

Patent Applications as Glimpses into the Sociotechnical Imaginary: Ethical Speculation on the Imagined Futures of Emotion AI for Mental Health Monitoring and Detection

NADIA KARIZAT, University of Michigan, USA

ALEXANDRA H. VINSON, University of Michigan, USA

SHOBITA PARTHASARATHY, University of Michigan, USA

NAZANIN ANDALIBI, University of Michigan, USA

Patent applications provide insight into how inventors imagine and legitimize uses of their imagined technologies; as part of this imagining they envision social worlds and produce sociotechnical imaginaries. Examining sociotechnical imaginaries is important for emerging technologies in high-stakes contexts such as the case of emotion AI to address mental health care. We analyzed emotion AI patent applications ($N=58$) filed in the U.S. concerned with monitoring and detecting emotions and/or mental health. We examined the described technologies' imagined uses and the problems they were positioned to address. We found that inventors justified emotion AI inventions as solutions to issues surrounding data accuracy, care provision and experience, patient-provider communication, emotion regulation, and preventing harms attributed to mental health causes. We then applied an *ethical speculation* lens to anticipate the potential implications of the promissory emotion AI-enabled futures described in patent applications. We argue that such a future is one filled with mental health conditions' (or 'non-expected' emotions') stigmatization, equating mental health with propensity for crime, and lack of data subjects' agency. By framing individuals with mental health conditions as unpredictable and not capable of exercising their own agency, emotion AI mental health patent applications propose solutions that intervene in this imagined future: intensive surveillance, an emphasis on individual responsibility over structural barriers, and decontextualized behavioral change interventions. Using ethical speculation, we articulate the consequences of these discourses, raising questions about the role of emotion AI as positive, inherent, or inevitable in health and care-related contexts. We discuss our findings' implications for patent review processes, and advocate for policy makers, researchers and technologists to refer to patent (applications) to access, evaluate and (re)consider potentially harmful sociotechnical imaginaries before they become our reality.

CCS Concepts: • **Human-centered computing**; • **Computing methodologies** → **Artificial intelligence**;

Additional Key Words and Phrases: Emotion Artificial Intelligence, Emotion AI, Emotion Recognition, Mental Health, Data Subjects, Healthcare, AI Ethics, Ethical Speculation, Sociotechnical Imaginary, Patents

ACM Reference Format:

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. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 106 (April 2024), 43 pages. <https://doi.org/10.1145/3637383>

Authors' addresses: Nadia Karizat, nkarizat@umich.edu, University of Michigan, Ann Arbor, Michigan, USA; Alexandra H. Vinson, ahvinson@med.umich.edu, University of Michigan, Ann Arbor, Michigan, USA; Shobita Parthasarathy, shobita@umich.edu, University of Michigan, Ann Arbor, Michigan, USA; Nazanin Andalibi, andalibi@umich.edu, University of Michigan, Ann Arbor, USA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

ACM 2573-0142/2024/4-ART106

<https://doi.org/10.1145/3637383>

1 INTRODUCTION

Affective Computing refers to “*computing that relates to, arises from, or influences emotions*” [161, p.1]. Emotion artificial intelligence (emotion AI), a type of affective computing, aims to infer and/or influence individuals’ emotions [67]. Emotion AI has been applied in contexts such as hiring, workplace, education, the automotive industry, healthcare, advertising and more [131]. While the universality of emotions [71] has been challenged by scholars [13, 93], and there are pervasive concerns about emotion AI with regards to accuracy, bias and invasiveness [10, 16, 110, 190], the global emotion AI market is expected to reach 56 billion \$USD by 2024 [210].

Aligned with the longer history of affective computing’s focus on mental health [161], one of the expanding domains for emotion AI application is indeed in mental health. Emotion AI technologies are increasingly (proposed to be) used in mental health, including for diagnosis, treatment, and communication [88]. It is important to consider the impacts of mental health-related emotion AI technologies because mental health is a crucial site for negotiating tensions about individual and group autonomy, privacy and safety, serving as an important factor in one’s quality of life [46, 61, 64, 188]. Mental health is also deeply intertwined with physical and emotional health [145, 163]. As a result, mental health is a central part of how a person experiences their life. Therefore, the implications of emotion AI, especially for mental health applications, are delicate and critical for society and the general public. Emerging technologies, including emotion AI, are often deployed in mental health without adequate critical attention to their consequences [31]. As such, the overarching research question we address in this paper is: *what are the ethical implications of applying emotion AI in mental health?*

To address this question, we turn to patent applications because they serve as a glimpse into the sociotechnical imaginaries [103] and promissory futures [24] of emotion AI. These imaginaries provide a site for speculation about the possible impacts of the technologies their inventors describe [198]. We qualitatively analyzed emotion AI U.S. patent applications ($N = 58$) describing technologies that aim to monitor and detect emotions and mental health conditions. We interrogate the imagined potential futures of emotion AI in mental health care by asking: *What problems do mental health patent applications imagine that emotion AI solves—and for whom? In other words, how do emotion AI patents legitimize [199] these technologies and the sociotechnical futures they seek to facilitate?*

Using patents to better understand sociotechnical futures is not without precedent. For example, Science and Technology Studies (STS) scholars use patents to make claims about the past, present and future of technologies and their implications [32, 49, 62, 99, 104, 148, 151, 182, 200], such as the case of Delfanti and Frey who reviewed Amazon’s patents in order to better understand imagined futures of work and speculate on the impacts of “humanly extended automation” [62]. Patents exist as the material form of an organization’s or individual’s aspirations for their technological futures [49, 62, 99, 182]. In Computer-Supported Cooperative Work (CSCW), scholars have begun to also use patent applications as a way to speculate on the future deployment of technologies, such as the application of emotion AI in the workplace [28]. Informed by this prior scholarship, we view patent applications as a window to examine how developers of new technologies envision the futures of a technology in order to discuss its implications. While patent applications do not necessarily reflect the present reality as they may not be granted nor implemented as described *if granted* [33], they provide insight into technologists’ imagination and the potential futures in which these technologies may one day exist [32, 49, 106].

As such, the language in mental health emotion AI patent applications can be viewed as legitimizing certain social practices [199]; acting as discourse in which inventors justify certain relations between actors (e.g. clinicians, those with mental health conditions) where their imagined technologies are desirable. Essentially, patent applications can help us understand, from the inventors’

perspectives, how they imagine their emotion AI should be used in the context of mental health care, and allow us to explore the proposed and implied uses that may be consequential for policy [144]. The language and descriptions found in patents exist as one piece of the puzzle composing what Hilgartner refers to as “vanguard visions”, ideas about the future produced by visionaries who have visions for progress not yet fully realized [95]. While US patent law does not require patent applicants to detail the evidence for their technologies, or prove they ‘work’ as described, or anticipate every use case for their inventions [154], we use patent applications to tap into vanguard visions of emotion AI to consider potential consequences and ways we might design or regulate emotion AI to prevent harmful outcomes. With regard to AI, these vanguard visions are in part realized through narratives of the future where AI is portrayed as the inevitable, society-altering, and desirable path for technological development [11].

Findings and Contributions. Our findings highlight that the imagined futures reflected in emotion AI mental health patent applications frame emotion AI as a solution to issues surrounding data accuracy, care provision and experience, patient-provider communication, emotion regulation, and preventing harms attributed to mental health causes. In interpreting our findings, we use an ethical speculation lens [78] to anticipate the potential implications of the imagined emotion AI-enabled futures described in patent applications. Ethical speculation uses imagination about technological futures as a way to speculate about the harms and ethical implications of a technology so that we may be proactive about preventing or mitigating such harms [78]. It is particularly helpful when the lack of transparency surrounding emerging technologies makes direct examination nearly impossible. Using this lens, we argue that these described technologies stigmatize mental health conditions (or ‘non-expected’ emotions), equating mental health with propensity for crime and a lack of agency. By framing individuals with mental health conditions as unpredictable and not capable of exercising their own agency, emotion AI mental health patent applications propose solutions that intervene in this imagined future: intensive surveillance, an emphasis on individual responsibility over structural barriers, and decontextualized behavioral change interventions. Using ethical speculation, we articulate the consequences of these framings, raising questions about the role of emotion AI as positive, inherent, or inevitable in health and care-related contexts. We also consider the implications of our work, including the potentials for an ethical speculation lens [78] to be incorporated into formal patent review processes by regulatory bodies like the U.S. Office of Science and Technology Policy, and how existing regulatory structures like the U.S. Patent and Trademark Office might bolster certain patent evaluation criteria to incentivize more just technologies and dynamic innovatory modes of technological development. Lastly, we discuss how policymakers, researchers and technologists may refer to patent applications to critically evaluate potentially harmful sociotechnical imaginaries before they are mainstream.

2 RELATED WORK

2.1 Sociotechnical Imaginaries, and Promissory Futures as articulated in Patent Applications

When inventors build their material objects, they are also imagining social worlds. Jasanoff and Kim [103, p.9] define these as sociotechnical imaginaries: “*collectively imagined forms of social life and social order reflected in the design and fulfillment of nation-specific scientific and/or technological projects*”. These imaginaries then shape both policies, and visions of technological progress [24, 103]. For example, expectations of scientific and technological innovation can trigger new technological domains with promised value, even when the deployment of such technologies may not be widespread or regularly used [129]. These visions of technological progress and promissory

discourses shape the expectations and subsequent action of different actors and groups by “*disclosing possibilities but also colonizing the future, imposing own accounts of the inevitable or the desirable*” [159, p.6]. In fact, Bareis and Katzenbach have found that the rhetoric in several countries’ AI policy documents “*establish[es] AI as a given and massively disrupting technical development that will change society and politics fundamentally*” [11, p.875]. Responses to technological futures invoked by promissory discourse may include actions like the mobilization of resources, the passing of policies to support development, as well as lead to the general public feeling they hold limited control over their own technological futures due to seemingly inevitable technological ‘progress’ [24].

Imaginations and discourses of promissory futures exist in both formal spaces such as legal or institutional contexts, as well as more general, public spaces inspired by innovation and emerging technology that increasingly speak to each other [172]. One place that these sociotechnical imaginations can be found is in patent applications. In patent applications, inventors describe the technology and its intended uses. These descriptions function as promissory discourses that invoke potential futures by constructing a sociotechnical imaginary [198]. Trends reflected in patents (applications and granted) can “*suggest [an] industry’s strategic direction and what developments are likely coming*” [198, p.100]. In this study, each patent application works to construct a sociotechnical imaginary of emotion AI that, when analyzed together, can allow larger patterns of promissory futures to emerge and be critically analyzed, such as patterns in the ways individuals are conceived.

For example, patent applications configure the rights, responsibilities, and appropriate behaviors of users [153, 211]. Studying how people are configured by patent applications helps to move from the notion of the data subject to the notion of data subjectification, or how data subjects¹ are produced by the technologies that these subjects interact with/use. Particularly in the realm of mental health care, understanding how subjects are produced, how subjectivity is shaped by technology, and the implications of technology use for agency, privacy, and control, are crucial for understanding the place that emotion AI has in contemporary US society. The articulations of sociotechnical futures contained in patent applications help us critically interrogate the relationships between patients and providers, and citizens and social institutions.

Patent applications give us insight into how inventors imagine applications of their technologies and, rather than granted patents, allow us access to the varied *imagined* futures of emotion AI for mental health – enabling researchers, regulators, and other actors to consider emotion AI’s implications and promissory futures in the mental health context. As we investigate the sociotechnical imaginations of emotion AI as articulated in patent applications in this work, we do not claim or intend to predict the future of emotion AI, but instead make possible ethical speculation [78] so that researchers, regulators, and other stakeholders may be made aware of general trends and inventors’ perspectives in the possibilities of emotion AI for mental health and consider its implications and promissory futures. Ethical speculation [78] has been used in the past to make sense of the impact emotion AI technologies might have in the workplace and what that might mean for the future of work [28], foreshadowing a future of work made possible by emotion AI and how it might require increased emotional labor from employees, and threaten worker autonomy [28]. Drawing on work reviewed here, in this paper, we recognize patent applications as artifacts that provide a window into the state of a technological domain (e.g., emotion AI technologies) and its projected futures [1, 60, 111].

¹In this paper, we refer to the person(s) who’s data is collected for purposes of detecting or monitoring emotion as data subject(s).

2.2 Emotion Theories and Tensions for their Broad Application in Emotion AI

Emotion has been of interest to psychologists, anthropologists, linguists, computer scientists, and more [175]. Theoretical approaches to emotions in emotion AI, largely informed by psychology, include: (1) the evolutionary approach [175] within which the Basic Emotion Theory (BET) suggests anger, disgust, fear, happiness, sadness, and surprise as the six basic emotions—this model is used in most emotion AI systems [36], despite being critiqued for scientific validity and assumptions [12, 14, 84, 147, 174] and being in contrast to emotions as socially constructed and dynamic [23]; (2) the appraisal approach includes theories that associate emotions with dimensions like arousal (excitement caused by a stimulus) and valence (liking toward a stimulus) [175]; (3) the constructionism approach considers emotions from a social-psychological stance [175]. Scholars have critiqued emotion AI systems for simplifying and reducing the complexities of human emotions and experiences into something easily scalable for consumption, inference and influence [57], encoding a limited set of emotions as a stand-in for the countless complex emotional experiences around the world [56]. These approaches to emotions incorporated in emotion AI systems renders emotion as something that is *explainable* and therefore, (supposedly) predictable [56].

