




Socially Interactive Agents as Peers

Justine Cassell



Author note:

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


Citation information:

J. Cassell (2022). Socially Interactive Agents as Peers. In B. Lugin, 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. 331-366). ACM.

DOI of the final chapter: [10.1145/3563659.3563670](https://doi.org/10.1145/3563659.3563670)

DOI of volume 2 of the handbook: [10.1145/3563659](https://doi.org/10.1145/3563659)

Correspondence concerning this article should be addressed to Justine Cassell, justine@cs.cmu.edu



22 Socially Interactive Agents as Peers

Justine Cassell

22.1 Introduction

Embodied conversational agents (ECAs) were originally implemented as experts in their respective domains—bankers, real estate agents or health care workers, for example. In this sense, those early ECAs, that provided information to their users, were similar to early dialogue systems, such as those developed to give information about trains [Ferguson and Allen 1998]. Even early chat systems, less concerned with accomplishing a task, still represented expertise in a domain, such as Rogerian psychology [Weizenbaum 1966]. More recently, smart speakers, such as Siri, Alexa, and Google Home, still give access to information but represent themselves as assistants rather than experts. It is an interesting twist (and beyond the scope of this paper to discuss); however, the relationship still embodies a power differential between the user and the system.

Intelligent tutoring systems and pedagogical agents (see Chapter 21 on “Pedagogical Agents” [Lane and Schroeder 2022] of this volume of this handbook) have followed much the same vein as early ECAs, focusing in vast majority on agents that act as *expert* tutors. These tutors pose questions, assess student knowledge, and selectively deliver tutorials for skills that students have not yet mastered. They have a representation of a subject area, of best practices in how to teach, and a model of the student that is updated as knowledge is acquired. Historically, many of these systems grew out of research into theories of how people learn, particularly by Herb Simon, John Anderson, and their respective students (e.g., Anderson et al. [1985]). Tutoring has been shown to lead to significant learning gains when carried out by human tutors and has also transferred well to AI-based intelligent tutoring systems that have been used with impressive results in thousands of K–12 and university classrooms worldwide [Koedinger and Corbett 2005]. Increasingly, these systems have amplified their performance by including the ability to recognize student emotional states. However, even those intelligent tutoring systems that include emotional awareness focus in large part on the emotions of the *students*, with the adult tutor monitoring and responding but not engaging at the level of an equal peer (e.g., Zakharov et al. [2008]).

While tutoring agents are derived from theories of how people learn, ECAs have been based on theories and descriptions of how people behave with one another [Cassell 2000]. As the paradigm of ECAs developed at the end of the last century, it was natural, then, that eventually researchers begin to include children’s behavior with one another as the model for their systems. This was particularly the case for researchers who had studied children’s interaction with other children through computational systems [Cassell 2002] or children’s interaction with non-embodied computational systems that listened and engaged without demonstrating expertise or assuming the stance of a teacher—what were called Story Listening Systems [Cassell 2004].

Children’s interaction with peers is enormously influential in their cognitive, social, and emotional development [Ladd 2005], and story listening systems evoked and supported children’s

stories as ways of scaffolding cognitive, social, and emotional development. These story listening systems eventually became embodied in child-like ECAs called *virtual peers* [Cassell et al. 2000].

In what follows we focus on those virtual peers—the kinds of ECAs and robots where the computer takes on the role of a peer, often communicating with age-appropriate language, and even looking like a child of the same age as the young person interacting with it. For the most part, the application of contemporary virtual peers is learning, and so in this chapter we narrow the focus to learning (as opposed to interactions to simply pass the time without other goals, for example) among students up through university, but we broaden the discussion beyond the classic school curriculum to informal learning outside the classroom. We also broaden the focus beyond what are sometimes called “core literacies”—reading, writing, and arithmetic—to include the learning of socio-emotional skills such as curiosity and establishing social bonds. In this way, we focus both on ECAs without social competencies and also specifically Socially Interactive Agents (SIAs) as peers.

While many topics could be covered in a chapter of this sort, we organize the remainder of the chapter around the nature of natural peer interaction in young people and its role in development and learning. We begin by outlining some core ways that communicative behaviors among peers play a role in children’s development and in learning. We then turn to the implementation of virtual peers and how they can engage with children along the same key dimensions. Finally, we turn to ethical considerations and a roadmap for future work.

22.2 Peer Interaction

Children in virtually all societies spend significant periods of time with their peers. The context may be school or home, overseen by adults or free of adult intervention, while playing or working, in dyads or in larger groups. In all of these contexts, peer interaction plays several important roles in children’s development. For example, while adults may provide notions of culturally appropriate narrative structure for storytelling, it is often with peers that children learn how to construct stories for an audience who cannot see what they are thinking [Vygotsky 1978, Christie and Stone 1999]. Likewise, while adults remain for the most part polite and encouraging when teaching children, children amongst themselves engage in argument and disagreement that in fact provides a unique resource for problem-solving and learning [Azmitia and Montgomery 1993, Pellegrini et al. 1998].

There has been increased focus on peer interaction in the developmental literature. This may well be because much of the developmental literature before a certain period considered that children spend their childhoods becoming adults, reproducing the best versions that they can of adult norms, and acquiring adult ways of interacting with peers. More recently, however, there has been a realization that children are members of a culture that is different from that of the adult world [Cook-Gumperz and Kyratzis 2001, Ladd 2005]. As Kyratzis [2004] notes,

Children are not merely unformed adults. . . they reformulate social categories (e.g., friendship, gender) appropriated from the adult culture in ways that are sensitive to context . . . and reflective of children’s personalities and momentary goals and agendas in the culture of peers (p. 626)

In this context, Zadunaisky and Blum-Kulka [2010] have identified peer interaction as a “double opportunity space,” on the one hand offering opportunities for children to construct their own kinds of childhood cultures, with important rules and roles for the group, and on the other hand providing opportunities for individual development in cognitive, linguistic, and social spheres. While action certainly plays an important role here (such as stomping away after losing a game), language and associated non-verbal (gesture, facial expression, posture shifts) and paraverbal (prosody and loudness) behavior is the primary way that both individual and social development takes place in the context of peers. For our purposes in this chapter, certain of these communicative behaviors with their peers are particularly important:

1. Assuming and switching roles and participation frameworks in pretend play and collaborative learning.
2. Taking positions in debate and disagreement using both language and embodied behaviors.
3. Constructing variants of speaking styles to demonstrate affiliation or allegiance to a particular cultural group and particular identity.
4. Different ways of communicating.
5. The role of social interaction in development and learning.

22.2.1 Roles and Relationships in Peer Interaction

Many of our daily activities rely on designated roles, and relationships among those roles, that set expectations, rights, and responsibilities for how we participate in an activity with a partner [Allwood 2000]. Asymmetric role-relationships like lecturer—audience, teacher—student, or mentor—mentee, often come with fairly rigid expectations about who should say what when, and how the other should respond. For example, in large lecture classes, it is rare that students challenge the choice of homework assigned. Some of these roles—such as CEO-employee—are based on status or power relationships of the participants and remain relatively stable across activities; however, others are based on expertise about the current topic or are locally allocated in other ways. For example, a student in my small conversational agents class may that same evening be my yoga instructor.

Peer-to-peer interaction, on the other hand, reflects a socially symmetric relationship that gives both participants the same rights and expectations to speak, even while roles may shift seamlessly over the course of an activity. Two children building a block tower in a kindergarten classroom may start on equal footing until one steps back and says, “you’re doing it wrong.” In this symmetric relationship, however, the first child has the right to disagree and continue building the tower as before. In the workplace, two colleagues may begin revising a document together until one notices that the other is deleting text and suggests the other begin working on the references. In this symmetric relationship, the colleague who is deleting text may defend her choice and continue to delete. The social equality of the relationship has benefits for the roles that each takes. Neuman and Roskos [1992] observed that children engaged in instructional conversation with a peer negotiate and coach each other’s literacy activities. Unlike the exchanges in adult—child conversation, children instructing one another often reverse roles and attribute the role of the

more capable peer according to the purpose of the play at hand. Based on studies such as these, taking roles has become an important part of the formal curriculum in many subjects. One striking example is in “pair programming” where students take turns being the driver and navigator in the writing of code [Campe et al. 2020].

22.2.2 Debate and Disagreement

In the absence of hierarchy, debate and disagreement can flourish, and these play an important role in learning and development. For adults, differences in life experience and perspective can lead to what has been termed “productive disagreement.” This concept has been popularized recently as a way of overcoming unproductive political conversations [Benson 2019]. However, the notion of productive disagreement in children goes back at least as far as Piaget. He described the benefits of cognitive conflict among peers as leading children to revise their views, unlike parent—child interaction where children are more likely to take as given the opinion of the adult [Piaget 1959, cited in Tudge and Winterhoff 1993].

There are in fact many positive functions of conflict in children’s peer interactions. Some of these can be traced to the need to deploy sophisticated linguistic strategies such as justification of one’s opinions [Eisenberg and Garvey 1981]. Others derive from children’s ability to maintain an argument as opposed to resolving it quickly [Genishi and DiPaolo 1982, Maynard 1985]. Recent studies have focused on the complex moment-to-moment interactional processes by which kinds of conflictual talk, such as gossiping about one another, teasing, and debates over who has won and who has played fairly in childhood games, are used among peers [Goodwin 1990]. Piaget traces disequilibrium in beliefs provoked by these kinds of conflict to the subsequent change in beliefs that leads to learning [Piaget 1947:1950]. A more contemporary interpretation of this approach, referred to as socio-cognitive conflict, demonstrates the ways that cognitive conflict, specifically embedded in social situations, can lead to better learning than when individuals learn alone [Mugny and Doise 1978]. Sinha et al. [2017] observed how children in a small group, working on building a Rube Goldberg machine (a contraption where the goal is to make the mechanism as complex as possible), challenged and disagreed with each other’s ideas. These challenges and disagreements were in fact the most important predictor of an increase in curiosity in those who were challenged, and also in the group as a whole.

22.2.3 Varying Speaking Styles

In addition to seamlessly shifting roles, children in interaction with one another also seamlessly shift their ways of speaking, adopting and adapting language that they hear spoken by various adults around them, as well as constructing their own variants of speaking styles to display allegiance to a broader youth culture. These speaking styles play an important role as children experiment with who they want to be, and how they want to be perceived. They also allow children growing up in situations where different dialects or languages are in contact to mark their affiliation to dominant and minority ethnic, racial, and gender identities [Rampton 1995]. A 2021 article in the New York Times newspaper described in detail the struggles of a young Black woman as she moved back and forth between her poor family in New York City, and the elite boarding school she attended in Hershey Pennsylvania. A large part of the girl’s struggle, as she described it, was around the decision of whether to adopt “speaking like a white person” full-time, and whether doing so would be a betrayal of her family [Elliott 2021]. Young people (and adults) who move between marginalized and mainstream communities often report this kind of code-switching

as a way to maintain their affiliation with their home community, while also making their way in a world where the standard dialect is associated with increased earning power and other kinds of success [Kallmeyer and Keim 2003]. However, the movement back and forth is not without the kind of stress that is reported in the New York Times article, and Ogbu [2008] has coined the term “oppositional culture” to describe the ways that school systems may inadvertently set up a situation where the student feels the need to define her identity *contra* the expectations of the school, and for that reason to refuse the dialect that the school insists on.

