

NeuroBridge: Using Generative AI to Bridge Cross-neurotype Communication Differences through Neurotypical Perspective-taking

Rukhshan Haroon
Computer Science
Tufts University
MA, United States
rukhsan.haroon@tufts.edu

Kyle Wigdor
Computer Science
Tufts University
MA, United States
kyle.wigdor@tufts.edu

Katie Yang
Computer Science
Tufts University
MA, United States
zihan.yang@tufts.edu

Nicole Toumanios
Child Studies and Human
Development, Cognitive Science
Tufts University
MA, United States
nicole.toumanios@tufts.edu

Eileen T. Crehan
Eunice Kennedy Shriver Center
UMass Chan Medical School
MA, United States
eileen.crehan@umassmed.edu

Fahad Dogar
Computer Science
Tufts University
MA, United States
fahad@cs.tufts.edu

Abstract

Communication challenges between autistic and neurotypical individuals stem from a *mutual* lack of understanding of each other's distinct, and often contrasting, communication styles. Yet, autistic individuals are expected to adapt to neurotypical norms, making interactions inauthentic and mentally exhausting for them. To help redress this imbalance, we build NeuroBridge, an online platform that utilizes large language models (LLMs) to simulate: a) an AI character that is direct and literal, a style common among many autistic individuals, and b) four cross-neurotype communication scenarios in a feedback-driven conversation between this character and a neurotypical user. Through NeuroBridge, neurotypical individuals gain a firsthand look at autistic communication, and reflect on *their* role in shaping cross-neurotype interactions. In a user study with 12 neurotypical participants, we find that NeuroBridge improved their understanding of how autistic people may interpret language differently, with all describing autism as a social difference that "needs understanding by others" after completing the simulation. Participants valued its personalized, interactive format and described AI-generated feedback as "constructive", "logical" and "non-judgmental". Most perceived the portrayal of autism in the simulation as accurate, suggesting that users may readily accept AI-generated (mis)representations of disabilities. To conclude, we discuss design implications for disability representation in AI, the need for making NeuroBridge more personalized, and LLMs' limitations in modeling complex social scenarios.

Author's Note: This paper has been conditionally accepted for publication in the Proceedings of the ACM SIGACCESS Conference on Computers and Accessibility (ASSETS 2025), to be held October 26–29, 2025, in Denver, CO, USA. This is a preprint version.

1 Introduction

Autism Spectrum Disorder (ASD) is a complex neurodevelopmental condition marked by differences in communication, cognition, sensory processing, and social behavior compared to neurotypical

development [1–3]. It is one of the most common neurodevelopmental conditions in the U.S., affecting an estimated 1 in 45 adults [4]. Key traits of autistic communication include a preference for a direct conversational style [5, 6], literal language [7, 8], and minimal use of social cues [9, 10]. These often contrast with neurotypical communication norms, which involve phatic exchanges, implied intent, and social nuance [5, 11–13]. Prior work shows that cross-neurotype communication breakdowns due to these differences can have severe consequences for autistic individuals, such as social exclusion in both online and physical social spaces [5, 6, 14, 15], professional setbacks [16, 17], and barriers to quality healthcare [18, 19].

Prior efforts to bridge this divide include technological, educational, and therapy-based interventions [6, 20–22]. However, these have predominantly targeted autistic individuals, often pressuring them to conform to neurotypical norms. The double empathy problem underscores that communication challenges between autistic and neurotypical individuals arise from reciprocal misunderstandings, necessitating efforts from both sides to work toward mutual understanding and acceptance [23]. Yet, interventions at the neurotypical end are nearly nonexistent, limited to passive, informational resources that offer no opportunities to practice learned concepts or incentives to get involved [24, 25]. This imbalance places the burden of adapting communication styles almost entirely on autistic individuals.

As large language model (LLM) powered chatbots like ChatGPT [26] and Character.AI [27] gain widespread traction, with tens of millions of users engaging with them daily, LLMs present a powerful, new avenue for designing immersive, scalable, and personalized human-AI interactions. Their ability to generate fluent, human-like text, interpret subtle linguistic cues, and adapt to diverse conversational styles makes them well-suited for simulating real-world communication scenarios, including those involving different neurotypes [28–30]. We believe this capability, if utilized responsibly, can be used to engage neurotypical individuals in interactive, personalized learning experiences that cultivate empathy and appreciation for autistic communication styles. While existing applications




Figure 1: NeuroBridge architecture and interaction flow. Users begin by entering a topic and then engage in a loop of sending messages, receiving responses, and getting feedback.

of LLMs in this space focus on providing communication support to autistic individuals [6, 31], we advocate for shared responsibility and shift the focus of intervention to the neurotypical end.

In this paper, we present NeuroBridge, an interactive platform designed to help neurotypical individuals better understand autistic forms of expression, and reflect on how their own behavior shapes cross-neurotype interactions. At its core, NeuroBridge utilizes LLMs to simulate: a) an AI character configured to be direct and literal, a style common among many autistic individuals, and b) four cross-neurotype communication scenarios in a feedback-driven conversation between the character and a neurotypical user. Informed by prior work and vetted by an advisory board of autistic individuals, these scenarios (outlined in Table 1) reflect common communication challenges faced by autistic individuals [5–8]. The character may request clarification from users when needed, and for each scenario, users are given tailored feedback to work through their differences with the character empathically. Through NeuroBridge, neurotypical individuals gain a firsthand look at autistic communication, and learn how to communicate more effectively with autistic individuals. Throughout development, members of the advisory board evaluated NeuroBridge at different stages and provided in-depth design feedback, which was incorporated to ensure the simulation reflected their lived experience as closely as possible.

Through an in-lab user study with 12 neurotypical participants recruited from a university setting, aged 18 to 34, we gather survey and in-depth qualitative data on the perceived usefulness of NeuroBridge, how it shaped participants’ perceptions of autism, their attitudes toward AI feedback, and LLMs’ ability to model complex communication scenarios. We find that NeuroBridge improved participants’ understanding of how autistic people may interpret language differently, with all describing autism as a social difference that “needs understanding by others” after completing the simulation. Participants valued the simulation’s personalized, interactive format and described AI-generated feedback as “constructive,” “logical,” and “non-judgmental.” On certain occasions, however, participants found the feedback instructional, which led to feelings of defensiveness. Most perceived the portrayal of autism

in the simulation as accurate, suggesting that users may readily accept AI-generated (mis)representations of disabilities. Despite strong overall performance, our findings show that LLMs are more adept at simulating certain social scenarios than others. To conclude our work, we present a discussion around the implications for representing disabilities through AI, the need and opportunities for making NeuroBridge more personalized, and the limitations of LLMs in modeling complex social scenarios.

To summarize, we make the following key contributions:

- Make a case to integrate the double empathy problem’s theoretical framework into practice for bridging cross-neurotype communication differences in a neurodiversity-affirming manner.
- Co-design and implement NeuroBridge with an advisory board of autistic individuals to help neurotypical people better understand autistic forms of expression, and reflect on their role in shaping cross-neurotype interactions.
- Evaluate NeuroBridge in an IRB-approved study with 12 neurotypical participants, collecting in-depth feedback on the simulation’s usefulness and impact on user perceptions of autistic abilities, user attitudes toward AI feedback, and the effect of LLM-driven personalization on user engagement.
- Present a discussion around design implications for representing disabilities through AI, the need and opportunities for making NeuroBridge more personalized, and the limitations of LLMs in modeling various, complex social scenarios.

2 Related Work

In this section, we review common characteristics of autistic communication, different types of interventions in autism, and the role of technology, particularly LLMs, in advancing them.

Scenario / Challenge	Description	Example	Interpretations
Indirect Speech Act	A statement with an implicit request or intent.	Can you open the window?	A literal question about the possibility of opening the window, or a polite request to open it.
Figurative Expression	A phrase whose meaning goes beyond the literal interpretation of words.	She has a chip on her shoulder.	A literal reference to something on one's shoulder, or as an idiom, one holds a grudge.
Emoji with Variable Interpretations	An emoji with fluid meaning, dependent on context, tone, and personal experience.	That presentation was on man... 🔥	The presentation was impressive, or as sarcasm, it was poor.
Being Misperceived as Blunt	A direct statement by an autistic person that unintentionally comes off as rude or blunt.	I don't like your idea at all.	A neutral expression of opinion, or harshly expressed criticism.

Table 1: List of communication challenges simulated in NeuroBridge, along with a description, example statement, and the different interpretations of the example that could cause misunderstanding in each scenario.

2.1 Characteristics of Autistic Communication

Numerous studies in disabilities and linguistics research show that key traits of autistic communication include a preference for a direct conversational style [5, 6], literal language [7, 8], and minimal use of social cues [9, 10]. These autistic norms are known to be rooted in Gricean maxims, which are unwritten rules that guide conversational cooperation by encouraging speakers to be truthful, clear, relevant, and concise [32]. For example, when asked, “Can you open the window?”, an autistic individual might respond with a literal “Yes,” interpreting it as a question about ability rather than a request. Similarly, it has been observed that autistic individuals may take figurative expressions such as sarcasm, metaphors, or sexual innuendos at face value [33]. Such literal interpretations can make it hard to infer others’ intentions or navigate the implicit nature of everyday conversation [32]. In addition, there is a common misconception that autistic people lack empathy, because their preference for directness may not align with socially accepted norms, and as a result, perceived as bluntness [34]. While these styles are common among many autistic individuals, it is important to note that autism is a spectrum, and they do not apply to all autistic individuals [2].

