Distinguishing Emotion AI: Factors Shaping Perceptions Including Input Data, Emotion Data Recipients, and Identity

Alexis Shore Ingber
University of Michigan
Ann Arbor, USA
ingber@umich.edu

Oliver L Haimson
University of Michigan
Ann Arbor, USA
haimson@umich.edu

Nazanin Andalibi
University of Michigan
Ann Arbor, USA
andalibi@umich.edu

Abstract

Emotion artificial intelligence (emotion AI), which claims to infer data subjects' emotions, is used across a range of high-stakes contexts in the U.S. Examining data subjects' perceptions of emotion AI is important given the impact its use can have on their access to opportunities, among other consequences. Such inquiry requires an empirical understanding of data subjects' attitudes about emotion AI compared to AI that does not infer emotions, and differences across potential 1) input data (e.g., facial data, speech data), and 2) recipients of its outputs (i.e., emotion data). To these ends, using a contextual integrity lens, we conducted a survey (n = 742) with a representative U.S. sample, oversampling for individuals of minority groups. Results demonstrate data subjects' attitudes toward emotion AI are significantly more negative than AI that does not detect emotion. Further, we find significant contrasts in perceptions of specific input data across gender, race, and disability status, particularly for neurodiverse individuals. We also find participants express concern with *all* emotion data recipients, with higher concern for their data being accessed by recipients with large power asymmetries. Based on these findings, we argue emotion AI should not be deployed without a comprehensive assessment of its implications across demographic groups, yielding implications for policymakers and developers.

CCS Concepts

• **Security and privacy → Privacy protections**; • **Social and professional topics → User characteristics**.

Keywords

emotion recognition, affective computing, input data, AI, contextual integrity, data recipients

ACM Reference Format:
Alexis Shore Ingber, Oliver L Haimson, and Nazanin Andalibi. 2025. Distinguishing Emotion AI: Factors Shaping Perceptions Including Input Data, Emotion Data Recipients, and Identity. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25), June 23–26, 2025, Athens, Greece*. ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3715275.3732034

FAccT '25, Athens, Greece

ACM ISBN 979-8-4007-1482-5/25/06
https://doi.org/10.1145/3715275.3732034

1 Introduction

As artificial intelligence (AI) proliferates across industries, developers and investors are exploring new ways to gain insights and make predictions about human behavior. A type of AI undergoing significant growth is that claiming to infer emotion and other affective qualities, which we refer to as emotion AI.[1] The market for emotion AI is estimated to be worth $42.9 billion USD by 2027 [68] and is already actively deployed across industries such as the workplace, hiring, healthcare, education, and social media [76]. Developers use a plethora of input data to train emotion AI systems such as facial expressions [33], voice patterns [59], spoken words [50], eye movement [58], gait [65] and electroencephalogram (EEG) data [6]. Of note, the surge of interest in emotion AI has been met with a wave of criticism regarding its accuracy [10], representativeness of training data [75], and underlying theoretical foundations [10, 107]. Additionally, empirical work has demonstrated emotion AI's harms related to privacy [4, 23, 96], mental integrity [76], and imposed emotional labor [14, 96].

Though AI developments have raised enthusiasm, the same should not be assumed for emotion AI.[2] In fact, recent work has demonstrated the U.S. public's discomfort with emotion AI across several contexts of deployment (e.g., healthcare, job interviews, social media) [5]. Distinguishing between AI and emotion AI is crucial given the unique sensitivity of emotion data and need for emotional privacy [4, 106]. Moreover, existing research on data subjects' perceptions of emotion AI has not examined 1) the sensitivity of *input data* used to train emotion AI, and 2) the relative negative consequences associated with *emotion data recipients* who may use emotion AI's outputs (i.e., emotion data). Nissenbaum's theory of contextual integrity [81] would argue these two factors contribute to defining the appropriateness of how emotion data flows. *Contextual integrity* is a framework invoking multiple dimensions of an information flow (e.g., data input, data receivers) to determine when such might be a violation of privacy [81]. As such, contextual integrity is a suitable framework to leverage in emotion AI's empirical analysis.

Once emotion AI systems process input data, the resulting emotion data can be delivered to a range of actors for future use (i.e., data recipients). These recipients may range from law enforcement officers [98] to healthcare professionals [79] who may see the benefit of capturing individuals' emotions. The variety in emotion AI's input data and recipients engenders a spectrum of implications

[1]Emotion AI, a type of affective computing systems [86], is also referred to as affect recognition, emotion sensing, affect sensing, and passive sensing, among others.
[2]For purposes of this paper, "AI" without a corresponding mention of "emotion" refers to AI as algorithmic systems that can perform tasks which would usually require human intelligence. We define emotion AI as algorithmic systems promising to automatically detect emotions and moods based on various kinds of data.

for the deployment of emotion AI, particularly across gender, race, and disability status given existing biases emotion AI may perpetuate along these dimensions [66, 91, 123]. Indeed, scholars have urged for a focus on marginalized individuals when undertaking a contextual approach to individuals' perceptions of information flows [118]. Nissenbaum's [81] theory serves as both the theoretical and methodological foundation of this paper, offering a robust framework for understanding the implications of emotion AI's deployment. Additionally, this paper expands upon existing typologies of information sensitivity (e.g., [78]). Empirical examination of the contextual variation across emotion AI deployment can help inform emotion AI regulation and development. This paper is guided by the following research questions:

RQ1: How might U.S. data subjects' attitudes towards AI differ from their attitudes towards emotion AI?

RQ2a: How might U.S. data subjects' perceptions of sensitivity differ across input data to emotion AI?

RQ2b: How might perceptions of sensitivity across input data to emotion AI differ across gender, race, and disability status?

RQ3: How might U.S. data subjects' perceptions of sensitivity differ across emotion data recipients?

To answer these research questions, we conducted a nationally representative survey of the U.S. population (stratified by gender, race, and age) in addition to oversampling transgender individuals, people with disabilities, and people of color (POC) (n = 742).

We find U.S. data subjects to have significantly more negative attitudes toward emotion AI than AI, which quantitatively distinguishes emotion AI from AI more broadly. Further, results reveal participants perceived a high degree of sensitivity across *all* input data types to emotion AI, with precise location, digital communications, and spoken words being the most sensitive. Moreover, disabled individuals and gender minorities (i.e., transgender and non-binary people) perceived particular input data as significantly more sensitive than individuals who are not disabled and cisgender men, respectively. Finally, we find U.S. data subjects believe there could be negative consequences when emotion data is received by *any* third-party actor. That being said, U.S. data subjects saw the least risk with people to whom they have personal relationships (e.g., friends, family) and were most concerned with data brokers (i.e., companies that collect and sell individuals' personal information) receiving their emotion data, followed by actors who may have a high stake in their future and access to opportunities (e.g., law enforcement, future employers).

Our results have implications for emotion AI policy and development in the U.S.. We argue for U.S. policymakers to approach future AI governance with recognition of *emotion AI* as distinct from AI and to create rules and restrictions around data *inputs* to emotion AI and *recipients* of emotion data. We also urge emotion AI developers to take action recognizing particular types of emotion AI input data as distinctly sensitive, particularly when their systems are being used by a diverse set of users. That disabled individuals and gender minorities perceived particular emotion AI data inputs as significantly more sensitive than their majority counterparts indicates a tension in including their data for more "inclusive" AI systems. Finally, participants' belief that negative consequences may arise upon *any* third party receipt of emotion data underscore the need to rethink permissible emotion data flows. Zooming out, our results foster an argument that future emotion AI research and development must foreground the perspectives and experiences of data subjects, prioritizing the contextual integrity of their information flows a priori.

2 Prior Work

2.1 Defining and Distinguishing Emotion AI

Although emotion AI has received scholarly critique from social scientists, these scholars have yet to formally corroborate the public's perceived differences between AI and emotion AI with the exception of a study conducted in Japan [67]. This study focused on emotion AI use in healthcare, finding Japanese people's attitudes related to emotion AI to be more negative than AI [67]. Moreover, Mantello et al. showed people do not trust and have privacy concerns over emotion AI's use in this context [67]. Other past work provides a series of independent analyses of attitudes toward AI and emotion AI, respectively. For instance, the U.S. public has been found to be more concerned than excited about the use of AI—as defined generally—in daily life [26], with these comfort levels differing across contexts of deployment [82]. These attitudes are similar to those reported about emotion AI [5]. U.S.-centric qualitative research has reported a negative view of emotion AI deployment in contexts such as social media (e.g., to inform well-being interventions) [4, 94] the workplace (e.g., to monitor worker productivity) [14, 23, 96], and hiring (e.g., to evaluate candidate fit during a job interview) [88]. Across contexts, data subjects' concern about the deployment of emotion AI has been rooted in its potential for inaccuracy and bias along with privacy and emotional labor harms [23, 38, 63, 95].

Despite studies of AI and emotion AI attitudes predominately existing independently, scholars have *conceptually* differentiated emotion AI from other forms of AI given the unique sensitivity of emotion data [4, 106]. In fact, a lack of ability to conceal emotions from technology runs counter to Erving Goffman's [35] argument that people alter their outward self-presentation in an effort to protect their privacy. Wachter [113] notes the deployment of emotion AI may inadvertently encourage individuals to alter their emotional displays in an attempt to give the system what they think it wants, as similarly noted by workers reflecting on the deployment of emotion AI [23]. In fact, Pyle et al. [88] found job-seekers to feel the use of emotion AI for interview evaluation was not only unjust, but would require the exertion of additional emotional labor. Part of this paper seeks to examine if the U.S. public maintains significantly different attitudes between emotion AI and AI more broadly defined. Such an examination lays the foundation for a contextualized analysis and regulation of emotion AI, as well as a higher level discussions of emotion AI's implications.