In addition to code-switching of this kind, children may engage in pretend play where they assume the voices of the characters they are enacting. In this context, even quite young children are capable of “playing teacher” or “playing mother,” both roles they have participated in only as observer and interlocutor [Goodwin 1993, cited in Kyratzis 2004; Cekaite and Aronsson 2005].

22.2.4 Difference in Peer Social Interaction

Thus far we have described peer interaction and peer learning in neurotypical populations. However, non-neurotypical individuals, such as those diagnosed with autism spectrum disorders (ASD), tend to exhibit social-emotional skills that differ from their neurotypical peers, and that impact their peer interactions in fundamental ways, such as difficulties with integrated verbal and non-verbal communication and with interpersonal relationship development, and insistence on behavioral and environmental sameness [American Psychiatric Association, 2013]. ASD is called a spectrum disorder because each individual with ASD may demonstrate the above traits to a greater or lesser extent. Nevertheless, the research and clinical communities have identified what they refer to as “high-functioning autism” or Asperger’s, where individuals tend to exhibit different socio-emotional skills from their typically developing neurotypical peers, but to a lesser extent. For this reason, individuals from this population are likely to be mainstreamed into schools where they will interact with neurotypical peers. In contemporary classrooms, where group learning is the norm in many countries, these individuals may have difficulties in benefiting from the learning context if they cannot engage in productive interaction with neurotypical peers, and their neurotypical peers likewise have difficulties engaging with them.

In reflecting on the role that virtual peers can play in the learning of socio-emotional skills as well as disciplinary topics, it is therefore important to understand the nature of social interaction and friendship among individuals with ASD, and in their interaction with neurotypical peers [Bauminger 2002, Bauminger et al. 2003]. Diagnostic criteria for autism define it in terms of “restrictive interests and behaviors” and “deficits in social interaction and communication” [DSM IV 1994] The mechanism underlying these behaviors is often described as an “impaired theory of mind” [Senju 2012] whereby the individuals in question may have difficulty “de-centering” or imagining a perspective on the world different than their own. In the lay literature, this is often described as difficulty in imagining the thoughts and feelings of others, although this does not align perfectly with the technical definition of the cognitive capacity to infer other’s mental states. It is true that individuals with ASD may find it difficult to understand the behavior, and intentions behind that behavior, of others in social interaction. On the other hand, the same is true in the other direction, where neurotypical individuals find it difficult to understand the intention of individuals with ASD in social interaction [Humphrey and Symes 2011]. This has been referred to as the “double empathy problem” [Milton 2012]—a breakdown of empathy and mutual understanding between people with differing ways of experiencing the world, or the “cross-neurological theory of mind” [Beardon 2017].

Perspectives such as these are important not just for improving our understanding of individuals with ASD but also for our understanding of social interaction in general and peer social interaction in particular. These theories point out the kinds of obstacles posed by different experiences (including language or dialect spoken, socio-economic status, as well as other life experiences) in creating a social bond that can productively support learning and other aspects of development.

22.2.5 Social Interaction during Task Behavior

It would be inappropriate for a teacher to share his problems with his spouse during his high school class on linear algebra. The same teacher, however, while grading papers with a colleague, may well share those marital difficulties. Similarly, when friends engage in a task, regardless of its nature, they often refer to past shared experience, disclose their feelings about what they are doing, and, depending on the cultural context, they may engage in mutual teasing. They also engage in more mutual eye gaze and greater alignment in their speech rate. These behaviors serve to build and maintain rapport between the participants [Cassell et al. 2007]. Cappella [1990] goes so far as to argue that the “construct of rapport is arguably one of the central, if not the central, construct necessary to understanding successful helping relationships.” In fact, such behaviors between friends, apparently unrelated to task, are associated with higher learning gains in a peer tutoring task, while the same behaviors between strangers are negatively correlated with learning [Finkelstein et al. 2012, Wang et al. 2012].

The reciprocal relationship between peers is in part what allows this kind of off-task behavior, whether the peers are children or adults. Sociality and the interweaving of social and task talk is perhaps the most representative behavior of peer interaction, and it plays an essential role in cognitive development and learning [Hartup 1996]. Interleaving the two is a primary way in which peers bond and create solidarity and rapport with one another, manifest their alliances, and demonstrate that their relationship is special and not subject to the same politeness rules of the wider culture [Kyratzkis 2004, Zhao et al. 2014]. As noted above, the off-task talk that achieves these goals may on its surface appear quite negative, such as joking at the expense of the other or teasing [Corsaro 1997, Kyratzkis 2004]. As a clue to the affiliative function of these apparently disruptive behaviors, researchers have observed that children justify their actions during arguments but do not try to resolve their disagreement, as the conflict remains an important part of healthy peer interaction [Genishi and DiPaolo 1982, Maynard 1985].

22.2.6 Peer-based Learning

In peer-based learning, students learn both from and with one another, in dyads or small groups. They learn by explaining their ideas to others and by participating in activities where they can learn from their peers. Not only do the learners develop skills related to the material being discussed, they also develop invaluable interpersonal skills as they work with others, give and receive feedback, and evaluate their own performance [Blum-Kulka and Dvir-Gvirsman 2010, Sin et al. 2019]. Considerable convincing evidence has accumulated demonstrating that if students are asked to discuss their answers with other students, their understanding of the material increases more than if they did an active learning component on their own [Bonwell and Eisen 1991, Johnson et al. 1991]. Similarly, structured group work can promote problem solving at a higher level than possible with individual effort alone [Millis and Rhem 2010]. Additionally, there is evidence that peer learning may help reduce attrition rates and increase engagement [Crouch et al. 2007, Porter

et al. 2013]. Peer-based learning is particularly effective for underperforming students [Robinson et al. 2005] and those in low-resource environments [Jacobson et al. 2001].

These studies demonstrate that learning does not happen in a cultural or social void. When, students learn together, a number of cognitive advantages accrue, linked in large part to the relative social equality of their relationships [Webb 1989]. It is not just the tutee who learns, however. The “tutor effect” refers to the fact that explaining a subject to somebody else can lead to learning [Sharpley et al. 1983]. Self-explanation, in the absence of a learning partner, has been linked to learning in many studies. However, the tutee’s challenges and questions also may encourage deeper reflection about the topic on the part of the tutor [Webb 1989]. Nevertheless, simply tutoring another is not sufficient as the experience of being both tutor and tutee may also encourage students to view learning as socially desirable [Rohrbeck et al. 2003] and to better understand how to learn. In fact, reciprocal tutoring [Palincsar and Brown 1984], where the students take turns as tutor and tutee, has been shown to be an important tool in classrooms, and forms the basis for many of the SIAs we will discuss below.

Learning how to read and write is particularly facilitated by peer talk [Teale and Sulzby 1986, Fuchs and Fuchs 2005], as is the learning of math and science [Newcomb and Brady 1982]. While these subjects require cognitive skills, they also rely on self-confidence and self-efficacy and perhaps for that reason, peers play a particularly important role [Rohrbeck et al. 2003]. In this context, as mentioned above, in young adults peer tutoring has been formalized in what is called “peer (or pair) programming,” a paradigm that has been shown to be a particularly effective approach to learning for those students traditionally marginalized in STEM (Science, Technology, Engineering, and Math), such as women or underrepresented minorities. Here, correlations have been found between the strength of the bond between peers and their learning gains [Zhong et al. 2016].

It is not just these school subjects that benefit from peer learning, however. Peer interaction plays a key role in the learning of a first and second language [Sato and Ballinger 2016]. In fact, as any parent who has moved to a country with a different language can attest, children most rapidly learn the second language in their interactions with peers, sometimes seeming to learn to speak a new language overnight!

22.3 Research on Virtual Peers and SIAs as Peers

In the first sections of this chapter, we discussed the value of peer interaction for learning and development. One might ask, then, what the value is of virtual peers—why not just stick to human peer contact? While the kinds of interactions discussed above are extremely valuable, peer-to-peer learning is not always possible. When peers cannot be found, when scheduling or distance or a pandemic makes assembling a dyad or group impossible, or when the student doesn’t get along with the only available peers, or is not understood by them, peer-based learning may not be viable. In these instances, SIAs as peers can make peer-based learning accessible by providing interactions and content matched to the learner’s (or learners’) stage in the learning process. Matched to the personal and interpersonal abilities of individuals or groups of learners, virtual peers can be capable of establishing peer-like bonds that sustain learning. Virtual peers, too, can maintain user privacy, which may allow learners to be more vulnerable about not understanding a particular subject matter.

Before we go any further, a note about terminology and its history. The first work on virtual peers was published in 2000 [Cassell et al. 2000]. In the same year the first paper was published

on imbuing ECAs with social competencies [Bickmore and Cassell 2000]. Both strands of research continued independently; however, it would take 12 years before the two topics were joined, in a study of rapport in peer tutoring [Finkelstein et al. 2012]. Since virtual peers existed before they were given social competencies, “virtual peer” is used as the generic term, and “SIA as peer” is used to refer to a virtual peer with social interaction skills. A similar trajectory was followed by robots. Breazeal’s seminal book *Designing Social Robots* was published in 2002 [Breazeal 2002]. In 2013, Kory et al. [2013] published the first paper on robot peers. In what follows, then, we use the term “virtual peer” to refer to both physical peer robots and peer agents on a screen (the virtual, in this context, refers to the fact that it is not a flesh-and-blood peer) and “SIA as peer” to refer to both socially interactive peer robots and socially interactive virtual peers.

In terms of the relative merits of the two kinds of embodiment—robot and graphics—for peer agents, the jury is still out. A number of studies have compared physical to graphical peer agents. However, in vast majority the studies have compared a physical robot to an image of that same robot on the screen [e.g., Kennedy et al. 2015]. This comparison does not do justice to the strengths of each kind of embodiment—primarily, the physical presence of the robot and the more natural lifelike movement of the virtual agent—and so the comparison does not tell us much about the subject of this chapter. For that reason, in the remainder of the chapter we lay out studies on both virtual peers with a physical instantiation and virtual peers on a screen in those places where each has made important advances in the use of peer agents to support children’s learning and development.

While contemporary virtual peers are graphical virtual agents displayed on a screen or physical robots, the very first virtual peer system was also one of the earliest intelligent tutoring systems, the text-based Learning Companion System described in the visionary 1988 article “Studying with the Prince: The computer as learning companion” [Chan and Baskin 1988]. Here an artificial student interacted with the real student while both learned under the guidance of an intelligent tutoring system. By including two tasks, learning by being tutored and tutoring, this system was based on the effective “reciprocal teaching” paradigm [Palincsar and Brown 1984] described above in which children take both the teacher’s and learner’s role.

