

“I Hadn’t Thought About That”: Creators of Human-like AI Weigh in on Ethics & Neurodivergence

Naba Rizvi

University of California, San Diego
La Jolla, USA
nabarizvi99@gmail.com

Taggert Smith

University of California, San Diego
La Jolla, USA
tds002@ucsd.edu

Tanvi Vidyala

University of California, San Diego
La Jolla, USA
tvidyala@ucsd.edu

Mya Bolds

University of California, San Diego
La Jolla, USA
mbolds@ucsd.edu

Harper Strickland

University of California, San Diego
La Jolla, USA
hstrickland@gmail.com

Andrew Begel

Carnegie Mellon University
Pittsburgh, USA
abegel@andrew.cmu.edu

Rua Williams

Purdue University
West Lafayette, USA
rmwilliams@purdue.edu

Imani Munyaka

University of California, San Diego
La Jolla, USA
imunyaka@ucsd.edu

Abstract

Human-like AI agents such as robots and chatbots are becoming increasingly popular, but they present a variety of ethical concerns. The first concern is in how we define humanness, and how our definition impacts communities historically dehumanized by scientific research. Autistic people in particular have been dehumanized by being compared to robots, making it even more important to ensure this marginalization is not reproduced by AI that may promote neuronormative social behaviors. Second, the ubiquitous use of these agents raises concerns surrounding model biases and accessibility. In our work, we investigate the experiences of the people who build and design these technologies to gain insights into their understanding and acceptance of neurodivergence, and the challenges in making their work more accessible to users with diverse needs. Even though neurodivergent individuals are often marginalized for their unique communication styles, nearly all participants overlooked the conclusions their end-users and other AI system makers may draw about communication norms from the implementation and interpretation of humanness applied in participants’ work. This highlights a major gap in their broader ethical considerations, compounded by some participants’ neuronormative assumptions about the behaviors and traits that distinguish “humans” from “bots” and the replication of these assumptions in their work. We examine the impact this may have on autism inclusion in society and provide recommendations for additional systemic changes towards more ethical research directions.

CCS Concepts

• **Human-centered computing** → **Accessibility**; • **Computing methodologies** → **Artificial intelligence**; • **Social and professional topics** → *Computing profession*.

Keywords

neuronormativity, neuroinclusivity, humanization, humanness in AI, human-like AI

ACM Reference Format:

Naba Rizvi, Taggert Smith, Tanvi Vidyala, Mya Bolds, Harper Strickland, Andrew Begel, Rua Williams, and Imani Munyaka. 2025. “I Hadn’t Thought About That”: Creators of Human-like AI Weigh in on Ethics & Neurodivergence. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT ’25)*, June 23–26, 2025, Athens, Greece. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3715275.3732218>

1 Introduction

As human-like AI-powered communication agents such as robots and chatbots become increasingly popular, it is important to examine the ethical concerns surrounding the ways human-like behavior and characteristics or *humanness* is defined and implemented in such technologies. Our approach to this follows the recommendations of prior research and applies relational ethics in combating algorithmic injustices [5]. To this end, we analyze foundational beliefs informing algorithmic design, question what is considered “normal”, and how these definitions may marginalize certain groups [5]. In particular, we examine how perceptions of humanness may marginalize autistic people, as prior work has shown media representations of autism impact public perceptions, and more positive and accurate representations may even help combat the explicit biases others’ hold toward them [31, 45].

The **double empathy problem** suggests the misunderstandings in communication among autistic and non-autistic people is a two-way issue and not a result of a deficits among autistic people [44]. In contrast, **neuronormativity** is the positioning and privileging of neurotypical social behaviors as the “norm”, thus introducing a



This work is licensed under a Creative Commons Attribution 4.0 International License. FAccT ’25, Athens, Greece

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1482-5/25/06
<https://doi.org/10.1145/3715275.3732218>

power imbalance between people with different neurotypes, such as those who are autistic [38]. Through a neuronormative perspective, autistic people are marginalized for lacking “social skills”. However, in alignment with the double empathy problem, prior work has shown that in neurotype-matched interactions, autistic people have high interactional rapport with other autistic individuals, suggesting that they have their own communication preferences instead of having a deficit of the communication behaviors commonly found in neurotypicals [12]. Although neuronormative beliefs lead to the dehumanization and marginalization of autistic people, punitive measures, and other epistemic injustices, they continue to be pervasive in our society [4, 9, 38]. For autistic individuals, neuronormativity has often resulted in dehumanizing comparisons to robots, animals, and other non-human entities, which compounds the importance of ensuring technological agents do not reproduce these stereotypes [56, 77].

Prior work in computing research has examined the ways in which technologies may fail to align with the needs of autistic and other neurodivergent (ND) individuals, or contribute to their marginalization. Autistic people have been notably excluded from the design of robots, which has resulted in robots often replicating dehumanizing stereotypes in their interactions with autistic people [56]. While researchers have identified the unique needs, exclusion, and marginalization of autistic people in other areas of computing research and highlighted the ways in which technology can be made more accessible for them [2, 3, 23, 24, 26, 56, 65, 69, 74, 83], there is a notable gap in research exploring the perspectives of the people who create such technologies, particularly their intentionality, perceived and intended impact, and other ethical considerations of their work. Additionally, it remains unclear how the creators of these AI systems conceptualize and operationalize humanness in their work, and whether these efforts perpetuate, support, or challenge neuronormativity.

To address this gap, our work examines the beliefs and experiences of researchers, designers, and engineers who work on human-like AI technologies (i.e. robots and chatbots). We conduct a qualitative study using an interpretative phenomenological analysis of two interviews and two surveys from each individual participant in our study. Our work also examines the alignment of their work with neuronormative standards of communication, and the barriers they face in making their work more accessible and inclusive.

To this end, we investigate the following:

- How do AI system makers conceptualize and implement “humanness” in their systems?
- From the developer perspective, what are the potential ethical and societal impacts of designing AI systems to mimic human communication?
- How do AI makers’ considerations toward replicating human behavior reinforce neuronormative standards and marginalize autistic individuals?
- What challenges prevent AI developers from incorporating neurodiversity and accessibility into their design processes, and how can these be addressed?

2 Related Work

Advancements in AI increasingly aim to replicate human qualities, yet ethical concerns arise from how “humanness” is defined, often overlooking the experiences of historically dehumanized groups such as autistic individuals. In this section, we explore how humanization is implemented in human-like AI agents such as robots and chatbots, the anti-autistic biases prevalent in AI, and the ways in which such technologies may perpetuate stereotypes or fail to align with the needs of autistic people.

2.1 Humanizing AI

2.1.1 Origins. Alan Turing proposed a simple and observable test to evaluate whether a machine can exhibit human-like intelligence through conversation [71]. The Turing Test, which assesses whether or not a judge can reliably distinguish between answers from a human and a machine [71], is considered a benchmark for assessing AI’s humanness. However, its emphasis on imitating human behavior raises ethical concerns about which set of humans are being treated as the gold standard for ‘humanness’, especially as the test considers it to be a key benchmark for intelligence [30, 47]. This is particularly relevant for marginalized groups, such as autistic individuals, who may be unfairly measured against neuronormative social expectations that AI is also trained to replicate. Despite these concerns, the test has continued to influence decades of AI research on natural language and cognition, from early chatbots to modern dialog models [30, 47, 62, 81].

2.1.2 Contemporary Work. From the Turing Test [70] through Weizenbaum’s ELIZA [76] and onward to Breazeal’s sociable robots [6], the drive to humanize technology has spanned domains from natural language processing to embodied robotics. It has shaped contemporary research on how robotic appearances and gestures evoke emotional responses, and continues to inform contemporary HCI work aimed at cultivating trust through empathetic yet transparent AI [13, 25]. Designers now consider anthropomorphic (i.e. human-like) behaviors such as voice, facial expressions, or gestures with ethical frameworks emphasizing user autonomy, cultural sensitivity, and disclosure of AI’s limitations [11, 20]. While researchers are increasingly moving toward responsible AI research, such work continues to overlook people with disabilities, particularly those who are neurodivergent.

2.1.3 Limitations. Recent works continue to build upon Turing’s foundational ideas. For example, one study on conversational agents highlighted that incorporating cognitive, relational, and emotional competencies can enhance user engagement and lead to more human-like capabilities in these agents [10]. Similarly, another paper highlighted the modalities (verbal, non-verbal, appearance) and footing (similarity and responsiveness) that can help optimize interactions with end-users by moving toward humanness [75]. However, a systematic literature review highlighted many ethical concerns with humanizing AI, such as the usage of AI to manipulate end-users, for example, by influencing their voting decisions [1]. Importantly, this work does not address the ethical implications of how humanness is defined and implemented in AI. This issue is particularly significant for autistic people who have been historically

dehumanized due to their cognitive, relational, or emotional differences [50] and have been unfairly compared to robots, with robots frequently and paradoxically being used as mentors to guide them toward conforming to a constructed notion of “humanness” [56, 77].