Communication breakdowns caused by these differences can lead to adverse consequences for autistic individuals, such as social exclusion in online and physical social spaces [5, 6, 14, 15], professional setbacks [16, 17], and barriers to quality healthcare [18, 19]. For example, autistic users have reported struggling to navigate innuendos in conversations with potential dates on dating applications, and facing harsh reactions on online public forums for being perceived as rude, as opposed to direct and factual, by others [5, 6]. Similarly, doctors may find it difficult to fully understand an autistic patient’s symptoms if they don’t express themselves in a way that aligns with their expectations [18]; in workplace environments, where traits such as diplomacy and politeness are valued, being overly direct can impact relationships with colleagues and slow down career advancement [35, 36]. Therefore, bridging these

differences is crucial to improving the day-to-day lives of autistic individuals.

2.2 Interventions in Autism

A number of educational, therapeutic, and technological interventions have been developed to support social skills development in autistic individuals. For example, peer-mediated interventions involve typically developing peers to support social interaction and communication development in classroom settings [37, 38]. Applied Behavior Analysis (ABA), though controversial in some communities, is commonly used to teach social skills through reinforcement [39, 40]. Additionally, research in Human-Computer Interaction (HCI) has advanced support through multiple technology-driven interventions [41–48]. For example, Park et al. combined augmented reality (AR) with drama therapy to facilitate accessible and adaptable language therapy for autistic children [46], while Ringland et al. built a whole-body interface to augment dance therapy for autistic children with sensory sensitivities [48]. Prior studies highlight the benefits of incorporating technology into interventions, such as greater user engagement [49], access to support [44], and customization for catering individual needs [50]. Broadly, these approaches align with the interventionist and medical models of disability, which view disability as an impairment to be managed or treated through targeted support [51, 52].

On the other hand, the social model of disability emphasizes that disabilities arise not from individual deficits, but from the mismatch between individuals and their social environments [53, 54]. In the specific context of autism, the ‘double empathy problem’ is a concept grounded in neurodiversity theory which posits that communication breakdowns between autistic and non-autistic individuals are bidirectional, stemming from differences in conversational norms and emotional expression [55]; thus, they are the result of mutual misunderstandings, rather than lack of empathy or deficits on part of autistic individuals. As such, interventions should support bidirectional accommodations, rather than focusing solely

on training autistic individuals to conform to neurotypical norms [56, 57]. Yet, interventions at the neurotypical end remain scarce, often limited to passive, informational resources that provide little opportunity for practicing learned concepts or incentive to get involved [24, 25, 58]. Notable efforts in this space include Autismity – The Autism Simulator [59], a VR-based awareness tool, and The Autism Reality Experience [60], a mobile sensory van. However, these initiatives are costly, difficult to scale, and primarily focus on the physical and sensory experiences of autistic people. While disability simulation has been critiqued for reinforcing stereotypes and erasing the everyday, systemic experiences of disabled people in the past [61, 62], recent scholarship suggests that framing it as an educational tool through the social model can improve its effectiveness [63]. Our work builds on this line of work by engaging participants as active learners rather than passive observers, providing personalized, targeted feedback to encourage accountability, self-reflection, and change, while remaining scalable and cost-effective.

2.3 LLMs, Communication, and Accessibility

Recent advances in generative AI have led to the emergence of large language models (LLMs) such as GPT-4 [26] and LLaMa [64]. LLMs are capable of generating fluent, human-like text, interpreting subtle linguistic cues, and adapting to a variety of conversational styles [28, 29]. These capabilities have opened up new possibilities for designing communication support tools for people with diverse needs, including those who are dyslexic, hard-of-hearing, and use augmentative and alternative communication devices [6, 31, 65–69]. Specifically in the context of autism, Jang et al. examined the use of LLMs for communication assistance at the workplace, finding that autistic individuals prefer LLMs over human colleagues due to greater convenience/availability, neutrality, and privacy [31]. Haroon et al. integrated LLMs into an instant messaging application to provide autistic users with in-situ communication assistance, and found that LLMs offer a convenient way for them to seek clarifications, provide a better alternative to tone indicators, and facilitate constructive reflection on writing technique and style [6]. Similarly, Barros et al. conducted participatory workshops with autistic social media users to identify their design needs and develop new features to address them; LLMs show promise to power many of the envisioned features [5].

However, most of these approaches reinforce a deficit-oriented model of disability by promoting adaptation to dominant social norms. In contrast, our work aims to use LLMs to help neurotypical individuals understand autistic forms of expression, and how their own behaviors shape cross-neurotype communication. Our approach directly aligns with Boyd’s concept of celebratory technologies, which highlights the value of neurodivergent ways of being and advocates for interventions that promote dignity, agency, and social inclusion, rather than focusing on remediation [70]. Researchers have also worked on identifying and mitigating risks, biases, and ethical concerns related to LLMs and disability [71–74]. Beyond accessibility, LLMs have been used to enhance efficiency, safety, and decision-making across several, critical domains [75–78].




Figure 2: The main interface of NeuroBridge is designed to replicate regular messaging apps, making it feel familiar to users. The message in the blue bubble was sent by the user, while the message in the gray bubble was sent by Julia, the AI character.

3 Overview of NeuroBridge

In this section, we outline the key components, design and implementation of NeuroBridge.

3.1 Components of NeuroBridge

All components of NeuroBridge (described below) are powered by an LLM. Figure 1 captures how these components interact with one another. For each component, we provide the LLM with a unique ‘prompt’ – a carefully crafted instruction given as input to guide the model’s output. A detailed description of our prompting strategy is provided in Section 3.2.3. As shown in Figure 2, the main interface of NeuroBridge resembles a standard chatting application.

Scenario Generator. The Scenario Generator creates a conversation scenario tailored for each user based on information they provide about themselves. Figures 3a and 3b show the interface for collecting this information, and an example scenario, respectively. The goal is to center the conversation with the AI character around a topic that is both interesting and relatable for the participant.

Message Options Generator. The Message Options Generator takes in a user message, and creates three different versions of it, which we call ‘message options’. This process is shown in Figures 4a and 4b. The message options are similar in meaning to the user’s initial message but vary in tone, clarity, or phrasing based on the given scenario (the scenarios are listed in Table 1). For instance, in the scenario involving indirect speech acts, one option may ask a question directly (“What methods...”), while others phrase the same question ambiguously (“Is there a way...”), as exemplified in Figure 4b. The user can then select and send one of the three message options. Similarly, for scenarios involving figurative expressions and emojis with variable interpretations, one message option uses literal language or a straightforward emoji, while others express the same idea figuratively. In the scenario involving misperceived bluntness, two options suggest the user found the character’s message blunt, while the third is a neutral response that shows understanding and acceptance of the character’s direct style. Interaction flows




Figure 3: The user registration screen, shown in (a), first gathers information about the user. Based on this information, NeuroBridge generates an AI character and a social scenario for the upcoming conversation, as shown in the scenario description screen in (b).

for these three scenarios are provided in Appendix A. In this way, the message options allow us to trigger different scenarios, while having the user craft the initial message allows for personalizing the simulation experience.

Response Generator. The Response Generator generates all messages sent by the character. If the user message is unclear, the character’s response is a request for clarification. This is shown in Figure 5. If the message is clear, the conversation is continued as usual. This is shown in Figure 6. In the scenario involving misperceived bluntness, if the user indicates that the character’s response was blunt, the character follows up in the subsequent message to explain that it was not meant that way.

Feedback Generator. The Feedback Generator generates scenario-specific feedback for the user. After getting the character’s response, users receive feedback through a dedicated panel in the chat interface. The feedback varies depending on the message option sent. If the message option sent is ambiguous or suggests the character’s response was rude, the user receives constructive feedback. This is shown in the gray panel in the center of Figure 5. Constructive feedback is structured such that it first highlights the difference in interpretation/intent between the user and the character, identifies the most appropriate message option, and then explains why it is most appropriate. The user is also provided with a message that they can send to continue the conversation empathically, as shown at the bottom of Figure 5. If the user sends the most appropriate message option, positive feedback is provided to the user, as shown in the gray panel at the bottom of Figure 6. Positive feedback serves as encouragement, and explains why the other message options might lead to a misunderstanding.

3.2 Development Process

3.2.1 Advisory Board. An advisory board of three autistic volunteers provided feedback on the design of NeuroBridge. All board members were enrolled as graduate and/or undergraduate students at Tufts University, and were identified through the authors’ existing network of autistic individuals who had previously participated in or expressed interest in serving on a board. Board members reviewed the prototype in three one-hour meetings held at the elementary, intermediate, and final stages of development. Each member evaluated NeuroBridge as a mock user and reviewed the AI-generated simulation, responses, and feedback, going through each simulated scenario at least twice. Each meeting was conducted by two researchers and recorded for later viewing. As detailed next, in-depth, open-ended feedback was obtained from the board and incorporated during development.