2.2 A Contextual Integrity Approach to Examining Emotion AI's Implications

Existing research on the public's attitudes toward AI has been critiqued for being overly general and lacking consideration of context [120]. A contextual integrity (CI) approach to emotion AI allows for a more detailed understanding of how people perceive the appropriateness of data that goes into emotion AI systems and to whom such data flows, as summarized in Table 1. CI argues information flows' appropriateness is dependent on the actors involved (i.e., senders, recipients), the information type/attribute, and the transmission principles established [81]. The transmission principle concerns the constraint surrounding the information flow (e.g., confidentiality). For example, while one may deem it appropriate to share health records with a doctor, they may feel less comfortable with this information flowing to a loan lender or romantic partner. In addition to perceiving a doctor as a more appropriate recipient of health data, one may also feel more comfortable with the stringent transmission principles to which healthcare professionals are obliged. A violation of contextual integrity happens when personal data moves against appropriate flows of information [81].

Table 1: Contextual Integrity Parameters Defining Emotion Data Flows

Parameter	Values
Sender	Emotion AI system
Subject	Data subject
Attribute	Precise location, Spoken words, Digital communications, Facial expression, Voice, Written words, Eye movement, Device usage, Non-facial biometrics, Body movement, Environment
Recipient	Employer, Data brokers, Future employers, Governmental actors, Law enforcement, Strangers on the Internet, Health insurance companies, Loan lenders, Current employers, Social media companies, Advertisers, Other tech companies (e.g., search engines, streaming services), Car insurance companies, Potential romantic partners, Retailers, Healthcare providers, Customer service representatives, Automotive companies, Educational instructors, My family, My romantic partner, My friends

Researchers have applied CI to understand the normative perspectives across a variety of information management scenarios. For instance, research has found older adults' information sharing practices were not only dependent on the perceived sensitivity of information—which differed across data types—but also *who* was going to receive the data and what they were going to do with it [30]. Additionally, researchers have used CI to understand acceptance of data sharing through COVID-mitigation apps, finding a significant influence of data type (e.g., test results, exposure status) and recipient (e.g., healthcare providers, employers) [28]. Similar findings have been reported in the context of smart personal assistants [18] learning analytics systems [71], and proctoring software [110]. This study applies CI to examine factors shaping data subjects' perspectives on emotion AI. Prioritizing ecological validity, we focus on data attributes and recipients because the transmission principles of emotion AI systems are predominately unknown to data subjects. Rather than evaluating the appropriateness of information flows that integrates each of the CI parameters simultaneously, this paper concentrates on understanding perceptions about emotion AI input data and emotion data recipients independently. This approach facilitates a deeper exploration of specific aspects of data transmission before attempting to assess the more complex interactions across all of CI's parameters.

2.2.1 What Goes In: The Sensitivity of Input Data to Emotion AI Systems. People across 10 different countries have reported concerns about AI given the the mass amount of personal input data needed for AI's operation [53]. Emotion AI is no different than AI in that it is fueled by different types of input data, representing CI's *attributes* parameter. For instance, developers of emotion AI have leveraged facial data [33], speech—both as spoken words[59] and written text [57]— eye movement [58], gait [65], physiological data such as EEG [6] and pulse [11] as input to emotion AI systems. Emotion AI developers have also built systems claiming to detect emotion through assessing patterns in data subjects' environment (e.g., location) [48, 97] or device usage [109]. These research explorations have been corroborated by emotion AI patent applications which propose a range of input data for emotion AI as deployed in the workplace [14] and mental healthcare [49].

Echoing the CI framework, input data types may engender different perceptions about emotion AI. Kozyreva et al. [56] found over half of their sample of U.S. data subjects to view collection and use of emails and private messages for algorithmic personalization as more unacceptable than that of location and browsing behavior. Outside of CI, previous research has typologized data, which could serve as input to AI systems, based on perceptions of sensitivity [21, 70, 78, 87]. Perceptions of sensitivity are critical given their impact on technology adoption and trustworthiness, especially in instances where the technology would attempt to force data disclosure [21].

2.2.2 Identity-based Variation in Perceived Input Data Sensitivity. In their argument for a contextual approach to understanding privacy, Wu et al. [118] push for a focus on marginalized subpopulations. McDonald & Forte [74] echo this sentiment in their proposal of vulnerability as a core concept in privacy theories such as CI to more accurately understand the privacy risks that technologies pose to marginalized individuals. In fact, it is possible sensitivity perceptions about input data to emotion AI may differ across demographic traits, particularly gender, race and disability status. Researchers have underscored attitudes toward AI are shaped by gender [9, 37, 44, 67] income [16, 66] and education level [16], with marginalized individuals maintaining more negative attitudes about AI. Given existing biases in emotion AI along the dimensions of gender, race and disability [66, 91, 123] which may result in discrimination [25], it is not surprising that evidence suggests these individuals may have unique concerns [88, 95] and discomfort [5] with emotion AI.

Specifically, we consider differences in perceptions across gender. With respect to input data, Kozyreva et al. [56] found cisgender women to be more concerned about the use of personal information as algorithmic input data than cisgender men. That being said, one paper found no significant difference between cisgender women and cisgender men on attitudes toward AI [108]. It is possible that incorporating a more nuanced definition of gender, inclusive of non-binary and transgender individuals, would provide more notable contrasts in data input sensitivity perceptions given the concerns these individuals have about AI [93] and emotion AI [5, 88].

In addition, given existing bias in AI across individuals with disabilities [115] which some fear extends to emotion AI [79], we consider how perceptions of emotion AI input data differ across disability status. In an interview study, disabled individuals placed blame on algorithmic training data for the algorithms' perpetuation of disability stereotypes [31]. Similar concerns were expressed by deaf and hard of hearing learners who were worried about the accuracy of collected emotion data [19]. We are specifically interested in individuals who are neurodiverse, as they may express and perceive emotions in non-normative ways [122]. Further, researchers have pointed to evidence where neurodiversity (e.g., autism) has been exploited to advance emotion AI [80]. Consequently, neurodiversity may be a particularly important disability to consider when designing and evaluating the ethics of emotion AI systems.

Finally, this paper considers the impact of race on perceptions about emotion AI input data given their unique experiences with AI [77]. Although POC may experience more harm as a result of subjection to emotion AI [25], recent work suggests they may still have positive attitudes toward AI [5, 83]. This tension may be explained by "Black optimism" [102]. Jared Sexton engages with Black optimism as a counterpoint to Afro-pessimism, where Black individuals perform acts of resistance as a form of agency against oppressive systems [102]. This is reflected in recent research, where people with darker skin tones had more positive attitudes toward AI than white people [83]. Distinguishing perceptions across racial groups is critical, as minoritized individuals' belief that algorithmic technologies are not designed with them in mind may lead to behavioral adjustments thought to achieve more optimal outcomes [77].

2.2.3 What Comes Out and Where it Flows: Emotion Data Recipients. In addition to data input types, CI argues the appropriateness of a information flows is dependent on the information recipient. Indeed, people are concerned with AI's potential to release personal information to multiple connected sources [53]. Zhang et al. found data subjects' perceptions about video monitoring systems depend on with whom the information is shared [121]. Across variation in input data, people perceive information flows as less sensitive if received by a friend than a trusted or unknown marketer [70]. Moreover, recent work which assessed acceptability of brain data collection and use across data recipients found government agencies to be the least acceptable [46]. Since emotion AI is intended to be deployed across contexts, its outputs may flow to a variety of actors. For example, emotion AI's deployment in healthcare may result in outputs flowing not only to healthcare providers, but also health insurance companies, data brokers, or personal caretakers [49]. This is significant when considering the emotion inferences that could be made by these recipients, which could profoundly impact individual livelihoods. In one experimental study, researchers found people imposed stricter boundaries around their personal data when they learned the system could generate inferences about them to personalize their newsfeed [7]. Specific to emotion AI, research has found job seekers—who may be subject to emotion AI during the interview process—to be concerned about *who else* may have access to their emotion data beyond a future employer [88]. In fact, individuals subject to emotion AI in the workplace question the very relevance of emotion data flowing to their employer [23]. For purposes of this study, we adopt Wirth et al's [117] definition of information sensitivity which is framed in terms of perceived negative consequences upon receipt by third party actors.

3 Methodology

We conducted a survey (n = 742) to answer our research questions. On average, the survey took participants 25 minutes (SD = 18.85). Participants were compensated at a rate of $12 per hour. This study was reviewed by our institution's review board and was deemed exempt from oversight.

3.1 Participants

We recruited participants from Prolific, which has been empirically demonstrated to be a reliable recruitment platform [84]. We recruited a U.S. sample that was representative across gender, race, and age (n = 300) ranging from ages 18-82 (M = 39.71, SD = 14.54), and oversampled for transgender individuals (n = 100), POC (n = 200), transgender POC (n = 50) and disabled individuals (n = 100). Our sample was finalized after manually reviewing respondents who failed the attention check questions,[3] completed the survey too fast, or did not consent. In terms of gender, the sample consisted of cisgender men (n = 256), cisgender women (n = 348) or gender minorities (n = 138).[4] The sample included POC (n = 402), representing all participants except for those were solely white (n = 340). Many participants were disabled (n = 274) though the majority did not report any disability (n = 468). A detailed breakdown of participant characteristics is included in the Appendix. [5]

3.2 Procedure

Since participants may not have been familiar with emotion AI, we provided the following definition at the start of the survey to keep in mind as they answered survey questions, adapted from prior work [23]:
"*Algorithmic systems that promise to automatically detect your emotions and moods (e.g., stress, boredom, calmness, fear, fatigue, attentiveness, happiness, sadness, disgust, surprise, anger) based on various kinds of data (e.g., facial expressions, voice, text/language, biometric information).*"

3.3 Measures

All quantitative measures were evaluated on a 7-point Likert scale ranging from 1 ("Strongly disagree") to 7 ("Strongly agree") unless otherwise noted.