22.3.1 Roles for Virtual Peers

The earliest contemporary virtual peer played the role of a conversational partner in collaborative storytelling [Cassell et al. 2000]. Collaborative peer storytelling has been shown to have a positive impact on early reading and writing literacy [Teale and Sulzby 1986]. Based on a study of real children telling stories with one another, “Sam the CastleMate” was designed to be projected lifesize onto a screen behind a toy castle. A “magic tower” allowed Sam to seem to pass toys back and forth from the real to virtual world, and sensors in each room of the castle, as well as embedded in small figurines, allowed Sam to follow the child’s movements with its eyes, to give contextually appropriate feedback, and to tell stories that took place in the same room that the child had just played in. The graphics were intentionally cartoon-like in order to constrain the child’s expectations and to avoid any ambiguity about whether Sam was “real” or not (see section on ethics for further discussion). Most importantly, Sam was not photorealistic because the focus of the research was on the impact of the virtual peer on the child’s behavior rather than a focus on the most lifelike behavior possible for the virtual peer. In this context the evaluation focused on whether the system evoked natural social interaction behaviors in children and whether it improved their emergent literacy skills. Some children therefore were asked to play with Sam by themselves, and other children played with Sam and one other child, in a triad. These interactions were

compared to children playing with another child or telling stories by themselves. Results demonstrated that children's interaction with Sam was much like their interaction with other children. In fact, some children even coached Sam in how to tell stories, as in the case of one boy who told Sam "Try to make a longer story next time. It's like this: 'The little boy was outside'. . . ." [Cassell 2004]. The key question, however, is whether there was any benefit to Sam's presence over and above what accrued to children playing with one another.

This did appear to be the case. The dyads of children playing without Sam sometimes told complete stories with decontextualized emergent literacy language. However, they also sometimes told stories that devolved into arguments or breaking parts of the castle. The stories that single children told with Sam, on the other hand, were more likely to show the important decontextualized language that predicts later literacy, and the amount of decontextualized language increased with each subsequent story that they told in collaboration with Sam. Likewise, when dyads of children played with Sam, their stories also demonstrated more emergent literacy behaviors than in Sam's absence, and Sam's presence also led them to engage in more pro-social collaboration of the kind that allowed them to get maximal educational benefit from the storytelling (op. cit.).

More recently, virtual peers have played a wide variety of roles in learning contexts. In a 2013 review article, Mubin et al. [2013] discuss the different roles that robots have taken in education and concludes that younger children are more likely to see the robot as a companion while older children see it more as a teaching aid. In a more recent review article, Belpaeme et al. [2018] add the teachable agent as a third useful role that can be taken by robots. In the teachable agent paradigm [Biswas et al. 2005, Chase et al. 2009], the student is always the tutor, teaching an agent that is described as younger and/or less knowledgeable (as the little boy described above did spontaneously when teaching Sam how to tell a long story). This perspective takes advantage of the "tutor effect" described above whereby students learn by teaching as well as when they are the student. Interesting work by Dillenbourg and colleagues has shown that even when children teach physically based skills, such as handwriting, a robot that plays the role of teachable agent, and that makes mistakes, can be very effective [inter alia Hood et al. 2015]. Research comparing the role of tutor, tutee, and peer found that children preferred peer robots [Looije et al. 2008]; however, explicit judgements of this sort do not always translate into performance, and so further research is required.

Research by Baylor and Kim [2005] showed the success of having a virtual peer take a variety of instructional roles. Research by Chen et al. [2020] demonstrated an active role-switching policy trained using reinforcement learning, in which the agent was rewarded for adapting its tutor or tutee behavior to the child's knowledge mastery level. Results demonstrate that both tutor and tutee roles were important. The former had a greater effect on learning while the latter had a greater positive impact on the student's affect. From research such as this, it is clear that the most effective virtual peers would take not just one role but be able to switch among them (including tutee, collaborator, and tutor) during a session. Similarly, it has been demonstrated that virtual peers can recognize simple roles in some natural narrative collaboration contexts and successfully elicit a shift in roles in the human partner, as well as shifting roles themselves. For example, in the context of children's spontaneous collaborative storytelling, Wang and Cassell [2003] has shown that children reliably began collaborative stories with their peers by attributing roles to one another. For example, one child might say "OK, you be the princess and I'll be the dragon." Another might describe the content of each role by saying "OK, the princess kills the dragon and saves the prince who was locked up in the castle." Based on an automatic analysis of speech and non-verbal behavior, the virtual peers were able to take on the roles the child had attributed to them, and, in

return, to attribute roles to the child. While this worked for simple role attribution in limited contexts, for the most part it is currently beyond the ability of most virtual peers to recognize the need for a new attribution of roles, or to fluidly switch roles.

At perhaps the most basic level of role that a virtual peer can take, a large number of studies have focused on virtual peers as ways to motivate learners to persevere. For example, virtual peers can successfully motivate an individual to stick to a task during a one-on-one tutoring session [Lane et al. 2013]. Virtual peers that demonstrate low competency can similarly raise self-efficacy. However, the same study showed that it was the high-competency virtual peer that increased learning gains [Kim and Baylor 2006]. As described above, we know that children also motivate one another to persevere in searching for solutions through debate and disagreement when collaborating in a small group [Sinha et al. 2017]. However, engaging in this kind of curiosity-inspiring conflict is still beyond the natural language capabilities of SIA systems, as described in Paranjape et al. [2018].

Nevertheless, virtual peers and other ECAs have been shown to be able to leverage group processes such as group trust, group emotion, conformity, norms, or cohesion, and to exert influence on groups (e.g., Sebo et al. [2020], Traeger et al. [2020]). Virtual peers can also influence the learner's stance toward learning. This has been notably demonstrated in a study showing that children were able to recognize a growth mindset in a peer-like robot, and then themselves adopt such a mindset in their own approach to learning [Park et al. 2017].

22.3.2 Debate, Teasing, and Disagreement in Virtual Peers

Given the significant role played by debate and disagreement in children's cognitive and social development, as well as in learning among peers, it seems natural to wonder how such phenomena might be incorporated into virtual peers. However, as the results described above demonstrate, teasing and insults can backfire and reduce learning gains when they are used in the wrong context—for example, among strangers rather than friends [Finkelstein 2017]. These results appear to hold for unfamiliar robots as well, as shown by Roth et al. [2019] where fairly mild robot insults resulted in reduced task scores for the people collaborating with the robot. For this reason, liberally peppering an interaction with insults is not going to help SIAs to act as learning companions. On the other hand, while insults play a negative role among strangers, by the same token they can play a positive role among friends where they may serve to mark the relationship between the two interlocutors as a special one that exists outside of the bounds of everyday politeness [Zhao et al. 2014]. In the context of the current chapter, this means that debate and disagreement are more likely to have a positive impact once rapport has been established. Indeed, in a situation where the rapport between child and SIA has been rated as high by external annotators, teasing does seem to be appreciated by the child. In fact, it is sometimes the child that initiates the teasing. An example comes from Finkelstein's work where 8- to 9-year-old children collaborated with a virtual peer on a science task [Finkelstein 2017]:

Child: What do think? Is [the bridge] going to be high or low?

SIA: Well, maybe we should make it lower so it has less room to wiggle around

Child: Ah! You took my idea, Alex! That was my idea because—

SIA: Nuh uh.

Child: Yes it was!

In this corpus, both child and SIA were likely to initiate teasing and, as the example above demonstrates, the episodes were well received by the child as well as initiated by the child. However, in this experiment, interactions took place over several weeks, which gave time for the relationship between child and SIA to develop. In addition, the data was collected in Wizard of Oz mode, which would have allowed the experiment to cut short any teasing that seemed to be missing the mark and resulting in ill feelings.

Future research in this area clearly requires an adaptive model of how specific conversational strategies can be deployed, based on the stage of the relationship and/or the level of rapport, and the user's own prior conversational strategies. As of this writing, in unpublished research the adaptive conversational strategy model of Zhao et al. [2014], that includes teasing and other violations of social norms, has been implemented as a SIA for peer tutoring in algebra, and one hopes that the results of this research will further illuminate whether and how putatively negative behavior such as teasing and disagreement might play a positive role in SIAs as peers.

22.3.3 Varying Speaking Styles in Virtual Peers

As described above, even quite young children are capable of adapting how they speak. In fact, by age 4 children can produce baby talk when interacting with an infant, even when they don't have an infant sibling [Weeks 1971 cited in Labotka and Gelman 2020]. By early school age, children can simplify their speech when interacting with a foreigner [Labotka and Gelman 2020]. As described above, children growing up in communities where low prestige dialects are spoken, and where these dialects are forbidden in school, can often fluently switch back and forth from the low-prestige dialect to school talk, even when their teacher considers them incapable of using the school-ratified dialect [Rader et al. 2011]. This code switching between dialects can serve as an important way of maintaining a link to two cultures—the culture of home, family, and tradition, and the culture of school, upward mobility, and mainstream societally ratified success. While some children are able and wish to navigate these different expectations and their associated ways of speaking, other children are either incapable of code-switching or find the two sets of expectations incompatible, as anthropologist Ogbu has discussed at length [Ogbu 1992, 2008].

Given the importance of dialect and other ways of speaking in establishing identity and building social affiliation with others, one might expect that the use of dialect by virtual peers and SIAs as peers would also play an important role in their interactions with users. Finkelstein et al. [2013] used a clever “distant peer” paradigm where 8- to 9-year-old children who spoke in vast majority African American English (AAVE), a low-prestige dialect used in many parts of the United States, were told that they were collaborating with a child from another classroom elsewhere in the city and that they would be communicating via recorded messages. In fact, the 2 four-minute voicemail messages from the other classroom (one social and one science) had been recorded by a bidialectal bicultural voice actress using one of three dialect patterns: one group heard AAVE in the social message and in the science message. A second heard AAVE for the getting to know one another recording and Mainstream American English (MAE) for the science recording. A third group heard MAE for both recordings. In order to test the impact of the social and science message on the children's dialect and on their production of science content, the children first heard the social message from the distant peer. They then recorded a social message in return. Then they recorded a science message. Then they listened to the distant peer's science message, and finally they recorded their own second science message. The dependent variable was the difference between the number of scientifically valid arguments produced in the first science message (before hearing the virtual peer's science message) and the second science message (produced after hearing the distant peer's science message).

While all children produced more science arguments after hearing the distant peer's model, the authors found that students in the AAVE condition demonstrated greater gains in the use of this school-ratified science discourse than students in the MAE condition (with no difference for the code-switching condition from either of the other conditions).

A later experiment [Finkelstein 2017] replicated the results using actual virtual peers, both during a six-week experiment and in a one-shot experiment with a larger number of participants. The virtual peers collaborated with children also 8 to 9 years old who also spoke primarily AAVE. Here the virtual peers looked identical across both conditions but spoke either MAE only (monodialectal) or spoke AAVE for brainstorming about a science task and MAE for practicing a presentation of their work for the teacher (the AAVE-only condition was omitted due to strong resistance from teachers). The use of two registers—brainstorming and speaking to a teacher—was introduced to justify the use of the different dialects. Here too the authors found that children in both conditions showed gains in the use of school-ratified science discourse (by which is meant both well-reasoned arguments and the language of science, such as referring to hypotheses and evidence) from pre-test to post-test. However, children who worked with the virtual peer that spoke their dialect demonstrated significantly greater gains in school-ratified science discourse from pre- to post-test than children in the monodialectal MAE condition, and the effect was strongest for children who were reading below grade level. The children in the code-switching condition also increased their participation in the science activity over the period of the experiment, unlike children in the MAE condition who visibly became increasingly aggressive with the agent and less participative over time.