3.2.2 Iterative Development. Several improvements were made based on feedback from the advisory board. For instance, they recommended that when a user sends an unclear message, the AI character should ask a clarifying question like, ‘Do you mean X or Y?’ rather than assuming one of X or Y, reflecting how they typically process uncertainty in real life. They also emphasized the importance of sharing these interpretations in more detail with neurotypical users in the feedback so that they can understand exactly how an autistic person might interpret language differently. They verified that two out of three message options in the simulated scenarios could, in fact, lead to a misunderstanding, while the remaining one was most appropriate. Additionally, the board reviewed the AI character’s blunt responses and agreed that the tone reflected their past communication experiences, which sometimes led to negative reactions from others. Consistent with prior studies [6, 55], they emphasized the importance of encouraging neurotypical individuals to understand autistic communication styles and viewed NeuroBridge’s approach as effective.

Moreover, we also conducted five pilot studies with neurotypical users for preliminary testing and feedback. In the initial version of NeuroBridge, users had no control over the conversation topic or message composition; they were given a set topic and pre-determined message options to choose from. Based on feedback from pilot users, we added the ability for users to select a topic of interest and compose their own messages, which are then used to generate personalized scenarios, message options and feedback. Initially, we had also included filler turns, so that only every other message triggered a scenario, creating a more natural conversation flow. This nearly doubled the interaction time, so we eventually removed them. Additionally, we observed that our pilot users, including those knowledgeable about autistic communication, repeatedly re-read the message options and carefully reasoned through them, highlighting both the intended subtlety of the differences and the cognitive effort involved in making the correct selection.

3.2.3 Prompting Strategy. LLMs take in input in the form of natural language, provided through ‘prompts’. A prompt is a carefully crafted instruction that guides the model’s output. Following prior work, we iteratively refined the prompts for each task, such as generating message options, character responses and feedback [6, 65]. Through repeated testing, we optimized them for consistency and reliability. Note that we did not instruct the LLM to act autistic or generate feedback from the perspective of an autistic person. Instead, we provided carefully crafted examples of message options, interpretations, and feedback for each scenario as ‘sample outputs’ in the prompts. These examples helped ensure the LLM responded in the way we intended. Providing examples to improve output quality is an effective prompting technique and commonly known as few-shot learning [79]. We avoided referencing autism in any of the prompts to prevent perpetuating existing biases about autism in LLMs [71]. All prompts have been made available as Supplementary Material for reproducibility.

3.2.4 Implementation. The frontend of NeuroBridge was developed using React and shadcn/ui, while the backend was built with FastAPI, incorporating both REST and WebSockets to facilitate real-time chat functionality. GPT-4o (GPT-4o-2024-0513 Regional) was used for LLM generation in all tasks, except for generating user message options involving emojis with variable interpretations. We used Claude 3.5 Sonnet (us.anthropic.claude-3-5-sonnet-20240620-v1:0) for it, as it outperformed GPT-4o on this task. Both models were accessed through a deployment on Microsoft Azure. The front-end was deployed on Cloudflare Pages, and the back-end was containerized using Docker and deployed on Google Cloud Run, with data storage managed through MongoDB Atlas.

4 Methodology

In this section, we provide our positionality statement and an overview of the recruitment process, user study, and data collection and analysis methods. All study procedures were approved by the Institutional Review Board (IRB) at Tufts University.

4.1 Positionality Statement

We disclose that a majority of the authors identify as neurodivergent and have professional and/or personal relationships with neurodivergent people. Our approach is informed by the social model of disability, which frames disability as arising from societal structures rather than individual impairments. This perspective underpins our view that cross-neurotype communication breakdowns reflect differences in communication styles rather than deficits in social skills, shaping both the design and interpretation of this work.

4.2 Recruitment

Twelve neurotypical participants were recruited through flyers posted around Tufts University’s main campus in Medford, MA, USA. Interested individuals completed a screening survey to determine eligibility. The inclusion criteria were: a) aged 18 or older, b) fluency in English (reading and writing), and c) ability to perform basic computer tasks. Participants were also asked about their familiarity with autistic communication styles in the screening survey, and an equal number were selected from each familiarity group. All participants identified as non-autistic. Participant information is shown in Table 2.

4.3 User Study Overview

User study sessions were conducted in person, on-campus in a lab setting, where participants were provided with a secure personal computer and monitor. Each session lasted about ninety minutes. Participants started by reading and signing a consent form describing the purpose and procedures of the study. Then, they proceeded to enter their name, pronouns, and a topic of interest, which was used to generate a social setting for the conversation with the AI character. After receiving instructions on how to use the interface, participants sent the first two messages to familiarize themselves with the system. No scenarios or feedback were triggered during this phase, as the first two were configured as test messages. This introductory phase allowed them to ask questions and get comfortable with the interface. Participants were encouraged to think aloud about their reasoning for selecting a message option, as well as their thoughts on the AI character’s responses and on the feedback they received in the remainder of the study. Participants interacted with the character until they had completed two rounds of each of the four scenarios.

4.4 Data Collection and Analysis

Participants’ screen activity and audio were recorded during the user study. Upon completing the user study, participants took part in a semi-structured interview followed by a survey in the same session. The interviews were also audio-recorded. The 11-item Likert scale survey had statements rated on a 7-point scale from ‘Completely Disagree (1)’ to ‘Completely Agree (7)’. The survey and interview delved into the usefulness of the simulation, its impact on participants’ perceptions of autism, their attitudes toward AI feedback, and the effect of personalization on user engagement. The survey results provide an overview of self-reported user perceptions, while qualitative insights offer richer, explanatory insights about their experience using NeuroBridge.

(a) Message crafted by user.

(b) Rephrased message options based on message crafted by user.

Figure 4: The user is first prompted to input a message to send to the AI character, as shown in (a). Then, three unique variations are generated and displayed to the user, prompting them to select the best-phrased message, as shown in (b).

Given the small sample size ($N=12$), we report descriptive statistics (mean and standard deviation) for the survey results, along with verbatim survey statements and the cumulative percentage of responses indicating agreement (options ranging from 1 to 3, both inclusive) or disagreement (options ranging from 5 to 7, both inclusive) on the Likert scale in Figure 7. This approach is adapted from Goodman et al. and Adnin et al. [65, 73]. For qualitative analysis, we used Braun and Clarke’s thematic coding approach [80] with a deductive framework. Prior to the study, we developed the following set of deductive codes to categorize: perceptions of the simulation’s usefulness; trust in the AI-generated simulation; reactions to AI-generated feedback; understanding and perceptions of autistic communication styles; and suggestions for improvement. A member of the research team first transcribed the audio data and then contextualized them with observations from the screen recordings. After importing the transcripts into NVivo [81], they extracted relevant quotes by reading the transcripts line by line, grouped them into themes, discussed the themes with other team members, and reviewed and refined them. Another member of the research team, who was not part of the initial study team, independently validated the themes and the data associated with each theme. A similar approach was used by Ahsen et al. and Haroon et al. [6, 50].

5 Findings

In this section, we discuss and synthesize our findings, supported by participants’ quotations and relevant survey results.

5.1 Usefulness of the Simulation Experience

5.1.1 Helps Develop Communication Awareness. Several participants (P1, P3, P5, P6, P7, P8, P9, P11, P12) reported that taking part in the simulation helped them understand how an autistic person might interpret language differently. Many were surprised to see that these interpretations were plausible, and even obvious in hindsight, but never occurred to them. They highlighted that the AI character’s interpretation, included in the feedback, helped them understand exactly what part of their message could be received differently by an autistic person. For instance, P11 shared, “If the feedback just said ‘the figurative part in your message could cause confusion’, I might’ve thought, ‘Okay, but why?’ The example [interpretation] provided helps me understand what exactly Wendy [the AI character] is thinking when she is reading this.” Similarly, P3 reflected, “Explaining how the rocket emoji could be interpreted differently with an example [of an autistic interpretation] gave me a chance to see Jason’s [the AI Character] perspective.” Echoing these sentiments, P5 felt the feedback was useful for navigating future interactions. This was reflected in participants’ behavior as well. Upon encountering a similar scenario later in the simulation, most were able to identify the most appropriate response and referred to feedback from a previous turn to back their rationale. Overall, participants strongly agreed ($avg. = 5.83, s.d. = 1.19$) that the simulation helped them recognize the differences in communication styles of autistic and non-autistic individuals, as shown in row 1 of Figure 7.

Is there a way to speed up a vision model on an embedded device?

do you want to know if it's technically possible or are you looking for specific methods to do it?

Clarify Your Question for Julia 😊

The phrase 'Is there a way to speed up a vision model on an embedded device?' can be interpreted in two ways: Julia might think you're asking if it's technically possible, which could be answered with 'yes' or 'no', or she might think you're asking for specific methods on how to do it. Because people have different communication styles, Julia sought clarification. To avoid confusion, specify if you want a simple answer or detailed information.

As An Alternative, You Could Try:

What methods can be used to speed up a vision model on an embedded device?

The alternative message clearly asks for specific methods to speed up a vision model on an embedded device. This phrasing removes ambiguity by directly requesting detailed information, reducing the chance of misunderstanding. By specifically asking for methods, Julia understands that you are not just asking if it's technically possible but are looking for practical solutions, making the intended meaning clearer.