AI Attitudes were measured using four items adapted from previous literature [36] after being provided with a definition of

[3]Within one of the matrix questions, participants were asked "To ensure you are paying attention, please select Strongly agree." Participants who did not select "Strongly agree were removed from the sample. In another matrix question, we asked "How likely is it that dogs can read books and solve complex mathematical equations?" and removed participants who answered that they considered this likely.
[4]This subgroup included participants who were nonbinary or transgender. It also included participants who were unsure about their gender identity or self-described as a gender identity not listed in the survey.
[5]Racial and gender categories were mutually exclusive. For example, participants who identified as multiple races (e.g., East Asian and Hispanic or Latino) were categorized as such (i.e., Mixed race, multiracial, or biracial) as opposed to being coded as *both* East Asian **and** Hispanic or Latino

AI.[6] These items were as follows: "I believe that AI will improve my life," "I believe that AI will improve my work," "I think I will use AI technology in the future" and "I think AI technology is positive for humanity" (M = 4.96; SD = 1.45; a = .93).

Emotion AI Attitudes were measured using the same four items as AI attitudes [36], marginally edited to include emotion (e.g., "I believe that emotion AI will improve my life") (M = 3.47; SD = 1.70; a = .96).

Perceived Sensitivity of Input Data was assessed through asking participants to rate their perceptions about the emotion AI input data from 1 ("Not sensitive at all") to 7 ("Extremely sensitive"). These input data, representing the attributes parameter of CI, are listed in Table 1. We designed these items based on input data types that existing or proposed emotion AI systems use [14].

Perceived Sensitivity of Emotion Data Recipients was measured through asking participants to reflect on common recipients of emotion data using a prompt adapted from Wirth et al. [117]. As established by Wirth et al.[117], perceived negative consequences served as a proxy for perceived sensitivity. Participants were prompted: "There would be negative consequences for me if information about my emotions as inferred by emotion AI were available to…" followed by 21 emotion data recipients listed in Table 1. We included recipients based on available information about existing or proposed emotion AI systems [4, 23, 49, 88, 94, 96].

3.4 Limitations

This study has several limitations. First, we assessed participants' perceptions of input data to emotion AI and emotion data recipients independently. Future work may undertake a factorial design where participants would assess input data *across* different data receivers, following previous work [46]. That said, our study provides the groundwork for understanding emotion AI from a CI perspective. Additionally, previous research has indicated attitudes toward AI differ across cultures [89]; as such, our findings may not generalize outside of the U.S. context. Finally, our analysis maintained broader categories for race, gender, and disability status to inform comparisons between meaningfully sized groups. Future work may use qualitative methods to engage more deeply with smaller samples of marginalized identity groups.

4 Results

In this section, we address RQ1-3. Our analysis distinguishes emotion AI from AI and reveals U.S. data subjects' sensitivity perceptions about emotion AI input data and emotion data recipients.

4.1 Distinguishing Emotion AI from AI

RQ1 sought to examine if data subjects view emotion AI as distinct from AI by assessing their attitudes towards each. A paired samples t-test indicated that participants had significantly more negative attitudes toward emotion AI (M = 3.47, SD = 1.70) than AI (M = 4.96, SD = 1.45), $t(1445.1) = -18.11$, p<.001, 95% CI: -1.65,-1.33, Cohen's d = .34). This indicates data subjects classify emotion AI as distinct from AI.

[6]The definition of AI that we provided to participants was "algorithmic systems that can perform tasks that would usually require human intelligence"

4.2 Perceptions about Emotion AI Input Data

RQ2a aimed to address how perceptions of information sensitivity differ across emotion AI input data types. Across the sample, we found participants to perceive precise location as the most sensitive emotion AI input data (M = 5.68, SD = 1.62), followed by spoken words (M = 5.55, SD = 1.58), digital communication (M = 5.46, SD = 1.58), facial expressions (M = 5.44, SD = 1.63), voice (M = 5.38, SD = 1.63), written words (M = 5.32, SD = 1.63), eye movement (M = 5.23, SD = 1.72), device usage (M = 5.20, SD = 1.75), non-facial biometrics (M = 5.16, SD = 1.75), body movement (M = 5.13, SD = 1.69), and their more general environment (M = 4.78, SD = 1.66).

RQ2b asked if perceptions of sensitivity about emotion AI input data would differ across U.S. data subjects' gender, race, or disability status. Descriptive statistics for race, gender, and disability contrasts across input data to emotion AI can be found in Tables 4-6 in the Appendix. Data were transformed into a long format to allow for direct contrasts in perceived sensitivity across input data types within each demographic group of interest. Due to the exploratory nature of this study and the risk of missing important patterns, we chose not to apply Bonferroni corrections, which could increase Type II errors.

4.2.1 Race. . A Levene's test indicated unequal variances between groups ($p < 0.05$), prompting the use of a Welch's t-test to compare perceived sensitivity of each type of input data across racial groups (see Figure 1a). Results show a significant difference between white people and POC on sensitivity perceptions about non-facial biometric data (i.e., heart rate and skin temperature), $t(734.8)=2.08$, p= .038, Cohen's d = .05. White participants (M = 5.30, SD = 1.67) perceived non-facial biometric data as input to emotion AI as significantly more sensitive than POC participants did (M = 5.03 , SD = 1.80). All other perceptions of input data sensitivity did not significantly differ between POC and white participants.

4.2.2 Gender. . We also sought to understand how perceived sensitivity of emotion AI input data differed across gender (See Figure 1b). To examine this, we ran a two-way analysis of variance (ANOVA). ANOVA results reveal both gender ($F(2, 8868) = 85.52$, p < 0.001, partial eta^2 = .02) and input data (F(11, 8868) = 15.48, p < 0.001, partial eta^2 = .02) to have a significant effect on reported levels of sensitivity. Pairwise comparisons, conducted through Tukey HSD tests, demonstrate that compared to cisgender men, gender minorities reported a significantly higher level of sensitivity around the use of body movement (95% CI [.12, 1.45], p =.002), eye movement (95% CI [.23, 1.56], p <.001), facial expressions (95% CI [.03, 1.36], p =.03), precise location (95% CI [.01, 1.33], p =.04) and voice (95% CI [.04, 1.37], p =.02) as emotion AI input data. There were no significant contrasts observed between cisgender men and cisgender woman nor cisgender woman and gender minorities across any of the input data types.

4.2.3 Disability. . Finally, we wanted to understand how sensitivity perceptions of emotion AI input data differed across disability status (See Figure 1c). Welch's t-tests (the appropriate test because of unequal variances, Levene's test p < .05) revealed that compared to individuals without a disability, disabled individuals perceived a significantly higher level of sensitivity around the use of eye movement ($t(530.9)= 3.20$, p= .001, Cohen's d = 0.1) and non-facial

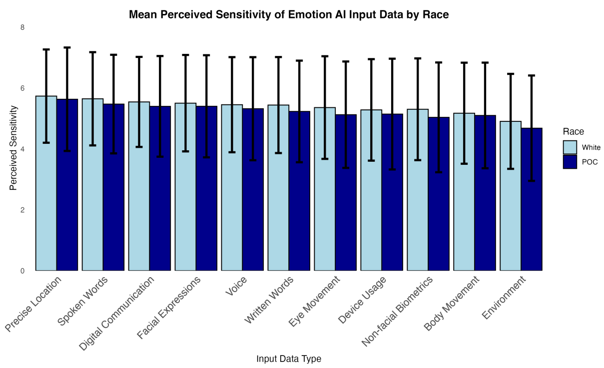

(a) Race

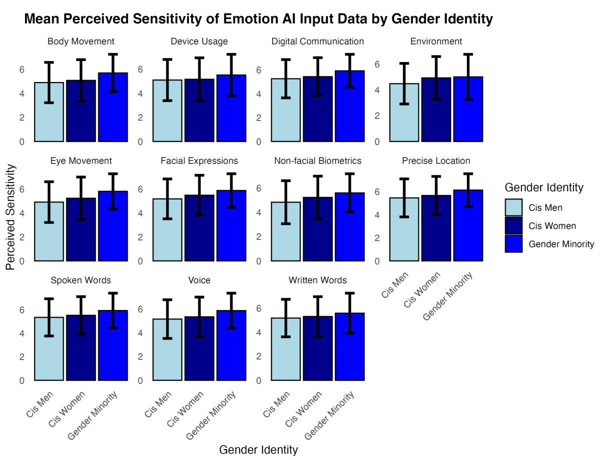

(b) Gender identity

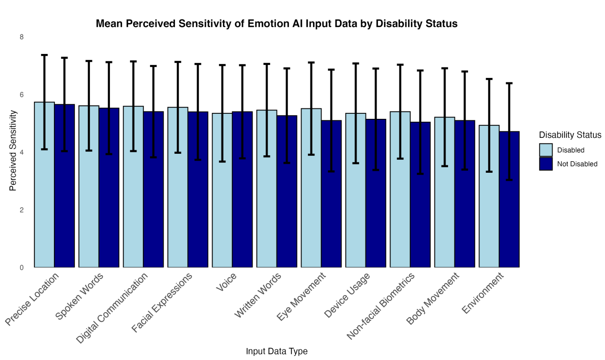

(c) Disability

Figure 1: Perceived sensitivity of input data across race, gender and disability.

biometrics ($t(529.3)=2.77$, $p= .006$, Cohen's $d = 0.1$) as emotion AI input data.

Given emotion AI's impact on neurodivergent individuals [80], we ran a post-hoc analysis comparing neurodivergent data subjects to the rest of the sample. We found neurodivergent participants perceived a significantly higher level of sensitivity across *all input data* to emotion AI than neurotypical participants (See Table 2).

