. Trinh, L. Ring, M. K. Paasche-Orlow, and J. W. Magnani. 2016. A smartphone-based virtual agent for atrial fibrillation education and counseling. In *International conference on intelligent virtual agents*, pp. 120–127. Springer.

A. Klapper, R. Ramsey, D. Wigboldus, and E. S. Cross. 2014. The Control of Automatic Imitation Based on Bottom–Up and Top–Down Cues to Animacy: Insights from Brain and Behavior. *Journal of Cognitive Neuroscience*, 26(11): 2503–2513. ISSN 0898-929X. DOI: 10.1162/jocn_a_00651.

J. Kory and C. Breazeal. 2014. Storytelling with robots: Learning companions for preschool children's language development. In *2014 RO-MAN: The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pp. 643–648. DOI: 10.1109/ROMAN.2014.6926325.

J. M. Kory-Westlund. 2019. *Relational AI: Creating Long-Term Interpersonal Interaction, Rapport, and Relationships with Social Robots.* PhD thesis, Massachusetts Institute of Technology.

J. M. Kory-Westlund and C. Breazeal. 2019a. Assessing children's perceptions and acceptance of a social robot. In *Proceedings of the 18th ACM International Conference on Interaction Design and Children*, IDC '19, p. 38–50. ACM. ISBN 978-1-4503-6690-8. http://doi.acm.org/10.1145/3311927.3323143. DOI: 10.1145/3311927.3323143. event-place: Boise, ID, USA.

J. M. Kory-Westlund and C. Breazeal. 2019b. Exploring the effects of a social robot's speech entrainment and backstory on young children's emotion, rapport, relationship, and learning. *Frontiers in Robotics and AI*, 6. ISSN 2296-9144. https://www.frontiersin.org/articles/10.3389/frobt.2019.00054/full. DOI: 10.3389/frobt.2019.00054.

J. M. Kory-Westlund, M. Martinez, M. Archie, M. Das, and C. Breazeal. 2016. Effects of Framing a Robot as a Social Agent or as a Machine on Children's Social Behavior. In *The 25rd IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 688–693. IEEE Press.

J. M. Kory-Westlund, L. Dickens, S. Jeong, P. L. Harris, D. DeSteno, and C. L. Breazeal. 2017a. Children use non-verbal cues to learn new words from robots as well as people. *International Journal of Child-Computer Interaction*. ISSN 2212-8689. DOI: 10.1016/j.ijcci.2017.04.001.

J. M. Kory-Westlund, S. Jeong, H. W. Park, S. Ronfard, A. Adhikari, P. L. Harris, D. DeSteno, and C. Breazeal. 2017b. Flat versus expressive storytelling: Young children's learning and retention of a social robot's narrative. *Frontiers in Human Neuroscience*, 11. ISSN 1662-5161. DOI: 10.3389/fnhum.2017.00295.

J. M. Kory-Westlund, H. W. Park, R. Williams, and C. Breazeal. 2018. Measuring young children's long-term relationships with social robots. In *Proceedings of the 17th ACM Conference on Interaction Design and Children*, pp. 207–218. ACM.

J. M. Kory-Westlund, C. Breazeal, and A. K. Ostrowski. in review. Physical or Virtual Robots—Which is better? A Survey of Embodiment and Presence in Human-Robot Interaction.

H. Kozima, M. P. Michalowski, and C. Nakagawa. 2009. Keepon. *International Journal of Social Robotics*, 1(1): 3–18.

I. Kruijff-Korbayová, E. Oleari, A. Bagherzadhalimi, F. Sacchitelli, B. Kiefer, S. Racioppa, C. Pozzi, and A. Sanna. Oct. 2015. Young Users' Perception of a Social Robot Displaying Familiarity and Eliciting Disclosure. In *Social Robotics*, pp. 380–389. Springer, Cham. DOI: 10.1007/978-3-319-25554-5_38.

N. Krämer and A. Manzeschke. 2021. *Social Reactions to Socially Interactive Agents and Their Ethical Implications*. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM, 1. DOI: http://dx.doi.org/10.1145/3477322.3477326.