Scholars have contested the broad application of the basic emotion theory, as well as the fundamentals behind emotion AI [93, 176, 196]. The idea of universal emotion and expression [71] is criticized with claims that there is little to no evidence that one is able to reliably deduce emotion from facial expression [14, 15, 100], and scholars suggest that Ekman's basic taxonomy of 6 emotions is insufficient to encapsulate the breadth of human emotion and its subsets (e.g. joy or pleasure as a subset of happiness) [55]. Other research has found that automated emotion recognition of facial expressions, audio speech, or text produces race, age, and gender biases through incorrect identification [86, 110, 169, 170, 195], and there has been increasing uncertainty regarding the validity of the data and modeling used in emotion AI systems [191]. The inference of individual and communities' emotions has been dubbed by some to be fundamentally invasive and manipulative [2, 82, 215]. Critics of emotion AI more broadly argue that AI systems encode emotions in ways that are inherently reductive and finite, and therefore inadequate for representing the infinite and varied emotional experiences that exist around the world [57, 93].

Instead of taking the claims made in patent applications about the capabilities of the described emotion AI technologies at face value, we are able to trace the premise of these technologies—the ability to categorize, classify and identify emotion—to prior theories that have been critiqued and contested. Through this understanding of the various emotion theories that may inform emotion AI *and* critiques identifying their limitations, we take a critical approach in our analysis of emotion AI mental health patent applications in this work, specifically as it pertains to emotion AI futures that imagine a world where emotions can be accurately detected and monitored and where these insights address myriad problems in health, safety and care-related contexts. We understand that emotions are complex and difficult to neatly define or explain the ways they manifest or present. By understanding theories undergirding emotion AI and their critiques, we are more prepared to question and critically make sense of the sociotechnical imaginaries of emotion AI for mental health without being swept up by AI hype and these technologies' grand promises, referred to by Narayanan as AI snake oil [136].

2.3 Emotion AI Applications in Mental Health

Emotion AI includes a variety of technologies using “*affective computing and artificial intelligence techniques to sense, learn about, and interact with human emotional life*” [132], and often uses a wide range of data (e.g. voice, text, online behavior, facial expressions, gait, biosignals, etc.) [131, 132, 179]. The convergence of AI and medical treatment has been dubbed by some to be “*inevitable*” [122]

and others to be fundamentally desirable for healthcare [123]. On the one hand, the sociotechnical imaginary of AI applications in healthcare radiates positive attitudes towards the deployment of AI technologies in health and care-related contexts [16, 204], with promises of reducing healthcare costs [141, 152], providing increasingly tailored medical care [16, 52, 152] and identifying rapid diagnoses of health conditions [16, 152]. On the other hand, the application of AI in health and care-related contexts may introduce or exacerbate issues for healthcare providers and patients. Scholars argue AI may further bias and unfairness in medical settings [3, 4, 40, 109, 152, 158], leading to unequally distributed harms and benefits [3, 4, 152, 158]. AI has the potential to intensify existing bias against disadvantaged groups in healthcare unless deliberate human choices are made to counteract and mitigate biases in AI training data, design and application [40].

Although AI technologies attempt to write their own futures, there is little evidence that they can be successfully deployed in clinical practice [109]. Emerging evidence, such as Wong et al.'s [209] study of the rapid adoption in the U.S. of a poorly-performing predictive algorithm for detecting sepsis, is exemplary of the dangers of adopting automated technologies and assuming they will work in practice as promised. Researchers have deduced values of transparency, privacy and trust as concerns of clinicians and patients with regards to AI use in the context of clinical chatbots [43], clinical decision support systems [142, 204], with broader implications for these values in the medical system overall [16, 40, 119, 120]. The application of AI in clinical practice has been shown to negatively impact patient adherence to treatment and care instructions [43], as well as adversely impact the patient-provider relationship [43, 119, 152, 185]. Actors who choose to deploy and implement AI in healthcare settings should reconcile issues of incompatibility with clinician and patient needs [150, 204], new usability challenges and other hurdles to implementation [109, 152, 185, 204], as well as an (in)ability to fit into a clinician's workflow [76, 141, 185, 204].

In the realm of mental health care, *emotion* AI, as a unique class of AI, is being envisioned and deployed for emotion (and other affective phenomena) monitoring and detection. Because they aim to associate certain 'objective' criteria (e.g. vocabulary, speech patterns) with varying mental health or emotional states, emotion AI technologies echo a legacy in early psychiatry, but at a breadth and level of granularity not possible previously [70]. For example, scholars have designed smart clothing [214], smart mirrors [189], video game systems [205], and audio analysis tools [164] that claim to be able to detect and monitor emotions, as well as mental health conditions based on myriad data collected (e.g., physiological, audio, facial expression data). This prior work highlights how a wide breadth of attributes of data subjects are viewed as collectable data points that can produce emotional, mental health and behavioral observation and insights [68].

Researchers have worked to understand individuals' experiences with chatbots for mental health support [25, 44, 114] or the reasons people decide to know or not know their risk predictions for developing various mental health conditions [126], noting mixed benefits and a myriad of concerns. Emotion AI has also been shown to feel invasive to individuals' sense of privacy, causing emotions such as fear and helplessness, as well as bringing about concern of emotional manipulation [8, 28, 54, 181]. These concerns in tandem with the growing interest in emotion AI applications for mental health make it pertinent to consider the ethical implications embedded within the promissory futures of emotion AI – a goal we undertake through our exploration of emotion AI patent applications that monitor and aim to detect individuals' emotions and mental health status.

3 METHODS

We assembled the dataset for this paper from a larger dataset of emotion AI patent applications we developed as part of a larger project. The initial dataset was narrowed down to a corpus of 58 mental health emotion AI patent applications concerned with monitoring and detecting emotions and/or mental health—the data our analysis draws from.

3.1 Establishing the Dataset

We collected patent applications rather than *granted* patents because we are interested in the futures their descriptions imagine; whether or not they are granted is immaterial. Furthermore, because emotion AI is relatively new, there are many more applications than patents (all applications are published within 18 months, whether or not they are granted). We scoped our initial search to United States patent applications using the InnovationQ Plus database. Throughout the process of establishing our dataset, we frequently consulted with a patent librarian for guidance.

We queried patent applications using a combination of Cooperative Patent Classification (CPC) codes (i.e. codes used to organize patent applications by their common subject matter) and keywords. We used a broad and more general selection of CPC and keywords in our initial query aiming to pull emotion AI patent applications. We applied the CPC parent code G06 (Computing, Calculating, Counting) to exclude patent applications about human emotion recognition². Next, we queried for specific keywords in patent applications' abstracts to compile patent applications concerned with emotion or related concepts: "emotion detection" OR "emotion recognition" OR "emotion prediction" OR "mood detection" OR "mood recognition" OR "mood prediction" OR "affect recognition" OR "affect detection" OR "affect prediction" OR "mental health" OR "mental illness" OR "digital phenotyping" OR "emotion" OR "mood". We chose to omit "affect" as a stand alone keyword due to it surfacing unrelated results to the query results due to its second meaning "to have an effect on". While this keyword and CPC code search may not have pulled every relevant patent from the database, it does provide a thorough dataset of patent applications where computational emotion recognition was salient enough to the applicants that they included these keywords in the abstract.

After collecting the larger dataset ($N=1163$ patent applications), we identified the smaller group focused on emotion AI. To do this, we manually checked the patent keywords and the sections of patent applications explicitly mentioning emotion. We included patent applications that detected and monitored emotions of 1) individual humans and groups of humans, 2) songs, movies and other content created by humans, and 3) patent applications that accepted self-reports of emotion data if they were later manipulated or analyzed (e.g. technology that analyzed individual reports to estimate the emotions of a group of people). We did not include patent applications that made use of human emotion in order to re-purpose it, such as systems that allowed for the individuals to communicate emotions to another without analyzing explicit emotion data (e.g., sending emoji over an IM system). This process resulted in 879 patent applications.

We then used the Partnership on AI's taxonomy from their "The Ethics of AI and Emotional Intelligence" report [88] to categorize the 879 patent applications (adding emergent categories as needed) for their domains (e.g., mental health, advertising), identifying 58 patent applications that




Fig. 1. Establishing our Final Dataset

²Human emotion recognition refers to *humans* recognizing or detecting other humans' emotions, not algorithms claiming to detect or infer emotions. As we were interested in emotion recognition and emotional AI technologies, we applied CPC parent code G06 to make sure that the patents pulled in our search directly involved computing technologies and *excluded* patents referring to *humans* being able to recognize humans' emotions.

describe technologies that monitor and claim to detect mental health conditions and emotions on behalf of care providers, personal networks, the data subject themselves and society more broadly—the data we used for this analysis.

3.2 Dataset

In establishing our dataset, we narrowed down a larger dataset of patent applications describing computational emotion recognition technologies to arrive at 58 patent applications describing specifically emotion AI technologies that monitor and detect mental health conditions and emotions. For a table that details the patent applications in our corpus, please refer to the the Appendix. Our dataset captures our goals of accessing the sociotechnical imaginary of emotion AI for the detection and/or monitoring of mental health conditions and emotions. We recognize that our dataset represents a *snapshot* in time of these sociotechnical imaginaries of emotion AI. We do not claim, and caution against, viewing this dataset and its subsequent analysis in this paper as a timeless capture of the promissory futures of emotion AI for mental health. Instead, we hope researchers and readers alike view this dataset as emerging from a specific space and time. And, from this snapshot in time, we are able to draw and discuss the ethical implications of emotion AI for mental health that may be relevant in other sociotechnical futures of emotion AI.

Our analysis and interpretation provides a partial, yet insightful and unique, view into emotion AI's sociotechnical imaginary that may indeed be more expansive than what is described in the patent applications of our dataset. We do not argue that ethical speculation as a lens, and patent applications as a source for accessing sociotechnical imaginaries can provide an all-inclusive, god's-eye view of the sociotechnical imaginaries of emotion AI for mental health. However, our approach does provide *a* view into a segment of the promissory discourses that invoke the sociotechnical futures for mental health near and far [198], and as such, warrant speculation.

3.3 Analysis

The first author first read all patent applications and extracted text relating to the described technologies' various applications and the problems their technologies propose to address; this was motivated by our goal to understand the ways inventors legitimize [199] emotion AI inventions in the mental health context. As mentioned in Section 2.1, patent applications are one space where sociotechnical imaginaries may be found, providing glimpses into the imagined futures of technologies like emotion AI for mental health. Inventors use patent applications not only to describe how their proposed invention works, but also to make their case for why the invention is useful. These justifications invoke sociotechnical imaginaries—visions of a potential future world that the invention is intended to create or prevent. We conceptualize patent applications as data sources with descriptive text providing insight into the broader sociotechnical imaginaries of emotion AI for mental health—aligned with prior work [32, 49, 62, 99, 104, 148, 151, 182, 200].

We then used an open-coding approach [51] to analyze these data. Using Dedoose, a computer-assisted qualitative data analysis software (CAQDAS), the first and last authors open-coded 20 patent excerpts to assemble a codebook with codes to better understand the proposed applications and solutions to problems the emotion AI technologies described in each patent. The codebook consisted of parent and child codes. For example, the parent code “Health and Well-Being Impacts” had child codes such as “Improve and Support Emotional State, Improve and Support Mental Health, Improve and Support Well-Being and Improve and Support Physical Health.” The two authors met to discuss their codes and observations. The first author then coded the remaining data and kept a written record of memos and notes to frequently discuss with the last author. To maintain consistency with codes and their application, the first and last author also regularly discussed the meaning and application of codes throughout the coding process, attending to the ways these

codes emerged and reflected the data in our patent corpus. Following this process, we applied axial coding[51] to group codes into larger categories and themes, noting connections between themes. The first and last authors met throughout this process to discuss and refine themes and relationships between the groups of codes. All authors engaged in discussions about these themes, refining our final interpretations of them. When coding the data, the first author coded direct and descriptive claims of usefulness within the patent applications, and later accounted for implied or underlying uses of these technologies when grouping codes into broader categories and themes surrounding the imagined futures of emotion AI for mental health. Altogether, these direct and implied claims of emotion AI's usefulness provided ample insights for us to speculate on the ethical implications of emotion AI technologies applied to mental health.

3.4 Limitations

Several aspects of the patent applications shape how we are able to use them to understand the sociotechnical imaginaries of emotion AI. For example, because these patent applications were submitted in the U.S., the imagined emotion AI-enabled mental health futures accessed via the patent applications may be more U.S.-centric, reflecting the geography and contexts in which inventors' envisioned their technologies' being implemented. Additionally, the format of patent applications encourages broad description in order to prevent limiting potential use-cases for a technology, only requires information on three evaluation criteria in the U.S. [157]: *usefulness, novelty and non-obviousness*, and does not mandate any commentary on harms or ways to mitigate potential harms. As a result, patent applications themselves are inherently not fully representative of the potential future applications of a described technology. Still, as noted in our review of the literature, patent applications do provide valuable insights on the promissory futures of emerging technologies including emotion AI for mental health, allowing us to better understand how emotion AI technologies are envisioned and legitimized by inventors, and the implications they have for emotion AI-enabled sociotechnical futures. Future work can conduct interviews or design workshops to gather perspectives on emotion AI uses for mental health from various actors such as technologists, healthcare providers who (may) choose to deploy emotion AI technologies, and patients who may encounter these technologies. For example, a design workshop may prompt participants to imagine technologies that address mental health while not contributing to the ethical concerns we have identified (e.g. stigmatization, surveillance), or to consider what trade-offs are warranted in order to center and prioritize the individuals a technology is intended to support. Additionally, our patent search to develop the dataset may have excluded patent applications relating to emotion AI technologies not described as part of one of the domains of technology in our initial search, but may be appropriated for mental health purposes. Finally, some of our interpretations and examples undergirding the ethical implications of emotion AI technologies are bound to the U.S. context, such as the technologies' potential connection to law enforcement and police.