Send this message to clarify and continue the conversation.

sorry for being unclear. i'd like to know what specific methods you might recommend to speed up a vision model on an embedded device.

Figure 5: The message option sent by the user is shown in the blue message bubble. After, the gray message bubble shows that the AI character asks the user to clarify what they meant since an incorrect message option was sent. Then, the user receives a two-part constructive feedback (shown in the center gray panel) explaining why their choice was incorrect and why the other option was more appropriate. The user is then prompted to send the provided follow-up message to clarify and continue the conversation.

5.1.2 Closest to a Real Interaction with an Autistic Person. Multiple participants (P5, P7, P9, P11) described the simulation as the closest they had come to interacting with an autistic person. They believed this was useful, as people often hold common misconceptions about autism that are unlikely to change without interacting with an autistic person in real-life. Since the simulation closely resembled such an experience, and because having open, exploratory conversations with an autistic person is not always possible, participants believed it served as an effective alternative. P7, who had in-depth knowledge of autistic communication through lived experience with their autistic sister, expressed "An interaction like this is probably the closest you can really get to emulating the experience of interacting with someone with autism." They described the platform as a safe, low-stress environment for learning, and contrasted it with real-life

interactions, "When interacting with someone with autism... things can kind of spiral out of control very quickly." In contrast, "[With the simulation] you're sort of on some guardrails..." Reflecting on their own experience, they added, "When I was growing up, this would have helped me a lot in understanding my autistic sister." P5 echoed these sentiments, noting that the AI character's responses allowed them to see how their message might have caused confusion for an autistic person if this was a real interaction, "You get to actually see what could happen if you say something that can cause confusion... it is very realistic, and prepares you to have a conversation [with an autistic person]." Overall, participants agreed (avg. = 5.50, s.d = 1.45) that the character's responses felt natural and realistic, as shown in row 11 of Figure 7.




Figure 6: The message option selected by the user is shown in the blue message bubble. In the gray message bubble, the AI character responds as usual because the user selected the correct option. Positive feedback (shown in the bottom gray panel) is provided to reinforce the user's choice and explains why the incorrect options may have caused confusion.

5.1.3 Enables Active Learning. Several participants (P1, P2, P3, P4, P5, P6, P10, P12) appreciated the interactive nature of the platform, highlighting that it allowed them to apply what they were learning hands on. Having grown up with an autistic father, P6 had in-depth experience/knowledge of autistic communication styles. Yet, they expressed simply knowing wasn't the same as applying that knowledge. The feedback they received on one of their responses revealed perspectives they hadn't considered, "*I'm pretty knowledgeable on how autistic people communicate, but I didn't even think about how the chicken emoji could be interpreted like that [as described in the feedback]. After looking at the feedback, I was like, oh... yeah, you're right.*" Similarly, P3 pointed out that conversing with the AI character exposed gaps in their knowledge, "*It's not until you actually try to have a conversation that you can really see what they might not understand in what you say.*" In addition, P1 noted that the process of crafting the response helped reinforce what they were learning, "*I had to actually think about what the response [by the character] was, and how to best word my response to continue the conversation.*" P3, P5, P6 and P10 echoed these sentiments. For P6, this was especially useful in moments of friction, "*They [the character] said something in not the nicest tone... and I had to think through the response.*" As shown in row 2 of Figure 7, participants generally agreed (avg. = 5.58, s.d. = 1.83) that the simulation would influence their future communication with autistic individuals.

5.1.4 Personalized Feedback. Several participants (P3, P5, P6, P9, P10, P11, P12) emphasized what made the simulation especially useful was the personalized nature of the feedback. Rather than presenting abstract or generic examples, the system provided feedback on messages that they had sent and were based on their original input. P11 explained, "*What really helps is having the confusing parts*

of your own speech specifically pointed out." In addition, participants found value in seeing how the message they had come up with could be easily rephrased to cause or prevent a communication breakdown. In one instance, after reading the message options, P11 exclaimed, "*A lot of the time when I was writing up my response, I didn't even consider other ways to say the same thing. It is interesting to see what those options were and think about which of those made the most sense.*" P12 echoed this sentiment, highlighting moments when the tool improved upon what they had tried to say, "*I was okay with the way I worded myself, but it wasn't perfect and then it would give me a better option that accomplished what I wanted to say in a very autistic-friendly way.*" In this way, personalized feedback encouraged participants to reflect on their own communication style and assumptions.

5.1.5 Engaging and Immersive. Several participants (P1, P4, P7, P8, P9, P10, P11, P12) described the simulation as more engaging than other, common ways of gaining awareness, such as awareness blogs or videos. P9, who was already interested in chatbots, appreciated how the experience "*replicates that feeling of talking to a real person*", adding that it was "*more engaging to have what feels like an actual conversation*" rather than passively consuming information. P11 echoed this sentiment, noting that reading felt "*a lot more educational*", whereas with the simulation, "*you learn on the way.*" P4 described traditional formats as "*passive*", and P8 shared, "*You read it, and then put it down and put it away, whereas this is a more memorable experience.*" Overall, as shown in row 3 of Figure 7, participants strongly preferred (avg. = 5.92, s.d. = 1.51) NeuroBridge over awareness blogs and videos. P12 further noted that reading or watching content can sometimes create a false sense of confidence, "*It actually kind of harms you because you think, 'Oh, I know what*

to avoid. I know what I need to do,' and you don't realize that just knowing about it doesn't mean you actually know how to apply it." They highlighted that in contrast, the simulation allows you to practice and test your understanding, making it easier to see what you truly grasp and where you might need to improve. For P1 and P7, the process of crafting their own responses kept them immersed throughout the simulation. As shown in rows 4 (*negatively-worded*, $avg. = 1.50, s.d. = 0.90$), 5 ($avg. = 6.92, s.d. = 0.29$), and 7 ($avg. = 6.58, s.d. = 0.67$) of Figure 7, users strongly valued the simulation's personalized and interactive format. In addition, participants generally did not feel the simulation was too time-consuming, as shown in row 6 (*negatively-worded*, $avg. = 2.33, s.d. = 1.23$) of Figure 7.

5.2 Feelings of Trust and Skepticism in AI

5.2.1 Factors Affecting Trust in AI. Several participants (P5, P6, P7, P9, P10) initially approached the simulation with skepticism because they were told it was powered by AI, but came to view it as trustworthy as they found their interaction with the AI character realistic, and the explanations provided by the system relatable and logically coherent. For instance, P5 described the AI character as "almost creepily realistic", and that it "easily could have been a real person." They found the reasoning provided in the feedback convincing, particularly because it systematically explained how their message could be interpreted in different ways. The feedback was structured to first recognize the sender's likely intent, then illustrate how and why the message might be received differently by the AI character, and finally suggest a better alternative along with a justification. This helped participants connect the dots between what they meant to say and how it could be misread. P5 reflected "It is the way it is explaining the phrases and the things I said... following a logical train of thought in its response." P6 shared a similar view, stating that they found the feedback provided to them during the simulation trustworthy because it also aligned with their own reasoning.

Participants (P7, P8, P9, P11) also highlighted that their trust in the system was shaped by their personal background, such as their prior exposure to autistic individuals and technology. For example, P11, who had limited experience interacting with autistic people, shared that they trusted and were open to receiving constructive feedback from the system because they did not consider themselves knowledgeable on the topic. P8 noted that while they personally trusted the chatbot, their grandmother would likely be much more skeptical of it. In their view, prior experience with technology played a key role in whether someone would take the simulation seriously, "My 75-year-old grandma would probably be very skeptical of it... whereas if she was talking to a professor, or someone with autism, she would believe them without hesitation." Similarly, P7, an engineering student with an autistic sibling, admitted to having an "intrinsic" bias against AI. They viewed it as a tool often misapplied to scientific problems beyond its limits, but found value in the simulation, "You kind of need to suspend disbelief. I know I'm talking to a machine, but it emulates it closely enough that I can get something out of it." Their personal connection to autism allowed them to look past their skepticism of AI.

5.2.2 Reactions to AI Feedback. At most occasions, participants described feeling curious, open, and motivated when they received

AI-generated feedback. For instance, P12 reflected, "It makes me curious, like, how can I, going back into real life, interacting with actual autistic people, tailor my language to make sure I'm communicating with them effectively?" In particular, participants appreciated that communication differences were framed constructively in the feedback, without labeling their response as "wrong". P3 echoed this sentiment and noted that even when they did not perform well, the system recognized that they were trying, "Even when I say something wrong, it isn't like, 'You're wrong.' Even the titles are 'Thoughtful Communication' and 'A Small Tweak to Make Your Message Clear'. They're very much acknowledging that you are trying." The use of emojis and a supportive tone contributed to making the feedback feel friendly and supportive, helping participants stay open and receptive. As P2 expressed, "I like the little star emoji [in the feedback]. It adds a nice little bit of flair and makes it feel like a little more celebratory." This sentiment was echoed by P5 and P9.