Importantly however, there was a strong mediating variable: children who collaborated with the bidialectal virtual peer demonstrated higher levels of rapport (as judged by external annotators), and it was the variable of rapport that predicted performance on the post-test [Finkelstein 2017]. This underscores the critical nature of social bonds in the success of virtual agents, as in the collaboration between human peers. It argues for social awareness as an essential part of virtual agents—a true argument for SIAs. Notably, not all studies with ECAs or robots have found the same results. Pazyzbekov et al. [2019] found no increased learning gains for students working with a robot that spoke their own dialect of ~~Kazakh~~ Kazakh. Similarly, a study on a pedagogical agent speaking high or low German also found no results on learning gains for the agent speaking the dialect of the student, although study participants gave higher likeability scores to the low German agent [Kühne et al. 2013]. Note that in these studies, however, the agent was not a peer but an authority figure, and the agents made no explicit attempts to knit social bonds, which may indicate that rapport is both more important and more easily built when virtual agents are peers than when they represent teachers or other authority figures. This is a key argument for SIAs as peers.

Nevertheless, integrating low-prestige dialects into SIAs presents its own challenges. In the longitudinal study reported above, children working with the agent that spoke their dialect produced more ratified science discourse, but they also produced increasing amounts of AAVE over time, a fact that may have caused difficulties for them in their classrooms, where their teachers were explicitly negative about this low-prestige dialect. And unfortunately, integrating AAVE into virtual peers did not solve one fundamental issue concerning low prestige dialects, which is internalized biases against low-prestige dialect speakers. Famously, sociologist Basil Bernstein found that people who spoke low-prestige dialects in the UK tended to find others who spoke similarly to be more likeable, but less competent and less capable of earning high salaries

[Bernstein 1961]. Similarly, Lugin et al. [2020] found that adults judged robots speaking their own low-prestige dialect to be more likeable and less competent. In Finkelstein's studies, children who interacted with Alex over several weeks did change their explicit language ideologies. That is, they were more likely to say that it was okay to speak AAVE in some contexts. However, across conditions, and even for those students who had collaborated with a code-switching virtual peer for 6 weeks, all of the children continued to rate people who spoke AAVE as significantly less smart than people who spoke mainstream English. This persistent issue highlights the fact that while virtual peers may play an important role in children's learning, they are not silver bullets. Eradicating bias is a multipronged societal issue.

22.3.4 Difference in SIAs as Peers

Speaking different dialects and in different registers is one important way that SIAs as peers can support children experiencing difference in learning and development. Another is to be sensitive to the needs of non-neurotypical individuals. Here virtual peers have been successfully deployed to assess cognitive skills in children with ASD [Zhang et al. 2020a] as well as to improve confidence in social skills among adolescents with ASD [Boccanfuso et al. 2016]. Some studies have demonstrated that children with ASD can deploy social skills in their interactions with virtual peers that they do not deploy with real human peers [Tartaro and Cassell 2008], suggesting that these skills are in some sense known but not deployed in the young people's everyday interactions. Based on this finding, some research has investigated whether it is possible to support adolescents with ASD in reflecting on social skills through programming virtual peers— systems that have been referred to as *authorable virtual peers* [Tartaro and Cassell 2006]. In a longitudinal study [Tartaro et al. 2015], teenagers with ASD were initially given a control panel that allowed them to choose behaviors for a virtual peer to perform in its interaction with another adolescent. Strikingly, the teenagers chose behaviors for the virtual peer to perform that they themselves did not use in interaction with their peers. Over a period of several weeks, the teenagers learned how to program novel behaviors for the virtual peer and even record language into the control panel for the virtual peer to utter. Once again, the teenagers developed social behaviors for the virtual peer that they themselves did not use. Results of the study demonstrated that programming social behaviors and controlling those behaviors for a virtual peer had a transfer effect such that the teenagers subsequently were better able to deploy some of the social skills that they had programmed in their interactions with their real peers [Tartaro et al. 2015]. The topic of autism and socially intelligent agents is further addressed in Chapter 25 on “Autism and Socially Interactive Agents” [Nadel et al. 2022] of this volume of this handbook.

22.3.5 Social Intelligence in Virtual Peers

One explanation for the success of peer-based learning comes from studies, some described above, that demonstrate that peers spend a fair amount of time managing social cohesion—the bonds that exist within social groups of various kinds. Even relatively young children can express their affiliation and bonds with others whom they resemble or wish to resemble, in quite sophisticated and effective ways [Kyratzis 2004]. Socially interactive peer agents can play similar roles. For example, a field trial of social peer robots in classrooms in Japan has shown that SIA robots can both detect friendship among the real children it interacts with and evoke friendship-like behaviors with children. This result was particularly the case for a sub-group of the children who specifically treated the robot as a peer (for example, asking it for advice about personal matters) but did not

want to know how it functioned. This suggests that those children needed to suspend disbelief about the robot's mechanical functioning in order to engage in friendship [Kanda et al. 2007].

In a closer examination of how rapport is built among human teenage reciprocal peer tutors working on linear algebra, Madaio et al. [2017] found that a number of different conversational strategies may play a role in raising levels of rapport. Some of the peer tutors, for example, couched negative feedback in indirect terms, and this indirectness negatively correlated with the level of rapport judged by external annotators, and positively correlated with the number of problems the peer tutees attempted, and the number they successfully completed. This use of indirectness was particularly the case for confident peer tutors, who may be better able to allocate some attention to the social context as well as to the tutoring content. In another study, teens working with SIAs as peers that manifested the same social interaction strategies as found in the human–human study demonstrated greater learning gains; however, this was only the case for the students who started with stronger knowledge of linear algebra. For those students whose prior knowledge was lower, a task only version of the agent was most successful in promoting learning gains [Guzman Garcia and Cassell, unpublished]. This suggests that social interaction with the SIA as peer may require some cognitive effort and as such is easier to maintain for those students for whom the task is less effortful. Evidence for this interpretation also comes from a neuroscientific study that showed that while interacting with a human interlocutor engaged the parts of the neural apparatus responsible for social interaction, interacting with a virtual human in an identical task engaged both the social and cognitive effort parts of the neural apparatus [Gayda et al. 2008]. Of course, as Tärning et al. [2020] explain, both developmental and neuroscientific evidence exists for the primacy of social stimuli. They link these results to the effectiveness of SIAs as peers whose non-verbal behaviors “prime a feeling of social partnership in the learner, which leads to deeper cognitive processing during learning, and results in a more meaningful learning outcome as reflected in transfer test performance” [Mayer and DaPra 2012, as quoted in Tärning et al. 2020]. Given the results above, however, it is clear that this social partnership may not work for every learner or perhaps only for particularly well-designed SIAs. In any case, it is clear that SIAs as peers can profitably use language and nonverbal behavior to create social cohesion in the service of task performance, as has been shown in a number of studies that have looked at the impact of social behaviors in virtual peers on learning. Baxter et al. [2017], for example, assessed the effects of a personalized robot on learning in the classroom, assessing performance on both a novel and familiar task. What is meant by personalized in their context is some overall non-verbal convergence to the child's movements, personability (friendliness, informal language, and lack of imperatives), and adaptation to task, such that the personalized peer robot allowed the child to repeat a task. While the ability to repeat a task might be considered as a confound, as it may derive more from good teaching technique than personalization, results did show a learning effect such that the personalized robot resulted in more learning on the novel topic (although not on the familiar topic). One important issue concerning social intelligence in virtual peers is how to maintain interest over the time it takes to build rapport, as virtual peers can be repetitive in their language and non-verbal behavior. Burger et al. [2017] implemented a module that was capable of engaging in mutual self-disclosure with children over a two-week period. A larger number of self-disclosures on the child's part was associated with higher rapport, and higher rapport led to more use of the application over time (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]).

Techniques such as these will be important for future work with SIAs as virtual peers. In this context, encouraging evidence about the feasibility and impact on learning of integrating rapport

building behaviors into SIAs comes from a parallel research tradition in *parasocial relationships* between children and characters that are familiar to them from mainstream media such as movies, computer games, and television shows. American television shows such as *Blues Clues* and *Dora the Explorer* have developed ways to deploy socially contingent parasocial interactions, defined as techniques whereby characters are programmed to create pseudo conversations with children through comments, questions, and well-placed pauses [Lauricella et al. 2011]. An increasing number of studies have shown the effectiveness of this approach, both for children’s feelings of friendship with the character and for their learning in interactive computer games. In one study, for example, Calvert et al. [2020] who have spearheaded much of this work, have shown that the strength of young children’s parasocial relationship with *Dora the Explorer* predicted their learning gains on an early math skills task, and that this learning transferred to other tasks. Similar results come from the kinds of child characters being integrated into interactive videos such as those derived from the American TV show *Elinor Wonders Why*. While social bonds have not been assessed, recent work by Xu et al. [2022] shows that children respond significantly more accurately to science questions posed right after watching the interactive version of the video than after watching the broadcast version without interaction between the child character (a curious bunny named Elinor) and the child viewer. Even more strikingly, engagement is higher for the interactive child character than for a parent when engaging in dialogic reading (engaging the child in dialogue about the book being read), and the child’s story comprehension is also higher in the interactive child character condition [Xu et al. 2021]. In all these cases, the studies rely on the power of learning with an interactive peer.

22.4 Models and Modeling

As we move toward the conclusion of this chapter, it is important to note that, in addition to being effective interventions, virtual peers and SIAs as virtual peers can also serve another important role, and that is as models of human behavior, in the sense of McClelland’s explorations of “ideas about the nature of cognitive processes” [McClelland 2009]. In this context, researchers can create different versions of a virtual peer, and observe which look natural and which seem unnatural. More helpfully still, researchers can relatively quickly change the virtual peer’s behaviors, have the virtual peers interact with children, and thereby assess—for example—what kinds of disagreements are productive and what kinds obstruct learning rather than supporting it. In addition, virtual peers can elicit peer-like behaviors in children when experiments among real peers are difficult to carry out. This is the case, for example, with the experiments on the role of low-prestige dialects in learning described above. It is extremely difficult, and perhaps impossible, to find a context in which every variable is kept constant except the dialect that one child speaks with another, or to find a child who varies the dialect in one task and not another. Issues of socio-economic status, ethnicity, low- versus high-resourced schools, and several other variables that covary, have rendered difficult important research on the role of low-prestige dialects in learning, and their impact on children’s language ideologies. Virtual peers have allowed this research, as described above. Similarly, experiments with SIAs as peer robots interacting with pairs of students were able to discover instances of engagement that were productive and instances of engagement—often held up as the holy grail of interactive learning environments—that were unproductive [Nasir et al. 2021]. The results of studies such as these, that identify productive sequences for all students, or for particular populations, in particular contexts, can be integrated into virtual peers that then

are optimally helpful to their human partners. This double role, of model and intervention, underlies many virtual peers—and ECAs before them.