2.2 Accessibility Considerations

A systematic review of chatbot accessibility organized considerations into five categories: content, user interface, integration with other web content, developer process and training, and testing [67]. This work highlighted the importance of ensuring chatbots use straightforward and literal language with a single-minded focus for broader accessibility [67]. Similarly, another review detailed accessibility concerns in chatbots, and how they arise from a lack of support for alternative input methods and assistive technologies, lack of clarity and consistency, and a lack of simplicity [82]. The recommendations provided in this study included providing navigational assistance, keyboard shortcuts, and feedback while reducing disruptive factors to improve accessibility [82]. A similar study on pre-existing accessibility guidelines for human-robot interaction research found that some researchers had considered ad hoc guidelines in their design practice, but none of them showed awareness of or applied the guidelines in their design practices [52]. Notably, these guidelines have a lot of overlap with the ones proposed for chatbots, for example, both encourage using multiple modals, providing support for assistive technologies, navigation assistance, and feedback [52, 82].

2.3 Anti-Autistic Biases in AI

Prior studies have examined several ways in which AI displays biases through applying neuronormative standards in their classification of emotions and behaviors, and stereotypical representations of autism. Notably, one study found training AI on neurotypical data leads to a mismatch with the needs of autistic users, highlighting the limitations of datasets that do not include accurate representations of autistic people [3]. Similarly, prior work has examined the ways in which AI-powered emotion recognition in speech may perpetuate ableist biases through their focus on pathologizing the communication behaviors of autistic people, especially as they build upon foundational work dehumanizing autistic people by comparing them to computers [33]. Additionally, diagnostic voice analysis AI may misclassify autistic people’s speech as “atypical” or “monotonous” which results in the system making ableist assumptions [40]. AI-powered talent acquisition systems may also misunderstand autistic people’s behaviors by judging them based on neurotypical standards, resulting in misclassifying their lack of eye contact as a lack of confidence, for example, and negatively impacting their overall perceptions of autistic candidates [8]. Even generative AI has a tendency to produce biased representations of autistic people through stereotypical emotions such as anger and sadness and engagement in solitary activities [80].

2.4 Robots, Chatbots, and Autism

Anti-autistic biases are prevalent even in chatbots, as one study uncovered that GPT-4 displays biases toward resumes that mention disabilities, including autism [22]. Additionally, a systematic review of human-robot interaction research found robots may marginalize

autistic people through a power imbalance in their user interactions by acting as mentors who help autistic people move toward humanness [56]. Interestingly, there is a notable contrast between the preferences of autistic people and others when it comes to robots and chatbots—prior studies found that autistic people have a preference for non-anthropomorphic robots and chatbots [35, 54, 56], while anthropomorphized chatbots and robots may increase satisfaction for other users [34], particularly those who have a higher desire for human interaction [61]. This suggests that even the decision to implement human-like traits in a bot may marginalize autistic users by overlooking their preference in favor of positive user experiences for other groups.

2.5 Creator Perspectives on Ethics and Limitations

Prior work exploring the perspectives of creators suggests that ethical concerns are not purposefully overlooked, but are rather the result of other practical limitations that can be addressed to improve the alignment of theoretical ethical beliefs with real-world practices. For example, one study investigated the ethical caveats of conversational user interfaces by interviewing both the designers and end-users of chatbots, and identified mismatches between the designers’ user-centered values, and the resulting user experiences shaped by technical constraints in the real-world [43]. Similarly, another study examined the thoughts of design leaders toward implementing ethics in commercial technology settings found that although the designers recognize the importance of inclusive design, they face limitations due to the pressure to deliver products quickly, which creates a gap between their aspirational ethical guidelines and their real-world projects [39]. Other works identifying the needs of machine learning practitioners, designers, and data practitioners has highlighted the importance of practical tools and guidance in helping bridge the gap between literature on fair ML research and real-world constraints [28], as unclear best practices makes them struggle with “ethical correctness” [15].

3 Methods

Our study consisted of a comprehensive methodological approach encompassing virtual semi-structured interviews, strategic participant recruitment, and rigorous data analysis to explore participants’ experiences with creating human-like AI, and their perspectives on intersectionality, accessibility, and neurodiversity accommodations.

3.1 IRB Approval and Other Ethical Concerns

Our IRB was approved by the appropriate Institutional Review Board (IRB). Prior to participating in each portion of our study, the participants filled out a consent form that provided details on the time commitment, participant and researcher expectations, the nature of the data being collected, the compensation being provided, and other information about the research team. Participants were explicitly asked for permission to record their interviews, and were made aware that they could quit the study at any time at their own discretion.

3.2 Participant Recruitment

Our recruitment strategy consisted of purposive sampling, snowball sampling, and outreach through social media. We recruited participants based on specific eligibility criteria, including being U.S.-based, at least 18 years old, and employed as researchers, designers, or engineers working in the development of human-like AI systems. We used purposive sampling to target a pool of 154 professionals from U.S.-based tech companies and universities specializing in AI and user experience (UX). These professionals were found through a manual search on the websites of their employers. We also reached out to community-based organizations to recruit participants, such as affinity groups supporting underrepresented communities in computing and AI, nonprofits advocating for disabled and queer individuals in technology fields, and employee resource groups at tech companies. Finally, we used snowball sampling by inviting eligible researchers from our networks and those recommended by other participants. We posted advertisements on social media platforms, including LinkedIn and Twitter. Ultimately, 16 individuals qualified for and completed the full study.

We define “human-like AI” as systems mimicking human communicative or behavioral cues (e.g., text, speech, gestures). We include participants working on self-described human-like technologies (e.g. robots, chatbots, digital avatars). We did not use the phrase ‘human-like’ during our recruitment to avoid recruitment bias. Here is an example of a passage used in our recruitment:

“We are interested in learning about the background and experiences of professionals in AI, Machine Learning, and/or Data Science. We are specifically interested in learning more about the processes used during the development of human-facing tools and the life experiences of those professionals.”

3.3 Study Design

We conducted semi-structured interviews virtually, with each session lasting between 30 and 45 minutes. The study was divided into two sections, with each section having a survey and an interview. The first section collected data on the participants’ demographics and dived deeper into their experiences with working on a human-like AI project. The second section focused more on their views of diversity, disability, and neurodiversity. Due to the nature of the questions presented in the second part, we alternated the order in which the survey and interview were presented to each participant to provide counterbalance. Participants were compensated \$60 for their participation, provided in two installments: \$30 after completing each section.

The interview and survey questions were carefully-designed to minimize response bias and gain a deeper understanding of participant perspectives. A collaborative team of researchers with expertise in software engineering, responsible AI, and security/privacy research developed the study materials. The interview and survey questions are available in the Appendix A.

3.3.1 Surveys. The study included two surveys, which can be found in our Appendix sections A.1.1 and A.1.2. The first survey gathered demographic information, including age, gender, socioeconomic status, education level, disability, race, and marital status

with questions taken from previous studies [19, 37, 58, 60, 66, 72, 78]. The second survey explored participants’ views on topics such as intersectionality, diversity, and workplace accommodations for neurodiversity.

3.3.2 Interviews. The first interview consisted of carefully designed questions on the design, development, and testing of a recent project completed by our participants focusing on human-like AI. This interview also included questions on general accessibility considerations, accessibility for neurodivergent individuals, and the broader impact their project may have on shaping the perceptions of communication norms and the humanness of both their end-users and other AI creators who build on their work. The questions for this section of the interview are available in our Appendix Section A.2.1.

The second interview focused on understanding each participant’s views of intersectionality, diversity, disability, and neurodiversity. In this interview, we also presented each participant with scenarios of workplace interactions between colleagues of diverse backgrounds including different neurotypes that were adapted from a previous study [55]. The identities of the people in the scenarios were not revealed to the participants to avoid response bias, but were indirectly communicated through their dialogue focusing on sensory and communicational differences common in individuals with ADHD and autism. These questions were designed to dive deeper into the participants’ knowledge and acceptance of neurodiverse identities. The questions for this section of the interview are in our Appendix Section A.2.2.

3.4 Analysis

We used an interpretative phenomenological analysis (IPA) approach to examine the data collected [17]. IPA is a qualitative method that explores how individuals perceive and interpret their lived experiences, typically involving iterative steps such as immersing oneself in the data, generating initial codes, searching for emergent themes, and synthesizing these into interpretative narratives [17]. The method also requires researchers to ‘bracket’ or set aside their biases at the beginning of the analysis, and have a data validation process to ensure the analysis results reflect the participants’ experiences and perspectives [63]. We used this method to gain a better understanding of each participant’s experiences as they were from diverse professional backgrounds and were working on different kinds of projects. Additionally, this analysis method benefits from quality and depth over quantity and breadth, which allows us to concentrate more deeply on the data collected in our multiple interviews and surveys from each individual [17].

The analysis was conducted by a team of four researchers (authors 1-4), with the findings validated by two additional researchers (authors 5 and 6) specializing in neurodiversity and accessibility. All interviews were recorded and transcribed with the participants’ knowledge and consent. To ensure an unbiased approach, each researcher engaged in a bracketing process by documenting their own perspectives and biases before beginning the analysis. This process was used to help mitigate potential personal biases on the findings. Each researcher independently reviewed the interview transcripts, taking verbatim notes with direct quotes to capture the participants’ experiences. The team then engaged in iterative

group discussions to summarize findings, identify recurring patterns, and generate codes. These codes and themes were further refined collaboratively and validated by external experts to ensure the accuracy and rigor of the analysis.