In addition, participants found it useful to receive feedback not only when they failed to identify the most appropriate message option, but also when they succeeded. P12 elaborated, "I definitely think you should continue to provide feedback when things are going well. I get so frustrated when I only get feedback for doing something wrong. I want to know what I did well so I can keep doing it in the future. I want to know exactly what part of my behavior was good, not just 'your behavior is good, keep doing it', because otherwise, I'm not really sure what to continue." Participants noted that positive feedback was not only encouraging, but also helpful for learning. This was particularly important for users unfamiliar with autism. While they might have selected the correct answer, they could have done so without fully understanding why. The feedback helped validate their reasoning and fill in any gaps in understanding. P4 reflected, "If I don't know much about autistic communication, I might pick the right option for the wrong reason. So it's helpful to hear, 'Yes, this is right and here's why.'"

However, on a few occasions, participants expressed feeling defensive, describing the feedback as instructive and diminishing their sense of agency. For example, P10 remarked, "The phrasing of the feedback should come off a bit more neutral. Some lines come off as almost an attack on how you talk, especially when some people... may go into this with no prior experience interacting with someone with autism." P7, who had lived experience supporting an autistic family member, expressed that frustration and defensiveness are natural in cross-neurotype communication. They emphasized the need for the feedback to not only offer constructive suggestions, but also to validate these emotions, "For me, a big part of it is validating those feelings... You should insert something like, 'It is okay to feel frustrated sometimes, you are human too'... and then go into, 'Here's how you can be better and kind of meet them halfway.'" Similarly, P12 described one instance where they felt sidelined in the interaction, "It [the feedback] really frustrates me because I feel like it puts too much focus on Autumn [the AI character] and takes agency away from me... It feels like you're just playing to Autumn's whims." Notably, all of these reactions were observed during the scenario around misperceived bluntness.

5.2.3 Cannot Substitute Real Interactions. Participants (P7, P8, P11) emphasized that while the AI-driven simulation was useful, it was still important to hear directly from autistic individuals, rather




Figure 7: Survey results with verbatim statements and statistics. The percentage on the left represents the number of participants who selected values between 1 and 3 (both inclusive), while the percentage on the right represents the number of participants who selected values between 5 and 7 (both inclusive). Responses of 4 (middle) are excluded from both percentages.

P#	Age	Gender	Knowledge of Autistic Communication
P1	18–24	Female	I have no prior knowledge
P2	18–24	Female	I have heard of it but don't know much
P3	18–24	Female	I have heard of it but don't know much
P4	18–24	Female	I have a very basic understanding.
P5	18–24	Male	I have heard of it but don't know much
P6	18–24	Female	I have in-depth knowledge and/or experience
P7	25–34	Male	I have in-depth knowledge and/or experience
P8	18–24	Female	I have a very basic understanding.
P9	18–24	Female	I have in-depth knowledge and/or experience
P10	18–24	Male	I have no prior knowledge
P11	18–24	Non-binary/third gender	I have no prior knowledge
P12	18–24	Male	I have a very basic understanding.

Table 2: Participant demographics and familiarity with autistic communication styles.

than relying solely on an AI to represent them. P8, who had limited personal experience with autistic people, felt that the chatbot helped illustrate key communication pitfalls and did a good job of showing how seemingly clear messages could be received differently, but ultimately concluded, “*as a whole, having an experience with a person is a better way for getting to know them.*” P7 echoed this sentiment, framing the tool as part of a larger learning journey, “*If you wanted to create a package of how to interact with autistic people one-on-one, this would be an element of that, but it wouldn't be the whole thing.*” They appreciated the simulation’s ability to model scenarios and spark reflection, but felt it could only approximate the complex dynamics involved in a real conversation.

5.3 Concerns and Improvements

5.3.1 Perceptions of Autism. We were particularly interested in how the simulation shaped participants’ perceptions of autistic abilities. Survey results show that participants strongly agreed ($avg. = 6.42, s.d. = 0.90$) that “autism can be viewed as a social difference that needs understanding by others” after the simulation, as shown in row 9 of Figure 7. However, participants also expressed agreement ($avg. = 5.33, s.d. = 1.83$) with the statement, “social skills training, including understanding nuanced language, can be beneficial for autistic individuals”, as shown in row 8 of Figure 7. Qualitative results help contextualize this; some of our participants came away with reinforced stereotypes about autism. For example, P10 remarked that the AI character’s responses made them feel its text comprehension abilities as “*a bit below average*”, especially when it took common metaphors too literally. Similarly, P9 expressed concern that some users might interpret this behavior as a sign of cognitive inferiority, and stressed these literal interpretations need to be framed as a difference (as opposed to a deficiency) more concretely in the feedback. Similarly, P2 and P6 wondered whether the AI was underestimating autistic people’s abilities related to symbolic understanding, as they felt emojis like a thumbs-up or fire icon didn’t seem inherently complex, yet were treated as such by the AI character. In contrast, P9 agreed that while these emojis could be confusing depending on the context in which they are

used, they acknowledged the risk that users unfamiliar with autism might misread these incidents as evidence of limited ability.

During our meetings with the advisory board, we had reviewed several examples that neurotypical participants found to be too simple to be misunderstood, such as one involving a basic emoji. Members of the board pointed out that things that appear simple on the surface can be confusing depending on the context in which they are used. This reveals how neurotypical individuals may struggle to recognize that expressions they consider straightforward can be confusing for autistic individuals. Nonetheless, P9 made an interesting observation; although the AI character was configured to be literal, it did end up using metaphors once or twice. P9 felt this challenged the assumption that autistic individuals have below-average language skills or cannot understand/use figurative language, because the character was shown using it a few times. In P9’s view, this prevented a stereotypical portrayal of autism from being reinforced, while still highlighting that figurative language may not always be the preferred option.

5.3.2 Generating Message Options. One of the LLM’s core tasks was to generate alternative versions of the user’s message that were semantically identical but phrased differently, depending on the given scenario. However, several participants raised concerns about the quality of the message options, particularly in the scenario related to emojis with variable interpretations. Participants (P1, P5, P6, P7, P8, P12) found that the emojis added by the LLM often felt random or disconnected from the content of the message. For example, P12 described the use of crystal ball and alien emojis as “*super, super, weird*”, stating they wouldn’t have understood the purpose of adding them without reading the explanations in the feedback. Similarly, P1 stated that some of the emojis “*felt out of place*” and would confuse neurotypical individuals as much as autistic people. P7 expressed frustration at being “*forced into a series of bad options*”, highlighting a mismatch between the emoji’s tone and the content of the messages. Participants acknowledged that eventually the feedback helped clarify why those emojis were added, but the feedback was revealed to them only after they had sent the message. P8 wondered whether such abstract associations would be ever be apparent to anyone without the feedback. Overall,

nearly 40% of participants expressed some degree of confusion (*negatively-worded*, $avg. = 3.42$, $s.d. = 2.19$) during the simulation, as shown in row 10 of Figure 7.

5.3.3 Modeling the Blunt Scenario. Another key task for the LLM was to craft a blunt message on behalf of the character that would serve as a turning point in the conversation. This message was intended to simulate a situation in which the character might be perceived as blunt by the participant, triggering a harsh or confrontational response from them. However, several participants (P1, P2, P8, P11, P12) stated that these trigger messages did not always come off as blunt. Participants described the tone of these messages as “neutral”, “factual” or “reasonable” depending on the context. P12, for example, stated, “*They do not seem to me to be blunt... it’s a simple statement. They’re not elaborating, but they’re also completely answering my question.*” P1 similarly downplayed any negative tone, saying, “*I wouldn’t have thought that he [the AI character] was being blunt, or, you know... rude in any way.*” P2 added that such directness felt familiar and unremarkable, “*I’m used to hearing people say things like that... it seems neutral. It seems factual.*” As a result, some participants were confused about why they were presented with confrontational message options. P8, for instance, felt that message options like “What’s with the attitude?” did not align with their interpretation of the AI character’s message. “*Those surprised me as being options,*” they explained, “*because I didn’t interpret that [the trigger message] at all as giving attitude or being dismissive in any way.*” P11 described a moment where the AI character seemed to contradict itself by first saying, “*Do you want to hear about my experiences? I think they’re interesting,*” and then suddenly following with, “*They’re not interesting. Why do you want to know?*” The inconsistency left P11 confused, “*It’s like almost contradicting the text they just sent.*” In this instance, the LLM struggled to maintain conversational flow and logical coherence

6 Discussion

In this section, we reflect on our findings and discuss design implications for representing disabilities through AI, the need for making NeuroBridge more personalized, and LLMs’ limitations in modeling complex social scenarios.

6.1 Representing Disabilities through AI

A key challenge in creating accurate and complete representations of disability through AI lies in capturing the diversity of disabled people’s lived experiences [61]. This is especially true for autism, which spans a broad spectrum characterized by nuanced, and often subtle, differences. While the scenarios in NeuroBridge reflect common challenges experienced by autistic individuals with a direct and literal communication style, not everyone with this style finds them difficult to navigate. Moreover, other scenarios, such as those involving sarcasm or sexual innuendos, could be incorporated too [33]. Our participants were observant of this limitation, and suggested incorporating multiple AI characters to represent a broader range of scenarios and communication styles – echoing prior work that highlights how single-perspective disability representations can unintentionally reinforce stereotypes [61] and misconceptions [62]. Additionally, participants recommended adding in-situ ‘citations’ with each scenario, such as links to Reddit threads or first-hand

accounts from autistic individuals. This would not only enhance the credibility, transparency, and grounding of the AI-generated simulation, but also expose users to everyday experiences beyond those represented in the simulation. Understanding this context can help neurotypical individuals better understand how disabled individuals truly feel and identify with their disability [61, 63]. While gaps remain and it may be difficult to capture every nuance, LLMs’ ability to simulate diverse communication styles is a meaningful step toward more accurate disability representations in AI.