3.5 Researchers' Positionality

Our values inform this work, analysis, and interpretations – as expected and celebrated in qualitative interpretivist work [20, 113]. We imagine and aspire for a sociotechnical future where individuals are understood as complex, nuanced beings who deserve respect, dignity and understanding, and freedom from experiencing harm—where humanity and its complexity is centered and prioritized in any claims of technological 'progress'. In the mental health context, we wish for systems where individuals are seen as authorities over their own health, feelings and well-being, and able to enact agency and autonomy over their health and healthcare. We hold expertise in science, technology and medicine, emotional AI's societal implications, patents' politics, and information and computer

sciences. This combination of backgrounds make us well situated to analyze and speculate on the ethical implications of the imagined futures of emotion AI as revealed through patent applications. The first author is a doctoral student with a background in health informatics and social computing. The second author is an STS scholar with expertise in professional cultures and the construction of medical knowledge. The third author is an STS scholar whose research focuses on innovation politics and policy, including the patent system. The fourth author is a Human-Computer Interaction (HCI) and social computing researcher whose work addresses technologies' (including emotion AI) societal and ethical implications, with a particular attention to technological harms especially to marginalized communities. They have a background in both computer and information science, enabling them to understand both the technical and social aspects of patent applications, which they have done in prior work [28] as well.

4 FINDINGS

We find that the emotion AI technologies described in our patent application corpus that monitor and claim to detect emotional and mental health states promise myriad solutions to health, safety and care-related problems experienced by the 1) individuals being monitored, 2) health care providers, and 3) the broader society. Further, we show how patent applicants envision a future where emotion AI technologies address these problems and improve or support an individual's physical and mental health, emotional state, and overall well-being.³ An outline of our main findings is available in Table 1 at the end of this section.

Emotion AI patent applications for mental health uses propose a range of improvements to mental health care, mental health status monitoring, and data collection on mental health status—improvements that are motivated by articulations of need and invocations of a future where this need is met by the technology. We describe these technologies' promised solutions to issues of data accuracy, care provision and experience, patient-provider communication, mental health and emotion self-management, and safety for data subjects and society at large.

4.1 Solving Data Accuracy Challenges

We begin by detailing how patent applications commonly frame data inaccuracy as a problem to be addressed by the described emotion AI technologies. In doing so, patent applications portray data quality as a necessary precursor to the additional solutions proposed by these technologies and detailed in the remainder of our findings. Inventors describe emotion AI technologies as tools able to quantify 'mental health' and emotions, implicitly positioning these technologies as more reliable and accurate than qualitative and experiential knowledge of mental health. Patent applications frequently articulated emotion AI solutions as improving data accuracy for clinical, safety and general health contexts. They attributed increased data accuracy to a variety of the described emotion AI inventions' features: 1) the ability to collect and use data from multiple sources, and 2) framing emotion AI as tapping into objective and reliable data sources, such as [heart rate, vocabulary, voice pitch, etc.], in contrast to self-reported or human collected data.

³It is important to note that many of the results focus on broader visions of the imagined futures of emotion AI for mental health, as opposed to concrete specifics such as the exact mental health conditions or types of care providers patent applications identify in their technologies' imagined uses. This is an intentional choice; our goal is for readers to sit with the broader visions of these sociotechnical futures and their implications, as opposed to implications of emotion AI for a *specific* mental health condition or type of care provider. Equally important is that many patent applications in our corpus included sentiments of the broad nature and many potentials of their described technologies', emphasizing how the technologies' implementation is not limited to the example embodiment described in the application. They often refer to emotions, mental health, and providers in broad terms, perhaps to not limit any future implementations of their described emotion AI technologies. We present the results in this section as a description of the *broader trends and imagined uses* of emotion AI technologies for mental health.

4.1.1 More Complex and Varied Data Sources. Many patent applications argue that emotion AI's capacity to integrate multiple data sources produces more accurate detection and monitoring of emotional status, or mental health. For example, P14 says its invention uses multiple data sources to improve the accuracy of information collected and derived from the data: *"In the emotion analysis device according to the embodiment of the present disclosure, collection and analysis of the facial image, the acceleration information, and various physiological information are integrated, and the emotional health state of the target object is comprehensively analyzed and determined, which overcomes misjudgment caused by emotion recognition according to a single factor, so that the emotional health analysis and management are more comprehensive and accurate"* (P14). P14 described using multiple data sources as a way to generate more accurate emotional health detection and subsequent management. Similarly, P44 links combining self-reported data with data collected by emotion AI as producing more accurate information for clinical care: *"For example, by combining the self-reported SRM with automated, passive sensing by the smartphone for a short period of time, the models can further be individualized to each patient. This may further improve the accuracy of clinical information."* (P44). By combining multiple data sources as opposed to relying solely on one source, patent applications argue their described emotion AI technologies will improve accuracy and generate more reliable results when detecting and monitoring emotions and cognitive- and emotion-related health outcomes (e.g. depression, anxiety, dementia.)

4.1.2 Increased Objectivity and Reliability Compared to Human-Reported and -Collected Data. Patent applications commonly reinforce the idea that data collected by emotion AI is more "objective" and thus, reliable, compared to data collected or self-reported by people experiencing emotions, arguing frequently that the described inventions are *"[...]human-independent, less subjective..."* (P25). In our corpus, this claimed increased objectivity is often discussed in the context of self-reported data, observational data and survey or interview data collection in health care settings. P57, a patent application that articulates an invention to help detect and assess for risk of mental health conditions, positions its emotion AI technology as a solution to the challenges of self-reported data and those collected in clinical interviews: *"Besides being a costly and time-consuming process, the main problem with such measures is the lack of accuracy and reliability due to their reliance on patient self-reporting, as well as the lack of attention put on observable behavior. These diagnostics can be strongly biased by subjective observation and the absence of real-time or more natural assessments"* (P57). The idea of data collected by emotion AI as more objective, reliable and accurate than human-collected or -reported data is reiterated in other patent applications that promise *"[...]removing the involuntary but inevitable bias introduced by the patient recollecting the events"* (P34) or that explicitly state a goal *"to move away from measurement subject to recall biases of both the patient... and the clinician, moving towards objective measurement that will shed light on the components of an individual's day that contribute to the health of the patient..."* (P41). In these patent applications, the patient is not deemed a reliable narrator of their own affective experiences, and healthcare providers' human observation and data collection are viewed as liabilities to the 'truth.' In other words, emotion AI is framed as a neutral tool to bring to light a supposedly objective reality, not prone to the biases or weaknesses of the human recollection or collection of a data subject's own experiences.

In a few instances, emotion AI is described as a tool to validate self-reported data, therefore allegedly increasing overall data reliability and accuracy. For example, P12, a patent that describes analyzing voice to detect emotional state, explains that the individual's state *"can further be corroborated and or improved, through cross-referencing the individual's self-reported data with other biometric data, such as heart rate data, etc., when a particular state is self-reported and detected and recorded by the system..."* (P12). While some patent applications mention incorporating and

validating self-reported data into emotion AI systems, it is unclear how the data collected is ‘weighed’ or ‘valued’ in determining the final output of one’s detected emotional or health status.

4.2 Improving Care Provision and Experience

Patent applications in our corpus frequently describe emotion AI technology as solving problems relating to the provision and experience of care. We detail the ways emotion AI technologies are imagined as solutions to health and care-related problems, envisioned to improve and mediate the provision and experience of care between health care providers and data subjects by 1) assisting or replacing care providers in their duties (e.g. diagnoses, treatment), 2) improving care efficiency, and 3) allowing for more effective care to achieve desired outcomes.

4.2.1 Emotion AI Assisting and/or [at times] Replacing Healthcare providers. Many patent applications describe emotion AI assisting or fully replacing clinical providers and caregivers in their duties by supporting or independently providing treatment, managing medications and prescription processes, as well as supporting diagnoses and screenings. The patent applications situate this imagined emotion AI-assistance within the healthcare context where care providers face challenges to providing adequate care with limited time and high costs of healthcare. Some patent applications highlight the ways carrying out screening and diagnosis is “*costly and time-consuming*” (P57), and others emphasize how little time exists for clinicians “*to ask a predetermined number of questions to or to discuss a predetermined number of matters with the patient*” (P27). By highlighting these challenges, and by acknowledging the growing burden placed on health care providers [168, 207], patent applications envision emotion AI assisting, and at times fully replacing, care providers in their duties as desirable and beneficial.

4.2.2 Diagnosing and Screening for Health Conditions. Patent applications in our corpus conjure up emotion AI as supporting care providers with their responsibilities of detecting and diagnosing issues in mental, emotional, physical or cognitive state by way of directly performing or facilitating the screening and diagnostic testing of patients under their care. For example, P53 describes how emotion AI performs textual analysis when monitoring text by a patient (e.g. social media posts, emails, etc.) to generate “*accurate, non-invasive early detection or diagnosis of cognitive deficit or mental illness*” (P53) that may be presented to care providers and, in some cases, the individual themselves. (We note that what constitutes invasiveness and who decides is up for debate and highly contextual.) Other patent applications propose emotion AI that assist care providers in confirming or reaching conclusions around diagnoses, such as P54 that will ask follow-up questions of the patient when it detects a health state “*to confirm this determination. Answers to such questions may be used by medical professionals to make a medical diagnosis*” (P54). The patent applications envision emotion AI as benefiting diagnostic procedures through a claimed ability to detect health concerns and conditions more accurately and quickly than existing means (e.g self-reported data, clinical observation, clinical interview) as discussed in section 4.1.

4.2.3 Providing Treatment and Managing Medications. In addition to supporting diagnostic procedures, patent applications in our corpus consistently highlighted the ways their described technology would assist care providers in the provision of treatment or directly provide treatment to the data subject without a care provider. For example, P8 articulates an emotion AI that monitors the data subject’s emotions to allow for a better understanding of one’s emotional triggers by monitoring trends of data subjects’ emotions for a period of time and sharing these with others (e.g. clinicians, therapists) for purposes including informing treatment. In describing this scenario between a therapist and their client, P8 explains how the emotion AI “*upon receiving instructions from the user to share a list of instances of the target emotion—send the list of instances of the target*

emotion to the licensed therapist's device running the companion application. The therapist's device running the companion application can then prompt the user and the licensed therapist concurrently to discuss particular instances of the target emotion and prompt the licensed therapist to select coaching activities for the companion application to serve to the user in response to the system detecting future instances of the target emotion..." (P8) In this example, the emotion AI's stated abilities of detecting and monitoring emotion result in recommended therapeutic interventions and coaching activities to assist with treatment. In doing so, the emotion AI is imagined to support the care provider in providing treatment to improve the well-being and emotional state of the data subject. Additionally, patent applications like P15 claim their technology will identify a medication beneficial for a user given detected emotions and send the medication's name "to a pharmacy for fulfillment of a prescription or to a medical health professional for verification and writing of a prescription" (P15). In both instances, patent applications describe emotion AI technologies that aim to assist care providers in treatment and care provision.

At times, however, patent applications described emotion AI inventions that would provide interventions and treatments *directly* to the data subject, rather than proxied through a provider. P43 articulates emotion AI directly administering an intervention to those identified at "risk [for mental illness] and begin giving trace amounts of a new or existing drug, such as an anti-psychotic drug, to forestall problems from occurring" (P43). In another example, P42's described emotion AI detects one's mental state and, if deemed necessary, "automatically select[s] a corrective, therapeutic and/or enhancing action and output the selected action... may include soothing the user by using ambient music, recommending relocating the user due to, for example, undesired people around or an unsuitable room or area" (P42). These examples highlight how some patent applications imagine emotion AI as helping to deliver treatment or intervention without directly involving the care provider or using care providers' limited resources (e.g. time, supplies).

4.2.4 Emotion AI Improving Technical and Productive Efficiency of Care. The patent applications celebrate technological futures where emotion AI automates burdensome tasks for care providers and increases clinical reach (e.g., allows for a clinician to observe multiple people at once, remotely monitoring patients outside of an office) with reduced resources (e.g. hospital costs, time). Patent applications we analyzed also repeatedly imagined their inventions would lead to *faster* intervention and diagnosis for data subjects (implicating both technical and productive efficiency⁴), leading to better health outcomes overall. For example, P29 directly links its described technology's early diagnosis benefits to patients' having better health outcomes: "Early detection and treatment significantly improves patients' response to treatment and could prevent a progression to full relapse by prompting adequate clinical intervention. The proposed solution is deliberately designed to be time effective and has the potential to readily provide assessment reports..." (P29). In this case, emotion AI technologies are justified as enabling faster diagnoses, making possible better health outcomes for data subjects.