B. Lakey and S. Cohen. 2000. Social support theory and measurement. In S. Cohen, L. G. Underwood, and B. H. Gottlieb, eds., *Social Support Measurement and Intervention: A Guide for Health and Social Scientists*, pp. 29–52. Oxford University Press, New York, NY, US. ISBN 978-0-19-512670-9.

J. L. Lakin, V. E. Jefferis, C. M. Cheng, and T. L. Chartrand. 2003. The chameleon effect as social glue: Evidence for the evolutionary significance of nonconscious mimicry. *Journal of Nonverbal Behavior*, 27(3): 145–162.

P. Lally, C. H. M. van Jaarsveld, H. W. W. Potts, and J. Wardle. 2010. How are habits formed: Modelling habit formation in the real world. *European Journal of Social Psychology*, 40(6): 998–1009. https://onlinelibrary.wiley.com/doi/abs/10.1002/ejsp.674. DOI: 10.1002/ejsp.674.

C. H. Lane and N. L. Schroeder. 2022. *Pedagogical Agents*, pp. 307–330. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563669.

G. W. Lane, D. Noronha, A. Rivera, K. Craig, C. Yee, B. Mills, and E. Villanueva. 2016. Effectiveness of a social robot, "Paro," in a VA long-term care setting. *Psychological Services*, 13(3): 292–299. ISSN 1939-148X 1541-1559. DOI: 10.1037/ser0000080.

H. R. Lee, J. Sung, S. Šabanović, and J. Han. 2012a. Cultural design of domestic robots: A study of user expectations in Korea and the United States. In *RO-MAN, 2012 IEEE*, pp. 803–808. IEEE.

M. K. Lee, J. Forlizzi, S. Kiesler, P. Rybski, J. Antanitis, and S. Savetsila. 2012b. Personalization in HRI: A longitudinal field experiment. In *Human-Robot Interaction (HRI), 2012 7th ACM/IEEE International Conference On*, pp. 319–326. IEEE.

S. Lee, H. Noh, J. Lee, K. Lee, G. G. Lee, S. Sagong, and M. Kim. 2011. On the effectiveness of robot-assisted language learning. *ReCALL*, 23(01): 25–58.

I. Leite, C. Martinho, A. Pereira, and A. Paiva. Sept. 2009. As Time goes by: Long-term evaluation of social presence in robotic companions. In *RO-MAN 2009 - The 18th IEEE International Symposium on Robot and Human Interactive Communication*, pp. 669–674. DOI: 10.1109/ROMAN.2009.5326256.

I. Leite, G. Castellano, A. A. Pereira, C. Martinho, and A. Paiva. 2012a. Long-term interactions with empathic robots: Evaluating perceived support in children. Social Robotics, pp. 298–307. Springer.

I. Leite, A. A. Pereira, G. Castellano, S. Mascarenhas, C. Martinho, and A. Paiva. 2012b. Modelling empathy in social robotic companions. Advances in User Modeling, pp. 135–147. Springer.

I. Leite, C. Martinho, and A. Paiva. Apr. 2013. Social Robots for Long-Term Interaction: A Survey. *International Journal of Social Robotics*, 5(2): 291–308. ISSN 1875-4791, 1875-4805. DOI: 10.1007/s12369-013-0178-y.

I. Leite, G. Castellano, A. Pereira, C. Martinho, and A. Paiva. Aug. 2014. Empathic Robots for Long-term Interaction. *International Journal of Social Robotics*, 6(3): 329–341. ISSN 1875-4791, 1875-4805. DOI: 10.1007/s12369-014-0227-1.

I. Leite, M. McCoy, M. Lohani, D. Ullman, N. Salomons, C. Stokes, S. Rivers, and B. Scassellati. 2015. Emotional Storytelling in the Classroom: Individual versus Group Interaction between Children and Robots. In *Proceedings of the 10th Annual ACM/IEEE International Conference on Human-Robot Interaction (HRI'15). ACM (to Appear)*.

