

A Virtual Conversational Agent for Teens with Autism: Experimental Results and Design Lessons

Mohammad Rafayet Ali, Zahra Razavi, Abdullah Al Mamun, Raina Langevin, Reza Rawassizadeh, Lenhart Schubert, M Ehsan Hoque
University of Rochester

ABSTRACT

We present the design of an online social skills development interface for teenagers with autism spectrum disorder (ASD), who often lack access to social skills training. The interface is intended to enable private conversation practice anywhere, anytime using a browser. Users converse informally with an on-screen persona, receiving feedback on nonverbal cues in real-time, and summary feedback. The prototype was developed in consultation with an expert UX designer, two psychologists, and a pediatrician. Using the data from 47 individuals, feedback and dialogue generation were automated using the hidden Markov model and a schema-driven dialogue manager capable of handling multi-topic conversations. We conducted a study with nine high-functioning ASD teenagers, and through thematic analysis of post-experiment interviews, identified several key design considerations, notably: 1) Users should be fully briefed at the outset about the purpose and limitations of the system, to avoid unrealistic expectations. 2) An interface should incorporate positive acknowledgment of behavior change. 3) Realistic appearance of a virtual agent and responsiveness are important in engaging users. 4) Conversation personalization, for instance in prompting laconic users for more input and reciprocal questions, would help the teenagers engage for longer terms and increase the system's utility.

Author Keywords

Animated conversational agent design; social skills training; real-time feedback; autism intervention; nonverbal behavior; schema-based dialogue management;

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous; See <http://acm.org/about/class/1998> for the full list of ACM classifiers. This section is required.

Paste the appropriate copyright/license statement here. ACM now supports three different publication options:

- ACM copyright: ACM holds the copyright on the work. This is the historical approach.
- License: The author(s) retain copyright, but ACM receives an exclusive publication license.
- Open Access: The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single-spaced in Times New Roman 8-point font. Please do not change or modify the size of this text box.

Each submission will be assigned a DOI string to be included here.



Figure 1. Teenagers with Autism Spectrum Disorder interacting with LISSA.

INTRODUCTION

Autism spectrum disorder (ASD) is a developmental disorder which affects one in 59 individuals in the US alone [55]. Almost all the individuals with ASD show deficits in nonverbal communication [19]. They often fail to make appropriate gestures, eye contact, and smile to make their verbal communication compelling. Thus they often face difficulty in expressing their feelings and act out their frustrations through vocal outbursts [34], [51], [28]. Current practices for treating the communication skills deficit involve therapy sessions with behavior experts. There is a significant shortage of behavioral experts resulting in therapy time being limited and inaccessible. A computer-mediated social skills intervention has the potential to enable individuals with ASD to practice conversation frequently with a standardized and repeatable stimulus. Additionally, computer-mediated tools can enable a therapist to serve more individuals by monitoring their progress remotely while continuing their weekly face-to-face therapy.

Although computers have many advantages, social skills intervention for individuals with ASD is very challenging. Dr. Stephen Shore, an autistic professor of special education famously said, “If you’ve met one person with autism, you’ve met one person with autism”. This indicates that a single solution does not exist, and if we want to design any intervention, it needs to be customized for individual needs. Thus, we need the active participation of the ASD individuals in the design process, so that the design will be appropriate both in typical cases and in terms of individual variation. Because of the deficit in communication skills a participatory design is often hard to achieve.

In this paper, we review the design of an online social skills development interface LISSA – Live Interactive Social Skills Assistance; we focus on design lessons learned, based in part on additional trials with autistic teens, and particularly

on analysis of post-session interviews with participants. LISSA features a virtual agent capable of engaging the user in a multi-topic conversation and giving real-time feedback on the nonverbal cues, by analyzing the spoken language and the facial expressions of the user in real-time. In addition, LISSA provides a post summary feedback after the conversation. The motivating idea is that users can interact with LISSA repeatedly, and potentially learn and observe their skills improvement from the summary feedback in multiple conversation sessions, held in a private and safe environment. The design of the interface was guided by an expert UX designer, psychologists, and a pediatrician. The initial design of the interface was evaluated in a Wizard of Oz setting through a speed-dating study with 47 college students [1]. Participants demonstrated improvements in several areas of nonverbal communication using the interface, which includes eye contact, head nods, and smile. The feedback system and the dialogue manager were automated using the data collected through the study.

Subsequently, we began to focus on adapting the system design to teenagers with ASD. We decided to conduct an initial study that would inform further development of a system that could be evaluated in controlled intervention experiments, and subsequently deployed broadly. The initial study, whose lessons we explore here, involved nine teenagers with ASD. Figure 1 shows two teenagers with ASD interacting with LISSA in the lab. LISSA can be accessed through a laptop or computer which has Internet access and a webcam/mic. With the help of professionals in developmental and behavioral pediatrics, we recruited teens with ASD for preliminary interaction with LISSA and interviewed them about the experience after their interaction. Through a thematic analysis of the interview transcripts we then identified several key design guidelines, including the following: 1) Users should be fully briefed about the purpose of an interface, and its capabilities and limitations. 2) An interface should incorporate positive acknowledgment of behavior change. 3) Realistic appearance of a virtual agent and responsiveness are significant factors in engaging users. 4) Conversation personalization, e.g., user-sensitive turn-taking and prompting laconic users, would help the teenagers engage for longer terms, and thus have them benefit from the interaction.

Our paper makes the following contributions –

- We motivate and explain the design of the LISSA system, with emphasis on its adaptation towards helping teens with ASD to improve their conversational skills.
- We describe the automation of the system that makes it suitable for private, ubiquitous use.
- Based on our experience with ASD subjects, including transcribed post-session interviews, we discuss key design guidelines that we identify as important for future interface design for individuals with ASD.

RELATED WORK

Individuals with autism spectrum disorder (ASD) are characterized in varying degrees by their social interaction, difficulties in verbal and nonverbal communication, and repetitive behaviors [11]. High-functioning ASD individuals are those diagnosed with autism but functioning cognitively at a relatively high level (e.g., IQ greater than 70) [47], [10]. However, high functioning ASD individuals may still demonstrate deficits in communication, emotion recognition, and social interaction [47]. The existing treatments for high functioning ASD do not address the condition as a whole, rather they focus on individual symptoms. Thus there is no single intervention for such individuals. In this section, we discuss the existing computer-based social skills interventions for individuals with ASD.

Social skills intervention for ASD

Hourcade et al. [25] developed tablet apps for social skills development intervention for children with ASD. The authors conducted a randomized control study [26] with eight children with ASD to find empirical evidence of the effectiveness of the tablet apps. They used four apps from their Open Autism Software [25] and found that, when using the app, children spoke more, had more verbal interactions, and demonstrated physical engagement with the app activities. Gal et al. [18] developed a tabletop application as an intervention for improving social skills among children with high functioning ASD. The application allows two to four children to collaborate and create a narrative of a story for a given scenario. A study with 14 high functioning ASD children over a 3-week time period revealed that children were more likely to initiate positive social interaction, and demonstrate collaborative behavior after the intervention. Piper et al. [43] also developed a shared interface on tabletop known as SIDES to improve social skills among ASD individuals. The authors showed that their intervention tool enabled a middle school therapy class to become more effective in group work. Boyd et al. [8] designed SayWAT, a wearable assistive technology for adults with autism, which provides real-time feedback on prosody during face-to-face interaction. They demonstrated that their tool could detect atypical prosody and deliver feedback in real time without disruption to the conversation. Hayes et al. [21] have provided three prototype systems addressing the design challenges with the use of large group displays, mobile personal devices, and personal recording technologies. Through a qualitative study with 13 children, they presented design guidance for visual support. Benton et al. [4] presented a methodology for incorporating children with ASD in the design process. They conducted a study with 20 participants with ASD aged between 11 and 14 and, using their design methodology, came up with ten design guidelines specific for game design and idea generation. Madsen et al. [32] conducted a study with seven teenagers

with ASD to understand the design concepts needed for developing interfaces for the particular target group. They came up with three design considerations – an adaptation of the form factor, customizability of graphical user interface, and adaptation of user experience. Kamaruzaman et al. [27] developed a touchscreen-assistive learning numeracy app, known as TaLNA for children with ASD. The design was driven by participatory design guidelines.