Data Input	*t*	Mean (NT)	Mean (ND)
Precise Location	-2.52*	5.61	5.97
Digital Comm.	-3.98***	5.36	5.90
Spoken Words	-2.64**	5.48	5.86
Facial Expressions	-2.81**	5.37	5.78
Non-facial Biometrics	-4.44***	5.03	5.69
Eye Movement	-4.03***	5.12	5.71
Device Usage	-2.65**	5.13	5.55
Body Movement	-2.95**	5.04	5.50
Environment	-2.27*	4.71	5.07

Table 2: t-test comparing "Neurotypical" (NT) and "Neurodivergent" (ND) groups for emotion AI input data. Cohen's d = 0.3 across data inputs. *** p < 0.001, ** p < 0.01, * p < 0.05

4.3 Perceived Sensitivity of Emotion Data Recipients

RQ3 explored if sensitivity perceptions differed depending on the recipient of emotion AI outputs. Overall, U.S. data subjects perceived relatively high negative consequences across data recipients (see Figure 2). To some extent, lower perceptions of negative consequences corresponded to more personal relationships, though some recipients diverged from this pattern. For instance, participants felt the highest negative consequences would occur if data brokers (M = 5.35, SD, = 1.73), future employers (M = 5.30, SD, = 1.75), or governmental actors (M = 5.27, SD, = 1.86) were the recipients of their emotion data. By contrast, they perceived the least negative consequences would arise if their friends (M = 4.26, SD, = 1.80), romantic partner (M = 4.29, SD = 1.82), or family (M = 4.37, SD, = 1.79) gained access to their emotion data.

Given the significant differences in perceptions of input data, with some notable findings with respect to gender minorities, POC, and neurodivergent individuals, we filtered RQ3 results along these axes of identity posthoc (See Figures 3-5 in the Appendix). Results reflect similar patterns to input data in that average perceived negative consequences about data recipients were descriptively lower for POC and higher for gender minorities and neurodivergent individuals compared to the rest of the sample. Of note, although data brokers were similarly identified as the most concerning recipients of emotion data by the POC and gender minorities, neurodivergent individuals perceived future employers in this regard.

5 Discussion

This paper provides several insights into U.S. data subjects' perceptions of emotion AI. We first contribute a quantitative distinction between attitudes towards emotion AI and AI more broadly, finding U.S. data subjects have significantly more negative attitudes toward the former. Second, our findings suggest CI [81] parameters which characterize information flows to and from emotion AI systems matter and might violate contextual integrity. We find individuals of minoritized genders and disability status—particularly those who are neurodivergent—perceive a significantly higher sensitivity toward some emotion AI input data than cisgender men and those not living with a disability, respectively. Third, this paper

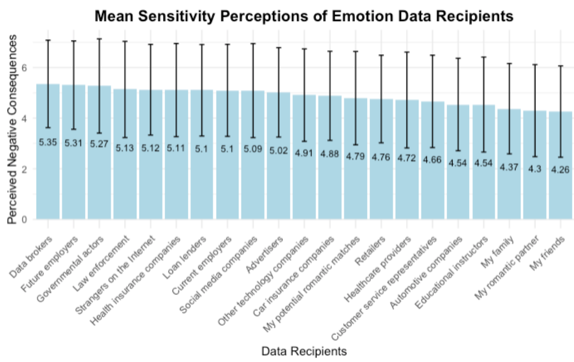

Figure 2: Average perceived sensitivity of emotion data recipients across the full sample. Mean values are displayed within each bar, and error bars represent standard deviations

contributes insights into data subjects' perceptions of emotion data recipients. Although data subjects see the potential for negative consequences across all recipients' access to emotion data, the most severe perceived outcomes are associated with those to whom data subjects have little personal connection and who hold a large power asymmetry. We discuss implications for AI governance and development including 1) maintaining contextual integrity and attention to data subjects' identities in future emotion AI regulation, development, and deployment, 2) recognizing sensitivity *across* variation in emotion AI input data, and 3) encouraging focused attention on emotion AI in broader AI regulation efforts.

5.1 With More Negative Attitudes, Data Subjects View Emotion AI as Distinct from AI

Our results reveal data subjects conceive of emotion AI as distinct from AI, with attitudes toward the former being significantly more negative. This finding extends past work which has built conceptual, but not empirical, arguments around the uniquely sensitive and intimate nature of emotion AI and emotion data [4, 106]. Indeed, emotion AI's promise to infer "emotion" implies an ability to obtain information which data subjects may prefer to remain concealed, or that they may not even be able to identify themselves [90]. Although scholars have critiqued emotion AI's potential to achieve its purported promise both on a theoretical [10, 62, 107] and technical level [52, 79], the market around emotion AI is projected to further grow [68]. Our results, paired with the evolving emotion AI market, affirm that policymakers should distinctly define emotion AI within broader AI governance initiatives, requiring specified oversight on its growth and implications. While the U.S. has made progress on AI governance both at the federal [69] and state (e.g., Colorado [22], California [17]) levels, none of these efforts have explicitly recognized emotion AI. The absence of a federal AI law in the U.S. presents an opportunity to ensure the explicit inclusion of emotion AI and emotion data in future policymaking. Such an approach would align with global leaders in Internet law such as the EU, whose AI Act has placed a ban on emotion AI use for workplace or education purposes unless its involves a healthcare or safety need [112].

5.2 Data Subjects Perceive *All* Emotion AI Input Data as Sensitive

Our results underscore U.S. data subjects perceive *all* inputs to emotion AI as sensitive. Some emotion AI developers would agree with these classifications, as demonstrated through efforts to address privacy issues arising from input data such as speech [27], eye-movement [58], and facial data [41]. In fact, some emotion AI developers have addressed emotion AI's privacy harms by using input data that is not connected to one's identity (e.g., relying on weather and traffic conditions as input data for in-car emotion AI systems) [12]. Others have remarked on perceiving facial expression data as more sensitive than body movement [13, 65], which aligns with data subjects' perceptions. The similarities between U.S. data subjects' perceptions and some emotion AI developers' privacy-protective efforts should motivate developers' prioritization of protecting sensitive information. This may mean not using certain types of input data at all due to their sensitivity. Further, policymakers can consult these findings as a guide for regulation which is informed directly by what data subjects perceive to be the most sensitive types of emotion AI input data.

Interestingly, participants reported precise location (e.g., a specific doctor's office) as the most sensitive emotion AI input data type, and the general environment (e.g., a generic doctor's office) as the least sensitive. The disparity between these types of emotion AI input data reflects data subjects' desire for location data *obscurity* (i.e., making data more difficult to find or understand [40]). Scholarship has noted the harms of collecting precise location explicitly in communities of color, as it not only allows for law enforcements' physical pursuit of individuals, but also permits connecting location data to other sensitive demographic attributes such as sexuality or health status [60]. In fact, facial emotion recognition systems may link its analytics with data subjects' locations to generate more targeted insights (e.g., emotions expressed at a retail store) [101]. Location data is indeed represented in 8% of emotion AI technologies proposed to be used in the workplace [14]. Although our results should not be misinterpreted as impetus to collect data subjects' general environment for emotion AI, they suggest data subjects' preference for location data obscurity. This aligns with previous research which identified data subjects' preference for location data obscurity in the context of a mobile application, specifically when other users chose to share an obscure as opposed to precise version of their location data [105]. Critically, our findings suggest that future emotion AI systems should not be able to connect emotion inferences with data subjects' location information.

Participants also saw spoken words and digital communications as highly sensitive emotion AI input data. This reveals data subjects' prioritization of their intellectual privacy, which includes the ability to engage in a private conversation [92]. Protecting against intrusion upon private conversations has long been an issue raised by lawmakers, yet has been exacerbated by digital tools that can surreptitiously capture them [45]. Compared to intrusions upon private conversations that are not supported by AI (e.g., eavesdropping on an offline conversation; taking a screenshot of a private text message), emotion AI represents a more invasive privacy violation by claiming the ability to deduce emotion and fuel subsequent inferences based on collected communications. Given the sensitivity data subjects perceive around spoken words and digital communications, emotion AI developers should take caution in using them as input data. This is particularly important as some developers view the use of spoken words as input data as *not* sensitive [20, 64]. Lastly, Ingber suggests law and policy must protect the privacy of offline and digital communications [45]. We propose that these protections expand to personal communications used as input to emotion AI.

5.3 Distinct Perceptions across Gender, Race, and Disability Status

We found participants' perceptions of particular input data significantly differ across race, gender, and disability status. This is not surprising given evidence surrounding emotion AI's bias along these axes [66, 91, 123]. At a high level, these results challenge efforts to use AI to mitigate identity-related bias [1, 8] because such technical approaches may not be fully effective nor able to overhaul data subjects' existing concerns. Further, our results corroborate and extend previous work suggesting identity-related characteristics have an influence on perceptions about AI [55].

In terms of race, we found white participants perceived non-facial biometrics (e.g., heart rate, skin temperature) as significantly more sensitive emotion AI input data than POC, aligning with recent work on emotion AI [5] and potentially reflecting the idea of Black optimism [102], which may extend to other POC. Black optimism contends individuals who have been historically mistreated as a result of their race may feel more optimistic about emerging technology as an act of resistance, which has been reported with respect to facial recognition technology [83]. It is possible that the normalization of surveillance POC experience as compared to white people [15] may have yielded a diminished sensitivity by the former. That being said, we suggest interpreting this finding with caution given its small effect size.

In terms of gender, minorities reported significantly higher levels of sensitivity around the use of body movement, eye movement, facial expressions, precise location, and voice as emotion AI input data compared to cisgender men. These findings reinforce broader concerns around the use of pervasive surveillance tools to fuel gender-based violence toward individuals who do not fall into a binary gender category [103]. Gender minorities' perceptions of several types of input data as significantly more sensitive may be due to the unique harms they face from misclassification [72, 116] which cisgender people do not normally experience. AI systems typically do not account for the complexities of trans and nonbinary identities and bodies, ultimately causing harm to gender minorities [32, 39, 42, 51, 54, 85, 100, 111].

Similarly, disabled participants perceived a significantly higher level of sensitivity around the use of eye movement and non-facial biometrics as emotion AI input data than individuals who were not disabled. These findings can be explained by some disabled individuals' non-normative means of self-expression [122] and emotion regulation [73]. It is likely that disabled participants were troubled by the idea of a system inferring emotions from subconscious biometric movement, as this leaves no opportunity to pass as able-bodyminded (i.e., concealing one's disability of body and/or mind) [104], which some may appreciate [114]. Thus, our findings resonate with arguments that emotion recognition systems inherently work against disability justice [47].