Patent applications articulate the ways emotion AI technologies might improve technical and productive efficiencies of care. Through automation and requiring fewer resources for equal or greater amounts of care, our corpus' patent applications articulate emotion AI technologies as improving technical and productive efficiency for care providers and their patients, the data subjects. These improved efficiencies imagined by patent application authors portray emotion AI with the ability to support health-care providers to deliver equal or additional amounts of care by optimizing resources.

⁴Technical efficiency in health care refers to achieving equal outcomes with less resources (e.g. cost, labor), and productive efficiency in health care refers to using different levels of resources to achieve more outcomes for the same cost (e.g. clinicians reaching multiple patients simultaneously)[149]

For example, P29 articulates emotion AI inventions as integral to improving technical efficiency by eliminating or reducing reliance on resources such as transportation to perform diagnostic procedures: *“Patients with mental illness may be unreliable. The patients are frequently late or not showing up on the scheduled time slot. The cost and logistics of transportation to and from the laboratory poses additional problems in the process. The present solution makes it possible for portable devices (e.g., laptops, tablets, smart devices, etc.) to host the entire diagnostic procedure”* (P29). The patent emphasizes challenges of time, cost, and transportation logistics as barriers to care that their envisioned portable emotion AI device helps navigate. In this example, technical efficiency is improved, by allowing for equal amounts of care (e.g. diagnostic procedures) with fewer resources (e.g. cost, time, transportation). Similar to when patent application authors frame data subjects’ own recollection of their experiences as barriers to the objective truth (as discussed in section 4.1.2), P29 frames the data subjects themselves as barriers to care that must be dealt with by asserting that patients with mental illness may be unreliable or careless with their time commitments.

Additionally, patent applications frame emotion AI technologies as helping healthcare providers deliver more care with similar levels of resources (e.g. time, labor, cost), increasing productive efficiency. For example, patent applications describe emotion AI technologies as allowing care providers to reach more patients, in more thorough ways, with less resources. P33 highlights this increased reach: *“Patient monitoring can also be done remotely...A doctor’s tablet or laptop can be used to monitor one or several patients and to perform more advanced comparisons and statistical analysis on them”* (P33). A doctor who would typically have to see and observe a patient in-person with limited time would now be able to remotely monitor the patient *and multiple other patients simultaneously*.

All in all, patent applications more commonly discuss what technical and productive efficiency means for clinicians, but not for patients (beyond requiring less frequent doctor visits). These patent applications frame the automation and increased assistance brought by their described inventions as positive for care providers to easily and rapidly provide equal or additional levels of care.

4.2.5 Emotion AI Boosting Effectiveness and Experiences of Care. In addition to increasing the efficiency of care delivery for clinicians, patent applications also imagine emotion AI that would help care providers improve the effectiveness of care experienced by data subjects and, thus, improve the desired health outcomes. In health care, effectiveness refers to whether an intervention achieves the intended or desirable results under normal circumstances [74]. These patent applications further the idea that emotion AI can support the healthcare provider or caregiver to deliver more effective treatments to the patient, while improving the individual’s experience and treatment compliance.

Monitoring and the large amounts of data collected by emotion AI are often credited with allowing for a *“continuity of treatment between the slew of social services, medical and mental health providers, and emergency and support personnel that is currently not available”* (P5). The continuous monitoring of data subjects celebrated in these described technologies is legitimized as allowing care providers to regularly monitor the ongoing experience of various treatments and medications to assess progress in the individual’s health and tweak personalized care as necessary for increased effectiveness of the care delivered. With increased understanding of how a patient is responding to treatment and medications, patent application authors argue that care providers would be able to improve efficiency of care by helping to personalize care treatment for the data subject. For example, patent P52 argued that its envisioned invention would allow for feedback from the emotion AI to care providers and data subjects to improve treatment efficiency: *“Feedback regarding which interventions are most helpful to which people could be provided directly through the intervention system (e.g., a prompt sent through phones telling a person that this exercise or technique is particularly helpful for you) or through therapist-provided feedback”* (P52). In this example, emotion

AI is imagined as helping individuals better understand which interventions are more beneficial for them, as well as assisting care providers in identifying treatments more effective for the individual.

Patent application authors also describe emotion AI as being able to improve the data subjects' overall care experiences by making personalized choices that reduce side effects to medication administered, decrease an individual's time commitment for treatment, etc. For example, P43 argues its invention has tremendous value for both care providers and data subjects with its claimed ability to facilitate *"better adjustment of medication, and possibly lower doses, having fewer side effects for the people, which would then make them more willing to stay on their medication, which is often a major problem..."* (P43) by emotion AI automatically detecting and monitoring reactions to mental health treatments. By emotion AI helping to shape treatments to reduce side effects and improve treatment experience, data subjects are imagined to benefit from this treatment personalization and thought of as more likely to adhere to treatment and experience better health outcomes.

Overall, we find that patent applications often imagine emotion AI technologies as providing value to health care providers through helping to assist, and at times fully replace, care providers in their responsibilities, boosting care efficiency, and improving care effectiveness and patient experience by making personalized care more easily achievable. Care providers are said to have less burden in care provision as well as more information to provide better care to data subjects. Data subjects, in this case patients, are the recipients of supposedly faster and improved personalized care that is thought to lead to improved health.

4.3 Helping Patient-Provider Communication

Patent application authors describe emotion AI technologies as solutions to improving communication between care providers and data subjects, and addressing challenges to empathy and a lack of understanding between patients and care providers. In most instances, emotion AI facilitating provider communication is justified as being able to positively impact an individual's emotional, mental or physical state, gather more thorough information from the individual, and to guide information delivery in ways that will be well-received by the individual.

Patent applications' described emotion AI technologies promise to strengthen the relationship between data subjects and care providers by facilitating healthier communication, and increasing understanding and empathy between the two; features of the patient-provider relationship some patent applications argue would not be possible with the time constraints of the current healthcare system. For example, beyond patent applications describing emotion AI technologies helping to personalize treatments and interventions (section 4.2.5), some imagined emotion AI supporting the mechanics (e.g. communication style) behind delivering care to meet the perceived needs and wants of the data subject. For example, P27 imagines emotion AI helping care providers to understand desired information and communication styles for a patient, and the patent application authors attribute this customization to improved patient experiences, arguing that *"knowledge about a patient's information needs can help the healthcare provider to tailor information and/or information-style to the patient and to optimize patient experience and patient adherence"* (P27). In this example, emotion AI technologies provide additional information about the desired information-style needs and preferences of the patient to deliver health-care related information that will be more positively received. By guiding communication around delivering care, P27 articulates emotion AI as improving patient-provider communication and supporting better patient experiences and adherence rates.

In framing emotion AI technologies as providing 'objective' data, patent applications envision care providers as using these technologies to discern a patient's authentic needs, enabling more informed decisions about the patient and their care plan. Patent application authors argue their described inventions can help provide more in-depth information about the patient, allowing a

“medical person [to] ask more intelligent questions, both of the caregiver and the patient allowing better decisions to be made” (43). In other words, more emotional and mental health data provides a larger breadth of information that clinicians and caregivers can ground their questions and conversations in; communication motivated and shaped by insights made possible by the described emotion AI technologies’ data outputs. Additionally, emotion AI was frequently reiterated as increasing emotional understanding, connected to improving care overall; for example, P16 states *“medical doctors, dentists, psychologist[s], psychiatrists, etc., may use the system to understand the real emotions felt by patients to enable better treatment, prescription, etc....”* (P16). By care providers understanding the ‘real, authentic emotions’ of their patients, increased empathy and overall understanding can occur between the two.

Altogether, emotion AI is imagined as facilitating the patient-care provider relationship through promoting better understanding of the patient, increasing information available to the care provider, and guiding communication to improve delivery of care and empathy overall.

4.4 Assisting Data Subjects with Mental Health and Emotion Regulation

While some proposed emotion AI technologies aim to improve the patient-provider relationship, others aim to assist data subjects with managing their ongoing mental health and emotion regulation. Our patent application corpus articulates these described emotion AI technologies as centering data subjects’ challenges for managing their mental health and emotions independently. Many patent application authors claim that their described technology either encourages, motivates or supports individuals to engage in behaviors and practices that improve their mental and physical health, emotional state or mood. These practices fall under self help and personal growth, addressing stimuli and triggers, and daily health-associated behaviors assumed to improve mental health, and health more broadly.

4.4.1 Self-Help and Personal Growth. Some patent application authors emphasized applications that can be categorized as relating to ‘self-help’ or personal growth. For example, many purported solutions to mental health and emotional self-management included emotion AI helping individuals have increased self-awareness in regards to their emotions, behaviors and mental health. This could include having access to their own mental health trends, support understanding how their relationships are functioning, etc. P1 describes a series of benefits made possible by its described emotion AI: *“To assist a user in general to understand his state of feeling that may result in physical and psychological changes, a type of expression is used to indicate to the user that his state of feeling may influence his logical thinking, wellbeing or his behavior emotional status....Users are able to analyze and react appropriately or make better decision in perceiving a certain situation. It can help to loosen the hold on negative emotions gained on an individual’s mind and body”* (P1). In this instance, emotion AI is imagined to detect the individual’s emotional state and provide awareness of how this feeling may shape their choices and well-being. As a result, the patent application authors argue this increase in self-awareness will improve decision making and support the handling of more negative emotions for the individual.

Some patent applications also claim their imagined technologies as increasing an individuals self-knowledge about their own health, claiming emotion AI can provide tailored information and personalized notifications to help individuals better manage and understand their health, as well as have improved overall health knowledge, such as understanding their health risks and how to prevent them. Increased knowledge about their emotions and mental health is framed as emotion AI supporting individuals being more informed of health in ways that leads to overall improvement in their health and well-being functioning. Patent P11 highlights the ways its articulated invention will detect health risks and provide tailored recommendations for certain actions to the data subject:

“The processor 204 may be further configured to recommend the first content or action to the user 112 to resolve the health-related issue based on the detect[ed] health status information. In such [a] case, the first content or action may indicate information about medicines or diagnostic centers or hospitals to the user” (P11). In this instance, the invention is imagined to provide tailored recommendations for various medications and care providers for the data subject to pursue to manage the detected health condition.

The “extra-clinical information” [34] collected by emotion AI inventions outside of a healthcare setting is framed as helping to provide more personalized information so that an individual can now gain knowledge about their health, as assessed by emotion AI, to better manage their mental and emotional states. Emotion AI technologies in our patent application corpus are often justified by framing the data subject’s lack of awareness about their own emotional or mental state as the problem their tools can solve via external support (e.g. prompts from an application).

4.4.2 Engaging with Stimuli and Triggers. Patent applications’ described emotion AI technologies promise to provide information and features to support a data subject in their day to day experiences by informing their reactions and engagements with the various stimuli and triggers they may encounter based on the types of emotions monitored and detected throughout the day. Through this increased understanding, data subjects are said to be better able to make decisions to support their mental, and emotional well-being. Some imagined emotion AI technologies promise the ability to provide a data subject alerts if a trigger or certain stimulus that may adversely impact them is nearby. Patent applications like P34 explain the ways its described invention detects reactions to various stimuli and how *“once a pattern had been identified: e.g. a positive or negative pattern with a particular person, for example, a fear of dogs or flying, the proximity of a bar for alcohol craving; the system could proactively suggest an intervention or use the knowledge of this trigger in a recommender system”* (P34). Through these alerts, inventors of emotion AI propose that individuals may be able to make informed changes to avoid or seek out certain triggers or stimuli in order to regulate their mental health status and manage mental health conditions. The ability to better operate in one’s environment and the circumstances an individual encounters is framed as supporting individuals in their day to day experiences, such as by directly altering *“the user’s itinerary to improve mood (e.g., change the default driving navigation plan or reschedule a typically stressful meeting to a better time of day for the user)”* (P48).

Emotion AI technologies described in our corpus also claimed abilities to change the content (e.g. text, media, social media content visible to a data subject) and the environment (e.g. user interface, virtual reality environment) a data subject was exposed to. Changes in what types of content was (or was not) visible, as well as changes to one’s immediate environment were presented as bringing about desirable changes (unclear desirable to whom) in a data subject’s mood or emotional state. Emotion AI technologies are described as positively changing the environment and content based on what it thinks are desirable stimuli and triggers for an individual’s overall mood, health and well-being, based on detected emotions. For example, patent P38 explains how, when its emotion AI detects a negative mood, the invention will begin *“sending one or more electronic messages to the client intended to mitigate the negative mood... and altering the client’s environment..”* (P38) in hopes that this exposure to certain content and changes to a person’s environment would have what is deemed as a positive impact. These instances highlight emotion AI independently making changes to one’s environment and the content they are exposed to for the alleged benefit of the data subjects’ well-being.

Some patent applications’ described emotion AI technologies that allowed others, not the data subject, to have control of the changes in one’s environment and consumed content to manipulate a data subject’s mood. For example, P19’s described invention heavily incorporates dynamics between

older adults and their family members who may face challenges staying connected or readily being present when articulating the technology's propensity for influencing mood. It explains how "a family member may, for example, with a click of the IMPROVE mood button on the mobile application, invoke automated functionality that seeks to improve the mood of the seniors... this feature is highly personalized in order to enable conversations and other interactions for each individual senior to achieve a desired mood" (P19). The family member, with a click of a button, is able to trigger content to be shared with the data subject to change their mood detected by the emotion AI. While this ability for another person to influence mood is framed as helping the individual by improving and supporting their mental, or emotional well-being, the patent applications rarely mention the data subjects as having any control, agency, and autonomy in these exposure modifications designed to influence their mood. And still, the emotion AI is imagined as making the ultimate determination of what stimuli or triggers to share or avoid in order to manipulate the mood of the data subject.