6.2 The Need for Situational Context

Our findings highlight the importance of LLM-powered interactivity, personalization, and realism in sustaining user engagement and active learning [82]. Informed by these insights, we propose incorporating ‘situational context’ into the simulation by situating communication scenarios within specific social roles or relationships [83], such as student-teacher or doctor-patient dyads. Note that appropriateness varies based on social dynamics; bluntness may be acceptable among friends but is generally less so between a teacher and a student. Incorporating situational context helps capture these nuances more accurately, while also raising the question of how individuals in authority roles, such as teachers or doctors, respond to feedback provided to them during the simulation. Our findings suggest factors like familiarity with autism and/or AI can affect user experience and attitude. Hence, it will be useful to examine how authority influences openness to critical feedback. Moreover, incorporating situational context could enhance transfer of knowledge and awareness from the simulation to real-world interactions. For example, simulating a disagreement between a student (role-played by AI) and teacher (role-played by a human), could be particularly beneficial for a real-life teacher, as they may face a similar situation at work [84]. Prior work shows that autistic individuals often use AI tools in hierarchical settings, where the risks and consequences of miscommunication are amplified [31]. Training neurotypical users in these scenarios will also help them recognize the higher stakes involved and the need for thoughtful, considerate communication.

6.3 The Fine Line in Trusting AI

Participants readily placed their trust in the AI-generated simulation and feedback, despite initially approaching it with skepticism. This was particularly observed among individuals with limited prior knowledge of autism. Given that LLMs have been shown to perpetuate biases against disabled individuals, including those on the autism spectrum [71, 72], it is crucial for users to calibrate [85] the amount of trust they place in AI-generated representations of disabilities. Some of our participants suggested the simulation should be paired with preparatory materials, such as a primer on autism, so that they feel more confident going into it and can view it from a critical lens. In addition, future iterations of NeuroBridge could consider incorporating features to facilitate structured and systematic reflection. These could include online discussion or chat features for engaging with other users or autistic/expert moderators. While LLM biases related to autism detection and demographics have been explored [71], exploring how LLMs simulate autistic communication styles with minimal prompting (as opposed to our

approach, which involved extensive instruction and no explicit mention of autism) warrants further investigation and could uncover additional biases.

6.4 Challenges of AI-driven Simulations

The LLM performed well for most tasks, but when failures occurred, they were often due to dependencies across tasks even though task decomposition has been shown to improve LLM performance [86]. For example, if the AI-generated message options (the first task) did not reflect the nuances of the scenario to be simulated very well, the LLM struggled to later provide a convincing explanation (the second task) for why those options might be perceived as confusing. This resulted in a trickle-down effect, with issues in the early stages undermining performance in later stages. For us, this posed a challenge as multiple components of NeuroBridge rely on each other to coherently scaffold the simulation. Interestingly, since much of the simulation’s content was open to interpretation, users often formed their own conclusions and were somewhat open to the AI’s different or even incorrect interpretations, thinking they might be valid as well. This observation aligns with prior work suggesting users may overly ascribe intent to AI, a phenomenon known as ‘algorithmic anthropomorphism’ [87]. In some cases, participants’ perceptions of autism were negatively influenced by their perceptions of the LLM’s capabilities [62]. For example, a few participants speculated that the AI had malfunctioned when they encountered a scenario they felt was too simple to be misunderstood by anyone. In this way, how users perceive AI may directly impact how they view the identities it represents.

6.5 Limitations

There are a number of limitations of our study. First, recruiting participants from a university setting limits the generalizability of our findings, as individuals from diverse age groups, backgrounds, and education levels may be less open to change, and as a result, react differently to feedback provided to them in the simulation. Hence, while our analysis reveals recurring themes, a broader demographic could uncover additional themes. Moreover, the study relied primarily on self-reported data, which may introduce bias as participants may not fully disclose their opinions. Incorporating more objective pre-post measures could offer deeper insights into the intervention’s effectiveness. Future research should also examine its long-term impact by exploring how it influences neurotypical individuals’ behavior in real-world interactions, particularly among family and friends of autistic people. Finally, we were only able to incorporate a limited set of communication scenarios, and a more comprehensive implementation would include a wider range.

7 Conclusion

In this paper, we argue for redressing the disproportionate burden autistic people bear in adapting to neurotypical communication norms. To this end, we present the design and evaluation of NeuroBridge, an LLM-powered tool that helps neurotypical individuals better understand autistic forms of expression, and reflect on their role in shaping cross-neurotype interactions. In a user study with 12 neurotypical participants, we find that NeuroBridge improved their understanding of how autistic people may interpret language

differently, with all describing autism as a social difference that “needs understanding by others” after completing the simulation. Participants valued the simulation’s personalized, interactive format, and even those familiar with autism through lived experience reported gaining new knowledge about autistic communication. While many found AI-generated feedback to be “constructive”, “logical”, and “non-judgmental”, participants also reported feeling defensive on a few occasions. Most perceived the portrayal of autism in the simulation as accurate, indicating that users may readily accept AI-generated (mis)representations of disabilities. Despite strong overall performance, our results show that LLMs are more adept at simulating certain social scenarios than others. We conclude by discussing design implications for disability representation in AI, the need for making NeuroBridge more personalized, and the limitations of LLMs in modeling complex social interactions.

8 Acknowledgments

We thank Ayesha Naeem, our advisory board, all members of the NAT lab and Crehan lab, and the anonymous ASSETS reviewers for their invaluable feedback that helped improve this work. This work was partially supported by a Tufts Springboard award.

References

- [1] American Psychiatric Association. *Diagnostic and statistical manual of mental disorders: DSM-5*. Autor, Washington, DC, 5th ed. edition, 2013.
- [2] Holly Hodges, Casey Fealko, and Neelkamal Soares. Autism spectrum disorder: definition, epidemiology, causes, and clinical evaluation. *Translational Pediatrics*, 9(S1):S55–S65, February 2020.
- [3] Enzo Grossi, Elisa Caminada, Michela Goffredo, Beatrice Vescovo, Tristana Castignano, Daniele Piscitelli, Giulio Valagussa, Marco Franceschini, and Franco Vanzulli. Patterns of restricted and repetitive behaviors in autism spectrum disorders: A cross-sectional video recording study. preliminary report. *Brain Sciences*, 11(6):678, May 2021.
- [4] CDCMMWR. Prevalence and Characteristics of Autism Spectrum Disorder ... — cdc.gov. https://www.cdc.gov/mmwr/volumes/72/ss/ss7202a1.htm?s_cid=ss7202a1_w. [Accessed 26-03-2025].
- [5] Belén Barros Pena, Nelya Koteyko, Martine Van Driel, Andrea Delgado, and John Vines. “my perfect platform would be telepathy” - reimagining the design of social media with autistic adults. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI ’23, New York, NY, USA, 2023. Association for Computing Machinery.
- [6] Rukhshan Haroon and Fahad Dogar. Twips: A large language model powered texting application to simplify conversational nuances for autistic users. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS ’24, New York, NY, USA, 2024. Association for Computing Machinery.
- [7] Agustín Vicente, Christian Michel, and Valentina Petrolini. Literalism in autistic people: a predictive processing proposal. *Review of Philosophy and Psychology*, 15(4):1133–1156, September 2023.
- [8] Tamar Kalandadze, Courtenay Norbury, Terje Nærland, and Kari-Anne B Næss. Figurative language comprehension in individuals with autism spectrum disorder: A meta-analytic review. *Autism*, 22(2):99–117, November 2016.
- [9] Joy Hirsch, Xian Zhang, J. Adam Noah, Swethasri Dravida, Adam Naples, Mark Tiede, Julie M. Wolf, and James C. McPartland. Neural correlates of eye contact and social function in autism spectrum disorder. *PLoS ONE*, 17(11):e0265798, November 2022.
- [10] Sarah Griffiths, Christopher Jarrold, Ian S. Penton-Voak, Andy T. Woods, Andy L. Skinner, and Marcus R. Munafo. Impaired recognition of basic emotions from facial expressions in young people with autism spectrum disorder: Assessing the importance of expression intensity. *Journal of Autism and Developmental Disorders*, 49(7):2768–2778, March 2017.
- [11] Vincent Miller. New media, networking and phatic culture. *Convergence the International Journal of Research Into New Media Technologies*, 14(4):387–400, October 2008.
- [12] Laura Sagliano, Marta Ponari, Massimiliano Conson, and Luigi Trojano. Editorial: The interpersonal effects of emotions: The influence of facial expressions on social interactions. *Frontiers in Psychology*, 13, November 2022.