Our findings problematize responsible AI efforts attempting to address the problem of AI bias by including more data from marginalized groups against whom systems are biased [2, 99, 119]. Because data subjects with disabilities and minoritized genders view various input data as sensitive and emotion data as associated with negative consequences, it follows that any effort to further "include" them by incorporating their data would not align with their values. We argue that because these individuals have significant concerns regarding various kinds of emotion AI input data, emotion AI systems should not collect their data. While this may foster systems that do not work for these groups well, we consider the possibility that this may not be an undesired outcome given their concerns. We concur with Hoffman's [43] argument that creating more "inclusive" technology disregards critiques against building particular technologies (e.g., emotion AI) or collecting sensitive data in the first place.

Results on disability were even more stark when focusing on neurodivergent U.S. data subjects. Compared to the rest of the sample, neurodivergent individuals perceived *all* of the potential emotion AI input data as significantly more sensitive. Much of affective computing broadly [86], and specifically emotion AI development, is motivated by the idea that disabled and neurodivergent people would benefit from emotion sensing [3]. Our results show neurodivergent individuals' perceptions fall in tension with this very idea, reporting all types of emotion AI input data to be significantly more sensitive than was reported by the broader sample. This is a critical finding that raises questions in line with Nagy [80] about assumptions embedded in the development of emotion AI for the benefit of neurodivergent people. It is possible that these individuals, assuming the technology works as claimed, do indeed experience benefits despite the perceptions we uncovered. However, given neurodivergent individuals' sensitivity toward both input and output data, we urge developers to question whether emotion AI development is the solution to the challenges neurodivergent individuals face, and what disparities and ethical concerns emotion AI may cause them.

5.4 All Recipients of Emotion Data Are Thought to Incur Negative Consequences

U.S. data subjects felt there could be negative consequences if their emotion data were to flow to *any* of the data recipients we asked about (See Table 1). That being said, we observed a loose pattern in that recipients to whom data subjects could have a personal relationship (i.e., friends, romantic partner, family, educational instructors) were perceived to have the least negative consequences. Although there are circumstances where, for example, a romantic partner may abuse access to intimate data as a means of control [61], it is sensible that data subjects would have a slightly higher level of trust if emotion data was received by a personal connection. This has been reflected in previous research which found data subjects to be significantly more comfortable sharing personally identifiable information with a friend than both trusted and unknown marketers [70]. By contrast, U.S. data subjects perceived a high likelihood of negative consequences associated with actors who have significant power. For instance, data brokers' primary role is to distribute personal information to other entities with little transparency or regulation [24]-including future employers, through measurements of "employability" [29]. Notably, while the most and least concerning emotion data recipients were consistent across the gender minority, POC, and general participant samples, neurodivergent individuals saw future employers as having the most negative consequences as recipients of emotion data. This corroborates evidence suggesting disabled individuals may be at a disadvantage in hiring that engages with predictive technologies [34], including emotion AI [79]. In sum, our results highlight U.S. data subjects' universal concern with the flow of their emotion data. We encourage policymakers to regulate the flow of emotion data given its innate sensitivity, as established in this study, with a particular focus on the expected negative consequences of such flows to powerful recipients.

6 Conclusion

This representative survey 1) distinguished emotion AI from AI more generally, and 2) leveraged CI as a theoretical framework to study the data flows both into and out of emotion AI systems. U.S. data subjects report a high level of perceived sensitivity across input data, with this being significantly higher for gender minorities and individuals with disabilities. We also demonstrate U.S. data subjects' belief that there would be negative consequences if their emotion data were to flow across a range of diverse data recipients, particularly those to which they have considerable asymmetrical power relationship. We urge policymakers to adopt a tailored approach to emotion AI which accounts for its impact based on data inputs, recipients, and the identity of its data subjects. Regulatory efforts may discourage emotion AI's development altogether, which, given the disparities across marginalized individuals, may be favorable.

Acknowledgments

We appreciate the reviewers' valuable feedback on this work. This work was sponsored by NSF CAREER award 2236674.

References

[1] Lavisha Aggarwal and Shruti Bhargava. 2023. Fairness in AI Systems: Mitigating gender bias from language-vision models. *arXiv preprint arXiv:2305.01888* (2023).

[2] Saba Akhyani, Mehryar Abbasi, Mo Chen, and Angelica Lim. 2022. Towards inclusive hri: Using sim2real to address underrepresentation in emotion expression recognition. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 9132–9139.

[3] Maram Fahaad Almufareh, Sumaira Kausar, Mamoona Humayun, and Samabia Tehsin. 2024. A conceptual model for inclusive technology: advancing disability inclusion through artificial intelligence. *Journal of Disability Research* 3, 1 (2024), 20230060.

[4] Nazanin Andalibi and Justin Buss. 2020. The human in emotion recognition on social media: Attitudes, outcomes, risks. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–16.

[5] Nazanin Andalibi and Alexis Shore Ingber. 2025. Public Perceptions About Emotion AI Use Across Contexts in the United States. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery. https://doi.org/10.1145/3706598.3713501

[6] Mohd Ayaan Anwar, Manan Agrawal, Neha Gahlan, Divyashikha Sethia, Gaurav Kumar Singh, and Rishabh Chaurasia. 2023. FedEmo: A privacy-preserving framework for emotion recognition using EEG physiological data. In *2023 15th International conference on communication systems & networks (COMSNETS)*. IEEE, 119–124.

[7] Sumit Asthana, Jane Im, Zhe Chen, and Nikola Banovic. 2024. " I know even if you don't tell me": Understanding Users' Privacy Preferences Regarding AI-based Inferences of Sensitive Information for Personalization. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–21.

[8] Edmond Awad, Loukas Balafoutas, Li Chen, Edwin Ip, and Joe Vecci. 2023. Artificial intelligence and debiasing in hiring: Impact on applicant quality and gender diversity. *Available at SSRN* (2023).

[9] Areej Babiker, Sameha Alshakhsi, Dena Al-Thani, Christian Montag, and Raian Ali. 2024. Attitude Towards AI: Potential Influence of Conspiracy Belief, XAI Experience and Locus of Control. *International Journal of Human–Computer Interaction* (2024), 1–13.

[10] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M Martinez, and Seth D Pollak. 2019. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological science in the public interest* 20, 1 (2019), 1–68.

[11] Mohamed Benouis, Yekta Said Can, and Elisabeth André. 2023. A Privacy-Preserving Multi-Task Learning Framework For Emotion and Identity Recognition from Multimodal Physiological Signals. In *2023 11th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*. IEEE, 1–8.

[12] David Bethge, Luis Falconeri Coelho, Thomas Kosch, Satiyabooshan Murugaboopathy, Ulrich von Zadow, Albrecht Schmidt, and Tobias Grosse-Puppendahl. 2023. Technical Design Space Analysis for Unobtrusive Driver Emotion Assessment Using Multi-Domain Context. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 4, Article 159 (Jan. 2023), 30 pages. https://doi.org/10.1145/3569466

[13] Carmen Bisogni, Lucia Cimmino, Michele Nappi, Toni Pannese, and Chiara Pero. 2024. Walk as you feel: Privacy preserving emotion recognition from gait patterns. *Engineering Applications of Artificial Intelligence* 128 (2024), 107565.

[14] Karen L Boyd and Nazanin Andalibi. 2023. Automated emotion recognition in the workplace: How proposed technologies reveal potential futures of work. *Proceedings of the ACM on human-computer interaction* 7, CSCW1 (2023), 1–37.

[15] Simone Browne. 2015. *Dark matters: On the surveillance of blackness*. Duke University Press.

[16] M Calice, L Bao, T Newman, D Scheufele, D Brossard, and M Xenos. 2022. US Public attitudes on artificial intelligence. *OSF Home* (2022).

[17] California State Legislature. 2024. California Assembly Bill No. 1008 (2023–2024). https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202320240AB1008 Accessed: 2024-10-25.

[18] Carolina Carreira, Cody Berger, Khushi Shah, Samridhi Agarwal, Yashasvi Thakur, McKenna McCall, Nicolas Christin, and Lorrie Faith Cranor. 2024. Who's Listening? Analyzing Privacy Preferences in Multi-User Smart Personal Assistants Settings. In *SOUPs Proceedings*.

[19] Si Chen, Jason Situ, Haocong Cheng, Suzy Su, Desiree Kirst, Lu Ming, Qi Wang, Lawrence Angrave, and Yun Huang. 2024. Inclusive Emotion Technologies: Addressing the Needs of d/Deaf and Hard of Hearing Learners in Video-Based Learning. *arXiv preprint arXiv:2410.00199* (2024).

[20] Shuaiqi Chen, Xiaofen Xing, Guodong Liang, and Xiangmin Xu. 2022. I Feel Stressed Out: A Mandarin Speech Stress Dataset with New Paradigm. In *2022 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. 583–589. https://doi.org/10.23919/APSIPAASC55919.2022.9980183

[21] Xi Chen, Pengxin Zheng, and Jian Mou. 2024. Understanding Chinese Internet users' information sensitivity in big data and artificial intelligence era. *Policy & Internet* (2024).

[22] Colorado General Assembly. 2024. Concerning the Implementation of an Automated Traffic Enforcement System. https://leg.colorado.gov/bills/sb24-205 Accessed: 2024-09-24.

[23] Shanley Corvite, Kat Roemmich, Tillie Ilana Rosenberg, and Nazanin Andalibi. 2023. Data Subjects' Perspectives on Emotion Artificial Intelligence Use in the Workplace: A Relational Ethics Lens. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–38.

[24] Matthew Crain. 2018. The limits of transparency: Data brokers and commodification. *New Media & Society* 20, 1 (2018), 88–104.