3.5 Positionality

This work was led by an autistic researcher with a neurodiverse research team. Although the researchers have diverse racial, ethnic, and gender backgrounds, we are all English-speaking and US-based. Thus, our work has limitations due to our Western and Anglo-centric perspective. We acknowledge that due to these limitations, our findings may not be applicable to cultures and languages, and encourage future work to explore other perspectives.

4 Results

We present insights from our interviews and surveys on the experiences and beliefs of researchers, designers, and engineers working on human-like AI. In particular, we focus on their views toward the accessibility and ethical implications of their work, and its alignment with neurodiversity. Examples of their projects include robots, chatbots, video games, and assistive technologies. Table 1 provides an overview of our participants' demographics while Figure 1 summarizes their knowledge and acceptance of the behaviors and traits common in neurodivergent individuals that may be considered atypical by neuronormative communication standards [55, 79]. The participants were classified as unaccepting, accepting, unaware, or aware based on their responses to our scenario-based questions in the second interview. For example, if a participant answered, "that's weird", when asked to imagine a colleague who wears headphones, but said "that's ok" when we specified the individual does it to avoid sensory overload, they were classified as being unaccepting but aware of neurodivergence as they initially displayed bias toward the behavior but understood it may just be a behavioral difference.

4.1 Desirable Traits

Through a qualitative thematic analysis of the participants' responses during our interviews, we uncovered 12 desirable traits they implement to make their technologies appear more human-like. Table ?? details the prevalence of each trait, with "personalized", and "uses natural language" being the most popular desirable traits. Notably, traits such as "ethical", and "compassionate" were less common.




Figure 1: An overview of our participants' knowledge and acceptance of neurodivergence.

4.2 Undesirable Traits

The participants also identified undesirable traits that caused communication breakdowns with their end-users and resulted in the users getting frustrated or upset with the system. These traits, shown in Figure 2, focused on the behaviors, utility, or appearance of the system. A bot would be considered "inefficient", for example, if it took too long to respond, understand the user, or complete the task specified. Similarly, "dysfunctional" bots were those that offered incorrect answers or experienced other technical difficulties. The participants also expected their bots to have appropriate tone based on the context of the communication, the identity of the user they were interacting with, and the system's assumed role in the interaction (e.g. as a boss, teacher, or peer). For example, the bots must not appear to be too 'child-like', which can be avoided by making them speak 'eloquently'. Another source of user frustration reported by our participants was a misalignment of their technologies with the participants' needs. These misalignments could be cultural, for example, if the bot struggles with a user's name, or simply due to their irrelevance to the user's needs. The most commonly reported undesirable trait was the bot appearing "uncanny". As shown in Figure 4, this was often the result of the system's communicative behaviors, emotional expression, or appearance. Figure 3 provides examples of the ways these behaviors and traits were described by our participants.

4.3 Accessibility

Accessibility considerations were mainly discussed by the participants for physical disabilities or globalization (n=12). Participants highlighted challenges such as language barriers and technological limitations, including slower devices. P3 emphasized, "It needs to be accessible to, you know, employees in Malaysia, Denmark, Europe, South America," while P6 noted issues such as users not owning personal devices and bots struggling to recognize Indian names, leading to user frustration. P1 acknowledged the need for accommodations in their app for medical professionals after drawing a personal connection to a sibling's vision impairment, noting that "he uses lots of accessibility tools for his study." However, this recognition only extended to physical disabilities. When asked about accommodations for neurodiverse individuals, the participant stated that they "don't have an answer for this." Similar conclusions were drawn by several other participants who considered and implemented accessibility for physical disabilities but not neurodiverse conditions (P4, P8, P11). Participant 11, who had implemented colorblind-friendly accommodations for their bot, mentioned that they were "not sure at all" and that "it's not something I consider" when asked about accessibility for neurodiverse individuals. P4 considered accessibility for blind and low-vision individuals, implementing a screen reader for their application, but stated that they were "not very sure how it would change their interaction with the tool" when it came to neurodiverse individuals.

4.4 Neurodivergence

Our findings reveal that while some participants demonstrated an awareness of the unique traits and differences of neurodivergent individuals (n=7), many had not considered their specific accessibility needs (n=12). For instance, P9 admitted, "I guess those populations

ID	Gender	Age	Race	Education	Disability	ND	Areas
1	M	31-40	White	PhD	None	Yes	AI
2	NB	18-22	2 or more races	Bachelors	None	Yes	AI
3	F	31-40	Other	No degree	Has disability (unspecified)	No	HRI
4	M	23-30	Asian	Bachelors	None	No	AI
5	M	51-60	Black	PhD	Blind	No	HCI
6	F	23-30	Asian	PhD	None	No	HCI
7	F	23-30	Asian	Bachelors	None	No	AI
8	M	23-30	Asian	PhD	Deaf	No	HCI
9	F	23-30	Other	PhD	None	No	HCI
10	M	31-40	White	PhD	None	No	HCI
11	M	23-30	Hispanic/Latino	PhD	None	No	AI
12	F	23-30	Asian	Bachelors	None	No	HCI
13	M	23-30	2 or more races	PhD	None	No	AI
14	M	51-60	White	PhD	None	No	HCI
15	M	31-40	Asian	PhD	None	No	AI
16	M	31-40	Hispanic/Latino	PhD	None	No	HRI

Table 1: The demographics of the participants in our study. ND refers to neurodivergence.

Desirable traits	Prevalence	Examples
Personalized	10	Mimics interactions with family members, peers, or friends
Uses natural language	7	Delayed responses, multimodal communication, organic conversations
Serious	6	Professional tone, not overly positive, not funny
Simple	6	Does not contain unnecessary features
Helpful	5	Offers clarifications, task-oriented, time-saving tool
Encouraging	3	Perky, has a rewards system
Compassionate	3	Friendly, empathetic, sympathetic
Therapeutic	3	Mimics warm nurse and interactions with therapists, calming
Emotive	2	Expresses emotions through vocal and facial cues
Teaches	2	Provides casual learning experience, teaches sensitive topics
Ethical	2	Safe, private, does not make users feel watched

Table 2: An overview of the different desirable traits and their prevalence as implemented by our participants in their bots to make them appear more human.

are not really my expertise,” while P13 acknowledged, “I actually don’t know enough, I guess, about autism to really know how this would affect them.” When prompted, this lack of understanding often led participants to default to design choices that either eliminated features rather than adapting them or made assumptions about neurodivergent users’ preferences. For example, P10 noted that humor could be misinterpreted by neurodivergent users, potentially leading to a poor user experience, and chose not to include it as a feature. Other participants (P7, P8, P16) acknowledged the value of customizing system features for various user groups, but also made assumptions about neurodivergent users’ interaction needs. P16 stated, “[neurodivergent people] need some kind of motivational behavior from the robot in order to encourage them to engage during the interaction,” and P7 suggested “some extra persuasive intervention tools” for neurodivergent people who “suffer from a lack of motivation.” Similarly, P3 assumed their current system would be “super, super friendly” for neurodivergent users, “even if they have like issues with spelling, because we have so

many ESL workers, there are some tolerances around [...] not exact matches,” and because the bot “responds with simple things.” None of these participants directly examined the needs of neurodivergent individuals in their studies. In contrast, P14 recognized the significant appreciation from neurodivergent communities for simply being considered during the design process, stating, “For the neurodiverse population [...] the fact that we were designing for them at all [...] there was an outsized appreciation for that. And I think that’s basically because they’re a commonly neglected or ignored community.”

4.5 Communication Biases

Participants also made assumptions about communication preferences and behaviors. Some of these assumptions involved the mode of communication. For example, P11 viewed text-based communication as “overly robotic”. P12 described switching from a numerical scale to using words in the bots’ interactions with the user to “humanize” the experience as numbers felt “unnatural”.

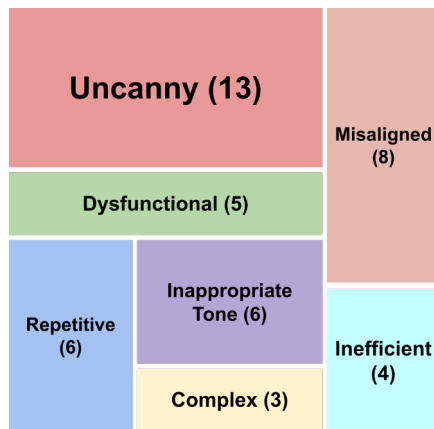


Figure 2: A list of undesirable traits discussed by our participants that may upset or frustrate their users.

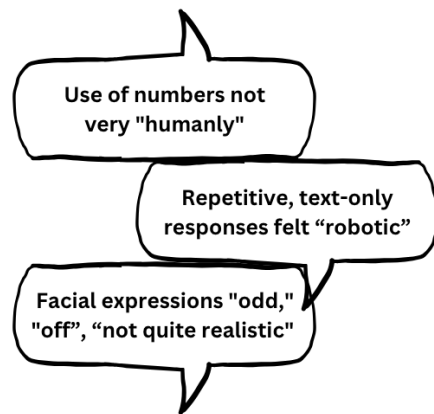


Figure 3: The words used by our participants to refer to undesirable traits and behaviors that frustrated their users.




Figure 4: A breakdown of the different aspects of a bot considered uncanny by our participants.