- [13] Sharice Clough and Melissa C. Duff. The role of gesture in communication and cognition: Implications for understanding and treating neurogenic communication disorders. *Frontiers in Human Neuroscience*, 14, August 2020.
- [14] Noah J. Sasson, Daniel J. Faso, Jack Nugent, Sarah Lovell, Daniel P. Kennedy, and Ruth B. Grossman. Neurotypical peers are less willing to interact with those with autism based on thin slice judgments. *Scientific Reports*, 7(1), February 2017.
- [15] Erinn H Finke, Jillian H McCarthy, and Natalie A Sarver. Self-perception of friendship style: Young adults with and without autism spectrum disorder. *Autism and Developmental Language Impairments*, 4, January 2019.
- [16] Christina H. Rebholz. Life in the uncanny valley: Workplace issues for knowledge workers on the autism spectrum.
- [17] Dokumen.pub. Asperger's syndrome workplace survival guide: A neurotypical's secrets for success. <https://dokumen.pub/aspergers-syndrome-workplace-survival-guide-a-neurotypicals-secrets-for-success-9781849059435-9780857008077.html>, n.d. Accessed April 11, 2025.
- [18] Christina Nicolaidis, Dora M Raymaker, Elesia Ashkenazy, Katherine E McDonald, Sebastian Dern, Amelia Ev Baggs, Steven K Kapp, Michael Weiner, and W Cody Boisclair. "respect the way i need to communicate with you": Healthcare experiences of adults on the autism spectrum. *Autism*, 19(7):824–831, April 2015.
- [19] Mary Doherty, Stuart Neilson, Jane O'Sullivan, Laura Carravallah, Mona Johnson, Walter Cullen, and Sebastian C K Shaw. Barriers to healthcare and self-reported adverse outcomes for autistic adults: a cross-sectional study. *BMJ Open*, 12(2):e056904, February 2022.
- [20] Kamilla Klefbeck. Educational approaches to improve communication skills of learners with autism spectrum disorder and comorbid intellectual disability: An integrative systematic review. *Scandinavian Journal of Educational Research*, 67(1):51–68, October 2021.
- [21] Ya-Chih Chang and Jill Locke. A systematic review of peer-mediated interventions for children with autism spectrum disorder. *Research in Autism Spectrum Disorders*, 27:1–10, March 2016.
- [22] Enid Gordon Wolf and Bertram A. Ruttenberg. Communication therapy for the autistic child. *Journal of Speech and Hearing Disorders*, 32(4):331–335, November 1967.
- [23] Damian Milton, Emine Gurbuz, and Beatriz López. The 'double empathy problem': Ten years on. *Autism*, 26(8):1901–1903, October 2022.
- [24] Autistic Self Advocacy Network. Awareness is not enough. <https://autisticadvocacy.org/2012/04/awareness-is-not-enough/>, April 2012. Accessed April 11, 2025.
- [25] Mark Kendall. Autism awareness still has a long way to go | opinion. *Newsweek*, March 2025.
- [26] OpenAI. *GPT-4 Technical Report*. March 2023.
- [27] Character.ai. Personalized ai for every moment of your day. <https://character.ai/>. Accessed April 12, 2025.
- [28] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST '23, New York, NY, USA, 2023. Association for Computing Machinery.
- [29] Murray Shanahan, Kyle McDonnell, and Laria Reynolds. Role-play with large language models, May 2023.
- [30] Yancheng Cao, Yangyang He, Yonglin Chen, Menghan Chen, Shanhe You, Yulin Qiu, Min Liu, Chuan Luo, Chen Zheng, Xin Tong, Jing Liang, and Jiangtao Gong. Designing llm-simulated immersive spaces to enhance autistic children's social affordances understanding in traffic settings. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*, IUI '25, page 519–537, New York, NY, USA, 2025. Association for Computing Machinery.
- [31] JiWoong Jang, Sanika Moharana, Patrick Carrington, and Andrew Begel. "it's the only thing i can trust": Envisioning large language model use by autistic workers for communication assistance. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA, 2024. Association for Computing Machinery.
- [32] Brian Reichow, Shawn Salamack, Rhea Paul, Fred R. Volkmar, and Ami Klin. Pragmatic assessment in autism spectrum disorders. *Communication Disorders Quarterly*, 29(3):169–176, May 2008.
- [33] Stella Lampri, Eleni Peristeri, Theodoros Marinis, and Maria Andreou. Figurative language processing in autism spectrum disorders: A review. *Autism Research*, 17(4):674–689, December 2023.
- [34] Lesley Kimber, Diarmuid Verrier, and Stephen Connolly. Autistic people's experience of empathy and the autistic empathy deficit narrative. *Autism in Adulthood*, 6(3):321–330, September 2023.
- [35] Kathryn A. Szechy, Pamela D. Turk, and Lisa A. O'Donnell. Autism and employment challenges: The double empathy problem and perceptions of an autistic employee in the workplace. *Autism in Adulthood*, 6(2):205–217, August 2023.
- [36] David B. Nicholas, Darren Hedley, Jena K. Randolph, Dora M. Raymaker, Scott M. Robertson, and Jonathan Vincent. An expert discussion on employment in autism. *Autism in Adulthood*, 1(3):162–169, September 2019.
- [37] Beihua Zhang, Shan Liang, Jingze Chen, Lin Chen, Weimin Chen, Shunshun Tu, Linyan Hu, Huimin Jin, and Lixi Chu. Effectiveness of peer-mediated intervention on social skills for children with autism spectrum disorder: a randomized controlled trial. *Translational Pediatrics*, 11(5):663–675, May 2022.
- [38] Jairo Rodríguez-Medina, Luis J. Martín-Antón, Miguel A. Carbonero, and Anastasio Ovejero. Peer-mediated intervention for the development of social interaction skills in high-functioning autism spectrum disorder: a pilot study. *Frontiers in Psychology*, 7, December 2016.
- [39] Theresa Eckes, Ulrike Buhlmann, Heinz-Dieter Holling, and Anne Möllmann. Comprehensive aba-based interventions in the treatment of children with autism spectrum disorder – a meta-analysis. *BMC Psychiatry*, 23(1), March 2023.
- [40] Kristen R. Choi, Bhumi Bhakta, Elizabeth A. Knight, Tracy A. Becerra-Culqui, Teri L. Gahre, Bonnie Zima, and Karen J. Coleman. Patient outcomes after applied behavior analysis for autism spectrum disorder. *Journal of Developmental and Behavioral Pediatrics*, 43(1):9–16, August 2021.
- [41] Peter Washington, Catalin Voss, Nick Haber, Serena Tanaka, Jena Daniels, Carl Feinstein, Terry Winograd, and Dennis Wall. A wearable social interaction aid for children with autism. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '16, page 2348–2354, New York, NY, USA, 2016. Association for Computing Machinery.
- [42] Carlos Duarte, Luis Carriço, David Costa, Daniel Costa, André Falcão, and Luís Tavares. Welcoming gesture recognition into autism therapy. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '14, page 1267–1272, New York, NY, USA, 2014. Association for Computing Machinery.
- [43] Minkyong Jeong, YoungTae Kim, Dongsun Yim, SeokJeong Yeon, Seokwoo Song, and John Kim. Lexical representation of emotions for high functioning autism(hfa) via emotional story intervention using smart media. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*, CHI EA '15, page 1983–1988, New York, NY, USA, 2015. Association for Computing Machinery.
- [44] Peter Washington, Catalin Voss, Aaron Kline, Nick Haber, Jena Daniels, Azar Fazel, Titas De, Carl Feinstein, Terry Winograd, and Dennis Wall. Superpowerglass: A wearable aid for the at-home therapy of children with autism. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 1(3), September 2017.
- [45] Andrea Tartaro. Authorable virtual peers for children with autism. In *CHI '07 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '07, page 1677–1680, New York, NY, USA, 2007. Association for Computing Machinery.
- [46] Jungin Park, Gahyeon Bae, Jueon Park, Seo Young Park, Yeon Soo Kim, and Sangsu Lee. Aedle: Designing drama therapy interface for improving pragmatic language skills of children with autism spectrum disorder using ar. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI EA '23, New York, NY, USA, 2023. Association for Computing Machinery.
- [47] LouAnne E. Boyd, Saumya Gupta, Sagar B. Vikmani, Carlos M. Gutierrez, Junxiang Yang, Erik Linstead, and Gillian R. Hayes. vrsocial: Toward immersive therapeutic vr systems for children with autism. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, page 1–12, New York, NY, USA, 2018. Association for Computing Machinery.
- [48] Kathryn E. Ringland, Christine T. Wolf, LouAnne Boyd, Jamie K. Brown, Andrew Palermo, Kimberley Lakes, and Gillian R. Hayes. Dancecraft: A whole-body interactive system for children with autism. In *Proceedings of the 21st International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '19, page 572–574, New York, NY, USA, 2019. Association for Computing Machinery.
- [49] Laura Tarantino, Margherita Attanasio, Tania Di Mascio, Giovanni De Gasperis, Marco Valenti, and Monica Mazza. On the evaluation of engagement in immersive applications when users are on the autism spectrum. *Sensors*, 23(4):2192, feb 2023.
- [50] Tooba Ahsen, Christina Yu, Amanda O'Brien, Ralf W Schlosser, Howard C. Shane, Dylan Oesch-Emmel, Eileen T. Crehan, and Fahad Dogar. Designing a customizable picture-based augmented reality application for therapists and educational professionals working in autistic contexts. In *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '22, New York, NY, USA, 2022. Association for Computing Machinery.
- [51] Simon Brisenden. Independent living and the medical model of disability. *Disability Handicap and Society*, 1(2):173–178, January 1986.
- [52] Deborah Marks. Models of disability. *Disability and Rehabilitation*, 19(3):85–91, January 1997.
- [53] Mike Oliver. The social model of disability: thirty years on. *Disability and Society*, 28(7):1024–1026, July 2013.
- [54] Tom Shakespeare. The social model of disability. *The Social Model of Disability*.
- [55] Damian Milton, Emine Gurbuz, and Beatriz López. The 'double empathy problem': Ten years on. *Autism*, 26(8):1901–1903, October 2022.
- [56] R. Edey, J. Cook, R. Brewer, M. H. Johnson, G. Bird, and C. Press. Interaction takes two: Typical adults exhibit mind-blindness towards those with autism spectrum disorder. *Journal of Abnormal Psychology*, 125(7):879–885, 2016.
- [57] L. Kimber, D. Verrier, and S. Connolly. Autistic people's experience of empathy and the autistic empathy deficit narrative. *Autism in Adulthood*, 2023.
- [58] IMEX Lab. Autism tmi virtual reality experience. <https://imex.psu.edu/project/autism-tmi-virtual-reality-experience/>, n.d. Accessed: 2025-04-16.