Virtual Agent-Based Social Skills Training

Nojavanasghari et al. [42] designed a virtual agent-based system to mediate human-to-human interaction for children with autism. The aim of their work was to design a tool to help improve the social skills of children with autism by providing visual support for the children and real-time feedback to the interactor about the children’s affective states, using a recommendation system. Foster et al. [17] designed ECHOES as a multimodal learning environment intended for children with ASD. The system features a virtual agent that engages a child in a collaborative learning activity and provides feedback based on sensed features including gaze direction and gesture. In a subsequent work with ECHOES, Bernardini et al. [5] designed a virtual agent as a credible social partner for children with ASD, that engages them in interactive learning activities. Milne et al. [36] designed a virtual agent as an educational tool for children with ASD, which helps improve their conversational skills and ability to deal with bullying. In a study with ten participants, the authors showed that participants who received the intervention had gained higher conversational skills and more knowledge about bullying than the control participants. Mower et al. [41] developed Rachel, an embodied conversational agent designed to elicit and analyze naturalistic interactions. This tool was designed for children with autism to encourage their affective and social behavior. Boujarwah et al. [7] presented a tool to enable non-expert humans to generate conversational scenarios, which can be used to teach children with ASD, appropriate behaviors in different social scenarios. DeVault et al. [14] developed SimSensei - a virtual agent in the context of the healthcare decision support system. The goal of this system is different than our work as it aims to identify psychological distress indicators through a conversation with a patient in which the patient feels comfortable sharing information. This system has both nonverbal sensing and a dialogue manager. The dialogue manager uses four classifiers to categorize the users’ speech, and hence to generate a relevant response. Hopkins et al. [24] developed a computer-based social skills training system for children with ASD. In a study with 49 individuals, the authors demonstrated that the program helped the participants improve their emotion recognition ability and social interactions. Razavi et al. [44] developed a conversational agent capable of conversing with teenagers with ASD. The authors employed a script-like schema [45] for guiding the dialogues, and generated appropriate responses using hierarchical pattern transductions. Tartaro et al. [54]

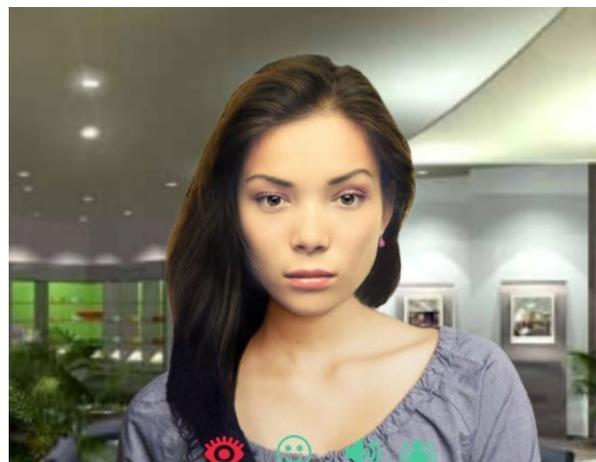


Figure 2. LISSA Interface. The virtual agent engages the users in a conversation and provides real-time feedback by flashing the icons at the bottom.

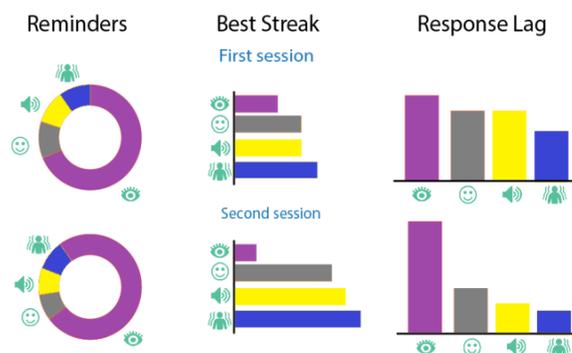


Figure 3. Post conversation summary feedback. Reminders: how often the users received feedback, Best Streak: how long the icons were green, Response Lag: how much time (on average) it took to turn a red icon green.

developed an authorable virtual peer for children with ASD. The virtual peer is capable of interacting with the children, sharing real toys, and responding to the children’s input.

Although interaction with conversational virtual agents proved to be effective in various applications, most state-of-the-art virtual agents are weak in terms of language understanding. Smartphone-based conversational agents can be impressive in their question-answering tasks but have been found to provide incomplete and inconsistent responses in areas like mental health, interpersonal violence, and physical health [37]. Woebot [16], a chat-bot recently developed at Stanford for delivering CBT to young adults with depression and anxiety, was able to significantly reduce symptoms of depression; however, its inability to hold a natural conversation was reported as the main issue users complained about.

Popular methods in dialogue managements include frame-based methods such as [16] and [46] where the user’s

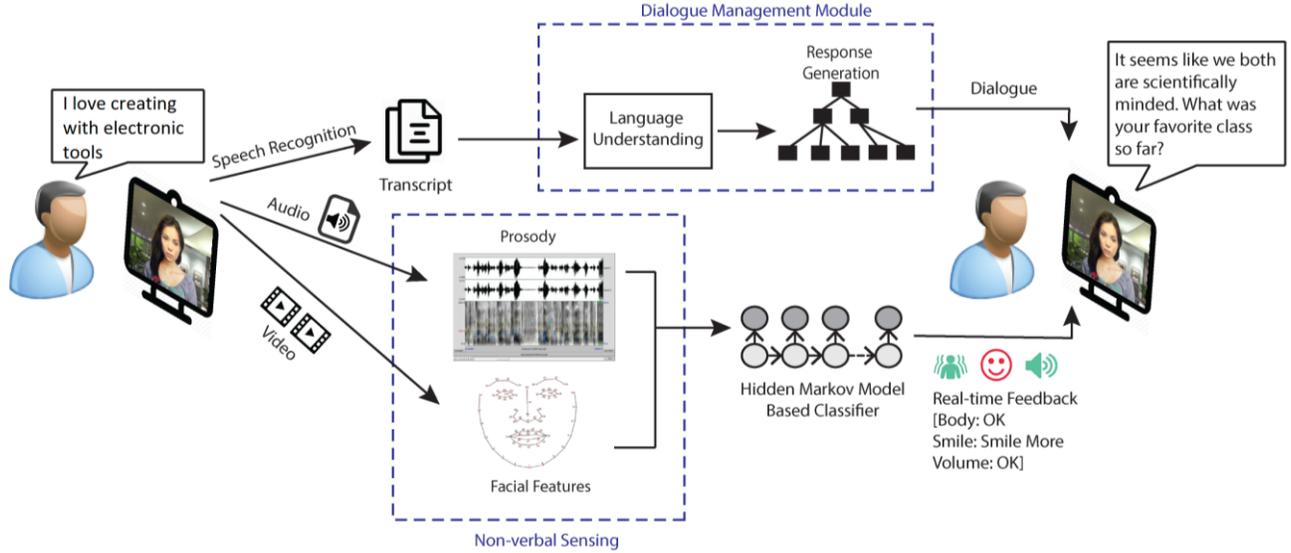


Figure 5. (From left to right) User having conversation with LISSA. System captures the speech, audio and video. Dialogue Management Module generates response. Nonverbal sensing extracts features and feeds to Hidden Markov Model based classifier, which produces real-time feedback in the form of flashing icon