4.4.3 Daily Health-Associated Behaviors. Patent application authors detailed the ways proposed emotion AI technologies would encourage daily behaviors thought to support one's mental health, health more broadly, and well-being. This included certain physical activity levels, food intake, technology usage and sleep levels thought to be beneficial to the health of the data subject. For example, some emotion AI patent applications promise the ability to "facilitate healthy eating habits by providing healthy recipes on demand based on groceries that user reports having available..." (P18) or promise to "remind a patient to engage and/or not engage in an activity...to remind a new mother to get out of the house or get more sleep..." (P54) due to the justification that its emotion AI could then monitor and detect for better emotional, physical and mental health outcomes. Other patent applications promise to monitor physical activity level and based on these data, encourage an individual to "improve their activity score" (P54). P11 explains how its described invention would assist with goal setting around physical activity and sleep to improve a person's well-being; explaining how the device sets "a timing goal of the physical activity and the sleep cycle for the user 112 to be achieved each day based on the user information received from the plurality of sensors 106. For example, if the user information of the user 112 indicates that the user 112 is overweight, the emotional storyboard generator 208 may set the timing goal as 2 hours of running for each day to promote [the] user's health and well-being" (P11). Among these examples, proposed emotion AI technologies are found to additionally monitor and detect other behaviors to encourage and promote those associated with better mental health outcomes, as well as health more broadly. These patent applications' described inventions supporting and encouraging various day to day habits or experiences are part of a larger imagined role for emotion AI technologies, one where emotion AI-promoted lifestyles are thought to improve one's physical, mental and emotional health.

4.5 Protecting Data Subjects' Well-Being and Preventing Harm

One recurring theme in our sample of patent applications was that inventors imagined their described emotion AI technologies as protecting well-being and preventing harms to data subjects and the broader society, framing these problems as caused by mental and emotional health problems. We also found that emotion AI technologies were legitimized as harm prevention solutions that work to protect the data subjects and society more broadly by way of alerts and notifications. These alerts could go to the data subjects themselves, but most commonly were sent to others in relative positions of power (e.g. health care providers, law enforcement) over data subjects.

4.5.1 Emotion AI Preventing Harmful Experiences and Circumstances Unwanted by data subject and/or society. A common theme among the patent applications was the notion that emotion AI could help prevent mental-health related conditions escalating towards experiences or circumstances that are framed as unpleasant or unwanted by the data subjects and/or society, such as

development of severe mental illness or committing crime. Patent applications describe emotion AI technologies as preventing the development or worsening of mental health for a person by monitoring and detecting a measurement of interest thought to be indicative of onset or worsened mental health, and responding with actions such as directly administering an intervention or alerting third parties to intervene. P43 claims the described invention can assist with early detection and diagnosis of a mental health condition, and argues this feature can benefit the individual by preventing worsened health: *“...It is felt that in some cases such early diagnosis can lead to a potential for eliminating the worst effects altogether, by allowing intervention before a first psychotic episode occurs and perhaps forestalling it entirely or at least mitigating the effects thereof”* (P43). Similarly, patent application authors attribute the monitoring of emotion AI as integral to helping target interventions for individuals to prevent mental health issues, such as P37 that explains how *“such a system can provide interventions, such as enabling preemptive care at the right moment and the right place for the subject using the system”* (P37). Patent applications also highlight their described technologies as preventing unwanted experiences beyond specific mental health conditions, such as conflicts in one’s relationship. For example, P52 explains how its invention can *“automatically detect and predict moods and events to send prompts to oneself or to other users in a social network. For example, if the algorithms detect that conflict is likely, the intervention system could be programmed to send a prompt that says ‘You are at risk for having conflict with your child/husband/friend. Would you like to try a relaxation exercise?’”* (P52). In these instances, emotion AI inventions are imagined to prevent escalation of conflict by predicting moods or events by alerting an individual that such conflict might occur. By proposing a just-in-time remedy to prevent or forestall a problem, these technologies raise questions of what interventions or remedies (both technical and non-technical) are valuable or ethical for helping people who might be in distress.

Patent applications commonly suggested mental health conditions *“represent growing risks and concerns for society”* (P57). Some patent applications frame emotion AI preventing an escalation of mental health symptoms or episodes as beneficial to others beyond the data subject by preventing *“a higher risk situation to the mental health client and in turn, their friends, loved ones, and the community at large”* (P5). P5 continues to explain how *“behavioral prediction of mental health clients is an important issue that needs a solution. Recent mass casualty incidents in the United States have inevitably pointed to individuals with known mental health issues that might have been avoided with early intervention and diversion”* (P5). By being able to detect and predict behavior, several patent applications claim to be able to prevent harm done by those with mental health conditions.

4.5.2 Helping Authorities Reach Data Subjects Vulnerable or Believed to Commit Harms. Patent applications argue their imagined emotion AI technologies can help authorities better reach those they are ‘responsible to help’ by producing alerts and notifications through detecting and predicting attributes like emotion, behavior or intent; these relationships include clinicians and their patients (as discussed in section 4.3) and police and their communities. Several patent applications describe their emotion AI technologies as supporting law enforcement in being more efficient and effective at performing their assigned duties, as well as helping care providers more readily reach their patients in perceived need of care. P17 explains how its described emotion AI can detect words of interest to alert third parties when deemed necessary, explaining how its invention can *“recognize high alert words, word combinations, phrases and sentences, and act accordingly... these are recognized as high alert phrases or words that will initiate notification by the robot to emergency contacts 519, for professional handling”* (P17). As another example, P5 highlighted the ways its invention could specifically support law enforcement: *“Law enforcement is often called upon to dedicate a minimum number of officers to respond to mental health clients who are at risk to themselves or others... The platform allows law enforcement to either respond earlier to the ‘Alert’ message or they can direct*

a mental health client to diversion community-based services where their unstable behavior can be deescalated before the call advances to an involuntary hospital admission. . .” (P5). Alerts and actions taken by authority figures in response to alerts produced by emotion AI technologies raise questions about potential consequences of inaccurate or failed emotion detection, as well as implications for individual privacy and autonomy. It is important to note that complex relationships exist between different individuals and their communities with institutions like law enforcement, such as Black, Indigenous, and People of Color (BIPOC) that have been subjected to invasive surveillance and police brutality [69, 83, 173]. Meanwhile, the analyzed patent applications uniformly position these powerful institutions as both harmless and desirable in their broader ‘community’.

Patent Applications construct Emotion AI to...
Solve Data Accuracy Challenges
<i>(a) Allow for more complex and varied data sources to be used, leading to more accurate detection and monitoring of data subjects’ emotional and/or mental health status(es)</i>
<i>(b) Make it possible to access more ‘objective’ and ‘reliable’ data sources as compared to human-reported and -collected data</i>
Improve Care Provision and Experience
<i>(a) Assist and/or replace healthcare providers already challenged by limited time and resources</i>
<i>(b) Diagnose and detect data subjects’ mental, emotional, physical or cognitive state more accurately and faster than currently possible</i>
<i>(c) Directly deliver or support the delivery of treatment, and manage medications to improve data subjects’ well-being and emotional state</i>
<i>(d) Increase the technical and productive efficiency of care by allowing healthcare providers to easily and quickly provide equal or additional amounts of care</i>
<i>(e) Improve the experience and effectiveness of care by allowing for faster and more personalized treatments that leads to better health outcomes</i>
Help with Patient-Provider Communication
<i>(a) Promote better understanding between a care provider and patient regarding the patient’s wants and needs</i>
<i>(b) Increasing the amount of information a care provider has access to</i>
<i>(c) Influencing communication so that care is delivered and positively received by patient(s)</i>
Assist Data Subjects with Mental Health and Emotion Regulation
<i>(a) Increase one’s self-awareness, leading to personal growth and self-help practices</i>
<i>(b) Assist data subjects’ navigating their surroundings, and helping them be mindful of certain stimuli or triggers that impact their emotions and mental health</i>
<i>(c) Promote changes to daily health-associated behaviors to support one’s mental health and health more broadly</i>
Protect Data Subjects’ Well-Being and Preventing Harm
<i>(a) Prevent harmful experiences and unwanted circumstances attributed to mental-health related conditions</i>
<i>(b) Help authorities more quickly identify and direct attention to data subjects who are deemed vulnerable and a risk to themselves and/or others</i>

Table 1. Breakdown of the imagined futures of Emotion AI for mental health, as described in patent applications in our corpus.

5 DISCUSSION

In this paper, we explored the sociotechnical imaginaries [103] of emotion AI through patent applications to understand the kinds of mental-health related problems emotion AI is imagined as solving, as well as how the discourse present in patent applications legitimize [199] these technologies. An outline of our key takeaways is available in Table 2.

Key Takeaways from the Ethical Speculation of...
<p>Emotion AI's Contributions to Mental Health Stigmatization</p> <p>(a) Framing mental health conditions as threats to public safety will lead to increased discrimination, surveillance and stigma on the basis of (claimed) detected mental health states.</p> <p>(b) Imagining patients as data subjects and passive recipients of care risks making people more reactive rather than proactive with their mental health.</p>
<p>Emotion AI's Part in Mental Health Surveillance Futures</p> <p>(a) Surveilling emotions and mental health states of data subjects (knowingly or unknowingly) can increase fear and paranoia, and further oppress communities on the basis of their (assumed) mental health, as well as increase the potential for privacy violations and paternalism in health care.</p> <p>(b) Envisioning medical institutions and law enforcement as trustworthy in emotion AI introduces new potential for harm.</p> <p>(c) Emotion AI that infers one's emotional states forces authenticity and removes the choice of being authentic, and emotionally vulnerable.</p>
<p>Emotion AI's Impact in Clinical Decision-Making Futures</p> <p>(a) Emotion AI can introduce additional confusion or uncertainty in healthcare providers' decision-making processes, and further decontextualize data subjects' from their health outcomes and care decisions.</p> <p>(b) Emotion AI is fundamentally reductive, simplifying complex and varied lived human experiences of data subjects into neat data points. As a result, critical aspects of human complexity will be excluded or dismissed in decision-making processes where data subjects' should be holistically and uniquely considered.</p>
<p>Emotion AI Shifting Focus From Systemic Health Problems and Barriers</p> <p>(a) Emotion AI perpetuating an understanding of health that centers the individual and focuses on personal responsibilities introduces the potential of further moralizing health by ignoring the multiplicity of health, and distracting from structural barriers to health that contribute to poor well-being and mental health in the first place.</p> <p>(b) Emotion AI focusing on issues of healthcare providers at the individual-level, as opposed to structural, systemic issues in healthcare can introduce new concerns while operating as short-term band-aids to deeply entrenched issues.</p> <p>(c) Emotion AI acts as a form of technosolutionism with grand promises that must be interrogated and questioned, specifically considering the ways these technologies may stigmatize, surveil, simplify or reduce complex issues and contexts critical for mental health and emotional well-being.</p>

Table 2. Key Takeaways from our Ethical Speculation on Emotion AI for Mental Health.

In this section, we use ethical speculation [78] as a lens to discuss the promissory sociotechnical futures of emotion AI in mental health contexts. Ethical speculation is useful to anticipate the

implications of technologies that “*inherently have unanticipated, if not unintended, consequences*” [78]. Prior work has used ethical speculation [28] to examine emotion AI’s implications by analyzing emotion AI patent applications in the workplace. Others have used speculation-like methods to discuss affective computing technologies’ future implications [19, 59]. By describing the futures envisioned by inventors and understanding some of the ways emotion AI might be deployed for mental health, actors like regulators and researchers can apply ethical speculation to general trends and inventors’ perspectives to consider the ethical implications of these possible futures *and* discuss interventions to prevent harms [77].

5.1 Mental Health Stigmatization Implications of Emotion AI Use

Emotion AI patent applications in our corpus depict mental health and those impacted in ways we argue may perpetuate mental health stigma. We speculate on how emotion AI in our corpus’ imagined futures may perpetuate the stigmatization of non-normative mental health and those with mental health conditions and its implications. Our findings demonstrate emotion AI patent applications envision data subjects ranging from unreliable patients to individuals capable of perceived positive change to both perpetrators and victims of harm due to their mental and emotional state(s). These beliefs around mental health are not unique to the patent applications in our corpus, but they do perpetuate commonly held ideas about mental health that contribute to mental health stigma [85, 187]. The ideas that certain mental health conditions are associated with helplessness, social ineptitude, unreliability, crime and harms can have consequences of increasing mental health stigma. Increased stigma can lead to discrimination and internalization of these ideas with potentials such as discouraging individuals from pursuing treatment and disclosing their mental health condition [53, 81, 184, 188]. When stigma becomes embedded within the sociotechnical imaginaries advanced in emotion AI patent applications, bias and discrimination toward people with mental health conditions can be perpetuated through the seemingly neutral technical interface of the emotion AI technology.