I. Leite, A. Pereira, and J. F. Lehman. 2017. Persistent Memory in Repeated Child-Robot Conversations. In *Proceedings of the 2017 Conference on Interaction Design and Children*, IDC '17, pp. 238–247. ACM, New York, NY, USA. ISBN 978-1-4503-4921-5. DOI: 10.1145/3078072.3079728.

R. Levitan, S. Benus, R. H. Gálvez, A. Gravano, F. Savoretti, M. Trnka, A. Weise, and J. Hirschberg. 2016. Implementing Acoustic-Prosodic Entrainment in a Conversational Avatar. In *INTERSPEECH*, pp. 1166–1170.

D. Leyzberg, S. Spaulding, and B. Scassellati. 2014. Personalizing robot tutors to individuals' learning differences. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-Robot Interaction*, pp. 423–430. ACM.

I. Lopatovska, K. Rink, I. Knight, K. Raines, K. Cosenza, H. Williams, P. Sorsche, D. Hirsch, Q. Li, and A. Martinez. Mar. 2018. Talk to me: Exploring user interactions with the Amazon Alexa. *Journal of Librarianship and Information Science*, p. 0961000618759414. ISSN 0961-0006. DOI: 10.1177/0961000618759414.

G. López, L. Quesada, and L. A. Guerrero. 2018. Alexa vs. Siri vs. Cortana vs. Google Assistant: A Comparison of Speech-Based Natural User Interfaces. In I. L. Nunes, ed., *Advances in Human Factors and Systems Interaction*, Advances in Intelligent Systems and Computing, pp. 241–250. Springer International Publishing. ISBN 978-3-319-60366-7.

N. Lubold. 2017. Building Rapport Through Dynamic Models of Acoustic-Prosodic Entrainment. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '17, pp. 297–300. ACM, New York, NY, USA. ISBN 978-1-4503-4656-6. DOI: 10.1145/3027063.3027132.

N. Lubold, E. Walker, and H. Pon-Barry. 2016. Effects of Voice-Adaptation and Social Dialogue on Perceptions of a Robotic Learning Companion. In *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*, HRI '16, pp. 255–262. IEEE Press, Piscataway, NJ, USA. ISBN 978-1-4673-8370-7.

N. Lubold, E. Walker, H. Pon-Barry, and A. Ogan. 2018. Automated Pitch Convergence Improves Learning in a Social, Teachable Robot for Middle School Mathematics. In C. Penstein Rosé, R. Martínez-Maldonado, H. U. Hoppe, R. Luckin, M. Mavrikis, K. Porayska-Pomsta, B. McLaren, and B. du Boulay, eds., *Artificial Intelligence in Education*, Lecture Notes in Computer Science, pp. 282–296. Springer International Publishing. ISBN 978-3-319-93843-1.

G. M. Lucas, J. Gratch, A. King, and L.-P. Morency. 2014. It's only a computer: Virtual humans increase willingness to disclose. *Computers in Human Behavior*, 37: 94–100. ISSN 0747-5632. DOI: 10.1016/j.chb.2014.04.043.

B. Lugrin, C. Pelachaud, and D. Traum. 2021. *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume*

1: Methods, Behavior, Cognition. ACM Press.

B. Mallinckrodt. 1989. Social support and the effectiveness of group therapy. *Journal of Counseling Psychology*, 36(2): 170.

B. L. McCornack. Dec. 1982. Effects of Peer Familiarity on Play Behavior in Preschool Children. *Journal of Genetic Psychology*, 141(2): 225. ISSN 00221325.

G. F. Melson, J. Kahn, Peter H., A. Beck, and B. Friedman. Sept. 2009. Robotic Pets in Human Lives: Implications for the Human–Animal Bond and for Human Relationships with Personified Technologies. *Journal of Social Issues*, 65(3): 545–567. ISSN 1540-4560. DOI: 10.1111/j.1540-4560.2009.01613.x.