responses are used to fill frame slots and to perform a task according to the contents of the frames. Another method, employed in SimCoach [40] and ASST [53] controls dialogue via transitions between predetermined states based on the user's input. While the method shows good performance in limited dialogues, scalability remains an issue for more open-ended conversation. Some other end-to-end systems achieve meaningful single-turn responses through data-driven methods (e.g., [49] and [57]), but cannot conduct a coherent longer dialogue; also, data-driven methods are susceptible to bias and privacy issues, among others [23].

Taken together, prior studies in computer-based social skills training and dialogue management suggest that dialogue-based systems hold considerable promise for helping individuals with ASD (or certain others with special needs). In our work, we have focused on improving users' nonverbal behavior in communication through real-time feedback. Through participatory studies, and post-session interviews, we have identified significant design guidelines for such virtual agent based intervention.

SYSTEM DESIGN

In this section, we outline the process of system design and implementation. We first describe the LISSA interface and then discuss the initial data collection process implemented with a Wizard-of-Oz prototype. Then we describe the machine learning techniques and dialogue management approach used to automate the system. The automation of the system, and its adaptation to the version aimed at ASD teens, prepared the way for the studies on which we focus here.

Training Real-time Feedback Model

Input: 46 video files $F = \{f_1, \dots, f_{45}\}$
 $\bar{v}_1, \dots, \bar{v}_T \leftarrow \text{ExtractFeatures}(F)$
 /* \bar{v}_i real valued vector with value range [0-1]*/
 $\bar{l}_1 \dots \bar{l}_T \leftarrow \text{GetLabels}(F)$
 /* \bar{l}_i is a four dimensional binary vector with value range [0, 1]*/
 For $i \in \{\text{eye_contact}, \text{smile}, \text{volume}, \text{body}\}$ do
 For $j \in \{\text{positive}, \text{negative}\}$ do
 Segment $\bar{v}_1, \dots, \bar{v}_T$ into $D_{i,j}$
 $D_{i,j} \leftarrow \bar{v}_k$ if $l_k^{(i)} = j$
 Train HMM $M_{i,j}$ using $D_{i,j}$

Generating Feedback

Input: Video Stream, V
 While ($\neg V.hasStopped()$) do
 $\bar{v}_t, \dots, \bar{v}_{t+\tau} \leftarrow \text{ExtractFeatures}(V)$
 For $i \in \{\text{eye_contact}, \text{smile}, \text{volume}, \text{body}\}$ do
 For $j \in \{\text{positive}, \text{negative}\}$ do
 $P_{i,j} \leftarrow P(\bar{v}_t, \dots, \bar{v}_{t+\tau} | M_{i,j})$
 If ($P_{i,\text{positive}} > P_{i,\text{negative}}$)
 Show positive feedback on i
 Else
 Show negative feedback on i

Figure 4. Feedback generation algorithm using HMM.

Interface

Our design team includes two psychologists, a psychiatrist, and a user experience designer. The online interface consists of two major parts: a conversation interface, and a post-conversation feedback interface. The LISSA interface features a virtual agent (see Figure 2), able to hold conversations with users, simultaneously providing real-time

feedback to them on their nonverbal behavior. The feedback is presented through four flashing icons placed at the bottom of the interface. The four icons represent eye-contact, smile, speaking volume, and body movement. The icons are green by default but turn to flashing red as a prompt to the user to adjust their behavior. We kept the simple red-green color scheme to reduce the cognitive load on the user, considering the fact that real-time feedback can be distracting and difficult to interpret [22]. The four nonverbal cues were selected because of their known applicability in improving social skills [50].

The post-conversation feedback interface (see Figure 3) shows a summary of the feedback provided during the conversation. This interface shows how many times the user received feedback (Reminder), how long the user kept the icons green (Best Streak), and how much time the user took to adjust nonverbal behavior (Response Lag).

Wizard of Oz Prototype

We first developed a Wizard of Oz prototype of the interface. There were two human operators for driving the system. One was responsible for the dialogue management and the other was responsible for giving feedback by flashing icons. Both operators were able to monitor the user remotely and control the interface through a web-based interface. This design allowed us to collect data from users and learn about their user experience without as yet applying machine learning techniques. Subsequently, we used the collected data to train machine learning models for automation.

Data Collection

To collect data and assess the viability of the interface we conducted a randomized control study with 47 college students and 8 female research assistants in the context of speed-dating. Speed-dating is gaining increasing popularity among researchers as a tool for studying social and communication skills [15]. The study results revealed that participants who used the LISSA interface improved their eye-contact and head nods [1]. From this study, we collected 46 videos of 23 individuals interacting with LISSA. The videos were captured through a camera attached on top of the computer monitor which displayed the LISSA program. We then employed six undergraduate research assistants majoring in Psychology to label the collected videos. The psychologists in our team provided multiple training sessions to the research assistants before the start of the labeling work. The research assistants watched the video recordings and marked those moments where the participants should receive feedback (i.e., red icon flash). We only considered those moments for feedback labels where more than two research assistants marked it as a feedback moment.

Automated System

We extracted facial and prosodic features from the recorded videos, including head pose, smile, facial action units, volume, and voice pitch. For this, we used off-the-shelf software tools, namely OpenFace [3], Praat [6], and SHORE [59]. We then trained a hidden Markov model (HMM) to

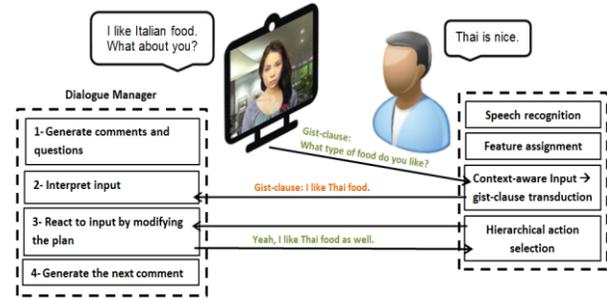


Figure 6. Overview of LISSA dialogue manager.

generate the flashing-icon feedback from the facial and prosodic features. In the past, HMMs have proven successful in modeling human behaviors and actions [58], [13]. Figure 4 shows the training and feedback generation technique.