As well as allowing researchers to tweak behavior and observe whether it is natural or not, and whether it is helpful or not, virtual peers can also help researchers better understand the very nature of dyadic interaction among peers. Here, in some sense virtual peers are acting as *simulations of theory of mind* [Decety and Grèzes 2006], both for the researcher and for the child interacting with the virtual peer. This is the case, for example, for children with ASD. As we described above, research shows that these children appear to find it easier to interact with virtual peers than with real peers [Tartaro and Cassell 2006, 2008]. However, when they are given the opportunity to program social behaviors into authorable virtual peers, they can then subsequently use some of the behaviors they have programmed in subsequent face-to-face interaction with their real peers [Tartaro et al. 2015]. This suggests that the virtual peer is a kind of “practice other”—an interlocutor who is easier to engage with than real peers, whose perspective is perhaps easier to assess (in the sense of a theory of mind), and who serves as a steppingstone to interaction with real peers.

Yet another role for virtual peers-as-model is to highlight places where we do not yet have adequate computational models of children’s language, particularly children’s language when speaking to their peers. And yet such models are required to build effective autonomous agents that truly speak like peers to their child users (we might envisage a future large language model called “Bertie” for example). Such computational models and associated corpora would also allow us to analyze peer communication more effectively. Failures in implementation are not usually published; however, one informative example comes from the work described above concerning the role of conflict in raising the level of curiosity in elementary school children. A series of papers published on the dynamics of curiosity in group learning (see Sinha et al. [2022] for the most complete discussion) describes the ultimate goal of implementing detailed models of curiosity and embedding them in a virtual peer or several virtual peers capable of engaging in conflict and thereby inspiring a rise in curiosity. However, attempts to build such systems (see Paranjape et al. [2018] for first steps) ran into issues due to the impossibility at the time of building a deep learning model of children’s peer talk that would automatically detect productive conflictual talk in children’s multiparty conversation, and autonomously generate utterances that would inspire such productive conflict.

22.5 Future Work: User Modeling and Conversational Strategies

Developing computational models of how conflict is displayed through verbal and nonverbal means, and how it can be evoked in children, remains an important goal for future work. However, the challenge of modifying automatically generated language to adapt to different populations—for example low-literacy adults (c.f. Martin et al. [2020])—is undergoing active research in the NLP community. One might image that such work would be applicable to generating child-like language. There is also recent work in open-domain chat-oriented systems, including using large-scale language models such as DialoGPT [Zhang et al. 2020b] and BlenderBot [Roller et al. 2021] to generate grammatical and locally relevant text in conversation; however, these systems are not yet able to generate appropriate text within a certain context beyond one or two turns. Those that learn in real time may also integrate biased or socially inappropriate user input for training data, and thus generate unsuitable language [Zhou et al 2020]. However, recent work has made strides in analyzing and generating conversational strategies of the kind that build rapport [Soni et al.

2021, Raphalen et al. 2022]. This latter work, too, will undoubtedly play a role in generating child-like language that builds a social bond with a real human peer.

In addition to the challenges highlighted in the sections above, user modeling remains an important challenge for SIAs as peers. User modeling has played a key role in tutoring systems, from their earliest days (cf. Sleeman & Brown [1982]). In user modeling in the tutoring context, the system keeps track of what a student knows and does not know, adapting the model as new knowledge is presented, and updating it when assessment shows that the student has acquired knowledge (or forgotten it). Some user models for tutoring systems differentiate between procedural knowledge (roughly, *how* to address a topic) and conceptual knowledge (roughly, a deeper knowledge of the domain, which allows the student to *generalize*) [Murray 1999, Rau et al. 2009]. In current user models for other domains (such as recommendation systems), user preferences are also recorded as well as, for some systems, an assessment of personality or other personal and interpersonal features of the user. Establishing and maintaining a user model for a SIA is particularly important as rapport and other social constructs change over time. And yet, user models for systems where the computer plays a peer are particularly difficult to design as the model must take into account the preferences and abilities of a student of a certain age and not those of the designer of the system. This may include vocabulary level, level of theory of mind, cultural references, size of gesture space (larger in younger children than in older), and other features of a student of a particular age group. Modeling a peer is important because prior research suggests peer discussions are successful in part because the discussion helps identify misunderstandings while still “speaking the students’ language” [Blum-Kulka and Dvir-Gvirsman 2010].

22.6 Ethics of Virtual Peers

Systemic racism and discrimination are embedded in our educational systems, including on the part of students toward their peers. For this reason, virtual peers must carefully consider notions of equity and inclusion from the moment of their conception [Perry and Lee 2019]. A powerful example of this concern comes from Finkelstein’s work on low prestige-dialects, described above. In a number of cases, the researchers were banned from the classroom for “advocating poor English.” The students themselves learned more when brainstorming with virtual peers that spoke as they did, but they still demonstrated internalized racism with respect to those agents. That is, while the children learned more with the agents who spoke as they did, when asked if those agents were smart, they replied that they were not, more than once specifying “because they speak ghetto” [Finkelstein 2017]. This reminds us that we must carefully navigate the ways in which AI systems can propagate bias and exclusion and must include the goal of reducing bias and increasing representation as part of the design criteria. Researchers must also pay attention to the composition of the datasets they use as training data. They must attend carefully to the appearance, voice, and behaviors designed for the agents to remove any unintended bias. For example, early pedagogical agents included negative stereotypical gendered characteristics, such as low-cut blouses and tiny waists for female agents, and negative stereotypical ethnic characteristics such as gold chains and backwards baseball hats for African American agents. Some research has turned to gender- and ethnicity-ambiguity to reduce the risk of reifying such stereotypes [Rader et al. 2011].

In addition to issues of bias, there are more general ethical issues to address. Voice assistants such as Alexa, Cortana, and Google Assistant have evoked a fair amount of fear among parents and teachers, and this fear also affects adult perceptions of SIAs as peers. Perhaps the most common of such responses by parents is the fear that children will come to believe that they need

not be any more polite with real people than they are with Alexa. Other worries include that children will no longer be able to distinguish between real and virtual playmates and will come to prefer the virtual versions, always accessible and always up for a game. These worries have been extensively covered by the press (e.g., Gonzalez [2018]), and even formed the subject of reports by governmental commissions (e.g., CNPEN [2021]). While there do not currently exist significant data to support these reservations, it is clear that general guidelines for the implementation of SIAs as peers should be followed, so as to ensure the emotional and physical safety of vulnerable users, particularly young children and children with special needs. For example, the CNPEN report suggests that SIAs as peers be clear about the fact that they are not real, that they not be photorealistic in appearance, and that it be possible to find out what data they are collecting and storing about child users. In some sense, these guidelines should hold for all AI systems, but when the users are children, they are perhaps particularly important to keep in mind. In addition, while the press has covered adult fears about SIAs as peers, it is just as important to understand children's own reactions and worries [Yip et al. 2019]. Finally, while guidelines for ethical use are important, it is also useful to remember that fears such as those described above have in fact accompanied the introduction of every new talking technology, from the radio to the television to videogames and today to virtual peers [Cassell 2020]. We might therefore wish to take a cue from a 1961 booklet on the dangers of television published by the US government and illustrated by famous cartoonist Walt Kelly:

there are few things to practice not doing. Don't be afraid of it. These things are probably here to stay. Don't be afraid of your child. He's not here to stay. He's a precious visitor. Do not wind your child up and set him to play with it unguided. Do not wind it up and set it to watch your child. A machine is a bad sole companion. It needs help. You can help it. Love your child. [Kelly 1961].

22.7 Conclusions

The development of virtual peers has built on work in intelligent tutoring systems and ECAs but goes beyond them. The 2020 (and beyond) pandemic has demonstrated the need for students to continue learning even when separated by distance and time from their peers and teachers. Like intelligent tutoring systems, virtual peers can be tuned to a student's capabilities and are available on demand at any time or place. In addition, virtual peers can take advantage of the ways in which the presence of peers potentiates learning and development, in terms of the productive cognitive conflict they can generate, the evocation of self-explanation, and the social bonds that underly much of learning and development. SIAs as peers, whether graphical agents or physical robots, are particularly well-placed to build learning-focused social bonds over time and to allow students to continue to make gains in learning as well as in the socio-emotional skills that will allow them to lead productive lives.

Bibliography

- J. Allwood. 2000. An activity based approach to pragmatics. In H. Bunt and W. Black (Eds.), *Abduction, Belief and Context in Dialogue*. John Benjamin's Publishing, 47–80. DOI: <https://doi.org/10.1075/nlp.1.02all>.
- American Psychiatric Association. 2013. *Diagnostic and Statistical Manual of Mental Disorders* (5th ed.)
- J. R. Anderson, C. F. Boyle, and B. J Reiser. 1985. Intelligent tutoring systems. *Science* 228, 4698, 456–462. DOI: <https://doi.org/10.1126/science.228.4698.456>.
- M. Azmitia and R. Montgomery. 1993. Friendship, transactive dialogues, and the development of scientific reasoning. *Soc. Dev.* 2, 3, 202–221. <https://doi.org/10.1111/j.1467.9507.1993.tb00014.x>.
- N. Bauminger. 2002. The facilitation of social-emotional understanding and social interaction in high-functioning children with autism: Intervention outcomes. *J. Autism Dev. Disord.* 32, 283–98. DOI: <https://doi.org/10.1023/a:1016378718278>.
- N. Bauminger, C. Shulman, and G. Agam. 2003. Peer interaction and loneliness in high-functioning children with autism. *J. Autism Dev. Disord.* 33, 489–507. DOI: <https://doi.org/10.1023/a:1025827427901>.
- P. Baxter, E. Ashurst, R. Read, J. Kennedy, and T. Belpaeme. 2017. Robot education peers in a situated primary school study: Personalisation promotes child learning. *PLoS One*, 12, 5, e0178126. DOI: <https://doi.org/10.1371/journal.pone.0178126>.
- A. Baylor, and Y. Kim. 2005. Simulating instructional roles through pedagogical agents. *Int. J. Artif. Intell. Educ.* 15, 2, 95–115.
- L. Beardon. 2017. *Autism and Asperger Syndrome in Adults*. Sheldon Press, London.
- B. Benson. 2019. *Why Are We Yelling? Learn the Life-Changing Art of Productive Disagreement*. Portfolio/Penguin Press.
- T. Belpaeme, J. Kennedy, A. Ramachandran, B. Scassellati, and F. Tanaka. 2018. Social robots for education: A review. *Sci. Robot.* 3, 21, eaat5954. DOI: <https://doi.org/10.1126/scirobotics.aat5954>.
- B. Bernstein. 1961. Social class and linguistic development: A theory of social learning. In A. H. Halsey, J. Floud, and A. Anderson (Eds.), *Education, Economy and Society*. New York.