5.1.1 Mental Health and Criminality. Our patent application corpus describes potential futures where individuals’ emotions and mental health states are able to be known, predicted and safeguarded against, especially when set to cause harm or conflict. Patent applications described data subjects with certain mental health states and emotions as dangerous and threats to the safety of themselves and others. While patent applications framing data subjects as potentially harmful to themselves and others may seem logical for patent applications articulating solutions for preventing harm, they *also* perpetuate the idea that individuals with mental health conditions are dangerous and predisposed to crime. Linking crime with mental health is a connection frequently made by mass media, politicians and the general public when violent events like mass shootings occur [79]. However, this connection is falsely drawn; few acts of violence are carried about by those with mental illness and individuals with mental health conditions are over 10 times more likely to be subjected to violent crimes than those without mental health conditions in the United States [63]. By framing mental health and poor emotional well-being as a threat to both the safety of the individual and society *and* directly implicating authorities (e.g. health care workers, law enforcement) in their solutions to ‘the mental health’ challenge, emotion AI patent applications portray individuals with certain mental health and emotional states as societal risks that need to be monitored and ‘controlled’—a continuation of past beliefs around mental health [177]. Due to emotion AI’s scalability to encompass a large number of data subjects and collect mass amounts of sensitive data related to mental and emotional health, this linkage of mental health with criminality raises alarms for normative futures where surveillance of a population already dealing with higher levels of paranoia [22] is deemed necessary for ‘public safety’.

Rationalizing emotion AI as technologies for safety requires an understanding of what safety is and *who* or *what* is a threat to *whose* safety; both subjective to the powers that might wield emotion AI in the name of public safety, *and* those data subjects being monitored. In the U.S., we have seen tensions emerge out of different understandings of safety and harms: police using teargas at anti-racism/anti-police brutality protests in the summer of 2020 [118], instances of Black people being arrested when trying to enter their own homes [201], a trend of book bans over hysteria around claims of ‘indoctrination’ on issues of Racial, Sexual and LGBTQ+ identity [135], and the U.S. government’s dismissal of employees thought to be LGBTQ+—and thus, ‘security risks’—during the Lavender Scare [105]. In all of these cases, we see understandings of harm and safety in tension with another’s identity, self-determination, or ability to resist perceived harms or abuses of power. Thinking speculatively raises concern that the technologies envisioned by these patent applications, which invoke emotion AI futures where people with mental health conditions are framed as threats to public safety or subject to criminality, will exacerbate these tensions and lead to increased discrimination, surveillance and stigma on the basis of (claimed) detected emotions and mental health states.

5.1.2 Mental Health and Agency. Patent application descriptions commonly imagined data subjects as passive recipients of care by deferring authority and decision making to others (e.g. care providers, police, family members), such as by automatically adjusting medication or proactively changing content shown to an individual to modify their emotions in ways thought desirable by the patent. We argue when technologies do not center the data subjects as agents of their own being by *deciding for data subjects* what desirable mental health is and the way to achieve it, they risk removing an individual’s ability to play an active role in deciding their well-being aspirations. This approach may also lead to internalization of stigmatized beliefs around mental health; that they are incompetent, or in need of fixing [85, 187]. We speculate by not positioning data subjects as owners of their own emotional experiences as ‘revealed’ or ‘deduced’ by these technologies, potential futures exist where people are made to be reactive rather than proactive with their mental health. For example, automated alerts to one’s doctor or family members remove an individual’s choice in deciding whom they want to be aware of their mental or emotional state, which may also vary based on circumstances for the same individuals involved. Individuals are made to react to the choices and subsequent impacts brought by emotion AI systems and the parameters the systems deem (in)appropriate when monitoring and detecting mental and emotional states. This may lead to individuals being placed in uncomfortable, awkward situations they never wanted to be in and perpetuate mental health stigma.

While the language and descriptions of mental health in patent applications may be stigmatizing, does that mean the tech *itself* contributes to stigmatization? We speculate, yes. Regardless of any potential failures of the systems, even *if* the technologies described in our patent application corpus operate ‘perfectly’ as intended, they posit and reduce data subjects to subordinate actors subject to the influence and decisions of authorities (e.g. medical, law enforcement) and/or emotion AI system outputs. Toyama’s law of amplification theory describes technology’s influence as amplifying pre-existing human beliefs and forces [197]. The stigmatizing beliefs reflected in patent applications about mental health risk amplifying and further entrenching these beliefs and their impacts on society at large [197], as opposed to contributing to equitable futures where individuals with different mental and emotional health states are imagined and supported as autonomous beings worthy of respect, self-determination and dignity.

5.2 Implications of Emotion AI Use in Mental Health Surveillance Futures

The patent applications in our corpus imagine sociotechnical futures where emotional and mental states are surveilled and the resulting data is made available for a wide range of “health” and “safety” uses. And yet, emotions and mental health conditions are, for numerous reasons, hard to truly know by those not directly experiencing them, rendered invisible or ambiguous [29, 121, 130]. Despite these qualities, emotion AI technologies in our corpus imagine a world where it is possible (and good) to surveil, monitor and detect affective phenomena for a variety of promised benefits relating to mental health, emotion regulation and harm prevention. This raises the opportunity to speculate about what it might be like to live in one of these potential future worlds. For some, the ubiquitousness of surveillance structures may lead to the development of attitudes commonly mistaken for indifference or compliance: surveillance-apatheia, “*an attitude that individuals learn... stems from the perception that there is little one can do to avoid [surveillance systems], so why concern oneself with deliberation and anxiety of them*”[72]. Others might instead practice resistance [128] or resilience [212] within surveillance structures. Here we speculate on the implications of emotion AI Systems in futures where they operate as a form of surveillance in mental health contexts.

Surveillance can lead to increased levels of fear, paranoia, and forced subjugation to webs of surveilling technologies by communities believed to be vulnerable to or with a propensity for committing harms [4, 90, 94, 117, 183]. For example, individuals with HIV are reluctant to seek help from health clinics and programs using surveillance technologies [108]. Emotion AI tools used for emotion monitoring and detection might similarly deter individuals from seeking mental health or emotional support from care providers due to a lack of trust or interest in engaging with surveilling tools. The paranoia of surveillance might also weaken an individual’s eagerness to openly engage in other daily contexts (e.g. stores, classrooms, work) due to concerns of an emotion AI technology detecting and subsequently ‘classifying’ them in some way (e.g. dangerous, at-risk to self) —which can be stigmatizing and consequential.

In contexts where conditions are associated with criminality, surveillance of identifying indicators (e.g. biometric data, emotions) “*requires disclosure of a status that, if known by the state, would open people to an elevated risk of harm*” [108]. As discussed in 5.1.1, our corpus links mental health with criminality and, similar to biometric surveillance for HIV, we speculate surveillance using emotion AI might enable the government, health entities or anyone with access to the emotion AI systems’ data to flag individuals as dangerous due to emotions and mental health states supposedly detected.

Technology-based monitoring is not new to mental health, and digital surveillance has raised ethical concerns. In 2017, the FDA approved Abilify MyCite (AMC), updating a popular medication for managing mental health conditions such as schizophrenia and depression with an ingestible tracker. It was the first medication to digitally track and monitor medication compliance [50]. AMC in practice is also paired with a “smart” application to record rest, number of steps, self-reported moods and rationale for abstaining from medication. Both sources of data are reported to the data subject’s healthcare providers, as well as select family and friends [50, 98]. AMC has raised ethical concerns for privacy, forced consent, increased paranoia among those taking the drug, and breakdowns of the patient-provider relationships [35, 186]. Additionally, coverage in the popular press describes AMC as fitting into a stigmatizing and paternalistic trend of mental health care, framing individuals as needing to be monitored, and controlled as unreliable agents for their health, despite mental health conditions having similar adherence rates for other chronic illnesses [186]. While AMC is not an emotion AI technology, we observe a similar trend in our corpus of technology increasing privacy violations, paternalism, and stigmatization.

5.2.1 Alerts. We found that alerts were an important surveillance method that may stigmatize and harm individuals. The potential consequences of inaccurate or failed emotion detection alerting authority has the potential for experiences like an individual having repeated unwanted and unnecessary encounters with police, or being continuously told they need to go see a doctor. What would happen to a person's sense of self or well-being if they were repeatedly flagged as being a threat to themselves and others or if they were incorrectly detected to be experiencing a health crisis? Emotion AI monitoring produces a form of surveillance where, by way of emotion AI alerts and inaccurate systems, the ability to further stigmatize and cause additional harm to data subjects exists.

Emotion AI technologies described in our corpus also reflect structural biases, particularly by assuming trusting relationships with authorities including law enforcement and healthcare professionals. Individuals' interactions and experiences with medical and law authorities are racialized [143, 208] and gendered [91]. For example, in clinical settings, clinicians are more likely to perceive men's symptoms as tangible medical conditions, as opposed to interpreting women's symptoms as a result of mental and social factors [91]. These gendered interpretations impact the kinds of care patients receive and contribute to gender disparities in the distribution of life-saving treatments and interventions [91]. In addition to gender disparities, racial disparities exist in medicine, such as worse reproductive health outcomes for Black women compared to white [7] or clinicians' disparate treatment for pain due to false beliefs of biological and pain tolerance differences by race [208]. Additionally, Black and other communities of color have been subjected to persistent surveillance and police brutality by law enforcement [66]; police violence and experiencing dehumanizing conduct by law enforcement have been shown to adversely impact the mental health and emotional well-being of Black communities [5, 143]. In these instances, we can understand how individuals' experiences with medical and law authorities are shaped by gender and race, and consequently, lead to adverse health experiences and outcomes and different relationships with these powerful actors. When technologies like those described in our corpus envision medical institutions and law enforcement as trustworthy and positive presences in a broader 'community,' they do not account for the lived realities of many and ignore potential futures of harm that contradict the justifications for their emotion AI invention.

While patent applicants in our corpus are not required by the U.S. Patent and Trademark Office to account for the social context of how their emotion AI inventions may be used by powerful actors, failing to do so contributes to the harmful idea that fields such as data science and technology involving AI are neutral and apolitical [87] while furthering systemic harms and inequalities. When inventors, technologists, and researchers (or any other positionality patent application authors might hold) are not required to reckon with the social context their inventions may operate in and with, it becomes possible to overlook any responsibility for the downstream effects of their creations, including those that inflict harm on populations. Reflecting on the potential futures of technologies and the actors involved in their creations, we view relevant actors as responsible for engaging in similar ethical speculations (as we do in this paper) to imagine potential harms and potential mitigation strategies; even if this speculation results in the conclusion that a technology is best to *not* design [18], despite external motivations like monetary profit or AI hype [216].

5.2.2 The "Right to be let alone" not promised in promissory Emotion AI Surveillance futures. Our corpus speaks to futures where our innermost feelings and emotions are no longer solely ours—one that alienates us from our emotions, rendering private feelings and emotions to be used for another's purpose (accurately or inaccurately). When awareness of one's mental health condition may lead to discrimination in employment, education, transportation, housing, etc. [38, 133, 167, 193], emotion AI and its surveillance uses bring to the forefront issues of a data subject's "right to be let alone"

[206], as well as their autonomy for deciding what to do or *whom* to tell about their mental health or emotional state. Emotion AI technologies that do not consider the data subjects' right to *meaningfully choose* to be monitored or to, for example, *choose* to be notified of a potential mental health diagnosis or emotional state, infringe on their personal agency and autonomy over their health and well-being—essentially forcing authenticity onto data subjects by forcing an obligation of disclosure. Taking the case of emotion AI promising benefits of predicting a mental health condition, prior work has demonstrated individuals seeking help from an early detection center for mental illness were resolute on whether or not they would choose to accept risk prediction for mental health conditions, believing this prediction would have large impacts on how they saw themselves *and* their quality of life for better or for worse [126]. Individuals in the study, regardless of being in favor of or against risk prediction, reported valuing control and self-determination [126].

It is unclear whether the technologies our patent applications describe would allow data subjects to decide whether or not to be alerted or for a health provider to be alerted and notify them of their mental or emotional state. However, such considerations are important for technologists hoping to build systems that respect one's autonomy, and thus, privacy [48]. As much of our corpus posits data subjects as passive recipients of care and not as authorities, as discussed in 5.1.2, we speculate the promissory future of emotion AI technologies is one filled with infringements on individuals' autonomy, rights to make decisions around care, and overall privacy – ones where even if the technology accurately infers one's emotional states, authenticity is forced upon them and the choice of being authentic, and emotionally vulnerable is taken away. To understand the potential ramifications of such infringements, future research may look at emotion AI systems already deployed and evaluate the ways in which they do or do not protect an individual's right to privacy, under what conditions people may be able to meaningfully consent to being subjected to emotion AI within the healthcare setting, and the impact this has on an individual's relationship to themselves and their relation to health overall.

5.3 Accounting for Emotion AI in Clinical Decision-Making Futures

One key feature of our corpus' technologies is algorithmic decision-making in the context of healthcare, and specifically, mental health care. It is clear patent applicants see emotion AI as improving health care, improving decisions and subsequent actions relating to care and treatment, diagnoses, health outcomes and patient-provider relations. And yet, algorithmic decision-making in healthcare may lead to higher levels of uncertainty for clinicians [89], decontextualization of data subjects [127] and further remove clinicians and patients from playing an active role in decision-making [39, 89]. For example, Jacobs found clinicians receiving machine learning treatment recommendations for depression did not experience better accuracy selecting treatments that matched with treatment recommendations by psychopharmacology experts [101]. At a fundamental level, AI involves finding patterns in external data to make predictions based on prior data collected [107]. In the case of our corpus, these predictions are used to direct clinicians in their decisions *or* for emotion AI systems to act directly in ways that influence health care for data subjects.