[25] Mateja Durovic and Tommaso Corno. 2024. The Privacy of Emotions: From the GDPR to the AI Act, an Overview of Emotional AI Regulation and the Protection of Privacy and Personal Data. In *Privacy, Data Protection, and Data-Driven Technologies*. Routledge, London, UK.

[26] Michelle Faverio and Alec Tyson. 2023. What the Data Says About Americans' Views of Artificial Intelligence. https://www.pewresearch.org/short-reads/2023/11/21/what-the-data-says-about-americans-views-of-artificial-intelligence/. Pew Research Center, accessed: 2024-12-13.

[27] Tiantian Feng and Shrikanth Narayanan. 2021. Privacy and Utility Preserving Data Transformation for Speech Emotion Recognition. In *2021 9th International Conference on Affective Computing and Intelligent Interaction (ACII)*. 1–7. https://doi.org/10.1109/ACII52823.2021.9597433

[28] Yuanyuan Feng, Brad Stenger, and Shikun Zhang. 2024. Contextual Acceptance of COVID-19 Mitigation Mobile Apps in the United States: Mixed Methods

Survey Study on Postpandemic Data Privacy. *Journal of Medical Internet Research* 26 (2024), e57309.

[29] Michael Fertik. 2012. Your future employer is watching you online. You should be, too. *Harvard Business Review. Retrieved from http://blogs. hbr. org/2012/04/your-future-employer-is-watchi* (2012).

[30] Alisa Frik, Julia Bernd, and Serge Egelman. 2023. A Model of Contextual Factors Affecting Older Adults' Information-Sharing Decisions in the US. *ACM Transactions on Computer-Human Interaction* 30, 1 (2023), 1–48.

[31] Vinitha Gadiraju, Shaun Kane, Sunipa Dev, Alex Taylor, Ding Wang, Emily Denton, and Robin Brewer. 2023. "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*. 205–216.

[32] Timnit Gebru. 2019. Oxford handbook on AI ethics book chapter on race and gender. *arXiv preprint arXiv:1908.06165* (2019).

[33] Tapotosh Ghosh, Md Hasan Al Banna, Md Jaber Al Nahian, M Shamim Kaiser, Mufti Mahmud, Shaobao Li, and Nelishia Pillay. 2022. A privacy-preserving federated-mobilenet for facial expression detection from images. In *International Conference on Applied Intelligence and Informatics*. Springer, 277–292.

[34] Kate Glazko, Yusuf Mohammed, Ben Kosa, Venkatesh Potluri, and Jennifer Mankoff. 2024. Identifying and Improving Disability Bias in GPT-Based Resume Screening. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 687–700.

[35] Erving Goffman. 1959. *The Presentation of Self in Everyday Life*. Anchor Books, New York.

[36] Simone Grassini. 2023. Development and validation of the AI attitude scale (AIAS-4): a brief measure of general attitude toward artificial intelligence. *Frontiers in psychology* 14 (2023), 1191628.

[37] Simone Grassini and Alexander Sævild Ree. 2023. Hope or Doom AI-ttitude? Examining the Impact of Gender, Age, and Cultural Differences on the Envisioned Future Impact of Artificial Intelligence on Humankind. In *Proceedings of the European Conference on Cognitive Ergonomics 2023*. 1–7.

[38] Gabriel Grill and Nazanin Andalibi. 2022. Attitudes and folk theories of data subjects on transparency and accuracy in emotion recognition. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (2022), 1–35.

[39] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M Branham. 2018. Gender recognition or gender reductionism? The social implications of embedded gender recognition systems. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–13.

[40] Woodrow Hartzog. 2018. *Privacy's Blueprint: The Battle to Control the Design of New Technologies*. Harvard University Press.

[41] Fabio Hellmann, Silvan Mertes, Mohamed Benouis, Alexander Hustinx, Tzung-Chien Hsieh, Cristina Conati, Peter Krawitz, and Elisabeth André. 2024. GANonymization: A GAN-based Face Anonymization Framework for Preserving Emotional Expressions. *ACM Trans. Multimedia Comput. Commun. Appl.* (Jan. 2024). https://doi.org/10.1145/3641107 Just Accepted.

[42] Mar Hicks. 2019. Hacking the Cis-tem. *IEEE Annals of the History of Computing* 41, 1 (2019), 20–33.

[43] Anna Lauren Hoffmann. 2021. Terms of inclusion: Data, discourse, violence. *New Media & Society* 23, 12 (2021), 3539–3556.

[44] Michael C Horowitz and Lauren Kahn. 2021. What influences attitudes about artificial intelligence adoption: Evidence from US local officials. *PLoS One* 16, 10 (2021), e0257732.

[45] Alexis Shore Ingber. 2024. Protecting Private Communications through Law and Policy: The Case of the Screenshot Feature. *Communication Law and Policy* 29, 1-2 (2024), 1–23.

[46] Emiram Kablo and Patricia Arias-Cabarcos. 2023. Privacy in the Age of Neurotechnology: Investigating Public Attitudes towards Brain Data Collection and Use. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 225–238.

[47] 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*. 455–466.

[48] Eiman Kanjo, Eman MG Younis, and Chee Siang Ang. 2019. Deep learning analysis of mobile physiological, environmental and location sensor data for emotion detection. *Information Fusion* 49 (2019), 46–56.

[49] Nadia Karizat, Alexandra H Vinson, Shobita Parthasarathy, and Nazanin Andalibi. 2024. Patent Applications as Glimpses into the Sociotechnical Imaginary: Ethical Speculation on the Imagined Futures of Emotion AI for Mental Health Monitoring and Detection. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–43.

[50] Mounika Karna, D Sujitha Juliet, and R Catherine Joy. 2020. Deep learning based text emotion recognition for chatbot applications. In *2020 4th International Conference on Trends in Electronics and Informatics (ICOEI)(48184)*. IEEE, 988–993.

[51] Sonia K Katyal and Jessica Y Jung. 2021. The gender panopticon: AI, gender, and design Justice. *UCLA L. Rev.* 68 (2021), 692.

[52] Harmanpreet Kaur, Daniel McDuff, Alex C Williams, Jaime Teevan, and Shamsi T Iqbal. 2022. "I didn't know I looked angry": Characterizing observed emotion and reported affect at work. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–18.

[53] Patrick Gage Kelley, Celestina Cornejo, Lisa Hayes, Ellie Shuo Jin, Aaron Sedley, Kurt Thomas, Yongwei Yang, and Allison Woodruff. 2023. " There will be less privacy, of course": How and why people in 10 countries expect {AI} will affect privacy in the future. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*. 579–603.

[54] Os Keyes. 2018. The misgendering machines: Trans/HCI implications of automatic gender recognition. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–22.

[55] Soojong Kim, Joomi Lee, and Poong Oh. 2024. Questioning Artificial Intelligence: How Racial Identity Shapes the Perceptions of Algorithmic Bias. *International journal of communication (Online)* 18 (2024), 677–700.

[56] Anastasia Kozyreva, Philipp Lorenz-Spreen, Ralph Hertwig, Stephan Lewandowsky, and Stefan M Herzog. 2021. Public attitudes towards algorithmic personalization and use of personal data online: Evidence from Germany, Great Britain, and the United States. *Humanities and Social Sciences Communications* 8, 1 (2021), 1–11.

[57] Sheetal Kusal, Shruti Patil, Ketan Kotecha, Rajanikanth Aluvalu, and Vijayakumar Varadarajan. 2021. AI based emotion detection for textual big data: Techniques and contribution. *Big Data and Cognitive Computing* 5, 3 (2021), 43.

[58] Guohao Lan, Tim Scargill, and Maria Gorlatova. 2022. EyeSyn: Psychology-inspired Eye Movement Synthesis for Gaze-based Activity Recognition. In *2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. 233–246. https://doi.org/10.1109/IPSN54338.2022.00026

[59] Siddique Latif, Hafiz Shehbaz Ali, Muhammad Usama, Rajib Rana, Björn Schuller, and Junaid Qadir. 2022. Ai-based emotion recognition: Promise, peril, and prescriptions for prosocial path. *arXiv preprint arXiv:2211.07290* (2022).

[60] Nicol Turner Lee and Caitlin Chin. 2022. Police surveillance and facial recognition: Why data privacy is imperative for communities of color. *The Brookings Institution. Erişim adresi: https://www. brookings. edu/research/police-surveillance-and-facial-recognition-why-data-privacy-is-an-imperative-for-communities-of-color* (2022).

[61] Karen EC Levy. 2014. Intimate surveillance. *Idaho L. Rev.* 51 (2014), 679.

[62] Ruth Leys. 2019. *The ascent of affect: Genealogy and critique*. University of Chicago Press.

[63] Bingjie Liu, Lewen Wei, Mu Wu, and Tianyi Luo. 2023. Speech production under uncertainty: how do job applicants experience and communicate with an AI interviewer? *Journal of Computer-Mediated Communication* 28, 4 (2023), zmad028.

[64] Daniel Isoso Machanje, Joseph Onderi Orero, and Christophe Marsala. 2019. Distress Recognition from Speech Analysis: A Pairwise Association Rules-Based Approach. In *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*. 842–849. https://doi.org/10.1109/SSCI44817.2019.9002972

[65] Matthew Malek–Podjaski and Fani Deligianni. 2021. Towards Explainable, Privacy-Preserved Human-Motion Affect Recognition. In *2021 IEEE Symposium Series on Computational Intelligence (SSCI)*. 01–09. https://doi.org/10.1109/SSCI50451.2021.9660129

[66] Peter Mantello, Manh-Tung Ho, Minh-Hoang Nguyen, and Quan-Hoang Vuong. 2023. Bosses without a heart: socio-demographic and cross-cultural determinants of attitude toward Emotional AI in the workplace. *AI & Society* 38, 1 (2023), 97–119.

[67] Peter A Mantello, Nader Ghotbi, Manh-Tung Ho, and Fuminobu Mizutani. 2024. Gauging public opinion of AI and emotionalized AI in healthcare: findings from a nationwide survey in Japan. *AI & Society* (2024), 1–15.