- T. Bickmore and J. Cassell. 2000. "How about this weather?" Social dialog with embodied conversational agents. In *Proceedings of the American Association for Artificial Intelligence (AAAI) Fall Symposium on Narrative Intelligence*. American Association for Artificial Intelligence, 4–8.
- G. Biswas, K. Leelawong, D. L. Schwartz, and N. Vye. 2005. Learning by teaching: A new agent paradigm for educational software. *Appl. Artif. Intell.* 19, 363–392. DOI: <https://doi.org/10.1080/08839510590910200>.
- S. Blum-Kulka and S. Dvir-Gvirsman. 2010. Peer interaction and learning. In P. McGaw (Ed.), *International Encyclopedia of Education* (3rd. ed.). Elsevier, Oxford, England, 444–449. DOI: <https://doi.org/10.1016/B978-0-08-0444894-7.00529-7>.
- L. Boccanfuso, E. Barney, C. Foster, Y. A. Ahn, K. Chawarska, B. Scassellati, and F. Shic. 2016. Emotional robot to examine different play patterns and affective responses of children with and without ASD. In *Proceedings of the International Conference on Human–Robot Interaction (HRI)*. IEEE Press, 19–26. DOI: <https://doi.org/10.1109/HRI.2016.7451729>.
- C. C. Bonwell and J. A. Eison. 1991. Active learning: Creating excitement in the classroom. ASHE-ERIC Higher Education Report No. 1. ERIC Number: ED336049.
- C. Breazeal. 2002. *Designing Sociable Robots*. MIT Press, Cambridge, MA.
- F. Burger, J. Broekens, and M. A. Neerinx. 2017. Fostering relatedness between children and virtual agents through reciprocal self-disclosure. In *BNAIC 2016: 28th Benelux Conference on Artificial Intelligence*. Springer, 137–154. DOI: https://doi.org/10.1007/978-3-319-67468-1_10.
- S. L. Calvert, M. M. Putnam, N. R. Aguiar, R. M. Ryan, C. A. Wright, Y. Liu, and E. Barba. 2020. Young children’s mathematical learning from intelligent characters. *Child Dev.* 91, 5, 1491–1508. DOI: <https://doi.org/10.1111/cdev.13341>.
- S. Campe, J. Denner, E. Green, and D. Torres. 2020. Pair programming in middle school: Variations in interactions and behaviors. *Comput. Sci. Educ.* 30, 1, 22–46. DOI: <https://doi.org/10.1080/08993408.2019.1648119>.
- J. N. Cappella. 1990. On defining conversational coordination and rapport. *Psychol. Inq.* 1, 4, 303–305. DOI: https://doi.org/10.1207/s15327965pli0104_5.
- J. Cassell. 2000. Nudge nudge wink wink: Elements of face-to-face conversation for embodied conversational agents. In J. Cassell, J. Sullivan, S. Prevost, and E. F. Churchill (Eds.), *Embodied Conversational Agents*. MIT Press, Cambridge, MA, 1–27. DOI: <https://doi.org/10.7551/mitpress/2697.003.0002>.

- J. Cassell. 2002. “We have these rules inside”: The effects of exercising voice in a children’s online forum. In S. Calvert, R. Cocking, and A. Jordan (Eds.), *Children in the Digital Age*. Praeger Press, New York, 123–144.
- J. Cassell. 2004. Towards a model of technology and literacy development: Story listening systems. *J. Appl. Dev. Psychol.* 25, 1, 75–105. DOI: <https://doi.org/10.1016/j.appdev.2003.11.003>.
- J. Cassell. 2020. The ties that bind: Social interaction in conversational agents. *Réseaux* 220–221, 2–3, 21–45. DOI: <https://doi.org/10.3917/res.220.0021>.
- J. Cassell, M. Ananny, A. Basu, T. Bickmore, P. Chong, D. Mellis, K. Ryokai, H. Vilhjálmsson, J. Smith, H. Yan. 2000. Shared reality: Physical collaboration with a virtual peer. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, 259–260. DOI: <https://doi.org/10.1145/633292.633443>.
- J. Cassell, A. Gill, and P. Tepper. 2007. Coordination in conversation and rapport. *Proceedings of the Workshop on Embodied Natural Language*. Association for Computational Linguistics, 41–50.
- A. Cekaite and K. Aronsson. 2005. Language play, a collaborative resource in children’s L2 learning. *Appl. Linguist.* 26, 2, 169–191. DOI: <https://doi.org/10.1093/applin/amh042>.
- T. W. Chan and A. Baskin. 1988. “Studying with the Prince”: The computer as a learning companion. In *Proceedings of the International Conference on Intelligent Tutoring Systems*. Montreal, Canada, 194–200.
- C. C. Chase, D. B. Chin, M. A. Opezzo, and D. L. Schwartz. 2009. Teachable agents and the protégé effect: Increasing the effort towards learning. *J. Sci. Educ. Technol.* 18, 334–352. DOI: <https://doi.org/10.1007/s10956-009-9180-4>.
- H. Chen, H. W. Park, and C. Breazeal. 2020. Teaching and learning with children: Impact of reciprocal peer learning with a social robot on children’s learning and emotive engagement. *Comput. Educ.* 150, 103836. DOI: <https://doi.org/10.1016/j.compedu.2020.103836>.
- J. Christie and S. Stone. 1999. Collaborative literacy activity in print-enriched play centers: Exploring the “zone” in same-age and multi-age groupings. *J. Lit. Res.* 31, 2, 109–131. DOI: <https://doi.org/10.1080/10862969909548042>.
- J. Cook-Gumperz and A. Kyratzis. 2001. Child discourse. In D. Schiffrin, D. Tannen, and H. Hamilton (Eds.), *The Handbook of Discourse Analysis*. Basil Blackwell, Oxford, 590–611. DOI: <https://doi.org/10.1002/9780470753460.ch31>.
- CNPEN: Comité National Pilote d’éthique Du Numérique, A. Grinbaum, L. Devillers, G. Adda, R. Chatila, et al. 2021. *Agents Conversationnels: Enjeux d’éthique*. [Rapport de recherche] Comité national pilote d’éthique du numérique; CCNE. 2021. ffcea-03432785v1f.

- W. A. Corsaro. 1997. *The Sociology of Childhood*. Pine Forge Press/Sage Publications.
- C. H. Crouch, J. Watkins, A. P. Fagen, and E. Mazur. 2007. Peer instruction: Engaging students one-on-one, all at once. *Research-Based Reform of University Physics* 1, 1, 40–95.
- J. Decety and J. Grèzes. 2006. The power of simulation: imagining one's own and other's behavior. *Brain Res.* 1079, 1, 4–14. DOI: <https://doi.org/10.1016/j.brainres.2005.12.115>.
- Diagnostic and Statistical Manual of Mental Disorders: DSM-IV*. 1994. American Psychiatric Association, Washington, DC.
- E. Eisenberg and C. Garvey. 1981. Children's use of verbal strategies in resolving conflicts. *Discourse Process.* 4, 2, 149–170, DOI: <https://doi.org/10.1080/01638538109544512>.
- A. Elliott. 2021. Dasani Leaves Home. *New York Times Magazine*, October 1, 2021, Retrieved from <https://www.nytimes.com>.
- G. Ferguson and J. Allen. 1998. TRIPS: The rochester interactive planning system. In *Proceedings of IAAI '99: Proceedings of the Sixteenth National Conference on Artificial Intelligence and the Eleventh Innovative Applications of Artificial Intelligence Conference Innovative Applications of Artificial Intelligence July 1999*. 906–907.
- S. Finkelstein. 2017. “Alex speaks with my voice!” *Rapport and Science Discourse with Bidialectal Virtual Peers* [Unpublished doctoral dissertation]. Carnegie Mellon University, CMU-HCII-17-109.
- S. Finkelstein, A. Ogan, E. Walker, R. Muller, and J. Cassell. 2012. Rudeness and rapport: Insults and learning gains in peer tutoring. In *International Conference on Intelligent Tutoring Systems*, Vol. 7315: Lecture Notes in Computer Science. Springer, Berlin, 11–21. DOI: https://doi.org/10.1007/978-3-642-30950-2_2.
- S. Finkelstein, E. Yarzebinski, C. Vaughn, A. Ogan, and J. Cassell. 2013. The effects of culturally-congruent educational technologies on student achievement. In *Proceedings of Artificial Intelligence in Education (AIED)*, Vol. 7926: Lecture Notes in Computer Science. Springer, 493–502. DOI: https://doi.org/10.1007/978-3-642-39112-5_50.
- D. Fuchs and L. S. Fuchs. 2005. Peer-assisted learning strategies: Promoting word recognition, fluency, and reading comprehension in young children. *J. Spec. Educ.* 39, 1, 34–44. DOI: <https://doi.org/10.1177/00224669050390010401>.
- J. Gayda, A. Tartaro, Z. Li, J. Cassell, and J. Y. Chiao. 2008. Neural Basis of Social Perception of a Human versus Virtual Human. Poster presented at the 15th Annual Meeting of the Cognitive Neuroscience Society (CNS), San Francisco, CA, April 2008.
- C. Genishi and M. DiPaolo. 1982. Learning through argument in a preschool. In L. C. Wilkinson (Ed.), *Communicating in the Classroom*. Academic, New York, 49–68.

- R. Gonzalez. May. 2018. Alexa, what are you doing to my kid's brain? *Wired Magazine*.
<https://www.wired.com/story/hey-alexa-what-are-you-doing-to-my-kids-brain/>.
- M. H. Goodwin. 1993. Accomplishing social organization in girls' play: Patterns of competition and cooperation in an African-American working-class girls' group. In S. T. Hollis, L. Pershing, and M. J. Young (Eds.), *Feminist Theory and the Study of Folklore*. University of Illinois Press, Urbana, IL, 149–165.
- J. Gratch and G. Lucas. 2021. Rapport between humans and socially interactive agents. In *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* (1st. ed.). Association for Computing Machinery, New York, NY, 433–462. DOI: <https://doi.org/10.1145/3477322.3477335>.
- A. Guzman Garcia and J. Cassell. Unpublished. The impact of rapport on tutoring by virtual peers.
- M. H. Goodwin. 1990. Tactical uses of stories: Participation frameworks within girls' and boys' disputes. *Discourse Process*. 13:1, 33–71. DOI: <https://doi.org/10.1080/01638539009544746>.
- W. W. Hartup. 1996. Cooperation, close relationships, and cognitive development. In W. M. Bukowski, A. F. Newcomb, and W. W. Hartup (Eds.), *The Company They Keep: Friendship in Childhood and Adolescence*. Cambridge University Press, 213–237.
- D. Hood, S. Lemaignan, and P. Dillenbourg. 2015. When children teach a robot to write: An autonomous teachable humanoid which uses simulated handwriting. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human–Robot Interaction (HRI'15)*. Association for Computing Machinery, New York, NY, 83–90. DOI: <https://doi.org/10.1145/2696454.2696479>.
- N. Humphrey and W. Symes. 2011. Peer interaction patterns among adolescents with autistic spectrum disorders (ASDs) in mainstream school settings. *Autism* 15, 4, 397–419. DOI: <https://doi.org/10.1177/1362361310387804> <https://journals.sagepub.com/doi/10.1177/1362361310387804#articleCitationDownloadContainer>.
- J. Jacobson, L. Thrope, D. Fisher, D. Lapp, N. Frey, and J. Flood. 2001. Cross-age tutoring: A literacy improvement approach for struggling adolescent readers. *J. Adolesc. Adult Lit.* 44, 6, 528–536. <http://www.jstor.org/stable/40013564>.
- D. Johnson, R. T. Johnson, and K. Smith. 1991. *Active Learning: Cooperation in the College Classroom*. Interaction Book Company, Edina, MN.
- W. Kallmeyer and I. Keim. 2003. Linguistic variation and the construction of social identity in a German-Turkish setting: A case study of an immigrant youth-group in Mannheim, Germany. In: J. Androutsopoulos and A. Georgakopoulou (Eds.), *Discourse Constructions of Youth*

Identities. John Benjamins Publishing Company, Amsterdam, 29–46. DOI: <https://doi.org/10.1075/pbns.110.03kal>.