W. Moyle, M. Bramble, C. Jones, and J. Murfield. Mar 2018. Care staff perceptions of a social robot called paro and a look-alike plush toy: a descriptive qualitative approach. *Aging & Mental Health*, 22(3): 330–335. ISSN 1360-7863. DOI: 10.1080/13607863.2016.1262820.

J. Nadel, O. Grynszpan, and J. C. Martin. 2022. *Autism and Socially Interactive Agents*, pp. 437–562. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563673.

A. F. Newcomb and C. L. Bagwell. 1995. Children's friendship relations: A meta-analytic review. *Psychological Bulletin*, 117(2): 306–347. ISSN 1939-1455 0033-2909. DOI: 10.1037/0033-2909.117.2.306.

J. Orkin and D. Roy. 2008. The restaurant game: Learning social behavior and language from thousands of players online. *J. Game Dev.*, 3(1): 39–60.

J. Orkin and D. Roy. 2009. Automatic learning and generation of social behavior from collective human gameplay. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, p. 385–392.

A. K. Ostrowski, D. DiPaola, E. Partridge, H. W. Park, and C. Breazeal. 2019. Older adults living with social robots: Promoting social connectedness in long-term communities. *IEEE Robotics & Automation Magazine*, 26(2): 59–70.

G. Palestra, G. Varni, M. Chetouani, and F. Esposito. 2016. A Multimodal and Multilevel System for Robotics Treatment of Autism in Children. In *Proceedings of the International Workshop on Social Learning and Multimodal Interaction for Designing Artificial Agents*, DAA '16, pp. 3:1–3:6. ACM, New York, NY, USA. ISBN 978-1-4503-4560-6. DOI: 10.1145/3005338.3005341.

H. W. Park, M. Gelsomini, J. J. Lee, T. Zhu, and C. Breazeal. May 2017a. Backchannel opportunity prediction for social robot listeners. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2308–2314. DOI: 10.1109/ICRA.2017.7989266.

H. W. Park, R. Rosenberg-Kima, M. Rosenberg, G. Gordon, and C. Breazeal. 2017b. Growing Growth Mindset with a Social Robot Peer. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '17, pp. 137–145. ACM, New York, NY, USA. ISBN 978-1-4503-4336-7. DOI: 10.1145/2909824.3020213.

H. W. Park, I. Grover, S. Spaulding, L. Gomez, and C. Breazeal. 2019. A Model-free Affective Reinforcement Learning Approach to Personalization of an Autonomous Social Robot Companion for Early Literacy Education. In *AAAI*.

R. H. Passman and J. S. Halonen. June 1979. A Developmental Survey of Young Children's Attachments to Inanimate Objects. *The Journal of Genetic Psychology*, 134(2): 165–178. ISSN 0022-1325. DOI: 10.1080/00221325.1979.10534051.

C. Pelachaud, C. Busso, and D. Heylen. 2021. *Multimodal Behavior Modeling for Socially Interactive Agents*, pp. 259–310. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM, 1. DOI: http://dx.doi.org/10.1145/3477322.3477331.

J. Perez-Osorio, E. Wiese, and A. Wykowska. 2021. *Theory of Mind and Joint Attention*, pp. 313–347. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM, 1. DOI: https://doi.org/10.1145/3477322.3477332.

R. W. Picard and J. Klein. 2002. Computers that recognise and respond to user emotion: Theoretical and practical implications. *Interacting with Computers*, 14(2): 141–169. ISSN 0953-5438. DOI: 10.1016/S0953-5438(01)00055-8.

R. Pieraccini. 2021. *Natural Language Understanding in Socially Interactive Agents*, pp. 147–172. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM, 1. DOI: http://dx.doi.org/10.1145/3477322.3477328.

M. Pontier and G. F. Siddiqui. 2008. A virtual therapist that responds empathically to your answers. In *Intelligent Virtual Agents*, pp. 417–425. Springer.

R. Prada and D. Rato. 2022. *Socially Interactive Agents in Games*, pp. 493–525. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: https://doi.org/10.1145/3563659.3563675.