Others focused on the tone or content of the communication. For example, P13 perceived overly positive responses as “unnatural”, especially when the bot is prompted to be angry or hateful toward certain communities. Additionally, some participants made assumptions surrounding the impact of certain behaviors. P11 believed making their chatbot repetitive and not providing a direct reward during interactions made it “boring” and “robotic” for the end-users and negatively impacted their engagement. Participants prioritized emotional responsiveness in context of their bots’ conversation and often equated this with humanness. P2’s bot was designed to be remain “empathetic” and “grounded” when discussing serious topics to maintain a calming environment for the user. P15 was particularly concerned with how empathy was expressed in the bot’s conversational style and tone, believing “robotic, monotonic advice would make the bot appear less human.

4.6 Neurotypical Bias in “Friendly” System Design

When participants discussed their design choices around making systems “friendly” or “personable,” they often defaulted to neurotypical social conventions and communication patterns [79]. For example, P3’s emphasis on making their system “cute, perky, and friendly” reflects assumptions about universal preference for social interaction styles that may not align with neurodivergent users’ needs who prefer more straightforward systems [56, 57]. Additionally, some participants’ systems (P3, P5) prominently featured an anthropomorphized face, a design choice likely to be less relevant to some autistic users for whom such conventional facial signifiers of emotion fail to resonate [57, 83]. P7 considered incorporating emojis and believed that emotional content was key to humanness, saying of their current system: “these kinds of agents don’t have emotions, so it would probably not be that fluid or, like, more human.” Our participants’ emphasis on emotional facial expression may cause users and researchers to associate a specific mode of communication via facial expressions with humanity, which misaligns with the findings of prior work showing autistic individuals prefer interacting with plain and featureless bots over humanized ones [57].

4.7 Ethics

The ethical concerns discussed by our participants were centered around privacy, model biases, and their perceived responsibilities in creating more ethical and accessible technologies. While participants were questioned about the broader impact of their work, the majority did not discuss any broader ethical considerations.

4.7.1 Privacy. Participants expressed varying perspectives surrounding privacy. Some participants were actively working on improving their systems’ privacy. For example, P6 described disabling certain features to avoid invading users’ privacy, while P12 highlighted efforts to avoid making users feel surveilled. However, other participants were less concerned with privacy. P10 did not prioritize privacy due to their focus on Gen Z users, observing generational differences in privacy attitudes, “Gen Z [are] less concerned about privacy than previous generations.”

4.7.2 Model Bias and Accessibility. Notably, ethical concerns related to AI model biases were raised by HCI-focused participants (P2, P6, P12) but were absent among AI-focused participants (P1, P4, P13), sometimes explicitly, as one participant remarked ‘we are machine learning people, not HCI people’. When discussing accessibility, AI-focused participants often downplayed its necessity. For instance, P1 stated, “I don’t know if [the results of] this model are necessary to be accessible because it’s more helpful to the doctors and biologists,” assuming that medical professionals do not have accessibility needs. Similarly, P11 remarked, “We weren’t UI people; we were ML people,” highlighting a lack of cross-disciplinary consideration in their designs. Such sentiments were also shared by other participants who believed accessibility is the responsibility of the product team (P4) or will inevitably occur when needed as a result of commercialization (P1). Overall, accessibility and ethics were noted as a future concern or beyond the scope of their responsibilities.

4.7.3 Broader Ethical Concerns. A significant concern was the lack of attention to broader ethical implications, particularly in relation to societal norms surrounding communication. Most participants (n=12) did not discuss the downstream effects of their work on either future researchers or end-users, particularly surrounding communication norms. This was especially troubling given the tendency of participants to mimic existing technologies like ChatGPT, as illustrated by P7: “[We] kept it really simple and tried to mimic other chat bots, for example, the ChatGPT interface.” Some of the participants more explicitly mentioned being inspired by communication behaviors that “humanized” other technologies, such as Duolingo’s bird icon which displays facial expressions corresponding to users’ engagement (P12). Yet, while the participants noted being influenced by the communicative behaviors of existing technologies, they did not discuss how their own work may similarly influence others.

4.8 Barriers

While the accessibility considerations mentioned by participants, such as having a simple interface, and providing navigational assistance were in alignment with prior work on accessibility standards for robots and chatbots [52, 82], the implementation of these features faced significant barriers, the most common among them being their organization’s priorities (P1, P3, P4, P7, P13, P15). As P3 explained, “So accessibility is probably not [our company’s] strong suit because it’s not a consumer-facing org,” overlooking the potential accessibility needs of employees and other end-users. Other barriers included funding constraints (P12), time limitations (P14), and technical challenges (P10, P11, P16).

5 Discussion

Our findings suggest that human-like AI continues to promote neuronormative standards of communication. We investigate the impact this may have on dehumanizing autistic people and recommend systemic changes to move toward more ethical research.

5.1 When “Uncanny” Meets Stigma: Parallels with Autistic Stereotypes

Our participants notably described their perceptions of uncanny qualities in their bots through stereotypes that are often also applied to autistic individuals. In particular, P11 described text-only communication as “too robotic,” due to a lack of human “warmth” or “natural flow” in the system’s interactions. Similarly, autistic people are also perceived as being “robotic” due to differences in their communication behaviors, which include a preference for text-based communication for clarity and reduced social pressure [29, 56, 77]. This highlights how our participants’ perceptions of “robot-like” communication may inadvertently reinforce problematic assumptions about neurodivergent behavior. Similarly, P15 pointed out that certain body movements and microexpressions seemed “off” or “not quite realistic,” contributing to an overall sense of the system being unnatural or inauthentic. These comments echo broader societal tendencies to dehumanize individuals who express emotions in ways deemed “atypical” [48], a common way autistic people are marginalized in our society and are often pressured to mask their natural behaviors despite the negative impact on their well-being [53].

Further highlighting the ways in which user interactions can be shaped by social stereotypes that correlate with disability stigma, P16 perceived the robot as more of a child than a caregiver. The disconnect between the system’s expected competence and its perceived immaturity caused some user frustration. This echoes the infantilization that autistic people frequently face [68], highlighting yet another way in which marginalization ties autistic identities to robot-like attributes [77]. Notably, P15 also drew attention to a broader distinction between humans and robots: the ability to infer another person’s internal state from subtle external cues. Similarly, autistic individuals are often accused of lacking a “theory of mind”—the assumedly universal ability to empathize by putting oneself in another’s position—according to dominant neuronormative standards [64]. These parallel perceptions of “deficient” empathy in both robots and autistic people reinforce the belief that atypical communication or emotional expression is inherently less human.

These participant insights not only emphasize how easily technology can be perceived as uncanny, but also how such perceptions are linked to normative expectations for communication, emotional expression, and social intuition. In doing so, they also highlight the risk that implementing such expectations in bot design in attempts to avoid the uncanny may reinforce harmful stereotypes about autistic people, who are frequently subjected to similar judgments and dehumanizing labels.

5.2 Humanizing Machines, Dehumanizing Humans

Dehumanization, whether subtle or overt, appears with alarming frequency in inter-group relations [36]. Subtler forms involve ascribing fewer human emotions or “complex” traits (e.g. maturity, civility or refinement) to outgroups, while more blatant instances compare marginalized groups to animals, machines, or primitive beings [36, 56]. Interestingly, in our study, the participants described frustrations with their bots being perceived as “childlike” or not

communicating “eloquently,” exemplifying the “complex traits” associated with machines that are also associated with dehumanized human groups. Additionally, our participants often associated behaviors preferred by autistic people, such as text-based communication, with robots, highlighting an explicit dehumanization of neurodivergence [29].

Such dehumanizing stereotypes can have a serious negative impact on the communities targeted by them, and thus it is important to address them. They may foster hostility, discrimination, and violence, both systemic and overt [36]. Those subjected to dehumanization may often experience negative emotions, develop more strained inter-group relationships, and may respond with reciprocal hostility [36]. To mitigate such dehumanization, prior work has suggested promoting inter-group contact, challenging hierarchical views of humanity, and emphasizing shared identities. Additionally, highlighting the similarities between groups can reduce subtle forms of dehumanization [36]. In fact, prior research has found employing these strategies in the virtual world can be effective in combating biased behaviors in the real world [7, 42, 46, 51, 59]. Thus, the incorporation of neurodiverse personas and behaviors into interactive AI agents can be a critical next step in combating dehumanization through the normalization of diverse communication styles.

5.3 Worlds Collide: How Virtual Interactions Impact Reality

Contact hypothesis theorizes that people’s prejudices toward particular social groups may be reduced through contact with the group, and a systematic review found it does typically reduce prejudice [49], though there are notable exceptions and limitations (such as self-segregation) which may impact its effectiveness in the real world [42]. However, prior work suggests contact with autistic people and knowledge of autism may improve autism acceptance among others. One study found familiarity with an autistic individual may decrease non-autistic people’s negative perceptions of autism [16], and another found providing autism acceptance training to non-autistic people reduces their explicit biases, improves their understanding of autism, and increases their interest in engaging with autistic people [31].

Researchers have examined the ways in which contact hypothesis occurring in digital spaces may impact people’s attitudes in real life, and have found diversity in video games may lead to more accepting attitudes amongst gamers towards diverse groups [46]. Similarly, other prior works have found that using racially diverse avatars led to less racially biased behavior in real world interactions [51], and even non-playable characters of diverse backgrounds may decrease users’ explicit biases [7]. In another study, researchers uncovered that conversational agents facilitating contact can improve inter-group attitudes even among groups that have a long history of conflict [59]. This also shows the important role conversational agents can play in addressing the common shortcomings of the contact hypothesis, which is strongly dependent on the nature of the contact as negative contact increases prejudice [42].