- T. Kanda, R. Sato, N. Saiwaki, and H. Ishiguro. 2007. A two-month field trial in an elementary school for long-term human–robot interaction. *IEEE Trans. Robot.* 23, 5, 962–971. DOI: <https://doi.org/10.1109/TRO.2007.904904>.
- W. Kelly. 1961. *Pogo Primer for Parents* (TV Division). Children’s Bureau Headliner Series, number 2. US Dept. of Health Education and Welfare, Social Security Administration, Children’s Bureau.
- J. Kennedy, P. Baxter, and T. Belpaeme. 2014. Comparing robot embodiments in a guided discovery learning interaction with children. *Int. J. Soc. Robot.* 7, 293–308. DOI: <https://doi.org/10.1007/s12369-014-0277-4>.
- Y. Kim and A. L. Baylor. 2006. Pedagogical agents as learning companions: The role of agent competency and type of interaction. *Educ. Technol. Res. Dev.* 54, 3, 223–243. Retrieved December 30, 2021, from <https://www.learntechlib.org/p/67659/>. DOI: <https://doi.org/10.1007/s11423-006-8805-z>.
- K. Koedinger and A. Corbett. 2005. Cognitive tutors: Technology bringing learning sciences to the classroom. In R. Sawyer (Ed.), *The Cambridge Handbook of the Learning Sciences (Cambridge Handbooks in Psychology)*. Cambridge University Press, Cambridge, 61–78. DOI: <https://doi.org/10.1017/CBO9780511816833.006>.
- J. Kory, S. Jeong, and C. Breazeal. 2013. Robotic learning companions for early language development. In *Proceedings of the 15th ACM on International Conference on Multimodal Interaction (ICMI’13)*. Association for Computing Machinery, New York, NY, 71–72. DOI: <https://doi.org/10.1145/2522848.2531750>.
- V. Kühne, A. M. Rosenthal-von der Pütten, and N. C. Krämer. August. 2013. Using linguistic alignment to enhance learning experience with pedagogical agents: The special case of dialect. In *Proceedings of the International Workshop on Intelligent Virtual Agents*, Vol. 8108: Lecture Notes in Computer Science. Springer, Berlin, 149–158. DOI: https://doi.org/10.1007/978-3-642-40415-3_13.
- A. Kyratzis. 2004. Talk and interaction among children and the co-construction of peer groups and peer culture. *Annu. Rev. Anthropol.* 33, 625–649. DOI: <https://doi.org/10.1146/annurev.anthro.33.070203.144008>.
- D. Labotka and S. A. Gelman. 2020. The development of children’s identification of foreigner talk. *Dev. Psychol.* 56, 9, 1657–1670. DOI: <https://doi.org/10.1037/dev0001078>.
- G. Ladd. 2005. *Children’s Peer Relations and Social Competence: A Century of Progress*. Yale University Press.

- H. C. Lane and N. Schroeder. 2022. Pedagogical agents. In B. Lugrin, C. Pelachaud, and D. Traum (Eds.), *Handbook on Socially Interactive Agents: 20 Years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics: Volume 2: Interactivity, Platforms, Applications*. ACM Press, 307–329. DOI: <http://dx.doi.org/10.1145/3563659>. 3563669.
- H. C. Lane, C. Cahill, S. Foutz, D. Auerbach, D. Noren, C. Lussenhop, and W. Swartout. 2013. The effects of a pedagogical agent for informal science education on learner behaviors and self-efficacy. In *Proceedings of 16th International Conference on Artificial Intelligence in Education, AIED 2013*, Vol. 7926: Lecture Notes in Computer Science. Springer, Berlin, 309–318. DOI: https://doi.org/10.1007/978-3-642-39112-5_32.
- A. Lauricella, A. A. Gola, and S. L. Calvert. 2011. Toddlers’ learning from socially meaningful video characters. *Media Psychol.* 14, 216–232. DOI: <https://doi.org/10.1080/15213269.2011.573465>.
- R. Looije, M. A. Neerinx, and V. de Lange. 2008. Children’s responses and opinion on three bots that motivate, educate and play. *J. Phys. Agents* 2, 2, 13–20. DOI: <https://doi.org/10.14198/JoPha.2008.2.2.03>.
- B. Lugrin, E. Ströle, D. Obremski, F. Schwab, and B. Lange. 2020. What if it speaks like it was from the village? Effects of a robot speaking in regional language variations on users’ evaluations. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 1315–1320. DOI: <https://doi.org/10.1109/ROMAN47096.2020.9223432>.
- B. Lugrin, C. Pelachaud, and D. Traum. (Eds.). 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, 538 pages. DOI: <https://doi.org/10.1145/3477322>.
- M. Madaio, K. Peng, A. Ogan, and J. Cassell. 2017. A climate of support: A process-oriented analysis of the impact of rapport on peer tutoring. In *Proceedings of the 12th International Conference of the Learning Sciences (ICLS)*. International Society of the Learning Sciences.
- L. Martin, E. Villemonte de la Clergerie, B. Sagot, and A. Bordes. 11–16 May. 2020. Controllable sentence simplification. In *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020)*. European Language Resources Association, 4689–4698.
- R. E. Mayer and C. S. DaPra. 2012. An embodiment effect in computer-based learning with animated pedagogical agents. *J. Exp. Psychol. Appl.* 18, 3, 239–252. DOI: <https://doi.org/10.1037/a0028616>.
- D. W. Maynard. 1985. On the functions of social conflict among children. *Am. Sociol. Rev.* 50, 2, 207–223. DOI: <https://doi.org/10.2307/2095410>.

- J. L. McClelland. 2009. The place of modeling in cognitive science. *Top. Cogn. Sci.* 1, 1, 11–38. DOI: <https://doi.org/10.1111/j.1756-8765.2008.01003.x>.
- B. J. Millis and J. L. Rhem. 2010. *Cooperative Learning in Higher Education: Across the Disciplines, Across the Academy*.
- D. Milton. 2012. On the ontological status of autism: The “double empathy problem.” *Disabil. Soc.* 27, 6, 883–887. DOI: <https://doi.org/10.1080/09687599.2012.710008>.
- O. Mubin, C. Stevens, S. Shahid, A. Mahmud, and J.-J. Dong. 2013. A review of the applicability of robots in education. *Technol. Educ. Learn.* 1. DOI: <https://doi.org/10.2316/Journal.209.2013.1.209-0015>.
- G. Mugny and W. Doise. 1978. Socio-cognitive conflict and structure of individual and collective performances. *Eur. J. Soc. Psychol.* 8, 181–192. DOI: <https://doi.org/10.1002/ejsp.2420080204>.
- T. Murray. 1999. Authoring intelligent tutoring systems: An analysis of the state of the art. *Int. J. Artif. Intell. Educ.* 10, 98–129.
- J. Nadel, O. Grynszpan, and J.-C. Martin. 2022. Autism and socially interactive agents. In B. Lugrin, C. Pelachaud, and D. Traum (Eds.), *Handbook on Socially Interactive Agents: 20 Years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics: Volume 2: Interactivity, Platforms, Applications*. ACM Press, 437–462. DOI: <https://doi.org/10.1145/3563659.3563673>.
- J. Nasir, B. Bruno, M. Chetouani, and P. Dillenbourg. 2021. What if social robots look for productive engagement? Automated assessment of goal-centric engagement in learning applications. *Int. J. Soc. Robotics* 14, 55–71. DOI: <https://doi.org/10.1007/s12369-021-00766-w>.
- S. Neuman and K. Roskos. 1992. Literacy objects as cultural tools: Effects on children’s literacy behaviors in play. *Read. Res. Q.* 27, 3, 202–226. DOI: <https://doi.org/10.2307/747792>.
- A. Newcomb and J. Brady. 1982. Mutuality in boys’ friendship relations. *Child Dev.* 53, 2, 392–395. DOI: <https://doi.org/10.2307/1128981>.
- J. U. Ogbu. 1992. Understanding cultural diversity and learning. *Educ. Res.* 21, 8, 5–24. DOI: <https://doi.org/10.2307/1176697>.
- J. U. Ogbu (Ed.). 2008. *Minority Status, Oppositional Culture, & Schooling*. Routledge, New York. DOI: <https://doi.org/10.4324/9780203931967>.
- A. S. Palincsar and A. L. Brown. 1984. Reciprocal teaching of comprehension-fostering and comprehension-monitoring activities. *Cogn. Instr.* 1, 2, 117–175. DOI: https://doi.org/10.1207/s1532690xci0102_1.

- B. Paranjape, Z. Bai, and J. Cassell. 2018. Predicting the temporal and social dynamics of curiosity in small group learning. In *Proceedings of the 19th International Conference on Artificial Intelligence in Education (AIED)*, Vol. 10947: Lecture Notes in Computer Science. Springer, Cham, 420–435. DOI: https://doi.org/10.1007/978-3-319-93843-1_31.
- H. W. Park, R. Rosenberg-Kima, M. Rosenberg, G. Gordon, and C. Breazeal. 2017. Growing growth mindset with a social robot peer. In *Proceedings of the 2017 ACM/IEEE International Conference on Human–Robot Interaction*. ACM, 137–145. DOI: <https://doi.org/10.1145/2909824.3020213>.
- A. Pazylbekov, D. Kalym, A. Otyunshin, and A. Sandygulova. 2019. Similarity attraction for robot’s dialect in language learning using social robots. In *2019 14th ACM/IEEE International Conference on Human–Robot Interaction (HRI)*. IEEE, 532–533. DOI: <https://doi.org/10.1109/HRI.2019.8673232>.
- A. Pellegrini, L. Galda, M. Bartini, and D. Charak. 1998. Oral language and literacy learning in context: The role of social relationships. *Merrill Palmer Q.* 44, 1, 38–54.
- A. M. Perry and N. Lee. September. 2019. AI is coming to schools, and if we’re not careful, so will its biases. *The Avenue*. <https://www.brookings.edu/blog/the-avenue/2019/09/26/ai-is-coming-to-schools-and-if-were-not-careful-so-will-its-biases/>.
- J. Piaget. (1947:1950). *The psychology of intelligence*. In M. Piercy and D. E. Berlyne (Trans.). Routledge & Kegan Paul, London (Original work published 1947).
- L. Porter, C. Bailey Lee, and B. Simon. March. 2013. Halving fail rates using peer instruction: A study of four computer science courses. In *Proceeding of the 44th ACM Technical Symposium on Computer Science Education*. ACM, 177–182. DOI: <https://doi.org/10.1145/2445196.2445250>.
- J. Piaget. 1959. *The Language and Thought of the Child* (3rd ed.). (M. Gabain and R. Gabain, Trans.). Routledge/Taylor & Francis Group.
- E. Rader, M. Echelbarger, and J. Cassell. May. 2011. Brick by brick: Iterating interventions to bridge the achievement gap with virtual peers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, 2971–2974. DOI: <https://doi.org/10.1145/1978942.1979382>.
- B. Rampton. 1995. Language crossing and the problematisation of ethnicity and socialisation. *Pragmatics* 5, 4, 485–513. DOI: <https://doi.org/10.1075/prag.5.4.04ram>.
- Y. Raphalen, C. Clavel, and J. Cassell. 2022. “You might think about slightly revising the title”: Identifying hedges in peer-tutoring interactions. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 2160–2174.