Dialogue Management

In order for LISSA to conduct a conversation, we developed an automated dialogue manager capable of handling multiple topics. Much of the content was based on our experience with the WOZ experiments. An outline of the functionality of the dialogue module is presented in Figure 6. LISSA leads the conversation by asking questions on different topics (often after a “personal” remark about the topic), and making relevant comments on the user’s responses -- comments intended to show actual understanding of the user. The dialogue manager can also handle some questions asked by user. It continually updates the conversation plan, based on the user’s responses. At the top level, LISSA uses a structure called a *schema*, which contains a list of expected successive events in a dialogue; it allows for actions by both interlocutors and can be dynamically modified by user responses. The schemas are hierarchically structured, allowing LISSA to insert subschemas into the dialogue plan, helping to make the conversation more spontaneous. In order to capture users’ inputs, we use the Nuance [60] speech recognizer. Automatic detection and control of turn-taking is still under development, and we required our users to indicate their turn taking by pressing a button on a wireless clicker

When LISSA’s dialogue manager receives input from the speech recognizer, it generates a high-level interpretation in the form of short, explicit, context-independent English sentences which we call “gist clauses”. Gist clauses are extracted by applying several pattern transduction trees to the user’s input, taking the current LISSA question as the context. In order to facilitate input matching, input words are automatically annotated with syntactic and semantic features before extraction of any gist clauses. Features are recursively attached to input words, such as GOODPRED for words like “happy”, and ones like SOCIAL-SCIENCE and (by recursion) ACADEMIC-SUBJECT for “linguistics”.

After gist clauses have been derived from the user’s input, they undergo a second stage of transduction, producing a

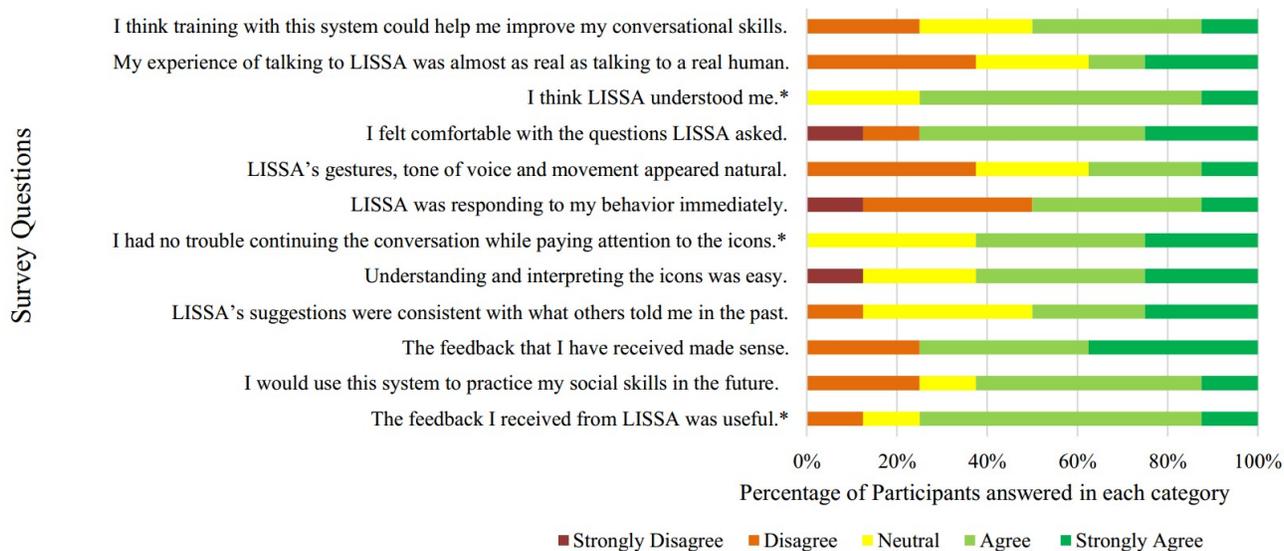


Figure 7. Survey questions and responses.

reaction for LISSA to output. If the user's input answers a question by LISSA, the reaction will usually be a relevant reaction to that answer; or, if there is a question among the extracted gist clauses (typically at the end of a user's input), LISSA is likely to answer the question. Every gist clause obtained is stored in LISSA's memory so that LISSA won't ask for information already provided by the user in previous turns. Also, the collected gist-clauses could be used for future enhancements that allow inference during the conversation and reference to previous contributions of the participants. Figure 5 provides an overview of the functioning of the automated system.

STUDY

Over the past three years, we have been recruiting participants for our studies through the developmental/behavioral pediatric research center at the local medical center. As already mentioned, our experimental sessions have been conducted with nine teenage participants (i.e., between 13 and 18 years old). The goal of this study has been to learn what aspects of LISSA are judged to be useful, and what adjustments need to be made in order to make LISSA a useful tool for iterative conversation training in the lives of teens with ASD anywhere. The sessions with the teens (whose parents were also invited) were scheduled on separate days. Each participant first interacted with LISSA for five minutes, then took a break for two minutes, followed by a second conversation with LISSA for another four minutes. As dialogue topics, we picked ones common in casual conversations such as "getting to know each other", "living in the current city", "crazy room" (aimed at eliciting imaginative responses), "city I want to move to in future", "free time", and "movies". During the conversations, the participants received real-time feedback through the flashing icons. After each conversation, the participants received

post-session summary feedback. We conducted an interview with both the accompanying parent and the participant right after the LISSA session. The interview included survey questions on LISSA's usability and capacity for open-ended discussion. The interview was audio recorded and then transcribed by professional transcribers.

SURVEY RESULTS

We presented 12 statements to the participants and asked them to specify their opinion ('strongly disagree' to 'strongly agree'). The questions were targeted to understand the usefulness of the feedback and the dialogues. The questions were inspired by the well-established system usability questionnaire [9] and modified to our specific needs. Figure 7 shows the specific questions and percentage of participants' answers in each category. The questions marked with a star (*) were answered significantly more ($p < 0.05$) with options 'agree' or 'strongly agree' compared to other options. We performed a single sampled non-parametric significance test [33] with Bonferroni correction [2] against the option 'neutral'.

Participants felt that they were being understood by LISSA. Participants also expressed that they could continue the conversation and pay attention to the icons without any trouble. This indicates that real-time feedback might be applicable to the teenage population. Additionally, participants felt that the feedback they received from LISSA was useful. During our interview session participants expanded on this perceived usefulness. The feedback was consistent and in accord with what their therapist said in the past. For instance, one participant had issues (i.e., slouching) with his posture and he received feedback through the 'body movement' icon. During the interview, the participant mentioned this and said that his parents often ask him to sit straight.

As can be seen in the figure, participants had mixed opinions about several questions, such as whether the conversational experience felt real, and whether LISSA's movements seemed natural. Three participants responded positively about the latter question, four responded negatively, and two were neutral. In the interview session, several mentioned that the lip movement and eye gaze were unnatural. Additionally, LISSA was not responding immediately. This was due to the fact that LISSA processes the dialogues and facial features in real-time and the processing takes place on a remote server. In our future versions of LISSA, we will make it more responsive by performing most of the computing locally.

QUALITATIVE ANALYSIS

We performed a thematic analysis [20] on the interview transcripts. In the past, thematic analysis was successfully used for user-centered design [35] and rapid online interface prototyping [31]. Additionally, thematic analysis was used for identifying the design guidelines for developing computer and phone-based technology to help improve the social skills of the children with ASD [39], [38]. In our analysis, three researchers performed thematic analysis and then the themes were merged to produce the final analysis report. As a basis for a qualitative thematic analysis of the interview transcripts, we considered individual interviews from the perspective of the following 14 labels: *usefulness, perceptiveness, related systems, accuracy, familiarity, curiosity, realism, speed, appearance, improvements, social, multitasking, uncertainty, and adult identity*. As a result, six themes (closely related to some of these labels) stood out to us as relevant to summarizing the experiences of the participants. These are elaborated in the following. The first theme briefly summarizes positively perceived aspects of LISSA, while the rest dwell mostly on aspects where further developments are desirable. We believe it is at least as important to focus on weaknesses as on strengths, as a guide to further development.