A. Pradhan, K. Mehta, and L. Findlater. 2018. "Accessibility Came by Accident": Use of Voice-Controlled Intelligent Personal Assistants by People with Disabilities. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pp. 459:1–459:13. ACM, New York, NY, USA. ISBN 978-1-4503-5620-6. DOI: 10.1145/3173574.3174033.

A. Pradhan, L. Findlater, and A. Lazar. Nov. 2019. "phantom friend" or "just a box with information": Personification and ontological categorization of smart speaker-based voice assistants by older adults. *Proc. ACM Hum.-Comput. Interact.*, 3(CSCW). https://doi.org/10.1145/3359316. DOI: 10.1145/3359316.

A. Ramachandran and B. Scassellati. 2015. Developing Adaptive Social Robot Tutors for Children. In *2015 AAAI Fall Symposium Series*.

N. Randall. 2019. A survey of robot-assisted language learning (rall). In *THRI*.

L. D. Riek and Don Howard. 2014. A code of ethics for the human-robot interaction profession. In *WeRobot 2014 Conference, University of Miami*.

L. Ring. 2017. *An Affect-Aware Dialogue System for Counseling*. PhD thesis, Northeastern University.

L. Ring, T. Bickmore, and D. Schulman. 2012. Longitudinal affective computing. In *International Conference on Intelligent Virtual Agents*, pp. 89–96. Springer.

L. Ring, T. Bickmore, and D. Schulman. 2015. Longitudinal Affective Computing. In *SpringerLink*, pp. 89–96. Springer Berlin Heidelberg. DOI: 10.1007/978-3-642-33197-8_9.

B. Robins, K. Dautenhahn, R. T. Boekhorst, and A. Billard. 2005. Robotic assistants in therapy and education of children with autism: Can a small humanoid robot help encourage social interaction skills? *Universal Access in the Information Society*, 4(2): 105–120.

K. J. Rotenberg, N. Eisenberg, C. Cumming, A. Smith, M. Singh, and E. Terlicher. Jan. 2003. The contribution of adults' nonverbal cues and children's shyness to the development of rapport between adults and preschool children. *International Journal of Behavioral Development*, 27(1): 21–30. ISSN 0165-0254. DOI: 10.1080/01650250143000571.

K. H. Rubin, W. M. Bukowski, and J. G. Parker. 1998. Peer Interactions, Relationships, and Groups. In *Handbook of Child Psychology*. John Wiley & Sons, Inc. ISBN 978-0-470-14765-8. DOI: 10.1002/9780470147658.chpsy0310.

T. Salter, K. Dautenhahn, and R. Bockhorst. Sept. 2004. Robots moving out of the laboratory - detecting interaction levels and human contact in noisy school environments. In *RO-MAN 2004. 13th IEEE International Workshop on Robot and Human Interactive Communication (IEEE Catalog No.04TH8759)*, pp. 563–568. DOI: 10.1109/ROMAN.2004.1374822.

C. Saund and S. Marsella. 2021. *Gesture Generation*, pp. 213–258. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition. ACM, 1. DOI: http://dx.doi.org/10.1145/3477322.3477330.

B. Scassellati, L. Boccanfuso, C.-M. Huang, M. Mademtzi, M. Qin, N. Salomons, P. Ventola, and F. Shic. Aug. 2018a. Improving social skills in children with ASD using a long-term, in-home social robot. *Science Robotics*, 3(21): eaat7544. ISSN 2470-9476. DOI: 10.1126/scirobotics.aat7544.

B. Scassellati, L. Boccanfuso, C.-M. Huang, M. Mademtzi, M. Qin, N. Salomons, P. Ventola, and F. Shic. 2018b. Improving social skills in children with asd using a long-term, in-home social robot. *Science Robotics*, 3(21). https://robotics.sciencemag.org/content/3/21/eaat7544. DOI: 10.1126/scirobotics.aat7544.