- M. A. Rau, V. Alevan, and Rummel, N. July. 2009. Intelligent tutoring systems with multiple representations and self-explanation prompts support learning of fractions. In *Proceedings of Artificial Intelligence and Education (AIED)*. IOS Press, 441–448. DOI: <https://doi.org/10.3233/978-1-60750-028-5-441>.
- D. Robinson, J. Schofield, and K. Steers-Wentzell. 2005. Peer and cross-age tutoring in math: Outcomes and their design implications. *Educ. Psychol. Rev.* 17, 4, 327–362. DOI: <https://doi.org/10.1007/s10648-005-8137-2>.
- C. Rohrbeck, M. D. Ginsburg-Block, J. Fantuzzo, and T. Miller. 2003. Peer-assisted learning interventions with elementary school students: A meta-analytic review. *J. Educ. Psychol.* 95, 2, 240–257. DOI: <https://doi.org/10.1037/0022-0663.95.2.240>.
- S. Roller, E. Dinan, N. Goyal, D. Ju, M. Williamson, Y. Liu, J. Xu, M. Ott, K. Shuster, E. M. Smith, Y. L. Boureau, and J. Weston. 2021. Recipes for building an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*. Association for Computational Linguistics, 300–325. DOI: <https://doi.org/10.18653/v1/2021.eacl-main.24>.
- A. M. Roth, S. Reig, U. Bhatt, J. Shulgach, T. Amin, A. Doryab, F. Fang, M. Veloso. October. 2019. A robot’s expressive language affects human strategy and perceptions in a competitive game. In *2019 28th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 1–8. DOI: <https://doi.org/10.1109/ROMAN46459.2019.8956412>.
- M. Sato and S. Ballinger (Eds.), 2016. *Peer Interaction and Second Language Learning: Pedagogical Potential and Research Agenda*, Vol. 45. John Benjamins Publishing Company. DOI: <https://doi.org/10.1075/llt.45>.
- S. S. Sebo, L. L. Dong, N. Chang, and B. Scassellati. March. 2020. Strategies for the inclusion of human members within human–robot teams. In *Proceedings of the 2020 ACM/IEEE International Conference on Human–Robot Interaction*. ACM, 309–317. DOI: <https://doi.org/10.1145/3319502.3374808>.
- A. Senju. 2012. Spontaneous theory of mind and its absence in autism spectrum disorders. *Neuroscientist* 18, 2, 108–113. DOI: <https://doi.org/10.1177/1073858410397208>.
- A. M. Sharpley, J. W. Irvine, and C. F. Sharpley. 1983. An examination of the effectiveness of a cross-age tutoring program in mathematics for elementary school children. *Am. Educ. Res. J.* 20, 1, 103–111. DOI: <https://doi.org/10.2307/1162677>.
- D. Y. E. Sin, T. C. T. Chew, T. K. Chia, J. S. Ser, A. Sayampanathan, and G. C. H. Koh. 2019. Evaluation of constructing care collaboration—Nurturing empathy and peer-to-peer learning in medical students who participate in voluntary structured service learning programmes for migrant workers. *BMC Med. Educ.* 19, 1, 1–13. DOI: <https://doi.org/10.1186/s12909-019-1740-6>.

- T. Sinha, Z. Bai, and J. Cassell. 2017. A new theoretical framework for curiosity for learning in social contexts. In E. Lavoué, H. Drachsler, K. Verbert, J. Broisin, and M. Pérez-Sanagustin (Eds.), *Proceedings of 12th European Conference on Technology Enhanced Learning (EC-TEL 2017)*, Vol. 10474: Lecture Notes in Computer Science. Springer, Cham, 254–269. DOI: https://doi.org/10.1007/978-3-319-66610-5_19.
- T. Sinha, J. Bai, and J. Cassell. 2022. A Novel Multimodal Approach for Studying the Dynamics of Curiosity in Small Group Learning. arXiv:2204.00545.
- D. Sleeman and J. Brown. 1982. *Intelligent Tutoring Systems*. Academic Press, New York.
- M. Soni, B. Cowan, and V. Wade. 2021. Enhancing self-disclosure in neural dialog models by candidate re-ranking. arXiv:2109.05090. Retrieved from <https://arxiv.org/abs/2109.05090>.
- B. Tärning, Y. Joo Lee, R. Andersson, K. Månsson, A. Gulz, and M. Haake. 2020. Assessing the black box of feedback neglect in a digital educational game for elementary school. *J. Learn. Sci.* 29, 4–5, 511–549. DOI: <https://doi.org/10.1080/10508406.2020.1770092>.
- A. Tartaro and J. Cassell. 2006. Using virtual peer technology as an intervention for children with autism. In *Universal Usability: Designing Computer Interfaces for Diverse User Populations*. John Wiley & Sons, New York, 231–262.
- A. Tartaro and J. Cassell. 2008. Playing with virtual peers: Bootstrapping contingent discourse in children with autism. In *Proceedings of International Conference of the Learning Sciences (ICLS 2008)*. International Society of the Learning Sciences, 382–389.
- A. Tartaro, J. Cassell, C. Ratz, J. Lira, V. Nanclares-Nogues. 2015. Accessing peer social interaction: Using authorable virtual peer technology as a component of a group social skills intervention program. *ACM Trans. Access. Comput.* 6, 1, 1–29. DOI: <https://doi.org/10.1145/2700434>.
- W. H. Teale, and E. Sulzby. 1986. Emergent literacy as a perspective for examining how young children become writers and readers. In *Emergent Literacy: Writing and Reading*. Ablex, Norwood, NJ.
- M. L. Traeger, S. S. Sebo, M. Jung, B. Scassellati, and N. A. Christakis. 2020. Vulnerable robots positively shape human conversational dynamics in a human–robot team. *Proc. Natl. Acad. Sci. U. S. A.* 117, 12, 6370–6375. DOI: <https://doi.org/10.1073/pnas.1910402117>.
- J. Tudge and P. Winterhoff. 1993. Can young children benefit from collaborative problem solving? Tracing the effects of partner competence and feedback. *Soc. Dev.* 2, 3, 242–259. DOI: <https://doi.org/10.1111/j.1467-9507.1993.tb00016.x>.
- L. S. Vygotsky. 1978. *Mind in Society: The Development of Psychological Processes*. Harvard University Press, Cambridge, MA.

- A. Wang and J. Cassell. 2003. Co-authoring, corroborating, criticizing: Collaborative storytelling between virtual and real children. In *Proceedings of the Workshop of Educational Agents: More than Virtual Tutors*.
- W. Wang, S. Finkelstein, A. Ogan, A. Black, and J. Cassell. 2012. “Love ya jerkface!”: Using sparse log-linear models to build positive (and impolite) relationships with teens. In *Proceedings of the 13th Annual Meeting of the Special Interest Group on Discourse and Dialogue*. Association for Computational Linguistics, 20–29.
- N. M. Webb. 1989. Peer interaction and learning in small groups. *Int. J. Educ. Res.* 13, 1, 21–39. DOI: [https://doi.org/10.1016/0883-0355\(89\)90014-1](https://doi.org/10.1016/0883-0355(89)90014-1).
- T. E. Weeks. October. 1971. Speech registers in young children. *Child Dev.* 42, 4, 119–131.
- J. Weizenbaum. 1966. ELIZA—A computer program for the study of natural language communication between man and machine. *Commun. ACM* 9, 1, 36–45. DOI: <https://doi.org/10.1145/365153.365168>.
- Y. Xu, J. Aubele, V. Vigil, A. S. Bustamante, Y.-S. Kim, and M. Warschauer. 2021. Dialogue with a conversational agent promotes children’s story comprehension via enhancing engagement. *Child Dev.* 93, 2, e149–e167. DOI: <https://doi.org/10.1111/cdev.13708>.
- Y. Xu, V. Vigil, A. Bustamante, and Warschauer, M. 2022. “Elinor’s talking to me!”: Integrating conversational AI into children’s narrative science programming. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. ACM, 1–16. DOI: <https://doi.org/10.1145/3491102.3502050>.
- J. Yip, K. Sobel, X. Gao, A. Hishikawa, A. Lim, L. Meng, R. F. Ofiana, J. Park, and A. Hiniker. 2019. Laughing is scary, but farting is cute: A conceptual model of children’s perspectives of creepy technologies. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI’19)*. ACM, 1–15. DOI: <https://doi.org/10.1145/3290605.3300303>.
- S. Zadunaisky Ehrlich and S. Blum-Kulka. 2010. Peer talk as a ‘double opportunity space’: The case of argumentative discourse. *Discourse Soc.* 21, 2, 211–233. DOI: <https://doi.org/10.1177/0957926509353847>.
- K. Zakharov, A. Mitrovic, and L. Johnston. 2008. Towards emotionally-intelligent pedagogical agents. In *Proceedings of the International Conference on Intelligent Tutoring Systems*, Vol. 5091: Lecture Notes in Computer Science. Springer, Berlin, 19–28. DOI: https://doi.org/10.1007/978-3-540-69132-7_7.
- L. Zhang, A. Z. Amat, H. Zhao, A. Swanson, A. S. Weitlauf, Z. Warren, and N. Sarkar. 2020a. Design of an intelligent agent to measure collaboration and verbal-communication skills of children with autism spectrum disorder in collaborative puzzle games. *IEEE Trans. Learn. Technol.* 14, 3, 338–352. DOI: <https://doi.org/10.1109/TLT.2020.3029223>.

- Y. Zhang, S. Sun, M. Galley, Y. C. Chen, C. Brockett, X. Gao, J. Gao, J. Liu, and W. B. Dolan. July. 2020b. DIALOGPT: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*. Association for Computational Linguistics, 270–278. DOI: <https://doi.org/10.18653/v1/2020.acl-demos.30>.
- R. Zhao, A. Papangelis, J. Cassell. 2014. Towards a dyadic computational model of rapport management for human–virtual agent interaction. In T. Bickmore, C. Sidner, and S. Marsella (Eds.), *Intelligent Virtual Agents (IVA) 2014*, Vol. 8637: Lecture Notes in Computer Science. Springer, Cham, 514–527. DOI: https://doi.org/10.1007/978-3-319-09767-1_62.
- B. Zhong, Q. Wang, and J. Chen. 2016. The impact of social factors on pair programming in a primary school. *Comput. Hum. Behav.* 64, 423–431. DOI: <https://doi.org/10.1016/j.chb.2016.07.017>.
- L. Zhou, J. Gao, D. Li, and H. Y. Shum. 2020. The design and implementation of Xiaolce, an empathetic social chatbot. *Comput. Linguist.* 46, 1, 53–93. DOI: https://doi.org/10.1162/coli_a_00368.