Even though there are benefits to promoting more diverse interactions online, our study evidences that neurotypical communication standards remain dominant in human-like robots and chatbots.

Consequently, neurotypical people rarely gain opportunities to see autistic traits and characteristics as inherently human, rather than “robotic.” Such exclusion reinforces the harmful perception that certain behaviors belong in the realm of machines, as they are not represented in systems designed to be human-like, compounding the marginalization of autistic individuals and other minority groups. For example, labeling particular behaviors as distinctly “human” versus “AI-like” can lead to the marginalization of individuals perceived to be using AI [27]. These issues can have dire consequences, such as job loss or false accusations of plagiarism, which disproportionately affect marginalized communities [21]. Thus, the absence of neurodiverse representations of humanness in such technologies is particularly concerning as they promote the neuronormative belief that there is only one right way to be a human [4]. As media representations play a major role in autism acceptance in society [45], it is important to ensure the technologies we create do not perpetuate biases regarding social norms that often dehumanize autistic people.

5.4 Beyond Conference Policies: Other Recommendations for Systemic Changes

Despite the recent push in publication standards at AI conferences to prioritize ethical considerations through means such as using ethics statements in submissions, we note that the broader ethical concerns of their work were not discussed by our participants. Even after exploring the conceptions of humanness implemented in their technologies, when asked about the ways these conceptions may influence the perceptions of human interactions in the real world or with AI held by others, the majority of participants ($n=12$) believed their work would have no impact on the way researchers building upon their work and their end-users view communication behaviors among humans. Similarly, the model biases and their impact were mentioned more frequently by our HCI-focused participants compared to our AI-focused ones, with the latter focusing solely on standard metrics such as precision to gauge the effectiveness of systems. This shows that despite the efforts to encourage more ethical work, AI researchers and engineers still struggle with understanding or explaining the broader ethical considerations that may arise from their work.

5.4.1 Not an Afterthought: Centering Ethics in AI Education. While AI technologies impact the lives of many humans either directly or indirectly, the creators of such systems have a tendency to view human concerns as beyond the scope of their responsibilities. In our study, participants frequently expressed sentiments such as “we are [machine learning] people, not HCI people” (P11), and referring to ethical concerns as the responsibility of others such as the product team (P1, P4). Such beliefs are mirrored even in our education system, with many AI educators having conflicting and contradictory thoughts toward ethics [32], and many universities separating those who work on “ethical” or “human-centered AI” from those who work on more “technical” projects through the creation of distinct departments or programs. This separation of “technical” and “human-centered” perspectives may result in the former group receiving inadequate training in identifying ethical issues. For example, prior work has found AI courses hosted on YouTube neglected discussing ethics in favor of more technical

content [18], and undergraduate computer science students believe ethics are not prioritized, valued, or reward in their training or job prospects [14].

In order to promote more ethical work, it is important for us to critically examine the separation of what we consider to be “ethical” or “responsible” AI from other forms of AI. Other engineering professions have standardized professional codes, legally binding standards, and license examinations, that prioritize safety and ethical concerns for all engineers. There is no such thing as an “ethical” structural engineer, for example, so why do we have that separation in AI? If we shrugged off ethical shortcomings as easily in other professions, we would need to exercise caution in only walking in buildings made by “ethical” engineers, for instance, to avoid having them spontaneously collapse under our feet. Yet, in the technical realm, the safety, well-being, and other ethical concerns of the broader general public are routinely sacrificed to be the “first” or “best” at releasing a product or a feature without any ethical examinations, despite the magnitude and scale of their impact on others’ lives sometimes being far larger than that of a single building. Perhaps, it is time we start a more thorough integration of ethics in our training of future makers, our evaluation of current makers and leaders, and the standards we all must uphold.

5.4.2 Organizational Support and Priorities. While our participants expressed an interest in making their technologies more accessible, particularly for users who are international or have physical disabilities, they often mentioned being limited by their organization’s support and priorities. For example, P3 noted accessibility was not a strong suit for their company, and reported having to “push” to implement their considerations to enhance user experiences. They described having limited control over the overall design of their products, even mentioning that they had more flexibility in their university projects. In contrast, P15 responded having more support such as established standards for accessibility at their organization, and on-going projects that focused specifically on accessibility. These experiences show the impact organizational priorities can have on promoting or hindering more inclusive design, as many participants reported technical features were a bigger priority than user experiences for their organizations. Yet, there is a notable lack of clearly defined standards for these organizations and, consequently, any accountability for failing to adhere to them. While researchers have attempted to standardize guidelines for improving the accessibility of systems such as robots and chatbots, as noted by prior work and reiterated by our own findings, these are not well-known even among the people who create these technologies [52, 67].

5.4.3 Fostering Diverse Teams for Socially Aware Perspectives. An example of an ethical rule from the National Society of Professional Engineers (NSPE) Code of Ethics states that engineers shall “perform services only in areas of their competence” [73]. Applying such a rule in computing would require us to pursue interdisciplinary collaborations more fervently, as our work directly impacts the lives and livelihoods of communities we may not be familiar with but who may be the area of expertise of other trained professionals.

Indeed, we found the ethical and accessibility considerations that were mentioned by our participants were often related to their own knowledge and familiarity with the groups they believed would

be impacted by their work. For example, P1 discussed accessibility considerations for blind/low vision users due to having a personal connection with a blind family member. Additionally, P2 mentioned implementing trauma-informed features in their work due to their experiences working with a minority community focusing on tasks that were often traumatizing. This highlights the importance of our other recommendation, which is fostering diversity in teams and, in support of the findings of prior work, encouraging greater community involvement in research studies [3, 5, 56]. We must give communities decision-making power while developing technologies that impact them, so they are not treated as token minorities and are able to shape our work [56]. This means ensuring they have the power to change the design, implementation, and outcome of our work instead of merely including their perspectives towards the end of the developmental cycle.

5.4.4 Changing How We Define Success. In our study, many participants linked the humanization of their systems to increased engagement and better user experiences overall, echoing the sentiments of prior work [10, 75]. However, the humanization of technological agents raises ethical concerns toward the trustworthiness of such systems. For example, humanized chatbots can blur the boundary between the real and the virtual worlds, prompting people to trust misinformation [41] or even to perceive these bots as genuinely human. This distortion has already culminated in tragedy, as in the case of a teenager who developed a fraught, “inappropriate” relationship with a chatbot, became increasingly isolated from his family, and ultimately took his own life ¹. Furthermore, this humanization may be unnecessary for certain groups of users such as autistic people, who may prefer interacting with simpler systems [57]. This leaves us with an important ethical consideration—do the benefits of humanizing technologies for certain users outweigh the harms this approach may cause to others? Perhaps we need to consider success metrics that go beyond user satisfaction, and think more deeply about the impact of our work on the lives of people we may inadvertently be marginalizing from solely focusing on the needs of the majority.

6 Conclusion

While there is a growing interest in humanizing AI agents such as robots and chatbots, our findings reveal this humanization is often done at the expense of autistic people’s preferences and broader inclusion in society. The creators of human-like AI who participated in our study displayed a clear preference for implementing neurotypical standards of communication in their work and often did not consider how this interpretation of humanness may impact their end-users perceptions of neurodivergent individuals. This is especially concerning as increasing diversity in the virtual world, and more positive representations of autism in the media have helped reduce explicit biases [31, 45, 51] among participants in previous studies, suggesting the impact more positive representations of diverse communication styles may have on autism inclusion in society. However, traits and behaviors commonly preferred by autistic people were often compared to non-human entities such as

¹<https://www.cnn.com/2024/10/30/tech/teen-suicide-character-ai-lawsuit/index.html>

robots, illustrating the ways in which communication norms explicitly dehumanize autistic people. We encourage a deeper inclusion of community perspectives, a more thorough integration of ethics, clearly defined standards, and accountability for organizations in upholding these standards to mitigate similar biases in future work.