- [59] Autismity. About the simulator - autismity, January 2025.
- [60] Training2Care. Autism reality experience (are) - training 2 care uk ltd.
- [61] Garreth W. Tigwell. Nuanced perspectives toward disability simulations from digital designers, blind, low vision, and color blind people. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, New York, NY, USA, 2021. Association for Computing Machinery.
- [62] Gary Kiger. Disability simulations: Logical, methodological and ethical issues. *Disability, Handicap and Society*, 7(1):71–78, 1992.
- [63] K. W. Barney. Disability simulations: Using the social model of disability to update an experiential educational practice. *SCHOLE: A Journal of Leisure Studies and Recreation Education*, 27(1):1–11, 2012.
- [64] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothee Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. LLaMA: Open and Efficient Foundation Language Models. arXiv preprint arXiv:2302.13971, 2023.
- [65] Steven M. Goodman, Erin Buehler, Patrick Clary, Andy Coenen, Aaron Donsbach, Tiffanie N. Horne, Michal Lahav, Robert MacDonald, Rain Breaw Michaels, Ajit Narayanan, Mahima Pushkarna, Joel Riley, Alex Santana, Lei Shi, Rachel Sweeney, Phil Weaver, Ann Yuan, and Meredith Ringel Morris. Lampost: Design and evaluation of an ai-assisted email writing prototype for adults with dyslexia. In *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '22, New York, NY, USA, 2022. Association for Computing Machinery.
- [66] Oliver Alonzo, Sooyeon Lee, Akhter Al Amin, Mounica Maddela, Wei Xu, and Matt Huenerfauth. Design and evaluation of an automatic text simplification prototype with deaf and hard-of-hearing readers. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '24, New York, NY, USA, 2024. Association for Computing Machinery.
- [67] HyungWoo Song, Minjeong Shin, Hyehyun Chu, Jin Hong, Jaechan Lee, Jinsu Eun, and Hajin Lim. Emobridge: Bridging the communication gap between students with disabilities and peer note-takers utilizing emojis and real-time sharing. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '24, New York, NY, USA, 2024. Association for Computing Machinery.
- [68] Ziming Li, Pinaki Prasanna Babar, Mike Barry, and Roshan I. Peiris. Exploring the use of large language model-driven chatbots in virtual reality to train autistic individuals in job communication skills. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI EA '24, New York, NY, USA, 2024. Association for Computing Machinery.
- [69] Stephanie Valencia, Richard Cave, Krystal Kallarackal, Katie Seaver, Michael Terry, and Shaun K. Kane. "the less i type, the better": How ai language models can enhance or impede communication for aac users. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, New York, NY, USA, 2023. Association for Computing Machinery.
- [70] LouAnne Boyd. Conceptualizing celebratory technologies for neurodiversity to reduce social stigma. In *Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '23, New York, NY, USA, 2023. Association for Computing Machinery.
- [71] Sohyeon Park, Aehong Min, Jesus Armando Beltran, and Gillian R Hayes. "as an autistic person myself:" the bias paradox around autism in llms. *CHI Conference on Human Factors in Computing Systems (CHI '25)*, page 17, 2025.
- [72] Vinitha Gadiraju, Shaun Kane, Sunipa Dev, Alex Taylor, Ding Wang, Emily Denton, and Robin Brewer. "i wouldn't say offensive but...": Disability-centered perspectives on large language models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, page 205–216, New York, NY, USA, 2023. Association for Computing Machinery.
- [73] Rudaiba Adnin and Maitraye Das. "i look at it as the king of knowledge": How blind people use and understand generative ai tools. In *Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '24, New York, NY, USA, 2024. Association for Computing Machinery.
- [74] Kate Glazko, Yusuf Mohammed, Ben Kosa, Venkatesh Potluri, and Jennifer Mankoff. Identifying and improving disability bias in gpt-based resume screening. *The 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)*, page 14, 2024.
- [75] Michael A. Hedderich, Natalie N. Bazarova, Wenting Zou, Ryun Shim, Xinda Ma, and Qian Yang. A piece of theatre: Investigating how teachers design llm chatbots to assist adolescent cyberbullying education. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA, 2024. Association for Computing Machinery.
- [76] Majeed Kazemitabaar, Runlong Ye, Xiaoning Wang, Austin Zachary Henley, Paul Denny, Michelle Craig, and Tovi Grossman. Codeaid: Evaluating a classroom deployment of an llm-based programming assistant that balances student and educator needs. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA, 2024. Association for Computing Machinery.
- [77] Muhammad Muzammil, Abisheka Pitumpe, Xigao Li, Amir Rahmati, and Nick Nikiforakis. The Poorest Man in Babylon: A Longitudinal Study of Cryptocurrency Investment Scams. In *Proceedings of The Web Conference (WWW)*, 2025.
- [78] Nur Yildirim, Hannah Richardson, Maria Teodora Wetscherek, Junaid Bajwa, Joseph Jacob, Mark Ames Pinnock, Stephen Harris, Daniel Coelho De Castro, Shruthi Bannur, Stephanie Hyland, Pratik Ghosh, Mercy Ranjit, Kenza Bouzid, Anton Schwaighofer, Fernando Pérez-García, Harshita Sharma, Ozan Oktay, Matthew Lungren, Javier Alvarez-Valle, Aditya Nori, and Anja Thieme. Multimodal healthcare ai: Identifying and designing clinically relevant vision-language applications for radiology. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA, 2024. Association for Computing Machinery.
- [79] Archit Parnami and Minwoo Lee. Learning from few examples: A summary of approaches to few-shot learning, March 2022.
- [80] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006.
- [81] K. Dhakal. Nvivo. *J Med Libr Assoc*, 110(2):270–272, April 2022.
- [82] Juhye Ha, Hyeon Jeon, Daeun Han, Jinwook Seo, and Changhoon Oh. Clochat: Understanding how people customize, interact, and experience personas in large language models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA, 2024. Association for Computing Machinery.
- [83] Judee Burgoon. *Expectancy Violations Theory*. 06 2015.
- [84] Michal Levi-Keren, Gabriella Landler-Pardo, Yehudith Weinberger, and Rinat Arviv Elyashiv. Simulation-based learning as a tool for assessing and fostering awareness of empathic patterns in teacher education. *Education Sciences*, 14(12):1338, December 2024.
- [85] Q.Vera Liao and S. Shyam Sundar. Designing for responsible trust in ai systems: A communication perspective. *2022 ACM Conference on Fairness, Accountability, and Transparency*, June 2022.
- [86] Tushar Khot, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal. Decomposed prompting: A modular approach for solving complex tasks. 10 2022.
- [87] Michelle Cohn, Mahima Pushkarna, Gbolahan O. Olanubi, Joseph M. Moran, Daniel Padgett, Zion Mengesha, and Courtney Heldreth. Believing anthropomorphism: Examining the role of anthropomorphic cues on trust in large language models. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, CHI EA '24, New York, NY, USA, 2024. Association for Computing Machinery.

A Simulation Flows

Interaction flows for the figurative expression, emoji with variable interpretations, and being misperceived as blunt scenarios are presented below. The interaction flow for the indirect speech acts scenario is discussed in Section 3.2 and Figures 4 and 5.

A.1 Figurative Expression




Figure 8: The simulation flow when an incorrect message option is sent in the "figurative expression" scenario. (1) shows the original message the user typed in; (2) shows the three message options generated and the user's choice; and (3) shows the AI character's response to that message and the feedback received.

A.2 Emoji with Variable Interpretation




Figure 9: The simulation flow when an incorrect message option is sent in the "emoji with variable interpretations" scenario. (1) shows the original message the user typed in; (2) shows the three message options generated and the user's choice; and (3) shows the AI character's response to that message and the feedback received.

A.3 Misperceived As Blunt




Figure 10: The simulation flow when an incorrect message option is sent in the "misperceived as blunt scenario". (1) shows the AI character's blunt message and the original message the user typed in; (2) shows the three message options generated and the user's choice; and (3) shows the AI character's response to that message and the feedback received.