[68] MarketsandMarkets. 2024. Emotion Detection and Recognition Market by Component, Technology, Application, and Region - Global Forecast to 2024. https://www.marketsandmarkets.com/Market-Reports/emotion-detection-recognition-market-23376176.html Accessed: 2024-09-05.

[69] Edward Markey. 2024. AI Civil Rights Act. https://www.markey.senate.gov/view/senator-markey-hosts-press-conference-to-introduce-his-ai-civil-rights-act Introduced in the U.S. Senate to address algorithmic discrimination..

[70] Ereni Markos, Lauren I Labrecque, and George R Milne. 2018. A new information lens: The self-concept and exchange context as a means to understand information sensitivity of anonymous and personal identifying information. *Journal of Interactive Marketing* 42, 1 (2018), 46–62.

[71] Marijn Martens, Ralf De Wolf, and Lieven De Marez. 2024. Datafication and algorithmization of education: How do parents and students evaluate the appropriateness of learning analytics? *Education and Information Technologies* 29, 7 (2024), 8151–8177.

[72] Samuel Mayworm, Kendra Albert, and Oliver L Haimson. 2024. Misgendered During Moderation: How Transgender Bodies Make Visible Cisnormative Content Moderation Policies and Enforcement in a Meta Oversight Board Case. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 301–312.

[73] Carla A Mazefsky and Susan W White. 2014. Emotion regulation: Concepts & practice in autism spectrum disorder. *Child and Adolescent Psychiatric Clinics* 23, 1 (2014), 15–24.

[74] Nora McDonald and Andrea Forte. 2020. The politics of privacy theories: Moving from norms to vulnerabilities. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–14.

[75] Andrew McStay. 2016. Empathic media and advertising: Industry, policy, legal and citizen perspectives (the case for intimacy). *Big data & society* 3, 2 (2016), 2053951716666868.

[76] Andrew McStay. 2024. *Automating empathy: Decoding technologies that gauge intimate life.* Oxford University Press.

[77] Zion Mengesha, Courtney Heldreth, Michal Lahav, Juliana Sublewski, and Elyse Tuennerman. 2021. "I don't think these devices are very culturally sensitive."—Impact of automated speech recognition errors on African Americans. *Frontiers in Artificial Intelligence* 4 (2021), 725911.

[78] George R Milne, George Pettinico, Fatima M Hajjat, and Ereni Markos. 2017. Information sensitivity typology: Mapping the degree and type of risk consumers perceive in personal data sharing. *Journal of Consumer Affairs* 51, 1 (2017), 133–161.

[79] Scott Monteith, Tasha Glenn, John Geddes, Peter C Whybrow, and Michael Bauer. 2022. Commercial use of emotion artificial intelligence (AI): implications for psychiatry. *Current Psychiatry Reports* 24, 3 (2022), 203–211.

[80] Jeff Nagy. 2024. Autism and the making of emotion AI: Disability as resource for surveillance capitalism. *New media & society* 26, 7 (2024), 3989–4007.

[81] Helen Nissenbaum. 2009. Privacy in Context: Technology, Policy, and the Integrity of Social Life.

[82] Ekaterina Novozhilova, Kate Mays, and James E Katz. 2024. Looking towards an automated future: US attitudes towards future artificial intelligence instantiations and their effect. *Humanities and Social Sciences Communications* 11, 1 (2024), 1–11.

[83] Hannah Overbye-Thompson, Kristy A. Hamilton, and Dana Mastro. 2024. Reinvention mediates impacts of skin tone bias in algorithms: implications for technology diffusion. *Journal of Computer-Mediated Communication* 29, 5 (2024), zmae016. https://doi.org/10.1093/jcmc/zmae016

[84] Eyal Peer, Laura Brandimarte, Sonam Samat, and Alessandro Acquisti. 2017. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology* 70 (2017), 153–163. https://doi.org/10.1016/j.jesp.2017.01.006

[85] Michel Perilo and George Valença. 2024. How Facial Recognition Technologies Affect the Transgender Community? A Systematic Mapping Study. In *Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing*. 1153–1160.

[86] Rosalind W Picard. 2000. *Affective computing.* MIT press.

[87] Gabriele Pizzi, Virginia Vannucci, Yupal Shukla, and Gaetano Aiello. 2022. Privacy concerns and justice perceptions with the disclosure of biometric versus behavioral data for personalized pricing tell me who you are, I'll tell you how much you pay. Consumers' fairness and privacy perceptions with personalized pricing. *Journal of Business Research* 148 (2022), 420–432.

[88] Cassidy Pyle, Kat Roemmich, and Nazanin Andalibi. 2024. US Job-Seekers' Organizational Justice Perceptions of Emotion AI-Enabled Asynchronous Interviews. *Proceedings of the ACM on Human-Computer Interaction* (2024).

[89] Mohammad Mominur Rahman, Areej Babiker, and Raian Ali. 2024. Motivation, Concerns, and Attitudes Towards AI: Differences by Gender, Age, and Culture. In *International Conference on Web Information Systems Engineering*. Springer, 375–391.

[90] Liran Razinsky. 2023. Better than they know themselves? Algorithms and subjectivity. *Subjectivity* 30, 4 (2023), 394–416.

[91] Lauren Rhue. 2019. Understanding the hidden bias in emotion-reading AIs.

[92] Neil Richards. 2015. *Intellectual privacy: Rethinking civil liberties in the digital age.* Oxford University Press, USA.

[93] Cami Rincón, Os Keyes, and Corinne Cath. 2021. Speaking from experience: Trans/non-binary requirements for voice-activated AI. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–27.

[94] Kat Roemmich and Nazanin Andalibi. 2021. Data subjects' conceptualizations of and attitudes toward automatic emotion recognition-enabled wellbeing interventions on social media. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–34.

[95] Kat Roemmich, Shanley Corvite, Cassidy Pyle, Nadia Karizat, and Nazanin Andalibi. 2024. Emotion AI Use in US Mental Healthcare: Potentially Unjust and Techno-Solutionist. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW1 (2024), 1–46.

[96] Kat Roemmich, Florian Schaub, and Nazanin Andalibi. 2023. Emotion AI at work: Implications for workplace surveillance, emotional labor, and emotional privacy. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20.

[97] Somaiieh Rokhsaritalemi, Abolghasem Sadeghi-Niaraki, and Soo-Mi Choi. 2023. Exploring emotion analysis using artificial intelligence, geospatial information systems, and extended reality for urban services. *IEEE Access* (2023).

[98] Javier Sánchez-Monedero and Lina Dencik. 2022. The politics of deceptive borders:'biomarkers of deceit'and the case of iBorderCtrl. *Information, Communication & Society* 25, 3 (2022), 413–430.

[99] Bernhard Schenkenfelder, Ulrich Brandstätter, Lukas Fischer, Rudolf Ramler, Dominik Laister, Martin Hartl, and Markus Wurm. 2024. Responsible AI Engineering: The Case of an Inclusive Image Annotation Team in a Global Technology Company. In *Proceedings of the 2nd International Workshop on Responsible AI Engineering*. 8–15.

[100] Morgan Klaus Scheuerman, Jacob M Paul, and Jed R Brubaker. 2019. How computers see gender: An evaluation of gender classification in commercial facial analysis services. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–33.

[101] Sovantharith Seng, Mahdi Nasrullah Al-Ameen, and Matthew Wright. 2021. A first look into users' perceptions of facial recognition in the physical world. *Computers & Security* 105 (2021), 102227.

[102] Jared Sexton. 2016. The social life of social death: On Afro-pessimism and Black optimism. In *Time, temporality and violence in international relations*. Routledge, 61–75.

[103] Jama Shelton, Kel Kroehle, Emilie K Clark, Kristie Seelman, and SJ Dodd. 2021. Digital technologies and the violent surveillance of nonbinary gender. *Journal of gender-based violence* 5, 3 (2021), 517–529.

[104] Emma Sheppard. 2018. Using pain, living with pain. *Feminist Review* 120, 1 (2018), 54–69.

[105] Alexis Shore and James J Cummings. 2023. An Exploratory Analysis of Interface Features Influencing Mobile Location Data Disclosure. *International Journal of Human–Computer Interaction* 39, 19 (2023), 3758–3773.

[106] Luke Stark. 2016. The emotional context of information privacy. *The Information Society* 32, 1 (2016), 14–27.

[107] Luke Stark and Jesse Hoey. 2021. The ethics of emotion in artificial intelligence systems. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 782–793.

[108] Jan-Philipp Stein, Tanja Messingschlager, Timo Gnambs, Fabian Hutmacher, and Markus Appel. 2024. Attitudes towards AI: measurement and associations with personality. *Scientific Reports* 14, 1 (2024), 2909.

[109] Boyuan Sun, Qiang Ma, Shanfeng Zhang, Kebin Liu, and Yunhao Liu. 2017. iSelf: Towards cold-start emotion labeling using transfer learning with smartphones. *ACM Transactions on Sensor Networks (TOSN)* 13, 4 (2017), 1–22.

[110] Arnout Terpstra, Alwin De Rooij, and Alexander Schouten. 2023. Online proctoring: Privacy invasion or study alleviation? Discovering acceptability using contextual integrity. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20.

[111] Eddie L Ungless, Björn Ross, and Anne Lauscher. 2023. Stereotypes and smut: The (mis) representation of non-cisgender identities by text-to-image models. *arXiv preprint arXiv:2305.17072* (2023).

[112] European Union. 2024. Artificial Intelligence Act. https://artificialintelligenceact.eu/recital/44/ Accessed: 2024-09-05.

[113] Sandra Wachter. 2021. 'Every smile you fake' — an AI emotion-recognition system can assess how 'happy' China's workers are in the office. https://cyber.harvard.edu/story/2021-06/every-smile-you-fake-ai-emotion-recognition-system-can-assess-how-happy-chinas Accessed: 2025-01-13.