B. Scassellati, J. Brawer, K. Tsui, S. Nasihati Gilani, M. Malzkuhn, B. Manini, A. Stone, G. Kartheiser, A. Merla, A. Shapiro, D. Traum, and L.-A. Petitto. 2018c. Teaching Language to Deaf Infants with a Robot and a Virtual Human. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pp. 553:1–553:13. ACM, New York, NY, USA. ISBN 978-1-4503-5620-6. DOI: 10.1145/3173574.3174127.

C. Schroeter, S. Mueller, M. Volkhardt, E. Einhorn, C. Huijnen, H. van den Heuvel, A. van Berlo, A. Bley, and H.-M. Gross. May 2013. Realization and user evaluation of a companion robot for people with mild cognitive impairments. In *2013 IEEE International Conference on Robotics and Automation*, p. 1153–1159. DOI: 10.1109/ICRA.2013.6630717.

A. Sciuto, A. Saini, J. Forlizzi, and J. I. Hong. 2018. "Hey Alexa, What's Up?": A Mixed-Methods Studies of In-Home Conversational Agent Usage. In *Proceedings of the 2018 Designing Interactive Systems Conference*, DIS '18, pp. 857–868. ACM, New York, NY, USA. ISBN 978-1-4503-5198-0. DOI: 10.1145/3196709.3196772.

G. R. Semin and J. T. Cacioppo. 2008. Grounding social cognition: Synchronization, coordination, and co-regulation.

S. Serholt and W. Barendregt. 2016. Robots Tutoring Children: Longitudinal Evaluation of Social Engagement in Child-Robot Interaction. In *Proceedings of the 9th Nordic Conference on Human-*

Computer Interaction, NordiCHI '16, pp. 64:1–64:10. ACM, New York, NY, USA. ISBN 978-1-4503-4763-1. DOI: 10.1145/2971485.2971536.

K. Severinson-Eklundh, A. Green, and H. Hüttenrauch. Mar. 2003. Social and collaborative aspects of interaction with a service robot. *Robotics and Autonomous Systems*, 42(3–4): 223–234. ISSN 0921-8890. DOI: 10.1016/S0921-8890(02)00377-9.

R. L. Severson and S. M. Carlson. 2010. Behaving as or behaving as if? Children's conceptions of personified robots and the emergence of a new ontological category. *Neural Networks*, 23(8): 1099–1103.

T. Shibata, K. Wada, Y. Ikeda, and S. Sabanovic. 2009. Cross-cultural studies on subjective evaluation of a seal robot. *Advanced Robotics*, 23(4): 443–458.

E. Short, K. Swift-Spong, J. Greczek, A. Ramachandran, A. Litoiu, E. C. Grigore, D. Feil-Seifer, S. Shuster, J. J. Lee, S. Huang, and others. 2014. How to train your dragonbot: Socially assistive robots for teaching children about nutrition through play. In *Robot and Human Interactive Communication, 2014 RO-MAN: The 23rd IEEE International Symposium On*, pp. 924–929. IEEE.

C. L. Sidner, T. Bickmore, B. Nooraie, C. Rich, L. Ring, M. Shayganfar, and L. Vardoulakis. July 2018. Creating New Technologies for Companionable Agents to Support Isolated Older Adults. *ACM Trans. Interact. Intell. Syst.*, 8(3): 17:1–17:27. ISSN 2160-6455. DOI: 10.1145/3213050.

N. Singh. 2018. *Talking Machines: Democratizing the Design of Voice-Based Agents for the Home.* Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA, USA.

T. Sinha and J. Cassell. 2015a. We Click, We Align, We Learn: Impact of Influence and Convergence Processes on Student Learning and Rapport Building. In *Proceedings of the 1st Workshop on Modeling INTERPERsonal SynchrONy And infLuence*, INTERPERSONAL '15, pp. 13–20. ACM, New York, NY, USA. ISBN 978-1-4503-3986-5. DOI: 10.1145/2823513.2823516.

T. Sinha and J. Cassell. June 2015b. Fine-Grained Analyses of Interpersonal Processes and Their Effect on Learning. In *Artificial Intelligence in Education*, pp. 781–785. Springer, Cham. DOI: 10.1007/978-3-319-19773-9_115.