5.3.1 Promissory Futures of Emotion AI Reducing Emotion in Health and Care Contexts. The patent applications in our study describe promissory futures where emotion AI technologies are poised as solutions to issues in health and care contexts. As a result, patent application descriptions act as discourses of legitimization [199]—discourses that rationalize futures where reducing complex emotions and mental health states and care-based decisions to algorithmic processes (using external data sources and physical attributes like facial expression) is desirable. This reduction of data subjects to a series of parameters for an AI system to compute a certain prediction has moral

implications [27] when considering who and what is (not) included or considered when making decisions for mental health care [202], especially critical when emotion AI may act *directly* on data subjects as was the case in some of our corpus. Villegas-Galaviz argues AI decision-making reduces decisions to data without accounting for individuals' unique and potentially vulnerable circumstances, highlighting how AI's reliance on patterns in past data leads to a technology that dismisses the complexities of data subjects' present [202]. As Bemme et al. warns, "the quest for holism through big data may thus lead to a re-emergence of the tyranny of reductionism [21]. By reducing a data subject and their care to specified metrics and data points, potentially crucial aspects of their personal experience are excluded from the decision-making process [6].

Our corpus imagines emotion AI futures where various data inputs (e.g. biometric, visual, text) implicate a data subject's emotional and mental health, for better or for worse, without first considering the individual's unique experiences or wants. Consider the implications of invention P43 from our corpus (described in section 4.2.3), which promises to automatically administer trace amounts of medication when the emotion AI concludes an individual is at risk for developing a certain mental illness. To be clear, this technology does not give data subjects a chance to make informed decisions about their health, fully understand the side effects of such medications, or consider alternative treatments before medicating them. Among the many imagined uses of the inventions featured in our corpus, it is critical for technologists to be proactive about giving the data subject meaningful authority and control over their emotion-related information and its uses.

The individuals emotion AI technologies monitor are experts in their own unique experiences and only they know the *actual reality* of their feelings[17, 165]. Emotion AI merely creates predictions and we argue these predictions should not be considered the objective, authoritative truth. While difficult due to the vagueness and sometimes convoluted descriptions of our corpus' computational methods for deducing emotion of data subjects, future research must work to evaluate and audit the ways in which emotion AI systems (are proposed to) operate in health care settings and act as authorities for decision-making, including but not limited to the perceptions of clinicians and patients in these decision processes (e.g., deciding a diagnosis, choosing a treatment). For example, one might explore clinicians' and patients' experiences delivering or receiving a diagnosis that was derived from an emotion AI system, focusing on features such as a sense of trust, feeling understood, or addressing concerns of failure, similar to [142] that examined how trust is experienced and valued by clinicians working with predictive modeling in healthcare. In doing so, we may understand not only the effectiveness of emotion AI systems in shaping healthcare itself, but also the impacts (positive or negative) that they may have on the human dynamics and relationships in care provision—contributing to centering and valuing the human experience over prioritizing the perceived 'solutions' promised by technological interventions like emotion AI.

5.4 Emotion AI Trends Divert Focus on Systemic Problems and Barriers for Health

Patents are useful in examining technological trends [1, 41, 42, 75], and our corpus reveals a trend of positioning the data subject as individually responsible for their mental health and well-being as opposed to focusing on structural barriers. Emotion AI patent applications posit their inventions as beneficial solutions to challenges impeding one's mental health and well-being, challenges involving individual behaviors and healthcare system constraints. We question whether emotion AI is the appropriate solution for the health-related problems our corpus aims to address. We turn to Baumer and Silberman's questions for researchers and practitioners to assess whether a technology may or may not be appropriate, "*Is there an equally viable low-tech or no-tech approach to the situation? Might deploying the technology result in more harm than the situation the technology is meant to address? Does the technology solve a computationally tractable problem rather than address an actual situation?*" [18]. We also take Toyama's law of amplification theory [197]—the notion

that technology's effects may be to amplify humans' existing beliefs and values—as a point of reflection on whether emotion AI is or can be the 'solution' to the many problems they are claimed as able to solve.

5.4.1 Individual Behaviors for Health. By imposing 'health' on data subjects with emotion AI, we argue emotion AI constructs health where the onus is on the individual to achieve health, without accounting for its fluctuating definitions and the systemic barriers to achieve them. Systemic barriers are structural forces (e.g. procedures, policies) that result in the exclusion or inability for a person to participate or engage in some activity (e.g. employment, healthcare) [178]. In regards to health, some examples of systemic barriers might be lack of transportation that limits one's ability to attend medical appointments [65, 194] or the lack of access to affordable fruits and vegetables due to living in a low-resourced community [26].

Our findings demonstrate emotion AI patent applications envision emotion AI as promoting behaviors thought to lead to overall improvements in mental, emotional and physical health [4.4], including behaviors such as eating "healthy", engaging in reflection, avoiding negative triggers, getting more physical activity, and refraining from extensive technology usage. Emotion AI systems and their inventors shape these daily healthy behaviors based on what their constructions of 'health' and 'healthy'. Lupton describes health as a social construct, arguing health is shaped by the historical, cultural and social contexts bodies have existed and operated in in the past and present [124]. As a result, the ways we understand health and our bodies' are continuously changing. For example, on illness, Brown argues instead of looking at illness as a biomedical absolute truth, we should view it as "*a set of understandings, relationships, and actions [about illness and diagnosis] that are shaped by diverse kinds of knowledge, experience, and power relationships...*" [30]. As health is often articulated in purposeful, goal-directed ways, health constructs embody values and moral judgements [116]. If one embraces the diversity of human experience, we must acknowledge health does not have a single correct definition, nor a single desirable end. The patent applications in our corpus do, however, prescribe a construction of health and healthy to data subjects.

For example, recall in section 4.4.3, P11 explained its technology would automatically set a goal for a data subject to complete two hours of running each day if they are thought to be 'overweight'. What is the experience of an individual who cannot engage in physical activities such as running due to time constraints, or an inability to get to a safe place to run? Likewise, as described in section 4.5.1, P18 intends to mediate healthy eating by providing healthy recipes according to what a data subject has available. What is to happen to an individual who cannot access 'healthy' foods in the first place because they live in a food desert without access to affordable, nutritious foods or lack the time to cook homemade meals? Systemic barriers, such as those to physical activity and 'healthy' eating behaviors, can impact an individual's ability to engage in 'health' as prescribed and monitored for by emotion AI systems [45, 58], like those described in P11 and P18.

If an emotion AI system repeatedly suggests and notifies an individual of 'healthy' actions without consideration of an individual's preferences or circumstances, they may experience something like information overload [166], resulting in fatigue from trying to sort what is relevant and useful to them, and cause them to miss potentially useful and relevant information and notifications as a consequence [37]. People might also experience stigmatization and discrimination due to being perceived and labeled as unhealthy [171]. Additionally, drawing from the law of amplification [197], we speculate those who already have access and/or the means to eat 'healthier', and engage in higher levels of physical activity will be more likely to continue to, however those without access will continue to lack access and the ability to engage in what our patent applications describe as daily healthy behaviors. By not accounting for or addressing systemic barriers to health, emotion AI use can then further amplify [197] the idea of 'personal responsibility' for health that ignores

and dismisses individuals' circumstances, and leads to judgements of one's moral character on the basis of health [192].

It is clear from our corpus that inventors imagine their systems' engagements with daily behaviors as facilitating and mediating 'healthy' behaviors and their subsequent outcomes. When recognizing the many systemic barriers to health *and* when reflecting on Baumer and Silberman's questions [18], we ask ourselves if emotion AI systems are inherently beneficial to individuals' health as they are described in our corpus and further speculated upon in this work. While it is impossible to predict all of the ways these emotion AI patent applications might be deployed, their descriptions related to the promotion and monitoring of daily 'health' behaviors raise concerns such as how they might exacerbate the already inequitable effects or pressures of participating in health and wellness-based incentives, promotions and benefits [80, 92, 97, 125, 180]. For example, an employer with the intention of improving the mental and emotional state of their workforce may choose to use emotion AI to monitor the daily behaviors of their employees, offering benefits and incentives to employees who engage in 'healthy' behaviors of the type and level deemed appropriate by the emotion AI system. As a result, employees unable to engage in those behaviors would be excluded from these benefits, may feel pressured to participate undermining meaningful consent, as well as have concerns around privacy, stigma, and discrimination [125]. Our corpus' described uses of promoting a narrow and individualistic conception of health raises concerns of how they may increase inequity for individuals with mental health conditions by not accounting for systemic barriers.

Fortunately, HCI scholars have begun considering health as being heavily shaped by multiple, intertwined systems. For example, Pendse et al. push for moving away from digital mental health tools that focus on treating symptoms and individual behaviors and instead towards tools that zoom out to account for structural and holistic healing [160]. Repeatedly, we found patent applicants promised to help improve individuals' treatment and diagnostic outcomes, as well as to help individuals manage their symptoms and behaviors with the intention of improving their mental and emotional health. Amidst these promises, the call for structural, holistic or communal understandings of healing were absent. We join scholars [160] arguing that the focus on individual treatments and symptoms ignores the structural barriers that often lead to or influence well-being and mental health in the first place.

We question if the implication for designing emotion AI for issues whose roots are far deeper than individual choice, such as individual behaviors, is to *not* design, particularly as band-aid or even harmful solutions to mental health challenges. Inspired by Toyama's law of amplification [197], we encourage future investigation into what existing human practices and forces instead exist—perhaps at the community level or with a focus on holistic healing—that may be supported and amplified by technology to result in improved mental health outcomes, while recognizing the diversity of definitions for health.

5.4.2 'Overcoming' Healthcare Challenges. Emotion AI patent applications describe myriad emotion AI solutions bypassing the time and resource constraints of clinicians, such as allowing clinicians to monitor and care for multiple patients simultaneously or lessen the frequency of doctors' visits required to deliver care. Emotion AI patent applications in our corpus reveal a trend in AI systems that focuses on problems like lack of time and resources for healthcare providers [section 4.2.1]. These problems are situated in a health care context experiencing systemic problems like the increased shortage of doctors [112], and growing burnout in healthcare professions [115, 168] that may result in reduced time and resources for healthcare providers. We urge technologists to consider whether the inventions they hope to create merely operate as bandages or pain medications, solutions to larger problems that camouflage or lessen the larger issues at hand.

When attempting to solve issues in healthcare, we must ask ourselves if the problem is at the individual healthcare worker level (e.g. lack of time or resources) or if the problem is broader. For example, is the solution to a doctor lacking time allowing them to be able to monitor multiple patients simultaneously *or* does the solution lie in addressing *why*, *for example*, there are fewer doctors available to provide care [112]. We argue emotion AI futures imagined in our corpus cling to technological solutionism, the idea that technology like algorithms can solve any problems efficiently and without additional trouble [134]. Morozov argues against solutionism, explaining that “*there are more fruitful, more humanistic, and more responsible ways to think about technology’s role in enabling human flourishing, but solutionists are unlikely to grasp them unless they complicate their dangerously reductionist account of the human condition*” [134, p.14]. By building emotion AI that can presumably detect emotions or mental health conditions, our corpus imagines the future of emotion AI as one that solves problems in health, care and safety contexts, by reducing their issues to individual-level problems (e.g. healthcare workers’ being inefficient [section 4.2.4]) as opposed to structural, systemic issues [112, 115, 168].

Additionally, the application of emotion AI technologies to address healthcare challenges can introduce new concerns of its own, such as the case of remote patient monitoring promised by several of our corpus’ inventions. While remotely monitoring patients may alleviate the burden of time for travel and receipt/provision of care, concerns exist around remote patient monitoring, including loss of interpersonal contact [203]; intrusion into the private and personal contexts of the data subject [9, 146]; misalignment between what patients and clinicians understand as important data for care [102]; and conflict between patients and clinicians as a result of unmet expectations for care [47]. Technologists must consider the collateral impacts of their inventions that on the surface may address one issue but introduce several new concerns.

It is important to note that many of the promised benefits found in these sociotechnical imaginaries of emotion AI for mental health, such as boosting efficiency of healthcare and healthcare resources, improving patient-provider communication, and supporting individuals in managing their own health are articulated as desirable in other discourses beyond emotion AI *and* mental health. For example, news releases from the National Institute of Health consistently announce projects they’re funding and new initiatives beginning by emphasizing the ways these efforts would improve efficiency and cost-effectiveness of healthcare [137–140]. As another example, improving patient-provider communication is discussed widely in medical trade journals, such as navigating health literacy challenges between patients and their clinicians [96], and comparing communication styles and its implications for patient self-management [162]. These discourses are part of the environment in which sociotechnical imaginaries emerge, and these imaginaries are aligned with a larger capitalist framework where technological developments are conflated with social progress and benefits [73]. In this way, the sociotechnical imaginaries of emotion AI for mental health are part of a larger social trend valuing qualities such as efficiency as desirable, and positioning technology as a way to achieve these aims.