References

- [1] Juby Abraham, George Joseph Cherian, and N Jayapandian. 2023. Systematic Review on Humanizing Machine Intelligence and Artificial Intelligence. In *2023 Second International Conference on Electronics and Renewable Systems (ICEARS)*. IEEE, 1092–1097.
- [2] Philip Baillargeon, Jina Yoon, and Amy Zhang. 2024. Who Puts the "Social" in "Social Computing"? Using A Neurodiversity Framing to Review Social Computing Research. *arXiv preprint arXiv:2410.15527* (2024).
- [3] Andrew Begel, John Tang, Sean Andrist, Michael Barnett, Tony Carbary, Piali Choudhury, Edward Cutrell, Alberto Fung, Sasa Junuzovic, Daniel McDuff, et al. 2020. Lessons learned in designing AI for autistic adults. In *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*. 1–6.
- [4] Katy Johanna Benson. 2023. Perplexing presentations: Compulsory neuronormativity and cognitive marginalisation in social work practice with autistic mothers of autistic children. *The British Journal of Social Work* 53, 3 (2023), 1445–1464.
- [5] Abeba Birhane and Fred Cummins. 2019. Algorithmic injustices: Towards a relational ethics. *arXiv preprint arXiv:1912.07376* (2019).
- [6] Cynthia Breazeal. 2003. Toward sociable robots. *Robotics and Autonomous Systems* 42, 3–4 (2003), 167–175. doi:10.1016/S0921-8890(02)00373-1
- [7] Priska Breves. 2020. Reducing outgroup bias through intergroup contact with non-playable video game characters in VR. *Presence* 27, 3 (2020), 257–273.
- [8] Maarten Buyl, Christina Cociancig, Cristina Frattone, and Nele Roekens. 2022. Tackling Algorithmic Disability Discrimination in the Hiring Process: An Ethical, Legal and Technical Analysis. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FACCT '22). Association for Computing Machinery, New York, NY, USA, 1071–1082. doi:10.1145/3531146.3533169
- [9] Amandine Catala, Luc Faucher, and Pierre Poirier. 2021. Autism, epistemic injustice, and epistemic disablement: A relational account of epistemic agency. *Synthese* 199, 3 (2021), 9013–9039.
- [10] Shalini Chandra, Anuragini Shirish, and Shirish C Srivastava. 2022. To be or not to be... human? Theorizing the role of human-like competencies in conversational artificial intelligence agents. *Journal of Management Information Systems* 39, 4 (2022), 969–1005.
- [11] Elizabeth F. Churchill and Mikael Wiberg. 2024. Speaking Human—Beyond Humanizing AI. *Interactions* 31, 3 (2024), 6–7. doi:10.1145/3658694
- [12] Catherine J Crompton, Martha Sharp, Harriet Axbey, Sue Fletcher-Watson, Emma G Flynn, and Danielle Ropar. 2020. Neurotype-matching, but not being autistic, influences self and observer ratings of interpersonal rapport. *Frontiers in psychology* 11 (2020), 586171.
- [13] Andrea Cuadra, Maria Wang, Lynn Andrea Stein, Malte F. Jung, Nicola Dell, Deborah Estrin, and James A. Landay. 2024. The Illusion of Empathy? Notes on Displays of Emotion in Human-Computer Interaction. In *CHI '24: Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. ACM, 446:1–446:18. doi:10.1145/3613904.3642336
- [14] Hana Darling-Wolf and Elizabeth Patitsas. 2024. "Not my Priority:" Ethics and the Boundaries of Computer Science Identities in Undergraduate CS Education. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–28.
- [15] Priya Dhawka and Sayamindu Dasgupta. 2025. The Social Construction of Visualizations: Practitioner Challenges and Experiences of Visualizing Race and Gender Demographic Data. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, New York, NY, USA. doi:10.1145/3706598.3713243
- [16] Cheryl L Dickter and Joshua A Burk. 2021. The effects of contact and labeling on attitudes towards individuals with autism. *Journal of Autism and Developmental Disorders* 51, 11 (2021), 3929–3936.
- [17] Virginia Eatough and Jonathan A Smith. 2017. Interpretative phenomenological analysis. *The Sage handbook of qualitative research in psychology* (2017), 193–209.
- [18] Severin Engelmann, Madiha Zahrah Choksi, Angelina Wang, and Casey Fiesler. 2024. Visions of a Discipline: Analyzing Introductory AI Courses on YouTube. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 2400–2420.
- [19] Soraya Fallah. 2017. *Giving a Voice to an Invisible Population: The Experiences of Families of Students with Disabilities from the Middle East, North Africa, and Southwest Asia (MENASWA) in the United States Special Education System*. Ph.D. Dissertation. California State University, Northridge.
- [20] Ali Fenwick and Gabor Molnar. 2022. The importance of humanizing AI: using a behavioral lens to bridge the gaps between humans and machines. *Discover Artificial Intelligence* 2 (2022), 14. doi:10.1007/s44163-022-00030-8
- [21] Louie Giray. 2024. The Problem with False Positives: AI Detection Unfairly Accuses Scholars of AI Plagiarism. *The Serials Librarian* (2024), 1–9.
- [22] Kate Glazko, Yusuf Mohammed, Ben Kosa, Venkatesh Potluri, and Jennifer Mankoff. 2024. Identifying and Improving Disability Bias in GPT-Based Resume Screening. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro, Brazil) (FACCT '24). Association for Computing Machinery, New York, NY, USA, 687–700. doi:10.1145/3630106.3658933
- [23] Josh Guberman. 2023. ActuallyAutistic Twitter as a Site for Epistemic Resistance and Crip Futurity. *ACM Trans. Comput.-Hum. Interact.* 30, 3, Article 38 (June 2023), 34 pages. doi:10.1145/3569891
- [24] Josh Guberman and Oliver Haimson. 2023. Not robots; Cyborgs—Furthering anti-ableist research in human-computer interaction. *First Monday* (2023).
- [25] Allyson I. Hauptman, Wen Duan, and Nathan J. McNeese. 2022. The Components of Trust for Collaborating With AI Colleagues. In *Companion Publication of the 2022 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '22 Companion)*. ACM, 72–75. doi:10.1145/3500868.3559457
- [26] Mohamad Hassan Fadi Hijab, Bilikis Banire, Joselia Neves, Marwa Qaraq, Achraf Othman, and Dena Al-Thani. 2024. Co-design of technology involving autistic children: A systematic literature review. *International Journal of Human-Computer Interaction* 40, 22 (2024), 7498–7516.
- [27] Jess Hohenstein, Rene F Kizilcec, Dominic DiFranzo, Zhila Aghajari, Hannah Mieczkowski, Karen Levy, Mor Naaman, Jeffrey Hancock, and Malte F Jung. 2023. Artificial intelligence in communication impacts language and social relationships. *Scientific Reports* 13, 1 (2023), 5487.
- [28] Kenneth Holstein, Jennifer Wortman Vaughan, Hal Daumé III, Miroslav Dudík, and Hanna Wallach. 2019. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need?. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM, New York, NY, USA. doi:10.1145/3290605.3300830
- [29] Philippa L Howard and Felicity Sedgewick. 2021. 'Anything but the phone!': Communication mode preferences in the autism community. *Autism* 25, 8 (2021), 2265–2278.
- [30] Cameron R. Jones and Benjamin K. Bergen. 2024. Does GPT-4 pass the Turing test?. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, 5183–5210. <https://aclanthology.org/2024.naacl-long.290>
- [31] Desiree R Jones, Kilee M DeBrabander, and Noah J Sasson. 2021. Effects of autism acceptance training on explicit and implicit biases toward autism. *Autism* 25, 5 (2021), 1246–1261.
- [32] Jaber Kamali, Muhammet Furkan Alpat, and Aras Bozkurt. 2024. AI ethics as a complex and multifaceted challenge: decoding educators' AI ethics alignment through the lens of activity theory. *International Journal of Educational Technology in Higher Education* 21, 1 (2024), 1–20.
- [33] Edward B. Kang. 2023. On the Praxes and Politics of AI Speech Emotion Recognition. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FACCT '23). Association for Computing Machinery, New York, NY, USA, 455–466. doi:10.1145/3593013.3594011
- [34] Katharina Klein and Luis F Martinez. 2023. The impact of anthropomorphism on customer satisfaction in chatbot commerce: an experimental study in the food sector. *Electronic commerce research* 23, 4 (2023), 2789–2825.
- [35] Kuan-Chou Ko, Chia-Wei Lin, and Zhi-Jun Yeh. 2024. Chatbot anthropomorphism might not be the design for all: examining responses to anthropomorphized chatbots by autistic individuals. *Marketing Letters* (2024), 1–13.
- [36] Nour Kteily and Emile Bruneau. 2017. Backlash: The politics and real-world consequences of minority group dehumanization. *Personality and Social Psychology Bulletin* 43, 1 (2017), 87–104.
- [37] Jennifer Lee and Karthick Ramakrishnan. 2020. Who counts as Asian. *Ethnic and Racial Studies* 43, 10 (2020), 1733–1756. doi:10.1080/01419870.2020.1734229
- [38] Mylène Legault, Amandine Catala, and Pierre Poirier. 2024. Breaking the stigma around autism: moving away from neuronormativity using epistemic justice and 4E cognition. *Synthese* 204, 3 (2024), 84.
- [39] Sharon Lindberg, Chiara Rossitto, Ola Knutsson, Petter Karlström, and Sirkku Männikkö Barbutiu. 2023. Doing Good Business? Design Leaders' Perspectives on Ethics in Design. *Proc. ACM Hum.-Comput. Interact.* 7, GROUP (2023), 2:1–2:22. doi:10.1145/3633067
- [40] Anna Ma, Elizabeth Patitsas, and Jonathan Sterne. 2023. You Sound Depressed: A Case Study on Sonde Health's Diagnostic Use of Voice Analysis AI. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FACCT '23). Association for Computing Machinery, New York, NY, USA, 639–650. doi:10.1145/3593013.3594032
- [41] Takuya Maeda and Anabel Quan-Haase. 2024. When Human-AI Interactions Become Parasocial: Agency and Anthropomorphism in Affective Design. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1068–1077.
- [42] Shelley McKeown and John Dixon. 2017. The "contact hypothesis": Critical reflections and future directions. *Social and Personality Psychology Compass* 11, 1 (2017), e12295.