H. Spencer-Oatey. 1996. Reconsidering power and distance. *Journal of Pragmatics*, 26(1): 1–24. ISSN 0378-2166. DOI: 10.1016/0378-2166(95)00047-X.

A. Stenzel, E. Chinellato, M. A. T. Bou, A. P. del Pobil, M. Lappe, and R. Liepelt. 2012. When humanoid robots become human-like interaction partners: Corepresentation of robotic actions. *Journal of Experimental Psychology: Human Perception and Performance*, 38(5): 1073.

H.-C. Sung, S.-M. Chang, M.-Y. Chin, and W.-L. Lee. Mar. 2015. Robot-assisted therapy for improving social interactions and activity participation among institutionalized older adults: A pilot study. *Asia-Pacific Psychiatry*, 7(1): 1–6. ISSN 1758-5872. DOI: 10.1111/appy.12131.

J. Sung, H. I. Christensen, and R. E. Grinter. 2009a. Robots in the Wild: Understanding Long-term Use. In *Proceedings of the 4th ACM/IEEE International Conference on Human Robot Interaction*, HRI '09, pp. 45–52. ACM, New York, NY, USA. ISBN 978-1-60558-404-1. DOI: 10.1145/1514095.1514106.

J. Sung, R. E. Grinter, and H. I. Christensen. 2009b. Pimp my roomba: Designing for personalization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 193–196. ACM.

J.-Y. Sung, L. Guo, R. E. Grinter, and H. I. Christensen. 2007. “My Roomba Is Rambo”: Intimate Home Appliances. In J. Krumm, G. D. Abowd, A. Seneviratne, and T. Strang, eds., *UbiComp 2007: Ubiquitous Computing*, Lecture Notes in Computer Science, pp. 145–162. Springer Berlin Heidelberg. ISBN 978-3-540-74853-3.

N. Suzuki and Y. Katagiri. 2007. Prosodic alignment in human-computer interaction. *Connection Science*, 19(2): 131–141.

F. Tanaka, A. Cicourel, and J. R. Movellan. Nov. 2007. Socialization between toddlers and robots at an early childhood education center. *Proceedings of the National Academy of Sciences*, 104(46): 17954–17958. ISSN 0027-8424, 1091-6490. DOI: 10.1073/pnas.0707769104.

S. Thrun, M. Bennewitz, W. Burgard, A. B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and others. 1999. MINERVA: A second-generation museum tour-guide robot. In *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference On*, volume 3. IEEE.

L. Tickle-Degnen and R. Rosenthal. Oct. 1990. The Nature of Rapport and Its Nonverbal Correlates. *Psychological Inquiry*, 1(4): 285–293. ISSN 1047-840X. DOI: 10.1207/s15327965pli0104_1.

E. Torta, F. Werner, D. O. Johnson, J. F. Juola, R. H. Cuijpers, M. Bazzani, J. Oberzaucher, J. Lemberger, H. Lewy, and J. Bregman. 2014. Evaluation of a small socially-assistive humanoid robot in intelligent homes for the care of the elderly. *Journal of Intelligent & Robotic Systems*, 76(1): 57–71.

D. Traum. 2022. *Socially Interactive Agent Dialogue*, pp. 45–76. The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application. ACM, 2. DOI: http://dx.doi.org/10.1145/3563659.3563663.

Y. Trope and N. Liberman. 2010. Construal-level theory of psychological distance. *Psychological Review*, 117(2): 440–463. ISSN 1939-1471(Electronic);0033-295X(Print). DOI: 10.1037/a0018963.

V. Turk. Dec 2016. Home invasion. *New Scientist*, 232(3104): 16–17. ISSN 0262-4079. DOI: 10.1016/S0262-4079(16)32318-1.

S. Turkle. Jan. 2007. Authenticity in the age of digital companions. *Interaction Studies*, 8(3): 501–517. DOI: 10.1075/is.8.3.11tur.