[114] Yvonne Wechuli. 2024. Passing as Able-Bodyminded, Disabled, or Supercrip: Rethinking Impression Management Strategies Through a Disability Lens. *Symbolic Interaction* (2024).

[115] Meredith Whittaker, Meryl Alper, Cynthia L Bennett, Sara Hendren, Liz Kaziunas, Mara Mills, Meredith Ringel Morris, Joy Rankin, Emily Rogers, Marcel Salas, et al. 2019. Disability, bias, and AI. *AI Now Institute* 8 (2019).

[116] Lauren Wilcox, Renee Shelby, Rajesh Veeraraghavan, Oliver L Haimson, Gabriela Cruz Erickson, Michael Turken, and Rebecca Gulotta. 2023. Infrastructuring care: How trans and non-binary people meet health and well-being needs through technology. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.

[117] Jakob Wirth, Christian Maier, Sven Laumer, and Tim Weitzel. 2019. Perceived information sensitivity and interdependent privacy protection: a quantitative study. *Electronic Markets* 29 (2019), 359–378.

[118] Philip Fei Wu, Jessica Vitak, and Michael T Zimmer. 2020. A contextual approach to information privacy research. *Journal of the Association for Information Science and Technology* 71, 4 (2020), 485–490.

[119] Claire Xu. 2023. HarmonyVisage: Ethical Facial Emotion Dataset Using Advanced Generative AI. In *2023 IEEE 14th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)*. IEEE, 0558–0563.

[120] Baobao Zhang. 2022. Public opinion toward artificial intelligence.

[121] Shikun Zhang, Yuanyuan Feng, Lujo Bauer, Lorrie Faith Cranor, Anupam Das, and Norman Sadeh. 2021. "Did you know this camera tracks your mood?": Understanding Privacy Expectations and Preferences in the Age of Video Analytics. *Proceedings on Privacy Enhancing Technologies* (2021).

[122] 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.

[123] Annuska Zolyomi and Jaime Snyder. 2024. An Emotion Translator: Speculative Design By Neurodiverse Dyads. In *Proceedings of the CHI Conference on Human*

Factors in Computing Systems. 1–18.

A Appendix

Table 3: Demographic Overview of Participant Sample

Variable	Category	Count
Race	American Indian or Alaska Native	15
	Black or African American	135
	East Asian	21
	Hispanic or Latino	59
	Middle Eastern	9
	Mixed race, multiracial, or biracial	122
	Native Hawaiian or Pacific Islander	3
	Prefer to self-describe	3
	South Asian	10
	Southeast Asian	25
	White	340
Gender	Cis Man	256
	Cis Woman	348
	Nonbinary (not trans-identified)	31
	Nonbinary Trans	62
	Trans (Unspecified)	3
	Trans Man	22
	Trans Woman	17
	Unsure/Additional	3
Disability	Deaf or serious difficulty hearing	11
	Blind or serious difficulty seeing	8
	Chronic illness	51
	Learning disability	3
	Mobility limitation	43
	Neurodiverse	65
	Speech or language impairment	3
	Mental health condition	71
	Motor limitation	8
	Prefer to self-describe	11
	No disability	468
Income Level	Less than $20,000	103
	$20,000 to $34,999	91
	$35,000 to $49,999	130
	$50,000 to $74,999	149
	$75,000 to $99,999	92
	$100,000 to $149,999	108
	$150,000 to $199,999	40
	$200,000 to $249,999	17
	$250,000 or more	12
Education Level	Less than high school	6
	High school or equivalent	108
	Some college	270
	Bachelor's degree	226
	Some graduate school	28
	Master's or professional degree	89
	Doctoral degree	15

Data Input	**White *M(SD)***	**POC *M(SD)***
Precise Location	5.73 (1.53)	5.63 (1.70)
Spoken Words	5.64 (1.53)	5.47 (1.62)
Digital Comm.	5.54 (1.48)	5.40 (1.65)
Facial Expressions	5.50 (1.58)	5.40 (1.68)
Voice	5.45 (1.56)	5.32 (1.69)
Body Movement	5.17 (1.66)	5.10 (1.73)
Device Usage	5.28 (1.67)	5.14 (1.82)
Written Words	5.44 (1.58)	5.23 (1.67)
Eye Movement *	5.36 (1.68)	5.12 (1.75)
Non-facial Biometrics**	5.30 (1.67)	5.03 (1.80)
Environment *	4.90 (1.56)	4.68 (1.73)

Table 4: Mean and Standard Deviation (SD) for Emotion AI Input Data by Race. $^{**}p < 0.05, ^{*} p < 0.1$.

Data Input	**Cis Men *M(SD)***	**Cis Women *M(SD)***	**Gender Minority *M(SD)***
Precise Location **	5.46 (1.65)	5.66 (1.65)	6.13 (1.41)
Spoken Words	5.36 (1.58)	5.54 (1.59)	5.94 (1.48)
Digital Comm.	5.26 (1.60)	5.43 (1.59)	5.92 (1.38)
Facial Expressions **	5.18 (1.66)	5.47 (1.66)	5.87 (1.40)
Voice **	5.16 (1.63)	5.34 (1.66)	5.87 (1.47)
Body Movement **	4.91 (1.68)	5.07 (1.72)	5.70 (1.55)
Device Usage	5.11 (1.71)	5.16 (1.78)	5.51 (1.72)
Written Words	5.19 (1.56)	5.32 (1.65)	5.59 (1.67)
Eye Movement **	4.91 (1.70)	5.23 (1.77)	5.81 (1.47)
Non-face Biometrics	4.82 (1.77)	5.23 (1.75)	5.59 (1.58)
Environment	4.47 (1.57)	4.93 (1.65)	5.00 (1.75)

Table 5: Mean and Standard Deviation (SD) for Emotion AI Input Data by Gender Identity. $^{**}p < 0.05, ^{*} p < 0.1$

Data Input	**No Disability *M(SD)***	**Disability *M(SD)***
Precise Location	5.67 (1.62)	5.73 (1.63)
Spoken Words	5.52 (1.59)	5.60 (1.55)
Digital Comm.	5.40 (1.58)	5.59 (1.55)
Facial Expressions	5.39 (1.66)	5.55 (1.57)
Voice	5.40 (1.61)	5.34 (1.67)
Body Movement	5.09 (1.70)	5.21 (1.70)
Device Usage	5.14 (1.76)	5.34 (1.72)
Written Words	5.26 (1.64)	5.45 (1.60)
Eye Movement**	5.09 (1.76)	5.51 (1.60)
Non-facial Biometrics**	5.04 (1.79)	5.40 (1.62)
Environment	4.71 (1.68)	4.92 (1.60)

Table 6: Mean and Standard Deviation (SD) for Emotion AI Input Data by Disability Status. $^{**}p < 0.05, ^{*} p < 0.1$.

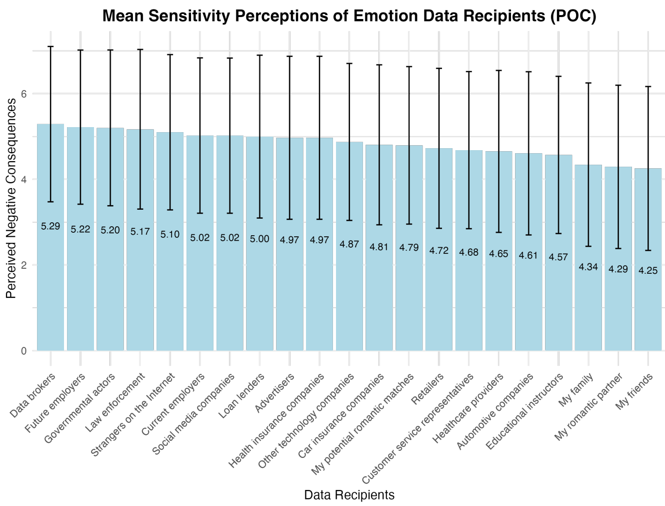

Figure 3: Average sensitivity perceptions across recipients of emotion data for people of color. Mean values are displayed within each bar, and error bars represent standard deviations

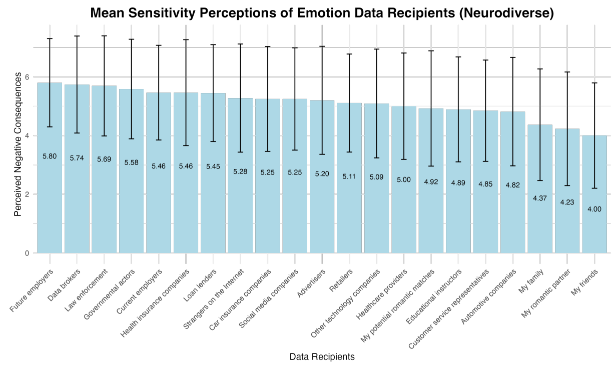

Figure 5: Average sensitivity perceptions across recipients of emotion data for neurodiverse individuals. Mean values are displayed within each bar, and error bars represent standard deviations

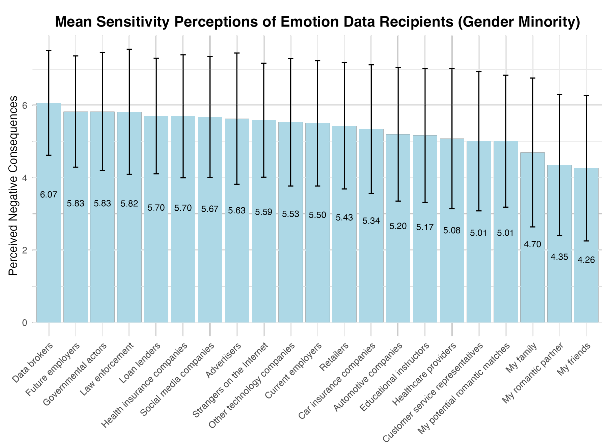

Figure 4: Average sensitivity perceptions across recipients of emotion data for nonbinary and transgender participants. Mean values are displayed within each bar, and error bars represent standard deviations