- [43] Thomas Mildner, Orla Cooney, Anna-Maria Meck, Marion Bartl, Gian Luca Savino, Philip R. Doyle, Diego Garaialde, Leigh Anne Clark, John Sloan, Nina Wenig, Rainer Malaka, and Jasmin Niess. 2024. Listening to the Voices: Describing Ethical Caveats of Conversational User Interfaces According to Experts and Frequent Users. In *CHI Conference on Human Factors in Computing Systems (CHI '24) Proceedings*. ACM, New York, NY, USA. doi:10.1145/3613904.3642542
- [44] Damian EM Milton. 2012. On the ontological status of autism: The 'double empathy problem'. *Disability & society* 27, 6 (2012), 883–887.
- [45] Gloria Mittmann, Beate Schrank, and Verena Steiner-Hofbauer. 2024. Portrayal of autism in mainstream media—a scoping review about representation, stigmatisation and effects on consumers in non-fiction and fiction media. *Current Psychology* 43, 9 (2024), 8008–8017.
- [46] Agnieszka Mulak and Mikołaj Henryk Winiewski. 2021. Virtual contact hypothesis: Preliminary evidence for intergroup contact hypothesis in interactions with characters in video games. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace* 15, 4 (2021).
- [47] San Murugesan. 2025. The Turing Test at 75: Its Legacy and Future Prospects. *IEEE Intelligent Systems* 40, 1 (2025), 20–24.
- [48] Cliodhna O'Connor, Judith Burke, and Brendan Rooney. 2020. Diagnostic disclosure and social marginalisation of adults with ASD: Is there a relationship and what mediates it? *Journal of Autism and Developmental Disorders* 50, 9 (2020), 3367–3379.
- [49] Elizabeth Levy Paluck, Seth A Green, and Donald P Green. 2019. The contact hypothesis re-evaluated. *Behavioural Public Policy* 3, 2 (2019), 129–158.
- [50] Amy Pearson and Kieran Rose. 2021. A conceptual analysis of autistic masking: Understanding the narrative of stigma and the illusion of choice. *Autism in Adulthood* 3, 1 (2021), 52–60.
- [51] Tabitha C Peck, Sofia Seinfeld, Salvatore M Aglioti, and Mel Slater. 2013. Putting yourself in the skin of a black avatar reduces implicit racial bias. *Consciousness and cognition* 22, 3 (2013), 779–787.
- [52] Malak Qbilat, Ana Iglesias, and Tony Belpaeme. 2021. A proposal of accessibility guidelines for human-robot interaction. *Electronics* 10, 5 (2021), 561.
- [53] Elizabeth M Radulski. 2022. Conceptualising autistic masking, camouflaging, and neurotypical privilege: Towards a minority group model of neurodiversity. *Human Development* 66, 2 (2022), 113–127.
- [54] Daniel J Ricks and Mark B Colton. 2010. Trends and considerations in robot-assisted autism therapy. In *2010 IEEE international conference on robotics and automation*. IEEE, 4354–4359.
- [55] Naba Rizvi, Andrew Begel, and Hala Annabi. 2021. Inclusive Interpersonal Communication Education for Technology Professionals.. In *AMCIS*.
- [56] Naba Rizvi, William Wu, Mya Bolds, Raunak Mondal, Andrew Begel, and Imani NS Munyaka. 2024. Are Robots Ready to Deliver Autism Inclusion?: A Critical Review. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–18.
- [57] Ben Robins, Kerstin Dautenhahn, and Janek Dubowski. 2006. Does appearance matter in the interaction of children with autism with a humanoid robot? *Interaction Studies. Social Behaviour and Communication in Biological and Artificial Systems* 7, 3 (2006), 479–512.
- [58] Bonnie Ruberg and Spencer Ruelos. 2020. Data for queer lives: How LGBTQ gender and sexuality identities challenge norms of demographics. *Big Data & Society* 7, 1 (2020), 2053951720933286. doi:10.1177/2053951720933286
- [59] Sofia Sahab, Jawad Haqbeen, Rafik Hadfi, Takayuki Ito, Richard Eke Imade, Susumu Ohnuma, and Takuya Hasegawa. 2024. E-contact facilitated by conversational agents reduces interethnic prejudice and anxiety in Afghanistan. *Communications Psychology* 2, 1 (2024), 22.
- [60] Ayden I. Scheim and Greta R. Bauer. 2019. The Intersectional Discrimination Index: Development and validation of measures of self-reported enacted and anticipated discrimination for intercategory analysis. *Social Science & Medicine* 226 (2019), 225–235. doi:10.1016/j.socscimed.2019.03.039
- [61] Ben Sheehan, Hyun Seung Jin, and Udo Gottlieb. 2020. Customer service chatbots: Anthropomorphism and adoption. *Journal of Business Research* 115 (2020), 14–24.
- [62] Heung-Yeung Shum, Xiao-Dong He, and Di Li. 2018. From Eliza to Xiaolce: Challenges and Opportunities with Social Chatbots. *Frontiers of Information Technology & Electronic Engineering* 19, 1 (2018), 10–26. doi:10.1631/FITEE.1700826
- [63] Jonathan A. Smith. 2004. Reflecting on the development of interpretative phenomenological analysis and its contribution to qualitative research in psychology. *Qualitative Research in Psychology* 1, 1 (2004), 39–54. doi:10.1191/1478088704qp0040a
- [64] David Smukler. 2005. Unauthorized minds: how "theory of mind" theory misrepresents autism. *Mental Retardation* 43, 1 (2005), 11–24.
- [65] Katta Spiel, Christopher Frauenberger, Os Keyes, and Geraldine Fitzpatrick. 2019. Agency of Autistic Children in Technology Research—A Critical Literature Review. *ACM Trans. Comput.-Hum. Interact.* 26, 6, Article 38 (Nov. 2019), 40 pages. doi:10.1145/3344919
- [66] Katta Spiel, Oliver L. Haimson, and Danielle Lottridge. 2019. How to do better with gender on surveys: a guide for HCI researchers. *Interactions* 26, 4 (2019), 62–65. doi:10.1145/3338283
- [67] Jeff Stanley, Ronna ten Brink, Alexandra Valiton, Trevor Bostic, and Becca Scollan. 2022. Chatbot accessibility guidance: a review and way forward. In *Proceedings of Sixth International Congress on Information and Communication Technology: ICICT 2021, London, Volume 3*. Springer, 919–942.
- [68] Jennifer L Stevenson, Bev Harp, and Morton Ann Gernsbacher. 2011. Infantilizing Autism. *Disability studies quarterly: DSQ* 31, 3 (2011).
- [69] Helen Taylor, Barry Ingham, David Mason, Tracy Finch, Colin Wilson, Clare Scarlett, Sebastian Moss, Carole Buckley, Anna Urbanowicz, Dora Raymaker, et al. 2023. Co-design of an NHS primary care health check for autistic adults. *Autism* 27, 4 (2023), 1079–1091.
- [70] Alan M. Turing. 1950. Computing Machinery and Intelligence. *Mind* 59, 236 (1950), 433–460. doi:10.1093/mind/LIX.236.433
- [71] Alan M Turing. 2009. *Computing machinery and intelligence*. Springer.
- [72] U.S. Department of Health and Human Services. 2018. *HHS Implementation Guidance on Data Collection Standards for Race, Ethnicity, Sex, Primary Language, and Disability Status*. Technical Report. <https://aspe.hhs.gov/datacncl/standards> Originally published in 2011.
- [73] Ibo Van de Poel and Lambèr Royakkers. 2023. *Ethics, technology, and engineering: An introduction*. John Wiley & Sons.
- [74] Martine Van Driel, John Vines, Belen Barros Pena, and Nelya Kotevko. 2023. Understanding autistic adults' use of social media. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW2 (2023), 1–23.
- [75] Michelle ME Van Pinxteren, Mark Pluymaekers, and Jos GAM Lemmink. 2020. Human-like communication in conversational agents: a literature review and research agenda. *Journal of Service Management* 31, 2 (2020), 203–225.
- [76] Joseph Weizenbaum. 1966. ELIZA—A Computer Program for the Study of Natural Language Communication Between Man and Machine. *Commun. ACM* 9, 1 (1966), 36–45. doi:10.1145/365153.365168
- [77] Rua M Williams. 2021. I, Misfit: Empty Fortresses, Social Robots, and Peculiar Relations in Autism Research. *Techné: Research in Philosophy & Technology* 25, 3 (2021).
- [78] Liliane Cambraia Windsor, Clay Shorkey, and DuWayne Battle. 2015. Measuring Student Learning in Social Justice Courses: The Diversity and Oppression Scale. *Journal of Social Work Education* 51, 1 (2015), 58–71. doi:10.1080/10437797.2015.977133
- [79] Sonny Jane Wise. 2023. *We're All Neurodiverse: How to Build a Neurodiversity-Affirming Future and Challenge Neuronormativity*. Jessica Kingsley Publishers.
- [80] Maciej Wodźniński, Marcin Rządęczka, Anastazja Szul, Marta Sokół, and Marcin Moskaiewicz. 2024. Visual Stereotypes of Autism Spectrum in DALL-E, Stable Diffusion, SDXL, and Midjourney. *arXiv preprint arXiv:2407.16292* (2024).
- [81] Jintang Xue, Yun-Cheng Wang, Chengwei Wei, Xiaofeng Liu, Jonghye Woo, and C.-C. Jay Kuo. 2024. Bias and Fairness in Chatbots: An Overview. *APSIPA Transactions on Signal and Information Processing* 13 (2024), e102. doi:10.1561/116.00000064
- [82] Theresa Zobel, Hendrik Steinbeck, and Christoph Meinel. 2023. Towards Inclusive Education: A Review of and Guide to Accessible Features in Chatbots for MOOCs. In *2023 IEEE Learning with MOOCs (LWMOOCs)*. 1–5. doi:10.1109/LWMOOCs58322.2023.10306062
- [83] Annuska Zolyomi and Jaime Snyder. 2021. Social-emotional-sensory design map for affective computing informed by neurodivergent experiences. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–37.