S. Turkle. 2017. *Alone Together: Why We Expect More from Technology and Less from Each Other*. Hachette UK.

S. Turkle, C. Breazeal, O. Dasté, and B. Scassellati. 2006. Encounters with kismet and cog: Children respond to relational artifacts. *Digital Media: Transformations in Human Communication*, pp. 1–20.

D. Utami, T. W. Bickmore, and L. J. Kruger. 2017. A robotic couples counselor for promoting positive communication. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 248–255. DOI: 10.1109/ROMAN.2017.8172310.

P. Valdesolo and D. DeSteno. 2011. Synchrony and the social tuning of compassion. *Emotion-APA*, 11(2): 262.

L. P. Vardoulakis, L. Ring, B. Barry, C. L. Sidner, and T. Bickmore. 2012. Designing Relational Agents as Long Term Social Companions for Older Adults. In *SpringerLink*, pp. 289–302. Springer Berlin Heidelberg. DOI: 10.1007/978-3-642-33197-8_30.

P.-P. Verbeek. 2006. Persuasive Technology and Moral Responsibility Toward an ethical framework for persuasive technologies. *Persuasive*, 6: 1–15.

P. Vogt, R. van den Berghe, M. de Haas, L. Hoffmann, J. Kanero, E. Mamus, J.-M. Montanier, O. Oudgenoeg-Paz, D. H. Garcia, and F. Papadopoulos. 2019. Second Language Tutoring using Social Robots. A Large-Scale Study. In *IEEE/ACM Int. Conf. on Human-Robot Interaction (HRI 2019)*.

K. Wada and T. Shibata. Oct. 2007. Living With Seal Robots—Its Sociopsychological and Physiological Influences on the Elderly at a Care House. *IEEE Transactions on Robotics*, 23(5): 972–980. ISSN 1552-3098. DOI: 10.1109/TRO.2007.906261.

B. E. Wampold. 2015. How important are the common factors in psychotherapy? an update. *World Psychiatry*, 14(3): 270–277.

A. Weiss and C. Bartneck. 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)*, pp. 381–388. IEEE.

A. Weiss, D. Wurhofer, and M. Tscheligi. June 2009. “I Love This Dog”—Children’s Emotional Attachment to the Robotic Dog AIBO. *International Journal of Social Robotics*, 1(3): 243–248. ISSN 1875-4791, 1875-4805. DOI: 10.1007/s12369-009-0024-4.

K. R. Wentzel. 1997. Student Motivation in Middle School: The Role of Perceived Pedagogical Caring. *Journal of Educational Psychology*, 89(3): 411–19. ISSN 0022-0663.

S. S. Wiltermuth and C. Heath. 2009. Synchrony and cooperation. *Psychological Science*, 20(1): 1–5.

D. Winnicott. 1953. Transitional objects and transitional phenomena; a study of the first not-me possession. *The International Journal of Psychoanalysis*, 34: 89–97. ISSN 1745-8315 0020-7578.

L. Yin and T. Bickmore. 2018. Culturally-aware healthcare systems. In *Advances in Culturally-Aware Intelligent Systems and in Cross-Cultural Psychological Studies*, pp. 97–110. Springer.

J. Yorke, C. Adams, and N. Coady. 2008. Therapeutic value of equine-human bonding in recovery from trauma. *Anthrozoös*, 21(1): 17–30.

K. Zawieska, B. R. Duffy, and A. Stronska. 2012. Understanding anthropomorphisation in social robotics. *Pomiary, Automatyka, Robotyka*, 16(11): 78–82.

R. Zhao, A. Papangelis, and J. Cassell. 2014. Towards a dyadic computational model of rapport management for human-virtual agent interaction. In *International Conference on Intelligent Virtual Agents*, pp. 514–527. Springer International Publishing.

R. Zhao, T. Sinha, A. W. Black, and J. Cassell. 2016. Socially-aware virtual agents: Automatically assessing dyadic rapport from temporal patterns of behavior. In *International Conference on Intelligent Virtual Agents*, pp. 218–233. Springer International Publishing.