Utility for practicing conversation in private

Participants generally found LISSA useful for practicing conversations. Additionally, they thought that the program was not hard to navigate. One participant said,

“Wasn't that hard to use. It could definitely be used for somebody who really needs help with conversations or for somebody who is not really social or for somebody who is not really the kind of person to be talkative.”

Participants liked the fact that the feedback was not coming from a human and they could use it in their private space without being observed. When we asked if they prefer human or a computer for giving feedback, some were ambivalent, some preferred the computer (*“I would rather have that (LISSA) for feedback,”*), and some were skeptical about LISSA.

Some participants with favorable reactions added that the automatic facial feature detection and the accuracy of the feedback were the main reason for endorsing LISSA for conversational skills training. One participant said,

“The fact that it was able to actually detect the facial features and everything being so accurate, I would consider that is good enough to actually train on.”

The caregivers, as well, liked the fact that LISSA allows users to have a conversation with a virtual agent instead of (for example) a stranger online. They felt that LISSA was realistic and appreciated that it provided quite complex, and private interactions.

Self-awareness

When conversing with LISSA, participants were often evaluating the experience in relation to real, social settings, for example, whether the feedback they received was truly appropriate. One participant noted that smiling broadly enough for LISSA to notice might be perceived as inappropriate in a public setting:

“If they can make it so that she can respond a little quicker and that she is able to pick up the smile little better because I was smiling a little bit she just didn't think it was good enough and if I go too big then people are gonna think I am creepy.”

Occasionally the self-aware evaluation of LISSA took the form of push-back against the presumption that they needed to improve their behavior, implicit in LISSA's feedback. One participant indicated that any inadequacy in their behavior with LISSA was due to the awkwardness of interacting with a virtual agent:

“Well, I am actually social. I do make good eye contact with other people but it's just kinda awkward you know.”

Another teenager commented similarly:

“I do try to be social with other people and I am good with that it's just that this was just a little new for me.”

From this it appears that some participants would need some time to become more familiar with LISSA. After multiple, unobserved interactions they might well feel more comfortable with the system. The comments also suggest that users might be more accepting of LISSA if the purpose of the system were more fully explained – that is, it is simply a tool for users who would like to improve technical aspects of their conversational ability, if they feel they could improve in that area.

Realism of the virtual human

Participants focused quite intently on the realism and precision of the LISSA persona. There were debates on how realistic the eyes, lips, face, voice, and overall movement were. Participants suggested various improvements, such as more flowing speech, immediate responses, and faster blinks:

“The other part I didn't like about her is because it took her a while to respond to my statements. So it was kind of confusing and irritating.”

“I would say blinking eyes was definitely a little bit slow, because if I blink it looks almost instant. But for that it was half a second total, it seems quite slow.”

“Nice sounding voice, nice face but the moving the lips thing was kind of a little creepy. Mainly because kind of it’s little too computer-ish maybe.”

When participants were asked about the usefulness of LISSA, their judgments were dependent on realism. For the other interview questions, participants had varying responses, but there always seemed to be an underlying fixation on realism.

“Well, it was good that she asked me what I liked to do. But since it wasn’t really real it just seemed all awkward for me. I would make it better by making the quality High Definition and High quality better with talking to her for real.”

Multi-tasking and feedback

When they were asked about multitasking between the conversations with LISSA and looking at the icons, there were different opinions: some found the icons helpful, and others wanted more explanation and bigger icons.

“Maybe the icons at the bottom make it a little bigger and maybe make it a little stand out a little bit more.”

And when asked about if they preferred feedback from a computer or human, the participants had mixed opinions. For example, one participant said,

“The feedback I received from LISSA was useful. Well it was kind of a bit choppy and kind of pretty much prefer something a bit more realistic.”

Interpreting feedback through LISSA’s behavior while conducting a conversation may be overwhelming for many participants. Perhaps if LISSA herself provided the feedback verbally during the conversations, rather than relying entirely on icons for feedback, it would feel more like a real interaction to users. Regardless of their preferences in feedback delivery, all participants expressed that the feedback was consistent with what other people had said to them in the past face-to-face training sessions. One participant said,

“They (feedback) were actually kind of handy and appeared to be pretty accurate.”

An important observation was that users broadly agreed on the need for real-time positive feedback. While the flashing red icons noticeably indicated the need for improvement, the reversion to static green icons was not sufficiently noticeable as positive feedback. This indicates both the need for changes in icon functioning and, once again, the desirability of verbal feedback (especially acknowledgement and praise) from LISSA.

Related systems

At the end of the interviews, participants tended to compare LISSA to other conversational agents. They were often

familiar with Siri and Alexa, and discussed the standard set by these virtual assistants in terms of speed and knowledgeability. They liked that LISSA could detect behavioral features, but hoped that LISSA could better understand and recognize them in the future. For example, a participant said,

“In its current state yes I might mess around with it some but I don’t believe I would use it as an actual social skills training and jump into actual conversations just yet.”

It seemed that they had high expectations of LISSA which in some cases led to impatience during the conversation. One participant said they would feel more comfortable talking with systems like Siri, that they believe understands them.

“That’s like my first time ever talking to a computer, well except for Siri that’s different though. That one is fine.”

The desire for adult identity

One teenager, after being urged by a caregiver to be honest, admitted that they probably wouldn’t choose to interact with LISSA, unless perhaps LISSA “popped up on their computer”. Part of the reason seemed to be that LISSA straddles the boundary between fantasy and reality; and while fantasy is fine for a kid, for real conversations they prefer actual humans:

“I don’t know if it was like a real person. Cause I do like fantasy things but I am also a kid who but when it comes to talking to people I like talking to real people.”

These participants were ambivalent about the future use of LISSA because they were inclined to regard it as a conversational tool for children. One participant said their schedule was too busy, but others could benefit from LISSA.

“Just so you know I am already 17 years old, I am growing up and some of these little kids’ things I have outgrown but not all of them.”

In one conversation, a participant felt momentarily uncomfortable with LISSA’s comment about a “crazy room,” (asking them to speculate what kind of crazy room they would enjoy) and said they wanted to feel like an average adult:

“Well, pretty much just the crazy room cause well I wanna be what your average adult is. Basically responsible, kind, but also a bit unique.”

These comments again suggest the need to make LISSA’s purpose clearer to users. It is not intended to be a surrogate human, but rather a tool for repetitive, private practice of conversational behavior. Also, being sensitive in the choice of terminology is important, and perhaps asking for a creative response is inappropriate for some participants.

DISCUSSION

Lessons about interface design for autistic teens

A majority of participants found that LISSA provided useful feedback and might well be helpful for practicing their

conversational skills. As we noted, they also liked the fact that LISSA would allow them to converse in private. Another point of interest is that the participants preferred real-time feedback to post-session feedback.

Our interviews with users helped to shed light on how the interface design could be further improved. Our qualitative analysis of these interviews provide evidence for the soundness of our design so far and grounds for optimism about our further development plans. The analysis also allowed us to formulate several important guidelines for the design of LISSA-like interfaces for conversation practice.

Appropriate Prior Briefing of Users about LISSA's Purpose

Our experience with users made clear that users' assessment of LISSA's behavior and potential utility for conversation practice depended very much on their expectations. They generally acknowledged that LISSA seemed to understand them and responded appropriately to inputs, yet was not genuinely human-like. The perceived shortcomings concerned LISSA's physical behavior, accuracy of perception of user behavior, and depth of knowledge. In part, this perception arose from comparisons with commercial systems such as Alexa and Siri, which have been optimized for smooth functioning in targeted information retrieval and other assistive functions.