A Appendix

Note: in attempts to avoid response bias, we have included extra questions in our surveys and interviews. These are denoted by “*” if included in the Appendix.

A.1 Surveys

A.1.1 Survey 1: Demographic Questions.

- (1) What is your job official title?

- (2) What industry do you work in?
 - Academia
 - Industry
 - Other: _____
- (3) Please specify the gender with which you most closely identify.
 - Woman
 - Man

- Non-binary
 - Prefer not to answer
 - Prefer to self-describe: _____
- (4) **What is your age?**
- 18–22
 - 23–30
 - 31–40
 - 41–50
 - 51–60
 - Over 60
 - Prefer not to answer
- (5) **Please specify your race/ethnicity.**
- White
 - Hispanic or Latino
 - Black or African American
 - Native American or American Indian
 - South Asian
 - Southwest Asian or North African
 - East Asian
 - Southeast Asian or Pacific Islander
 - Prefer to self-describe: _____
 - Prefer not to answer
- (6) **Parental Status ***
- I have or had children that I currently support and/or raise by myself, with a partner, or as part of a team or family
 - I do not have children that I currently support and/or raise by myself, with a partner, or as part of a team or family
 - Prefer to self-describe: _____
 - Prefer not to answer
- (7) **Marital status ***
- Single
 - Common Law Marriage
 - Married
 - Separated
 - Divorced
 - Widowed
 - Legal
 - Prefer to self-describe: _____
 - Prefer not to answer
- (8) **6-Question Disability Questionnaire**
- (1) Are you deaf, or do you have serious difficulty hearing?
 - (2) Are you blind, visually impaired, or do you have serious difficulty seeing, even when wearing glasses?
 - (3) Because of a physical, mental, or emotional condition, do you have serious difficulty concentrating, remembering, or making decisions?
 - (4) Do you have serious difficulty walking or climbing stairs?
 - (5) Do you have difficulty dressing or bathing?
 - (6) Because of a physical, mental, or emotional condition, do you have difficulty doing errands alone such as visiting a doctor's office or shopping? (15 years old or older)
- (9) **Sexual Orientation ***
- Asexual
 - Bisexual
 - Pansexual
 - Gay
 - Heterosexual

- Lesbian
 - Queer
 - Prefer to self-describe: _____
 - Prefer not to answer
- (10) **What is the highest degree or level of school you have completed?**
- Some high school credit, no diploma or equivalent
 - Less than high school degree
 - High school graduate (high school diploma or equivalent including GED)
 - Some college but no degree
 - Associate's degree
 - Bachelor's degree
 - Advanced degree (e.g., Master's, doctorate)
 - Prefer not to answer
- (11) **Which one of the following includes your total HOUSEHOLD income for last year, before taxes?**
- Less than \$10,000
 - \$10,000 to under \$20,000
 - \$20,000 to under \$30,000
 - \$30,000 to under \$40,000
 - \$40,000 to under \$50,000
 - \$50,000 to under \$65,000
 - \$65,000 to under \$80,000
 - \$80,000 to under \$100,000
 - \$100,000 to under \$125,000
 - \$125,000 to under \$150,000
 - \$150,000 to under \$200,000
 - \$200,000 or more
 - Prefer not to answer

A.1.2 Survey 2.

- (1) **Imagine your team switches to a remote-optional work environment and now allows your team members to have their videos off during meetings. How do you feel about this new policy?**
- Strongly Disagree
 - Somewhat Disagree
 - Neither agree nor disagree
 - Somewhat Agree
 - Strongly Agree
- (2) **Imagine your team does not allow people to have conversations in the open work spaces, and requires them to use designated spaces. How do you feel about this change in work policies?**
- Strongly Disagree
 - Somewhat Disagree
 - Neither agree nor disagree
 - Somewhat Agree
 - Strongly Agree

A.2 Interviews

A.2.1 Interview 1.

1. Tell me about a recent project you worked on where you built an AI system that others would potentially use.
 - a. When did you know you've reached the end of the design process for your system?

- i. When the current assignment is complete.
- b. When did you know you've reached the end of the design process for your Robot/chatbot/AI agent?
- c. What was the overall purpose or objective for your system?
- d. How did you decide the physical characteristics (including UI) of the bot?
 - i. What do you think would happen if you chose something else?
- e. How did you decide the behavioral and conversational characteristics of the bot?
 - i. What do you think would happen if you chose something else?
2. What are the metrics you use for evaluating the performance of your system?
3. How did you know that your work was effective? (ask if they created a solution to a problem, otherwise skip).
 - a. Can you tell me about a time when a user interacted with your system and got frustrated or upset?
 - b. What are some things that may make a user frustrated or not want to interact with your system?
4. Can you walk me through how or why you chose this specific mode of user interaction to reach your objective or solve your problem?
5. What real-world interactions can you think of that resemble your user interaction?
 - a. If they say no, ASK - Can you think of any other interaction that might be similar to the interaction you developed?
 - b. What are the identities of the people in these interactions?
 - c. Explain the interaction for me.
 - i. What would it look like if one of the people was autistic/ADHD/dyslexic?
 - ii. How would the interaction change?
6. What behavioral or communicational changes could you make to the bot that would impact the user's interaction with the bot?
 - a. How would that be different with a person who is funny?
 - b. How would that impact the user's interactions with folks who are and aren't funny/don't use humor in interactions?
7. What conclusions would researchers who are building off of your work reach or assume about human interactions, behavior, or AI? What about end-users? What conclusions might they reach?
 - a. What are the best practices in bot design you have learned in your work?
 - b. How do end-users' interactions differ from those with other bots?
8. In your opinion, should we design chatbots, robots, or other AI agents to meet the needs of various community groups?
 - a. Who are the people most likely to use your chatbot?
 - b. What would happen if someone from [specific group] used it?*
 - c. What would happen if you adjusted the bot to work with multiple user groups? Or specific user groups?*
 - d. What kind of accessibility features did you consider in your chatbot?*

- e. Should we design for people whose mental and neurological abilities may be considered atypical such as folks with autism, ADHD, and dyslexia?
 - i. Limitations?
 - f. Should we design for different age groups?*
9. Is there anything else you'd like to say or add in regards to your experience building AI entities?

A.2.2 Interview 2.

1. How would you describe your current professional position?
2. In your opinion, what experiences from your life greatly impacted or led you to your interest in your current career?
3. What made you stay or continue to pursue your current career?
4. Tell me about a time when your personal background overlapped with your current position or your work.

Intersectionality Questions

1. Has an ism (e.g., racism, sexism, ableism, homophobia) ever impacted you during your computer science journey?
 - a. Has an ism ever impacted you in your life?
 - b. If no, ask: At work, do you interact with a number of people who have a similar racial, ethnic, or gender background as you on a regular basis?
2. In your opinion, what is intersectionality?
 - a. Has it ever impacted your work or research? If so, how?
3. Tell me about a time when you used or experienced intersectionality in your work/research.*
 - a. If non-tech, let them provide it. If research-related, ask: Do you have any examples related to your current position?
4. Tell me about a time when you used or experienced a push for inclusivity in your work/research.*
5. Tell me a time when intersectionality has impacted you at work.*
6. Have you had any training in working with or building tools that support inclusivity?
 - a. If so, what did the training look like? Can you describe it?
 - b. Did it help you?
7. Have you had any support in working with or building tools that support inclusivity?
 - a. If so, what did the support look like? Can you describe it?
 - b. Did it help you?

Situational Awareness

1. Suppose your new colleague shows up wearing headphones to a research conference and continues wearing them even while presenting their work. How would you approach a conversation with them?
 - a. What if your coworker mentions they're wearing it because the room is too loud?
2. Imagine that your company celebrates its anniversary during Ramadan. Ramadan is a Muslim holiday where Muslims refrain from eating and drinking from sunrise to sunset. This year is their 50th anniversary, and the plan is to have a week full of free buffet-style lunch and brunch banquets over the weekend. However, about three coworkers have expressed concern that they won't be able to fully participate due to the way things are planned to happen.

- a. What is the problem from the coworkers' viewpoint?*
 - b. How might this problem be solved?*
3. Suppose your research assistant doodles during meetings, and your supervisor approaches you about the lack of professionalism exhibited by your RA. How would you address this situation?
4. Max and Alex are working on a research project and are expected to provide feedback to each other to ensure the project stays on track. [Share a conversation transcript.]
 - a. Why is Alex responding this way?
 - b. Why is Max responding this way?
 - c. How might you respond as their supervisor?