These reactions indicate the need for fuller preparation of users about LISSA's purpose: It is not a surrogate human, and it is not an app for access to useful knowledge or personal assistance. It is simply a tool for repeatedly practicing casual conversation for those who feel they could improve in that area. While LISSA has a range of verbal reactions to users depending on their particular inputs, and provides nonverbal feedback as a function of the user's behavior, the conversations are bound to be shallow, and to become more repetitive with multiple uses. Further, LISSA's physical behavior is not the focus; it is merely intended to be sufficiently human-like to make a casual conversation possible. All this should be made clear to potential users – along with a comment that there is no assumption that all users with ASD are lacking in the skills that LISSA is intended to help with. Such preparatory information could be provided both in advance of actual use of LISSA, and as part of LISSA's opening remarks (which already include “*I might sound a bit choppy, but ...*”). The post-session interviews of users should likewise focus on relevant aspects of LISSA's functioning. For instance, interview questions might include disclaimers, as in, “*We know that LISSA doesn't smile and blink very naturally, but was the content of her responses to you reasonable and natural?*”. In general, it is evidently important to prepare users not only for the capabilities of an interactive system, but also its limitations.

Positive Acknowledgment of Behavior Change

As noted in the previous section, the participants wanted to be made aware of positive changes. Perhaps flashing green could be used for behavioral improvements. Better yet, the virtual agent could say, for example, “*You have good eye*

contact now”. The efficacy of positive feedback and acknowledgment has been observed in past research [56], [48], [29], and our experience further confirms the desirability of positive acknowledgments for interventions aimed at social skills development.

Realistic appearance of a virtual agent and responsiveness

Notwithstanding disclaimers about LISSA's physical behavior, the issue deserves further attention. A possible reaction to users' comments about insufficiently realistic smiling, eye blinking, and reaction speed might be to back away from the “uncanny valley” (e.g., Chattopadhyay et al. [12]) by using a more cartoon-like avatar. However, this would risk reducing LISSA to a toy in the eyes of potential users. Instead, we interpret the users' comments as urging further development of the avatar towards greater realism. This is consistent with their age – they are approaching adulthood, and prefer a realistic to a childish avatar. In fact their comments suggest that the more life-like the character, the more likely they are to take it seriously. Realistic appearance of virtual characters has also been shown to be effective in other scenarios such as negotiation, tactical questioning etc. [30], [52]. The most important areas for improvement seem to be smiling and eye blinking. (Smiling is of course well-known to be very important in communication.) For example, smiling needs to be consistent with current feedback (one participant commented on co-occurrence of a smile by LISSA with negative feedback). Furthermore, smiles could be used directly to indicate improvements, or in support of positive verbal or icon feedback.

Conversation personalization would help the teenagers engage for longer terms.

Some participants expressed enthusiasm about home-use of LISSA. However, they varied in their opinions about the choices of topics. For example, while the “crazy room” topic struck a chord with some (e.g., they would fill it with video games), others objected to it, terming it childish. In addition, participants thought that it would be useful if LISSA could talk about topics of their own choosing. For example, one participant was very interested in computer programming and wanted to talk about it more. The LISSA program, at its current stage, is designed for initiating conversational topics, and treating a specialized topic like computer programming seriously would be a major challenge. However, adding further mundane topics is quite feasible, and in fact we have added many more in a related application currently under development. Thus we could personalize interactions to a considerable degree by having LISSA choose topics dynamically, skipping those that the user seems indifferent to. Also, choices could be made sensitive to the user's age or maturity. An immature 13-year-old may have quite different interests from a mature 18-year-old. (Some of our newly developed topics pertain only to seniors, just as some of LISSA's topics for ASD teens, such as bullying at school, pertain only to school-age users.)

Another opportunity for personalization lies in the verbosity or otherwise of the user. Our experiments showed that while some users provide expansive responses to LISSA, others respond tersely. As the goal of the system is to help users improve their communication skills, the system could gauge users' verbosity and provide helpful feedback where appropriate. For instance, LISSA might encourage laconic users elaborate their answers, or conversely, provide gentle suggestions about curtailing rambling or off-topic inputs.

Assessing users' verbosity throughout the conversation can also help improve turn-taking behavior. Users who tend towards longer answers should probably be allowed slightly longer silences before the turn is seized from them. The same applies to hesitant, slower speakers. Certainly humans adapt to such individual differences. This is an important research area -- to our knowledge, no available automatic turn-handling methods take into account individual speakers' verbosity or rate of speech.

One of the most important observations we made about teens with ASD in comparison with (neurotypical) college students, independently of verbosity, was that the ASD teens refrained from asking any reciprocal or other questions of LISSA (e.g., after telling LISSA about their favorite movie, asking "*and what's your favorite movie?*"). Whether this is due to less willingness to treat LISSA as human-like, or to limitations in social intuition, it is an area where verbal feedback by LISSA could be particularly useful; for instance LISSA might say, "*This would be a good point to ask me about my favorite movie. Would you like to try?*".

Limitations and Future Work

The current version of LISSA was not designed for immediate use in a randomized control intervention study. Rather, it is an exploratory system, which will enable a randomized control study after modification and enhancement based on the lessons learned from the trials with the initial set of ASD teens.

LISSA's dialogue manager was adapted from the initial version for college students to the anticipated needs of ASD teens, with advice from experts. It worked well, but the experiments have shown where improvements are most desirable, for example in topical adaptation to the user, inclusion of direct helpful hints in the verbal reactions to the user, and allowance for different turn-taking styles. Similarly, LISSA's nonverbal feedback system was trained on the data collected from college students, and although the participants perceived the feedback as useful, our experimental results indicate ways in which flashing icons, sensitivity to user smiles, and reaction speed could be improved. Also, while the post-session interviews indicated that the users liked the appearance and voice of the avatar, they saw a need for improvements in the naturalness of the avatar's behavior (especially smiles and blinking of eyelids). In our future work, we will design a customizable interface based on the knowledge we gathered through this study.

Data collected from teenagers with ASD using future versions of the system will help us to further improve the sensitivity and responsiveness of the system. In the current system, the dialogue and the feedback modules are independent. It clearly would be useful to tie the nonverbal feedback to the dialogue content, and to supplement nonverbal feedback signals with direct verbal ones. In the future, we will make the feedback dialogue aware.

CONCLUSION

In this paper, we described an interface capable of conducting a multi-topic conversation and provide feedback to help improve the user's overt behavior. The design benefited from the expertise of a pediatrician, psychologists, and a UX designer. We investigated further design desiderata through a study with nine teenagers with ASD. Using a thematic analysis we formulated several guidelines for improved interface design for teens with ASD. In future this knowledge will help guide interface design for this population as well as others with similar needs for social skills enhancement in casual dialogues.

REFERENCES

1. reference withheld to preserve Anonymity. No Title.
2. Richard A. Armstrong. 2014. When to use the Bonferroni correction. *Ophthalmic & physiological optics: the journal of the British College of Ophthalmic Opticians (Optometrists)*. <http://doi.org/10.1111/opo.12131>
3. Tadas Baltruš and Peter Robinson. OpenFace: an open source facial behavior analysis toolkit. <http://doi.org/10.1109/WACV.2016.7477553>
4. Laura Benton, Hilary Johnson, Mark Brosnan, Emma Ashwin, and Beate Grawemeyer. 2011. Ideas: An Interface Design Experience for Autistic Spectrum. *Proceedings of the 2011 annual conference extended abstracts on Human factors in computing systems - CHI EA '11*. <http://doi.org/10.1145/1979742.1979841>
5. Sara Bernardini, Kaska Porayska-Pomsta, and Harin Sampath. 2013. Designing an Intelligent Virtual Agent for Social Communication in Autism. *Proceedings of the Ninth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment Designing*.
6. Paul Boersma and David Weenink. Praat: doing phonetics by computer. Retrieved from <http://www.fon.hum.uva.nl/praat/>
7. Fatima Boujarwah, Gregory Abowd, and Rosa Arriaga. 2012. Socially computed scripts to support social problem solving skills. *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*, 1987. <http://doi.org/10.1145/2207676.2208343>

8. LouAnne E. Boyd, Alejandro Rangel, Helen Tomimbang, et al. 2016. SayWAT: Augmenting Face-to-Face Conversations for Adults with Autism. *CHI*. <http://doi.org/10.1145/2858036.2858215>
9. John Brooke. 2013. SUS : A Retrospective. *Journal of Usability Studies* 8, 2: 29–40. Retrieved from http://www.usabilityprofessionals.org/upa_publications/jus/2013february/brooke1.html
<http://www.usability.gov/how-to-and-tools/methods/system-usability-scale.html>
10. Latha Carpenter, Laura Arnstein Soorya and Danielle Halpern. 2009. Asperger's Syndrome and High-Functioning Autism. *Pediatric Annals*.
11. Centers Center for Disease Control, Prevention, and others. 2013. Learn the signs. Act early. *Program (www.cdc.gov/ActEarly)*.
12. Debaleena Chattopadhyay and Karl F. MacDorman. 2016. Familiar faces rendered strange: Why inconsistent realism drives characters into the uncanny valley. *Journal of Vision*. <http://doi.org/10.1167/16.11.7>
13. Pau Choo Chung and Chin De Liu. 2008. A daily behavior enabled hidden Markov model for human behavior understanding. *Pattern Recognition*. <http://doi.org/10.1016/j.patcog.2007.10.022>
14. David DeVault, Ron Artstein, Grace Benn, et al. 2014. SimSensei Kiosk: A Virtual Human Interviewer for Healthcare Decision Support. *International Conference on Autonomous Agents and Multi-Agent Systems*, 1061–1068.
15. Eli J Finkel, Paul W Eastwick, and Jacob Matthews. 2007. Speed-dating as an invaluable tool for studying romantic attraction: A methodological primer. *Personal Relationships* 14, 1: 149–166. <http://doi.org/10.1111/j.1475-6811.2006.00146.x>
16. Kathleen Kara Fitzpatrick, Alison Darcy, and Molly Vierhile. 2017. Delivering Cognitive Behavior Therapy to Young Adults With Symptoms of Depression and Anxiety Using a Fully Automated Conversational Agent (Woebot): A Randomized Controlled Trial. *JMIR Mental Health*. <http://doi.org/10.2196/mental.7785>
17. Mary Ellen Foster, Katerina Avramides, Sara Bernardini, et al. 2010. Supporting children's social communication skills through interactive narratives with virtual characters. *Proceedings of the international conference on Multimedia - MM '10*: 1111–1114. <http://doi.org/10.1145/1873951.1874163>
18. Eynat Gal, Liron Lamash, Nirit Bauminger-Zviely, Massimo Zancanaro, and Patrice L. (Tamar) Weiss. 2016. Using Multitouch Collaboration Technology to Enhance Social Interaction of Children with High-Functioning Autism. *Physical & Occupational Therapy In Pediatrics*. <http://doi.org/10.3109/01942638.2015.1040572>
19. Alexandra Livia Georgescu, Bojana Kuzmanovic, Daniel Roth, Gary Bente, and Kai Vogeley. 2014. The Use of Virtual Characters to Assess and Train Non-Verbal Communication in High-Functioning Autism. *Frontiers in Human Neuroscience*. <http://doi.org/10.3389/fnhum.2014.00807>
20. Gregory Guest, Kathleen MacQueen, and Emily Namey. 2012. Introduction to applied thematic analysis. *Applied Thematic Analysis*. <http://doi.org/http://dx.doi.org/10.4135/9781483384436>
21. Gillian R. Hayes, Sen Hirano, Gabriela Marcu, Mohamad Monibi, David H. Nguyen, and Michael Yeganyan. 2010. Interactive visual supports for children with autism. *Personal and Ubiquitous Computing*. <http://doi.org/10.1007/s00779-010-0294-8>
22. Richard G Heimberg. 2002. Cognitive-behavioral therapy for social anxiety disorder: Current status and future directions. *Biological Psychiatry* 51, 1: 101–108. [http://doi.org/10.1016/S0006-3223\(01\)01183-0](http://doi.org/10.1016/S0006-3223(01)01183-0)
23. Peter Henderson, Koustuv Sinha, Nicolas Angelard-Gontier, et al. 2017. Ethical Challenges in Data-Driven Dialogue Systems. Retrieved from <http://arxiv.org/abs/1711.09050>
24. Ingrid Maria Hopkins, Michael W Gower, Trista A Perez, et al. 2011. Avatar assistant: Improving social skills in students with an asd through a computer-based intervention. *Journal of Autism and Developmental Disorders* 41, 11: 1543–1555. <http://doi.org/10.1007/s10803-011-1179-z>
25. Juan Pablo Hourcade, Natasha E. Bullock-Rest, and Thomas E. Hansen. 2012. Multitouch tablet applications and activities to enhance the social skills of children with autism spectrum disorders. *Personal and Ubiquitous Computing*. <http://doi.org/10.1007/s00779-011-0383-3>
26. Juan Pablo Hourcade, Stacy R. Williams, Ellen A. Miller, Kelsey E. Huebner, and Lucas J. Liang. 2013. Evaluation of tablet apps to encourage social interaction in children with autism spectrum disorders. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*. <http://doi.org/10.1145/2470654.2466438>
27. Muhamad Fairus Kamaruzaman, Nurdalilah Mohd Rani, Harrinni Md Nor, and Mustafa Halabi Haji Azahari. 2016. Developing User Interface Design

- Application for Children with Autism. *Procedia - Social and Behavioral Sciences*. <http://doi.org/10.1016/j.sbspro.2016.02.022>
28. Connie Kasari, Nancy Brady, Catherine Lord, and Helen Tager-Flusberg. 2013. Assessing the minimally verbal school-aged child with autism spectrum disorder. *Autism Research*. <http://doi.org/10.1002/aur.1334>
 29. Connie Kasari, Marian Sigman, Peter Mundy, and Nurit Yirmiya. 1988. Caregiver interactions with autistic children. *Journal of Abnormal Child Psychology*. <http://doi.org/10.1007/BF00910499>
 30. Patrick Kenny, Arno Hartholt, Jonathan Gratch, et al. 2007. Building Interactive Virtual Humans for Training Environments. *Interservice/Industry Training, Simulation, and Education Conference (IITSEC) 2007*. <http://doi.org/10.1.1.116.662>
 31. M B Kinzie, W F Cohn, M F Julian, and W a Knaus. 2002. A user-centered model for web site design: needs assessment, user interface design, and rapid prototyping. *J Am Med Info Assoc*. <http://doi.org/10.1197/jamia.M0822.s>
 32. Miriam Madsen, Rana el Kaliouby, Micah Eckhardt, Mohammed E. Hoque, Matthew S. Goodwin, and Rosalind Picard. 2009. Lessons from participatory design with adolescents on the autism spectrum. *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems - CHI EA '09*. <http://doi.org/10.1145/1520340.1520580>
 33. H. B. Mann and D. R. Whitney. 1947. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics*. <http://doi.org/10.1214/aoms/1177730491>
 34. Morag Maskey, Frances Warnell, Jeremy R. Parr, Ann Le Couteur, and Helen McConachie. 2013. Emotional and behavioural problems in children with autism spectrum disorder. *Journal of Autism and Developmental Disorders*. <http://doi.org/10.1007/s10803-012-1622-9>
 35. Tara McCurdie, Svetlana Taneva, Mark Casselman, et al. 2012. mHealth consumer apps: the case for user-centered design. *Biomedical instrumentation & technology / Association for the Advancement of Medical Instrumentation*. <http://doi.org/10.2345/0899-8205-46.s2.49>
 36. Marissa Milne, David Powers, and Richard Leibbrandt. 2009. Development of a software-based social tutor for children with autism spectrum disorders. *Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group on Design: Open 24/7 - OZCHI '09*. <http://doi.org/10.1145/1738826.1738870>
 37. Adam S. Miner, Arnold Milstein, Stephen Schueller, Roshini Hegde, Christina Mangurian, and Eleni Linos. 2016. Smartphone-based conversational agents and responses to questions about mental health, interpersonal violence, and physical health. *JAMA Internal Medicine*. <http://doi.org/10.1001/jamainternmed.2016.0400>
 38. Joseph Mintz. 2013. Additional key factors mediating the use of a mobile technology tool designed to develop social and life skills in children with Autism Spectrum Disorders: Evaluation of the 2nd HANDS prototype. *Computers and Education*. <http://doi.org/10.1016/j.compedu.2012.11.006>
 39. Joseph Mintz, Corinne Branch, Caty March, and Stephen Lerman. 2012. Key factors mediating the use of a mobile technology tool designed to develop social and life skills in children with Autistic Spectrum Disorders. *Computers & Education*. <http://doi.org/10.1016/j.compedu.2011.07.013>
 40. Fabrizio Morbini, Eric Forbell, and D DeVault. 2012. A mixed-initiative conversational dialogue system for healthcare. ... *and Dialogue*.
 41. Emily Mower, Matthew P. Black, Elisa Flores, Marian Williams, and Shrikanth Narayanan. 2011. Rachel: Design of an emotionally targeted interactive agent for children with autism. *Proceedings - IEEE International Conference on Multimedia and Expo*. <http://doi.org/10.1109/ICME.2011.6011990>
 42. B. Nojavanasghari, C.E. Hughes, and L.-P. Morency. 2017. Exceptionally social: Design of an avatar-mediated interactive system for promoting social skills in children with autism. *Conference on Human Factors in Computing Systems - Proceedings*. <http://doi.org/10.1145/3027063.3053112>
 43. Anne Marie Piper, Eileen O'Brien, Meredith Ringel Morris, and Terry Winograd. 2006. SIDES: a cooperative tabletop computer game for social skills development. *Proceedings of the 2006 conference on computer supported cooperative work*. <http://doi.org/10.1145/1180875.1180877>
 44. Seyedeh Zahra Razavi, Mohammad Rafayet Ali, Tristram H. Smith, Lenhart K. Schubert, and Mohammed Ehsan Hoque. 2016. The LISSA virtual human and ASD teens: An overview of initial experiments. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. http://doi.org/10.1007/978-3-319-47665-0_55

45. Seyedeh Zahra Razavi, Lenhart K Schubert, Mohammad Rafayet Ali, and Mohammed Ehsan Hoque. Managing Casual Spoken Dialogue Using Flexible Schemas, Pattern Transduction Trees, and Gist Clauses. <http://doi.org/10.1.1.14.5879>
46. Hyekyun Rhee, James Allen, Jennifer Mammen, and Mary Swift. 2014. Mobile phone-based asthma self-management aid for adolescents (mASMAA): A feasibility study. *Patient Preference and Adherence*. <http://doi.org/10.2147/PPA.S53504>
47. James Ladell Sanders. 2009. Qualitative or quantitative differences between asperger's disorder and autism? Historical considerations. *Journal of Autism and Developmental Disorders*. <http://doi.org/10.1007/s10803-009-0798-0>
48. Ashley A. Scott-Van Zeeland, Mirella Dapretto, Dara G. Ghahremani, Russell A. Poldrack, and Susan Y. Bookheimer. 2010. Reward processing in autism. *Autism Research*. <http://doi.org/10.1002/aur.122>
49. Iulian Vlad Serban, Tim Klinger, Gerald Tesauro, et al. 2016. Multiresolution Recurrent Neural Networks: An Application to Dialogue Response Generation. 3288–3294. Retrieved from <http://arxiv.org/abs/1606.00776>
50. Stanley Siegman, Aron W.; Feldstein. 2014. *Nonverbal Behavior and Communication*. Psychology Press.
51. Emily Simonoff, Andrew Pickles, Tony Charman, Susie Chandler, Tom Loucas, and Gillian Baird. 2008. Psychiatric Disorders in Children With Autism Spectrum Disorders: Prevalence, Comorbidity, and Associated Factors in a Population-Derived Sample. *Journal of the American Academy of Child & Adolescent Psychiatry*. <http://doi.org/10.1097/CHI.0b013e318179964f>
52. M Tambe, W L Johnson, R M Jones, et al. 1995. Intelligent Agents for Interactive Simulation Environments. *AI Magazine*.
53. Hiroki Tanaka, Hideki Negoro, Hidemi Iwasaka, and Satoshi Nakamura. 2017. Embodied conversational agents for multimodal automated social skills training in people with autism spectrum disorders. *PLoS ONE*. <http://doi.org/10.1371/journal.pone.0182151>
54. Andrea Tartaro and Justine Cassell. 2006. Authorable virtual peers for autism spectrum disorders. *Proceedings of the Combined workshop on LanguageEnabled Educational Technology and Development and Evaluation for Robust Spoken Dialogue Systems at the 17th European Conference on Artificial Intelligence*. <http://doi.org/ECAI06>
55. U.S. Department of Health and Human Services. 2014. Prevalence of autism spectrum disorder among children aged 8 years - autism and developmental disabilities monitoring network, 11 sites, United States, 2010. *MMWR Surveill Summ*. <http://doi.org/24670961>
56. Stella Waterhouse. 2000. *A positive approach to autism*. Jessica Kingsley Publishers.
57. Jason D. Williams and Geoffrey Zweig. 2016. End-to-end LSTM-based dialog control optimized with supervised and reinforcement learning. <http://doi.org/arXiv:1504.01391>
58. Junji Yamato, Jun Ohya, and Kenichiro Ishii. 1992. Recognizing human action in time-sequential images using hidden Markov model. *Computer Vision and Pattern Recognition*. <http://doi.org/10.1109/CVPR.1992.223161>
59. SHORE™ - Object and Face Recognition. Retrieved from <http://www.iis.fraunhofer.de/en/ff/bsy/tech/bildanal/ysc/shore-gesichtsdetektion.html>
60. NUANCE: Speech Recognition Solutions. Retrieved from <http://www.nuance.com/for-individuals/by-solution/speech-recognition/index.htm>