In analyzing patent applications and, therefore, glimpsing into the sociotechnical imaginary [103] of emotion AI for mental health, we understand inventors justify emotion AI systems as solutions to problems in the health, care and safety space. We advocate that it is important to not get caught up in these vanguard visions of emotion AI [95] with their grand promises [11] and hype [216] without being critical of the stigmatizing, surveilling, oversimplifying and reductive potentials *obscured* in these promissory futures for our mental health and emotional well-being.

6 IMPLICATIONS

In this paper, we have demonstrated how accessing and characterizing the sociotechnical imaginaries of an emerging technology allows for researchers, such as those in the CSCW community, to

speculate on the ethical implications of potential sociotechnical futures. By accessing one source where sociotechnical imaginaries are located (e.g. patent applications), we were able to make sense of potential emotion AI futures for mental health and speculate on their ethical implications.

We recommend that researchers and regulatory bodies consider patents as early warning signals for potential ethical concerns surrounding emerging technologies—and ethical speculation [78] serves as one lens in which to illuminate these concerns. We suggest that ethical speculation be developed into a systematic process as part of patent review, ensuring that early warning signs of potentially problematic emerging technologies are considered and thus, awareness of these concerns made accessible to assist the U.S. Office of Science and Technology Policy (OSTP) and other regulatory agencies to develop appropriate public policies. We argue there is value in embedding ethical speculation into patent review processes and research on emerging technologies, aligning with recent calls for governments and regulatory agencies to use the patent system as a tool to regulate emerging technologies [155]. By considering the ethical implications of emerging technologies at the level of patent review, regulatory agencies are better prepared to consider and respond to the potential societal impacts of such technologies.

We argue more can also be done within existing regulatory structures as part of patent review and formal innovatory procedures to embed ethical evaluation and expand the range of relevant expertise when developing technologies. The U.S. Patent and Trademark Office may wish to strengthen their ‘utility’ criterion by specifically requiring patent applicants to provide substantial data about where, how and in what populations their proposed technologies have been evaluated and what the outcomes were. In this way, innovation processes—and subsequent research— may be incentivized to be more attuned and responsive to the data subjects actively involved in the development of such technologies [8]. Additionally, innovators, regulatory agencies and researchers must recognize the value of including non-technical communities and social sciences scholars into innovation processes; doing so “*foster[s] less linear, more dynamic innovation paths involving a more diverse array of experts*” [156] that can develop just and equitable technologies. We understand that our call for ethical evaluation (such as by increasing the ‘utility’ criterion) and requiring patent applicants to engage stakeholder populations and evaluate any proposed technology may raise concerns for slowing down innovation processes – which is not inherently bad. However, opportunities exist for patent applicants and inventors to collaborate with established groups of users/data subjects that may shorten the time it would take to produce an evaluation and satisfy a more robust utility criterion. We also urge the consideration of how engaging in anticipatory social or ethical analysis might produce alternative and improved innovation, moving away from a linear conception of innovation paths to one that embraces engaging with stakeholders and data subjects.

We believe that policymakers would be wise to look to patent applications and granted patents as one way to access and be aware of the very sociotechnical futures they might exist in and need to navigate as both members of society *and* as individuals in positions of power to regulate these futures. Organizations like OSTP are actively soliciting information from a diverse set of voices about the public and private sector uses of biometric technologies [144], recognizing the need to “*to understand the extent and variety of biometric technologies in past, current, or planned use; the domains in which these technologies are being used; the entities making use of them; current principles, practices, or policies governing their use; and the stakeholders that are, or may be, impacted by their use or regulation*” [144]. Our work helps to draw attention to the imagined scope and varied future uses of biometric technologies, specifically emotion AI for monitoring and detecting emotions and mental health.

Through highlighting the ethical implications of scope and varied planned uses of emotion AI for mental health, technologists may better approach problem formulation and solution development to not impose harms *and* providers may also be better prepared to consider the ethical implications of

emotion AI before deciding to deploy them in care, safety and health-related contexts. Researchers and technologists who develop these technologies may be proactive in working to evaluate emerging technologies for their ethical implications in myriad contexts, working to increase awareness of their impacts and get ahead of potential problematic applications *prior* to their deployment. Additionally, in the spirit of transparency *and* collaboration, technologists and researchers could be open with positive and concerning findings in evaluation processes—such as by responding to calls like [144], prompting collective efforts.

We expand on prior work [28] in questioning the inevitability of sociotechnical futures where emotion AI technologies are deployed for monitoring and detecting emotions. While what some have described as AI snake oil [136] may make emotion AI futures seem prewritten and at our doorstep, we challenge this technological deterministic view that any technological ‘progress’ and development is inevitable or independently occurring outside the social, political, and regulatory systems in which we play a part. This deterministic view *“leaves no space for human choice or intervention and, moreover, absolves us from responsibility for the technologies we make and use”* [213] and leaves us prey to the fallacy of technological progress equating to social progress. We have a responsibility—as (CSCW) scholars, as innovators, as regulators, as ordinary citizens—over the technologies we produce and use. It is crucial for the true enactment of this responsibility, however, that we are *all* included in technological development processes in a variety of ways (e.g. advisory boards, inventors, community members) so *all* voices and varied experiences are considered equitably before the claim of true progress is made.

7 CONCLUSION

This paper provides a partial, though unique and insightful, view into the sociotechnical futures of emotion AI for mental health as presented in the discourses found in patent applications. Through our analysis of 58 mental health patent applications filed in the U.S., we access the sociotechnical imaginary of emotion AI and discover that emotion AI technologies are legitimized by claiming to improve data accuracy, care provision and experience, patient-provider communication, emotion regulation, and harms attributed to mental health. Using an ethical speculation lens, we then unpack our findings to identify several ethical implications for emotion AI in the mental health domain. We argue that these patent applications justify emotion AI by framing individuals with mental health conditions (or unanticipated emotions) in stigmatizing ways, associating mental health with crime and passivity. We explain how emotion AI’s imagined futures, as seen through patent applications, portray a commitment to invasive surveillance, and highlight harms that may result. We articulate how these imagined emotion AI technologies recommend behavioral changes based on prescribed totalizing definitions of health that place the onus on the individual while dismissing systemic barriers to well-being and health. We call into question emotion AI as an inevitable or desirable part of health and care-related contexts.

8 ACKNOWLEDGMENTS

We are grateful to the AC and anonymous reviewers for their detailed feedback and suggestions that made this work and its contributions stronger. This work was supported by the National Science Foundation (award 2020872 and CAREER award 2236674). We would like to thank Karen Boyd for her role establishing the dataset of 58 patent applications as part of a larger project and documenting its creation process. The first author would also like to thank her two cats, Mishmish and Yeimy, who diligently supervised data analysis despite not understanding what patent applications and emotion AI are.

REFERENCES

- [1] Assad Abbas, Limin Zhang, and Samee U. Khan. 2014. A literature review on the state-of-the-art in patent analysis. *World Patent Information* 37 (June 2014), 3–13. <https://doi.org/10.1016/j.wpi.2013.12.006>
- [2] Access Now. 2022. Prohibit emotion recognition in the Artificial Intelligence Act. <https://www.accessnow.org/cms/assets/uploads/2022/05/Prohibit-emotion-recognition-in-the-Artificial-Intelligence-Act.pdf>
- [3] Muhammad Aurangzeb Ahmad, Steve Overman, Christine Allen, Vikas Kumar, Ankur Teredesai, and Carly Eckert. 2021. Software as a Medical Device: Regulating AI in Healthcare via Responsible AI. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. ACM, Virtual Event Singapore, 4023–4024. <https://doi.org/10.1145/3447548.3470823>
- [4] Muhammad Aurangzeb Ahmad, Arpit Patel, Carly Eckert, Vikas Kumar, and Ankur Teredesai. 2020. Fairness in Machine Learning for Healthcare. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. ACM, Virtual Event CA USA, 3529–3530. <https://doi.org/10.1145/3394486.3406461>
- [5] Sirry Alang. 2020. Police Brutality and the Institutional Patterning of Stressors. *American Journal of Public Health* 110, 11 (Nov. 2020), 1597–1598. <https://doi.org/10.2105/AJPH.2020.305937>
- [6] Ali Alkhatib. 2021. To Live in Their Utopia: Why Algorithmic Systems Create Absurd Outcomes. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–9. <https://doi.org/10.1145/3411764.3445740>
- [7] Ngozi F. Anachebe and Madeline Y. Sutton. 2003. Racial disparities in reproductive health outcomes. *American Journal of Obstetrics and Gynecology* 188, 4 (April 2003), S37–S42. <https://doi.org/10.1067/mob.2003.245>
- [8] 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*. ACM, Honolulu HI USA, 1–16. <https://doi.org/10.1145/3313831.3376680>
- [9] Tariq Osman Andersen, Jørgen Peter Bansler, Finn Kensing, Jonas Moll, Troels Mønsted, Karen Dam Nielsen, Olav Wendelboe Nielsen, Helen Høgh Petersen, and Jesper Hastrup Svendsen. 2019. Aligning Concerns in Telecare: Three Concepts to Guide the Design of Patient-Centred E-Health. *Computer Supported Cooperative Work (CSCW)* 28, 6 (Oct. 2019), 1039–1072. <https://doi.org/10.1007/s10606-018-9309-1>
- [10] K. S. Anthony. 2019. Company With Creepy AI That Recognizes Emotions Aims to 'Understand All Things Human'. <https://www.outerplaces.com/science/item/19329-affectiva-ai-emotional-recognition>
- [11] Jascha Bareis and Christian Katzenbach. 2022. Talking AI into Being: The Narratives and Imaginaries of National AI Strategies and Their Performative Politics. *Science, Technology, & Human Values* 47, 5 (Sept. 2022), 855–881. <https://doi.org/10.1177/01622439211030007> Publisher: SAGE Publications Inc.
- [12] Lisa Feldman Barrett. 2006. Are Emotions Natural Kinds? *Perspectives on Psychological Science* 1, 1 (March 2006), 28–58. <https://doi.org/10.1111/j.1745-6916.2006.00003.x>
- [13] Lisa Feldman Barrett. 2014. Opinion | What Faces Can't Tell Us. *The New York Times* (Feb. 2014). <https://www.nytimes.com/2014/03/02/opinion/sunday/what-faces-cant-tell-us.html>
- [14] 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 (July 2019), 1–68. <https://doi.org/10.1177/1529100619832930>
- [15] Lisa Feldman Barrett, James Gross, Tamlin Conner Christensen, and Michael Benvenuto. 2001. Knowing what you're feeling and knowing what to do about it: Mapping the relation between emotion differentiation and emotion regulation. *Cognition & Emotion* 15, 6 (Nov. 2001), 713–724. <https://doi.org/10.1080/02699930143000239>
- [16] Ivana Bartoletti. 2019. AI in Healthcare: Ethical and Privacy Challenges. In *Artificial Intelligence in Medicine*, David Riaño, Szymon Wilk, and Annette ten Teije (Eds.). Vol. 11526. Springer International Publishing, Cham, 7–10. https://doi.org/10.1007/978-3-030-21642-9_2 Series Title: Lecture Notes in Computer Science.
- [17] Thurstine Basset, Alison Faulkner, Julie Repper, and Elina Stamou. 2010. Lived experience leading the way: Peer support in mental health. *London: Together UK* (2010).
- [18] Eric P.S. Baumer and M. Six Silberman. 2011. When the implication is not to design (technology). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Vancouver BC Canada, 2271–2274. <https://doi.org/10.1145/1978942.1979275>
- [19] Anthony F. Beavers and Justin P. Slattery. 2017. On the moral implications and restrictions surrounding affective computing. In *Emotions and Affect in Human Factors and Human-Computer Interaction*. Elsevier, 143–161.
- [20] Ruth Behar. 2022. *The vulnerable observer: Anthropology that breaks your heart*. beacon press.
- [21] Dörte Bemme, Natassia F. Brenman, and Beth Semel. 2020. The subjects of digital psychiatry. (2020).
- [22] Jessica C. Bird, Felicity Waite, Eleanor Rowsell, Emma C. Fergusson, and Daniel Freeman. 2017. Cognitive, affective, and social factors maintaining paranoia in adolescents with mental health problems: A longitudinal study. *Psychiatry research* 257 (2017), 34–39. ISBN: 0165-1781 Publisher: Elsevier.

- [23] Kirsten Boehner, Rogério DePaula, Paul Dourish, and Phoebe Sengers. 2005. Affect: from information to interaction. In *Proceedings of the 4th decennial conference on Critical computing between sense and sensibility - CC '05*. ACM Press, Aarhus, Denmark, 59. <https://doi.org/10.1145/1094562.1094570>
- [24] Mads Borup, Nik Brown, Kornelia Konrad, and Harro Van Lente. 2006. The sociology of expectations in science and technology. *Technology Analysis & Strategic Management* 18, 3-4 (July 2006), 285–298. <https://doi.org/10.1080/09537320600777002>
- [25] Eliane M. Boucher, Nicole R. Harake, Haley E. Ward, Sarah Elizabeth Stoeckl, Junielly Vargas, Jared Minkel, Acacia C. Parks, and Ran Zilca. 2021. Artificially intelligent chatbots in digital mental health interventions: a review. *Expert Review of Medical Devices* 18, sup1 (Dec. 2021), 37–49. <https://doi.org/10.1080/17434440.2021.2013200> Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/17434440.2021.2013200>.
- [26] A.C. Bovell-Benjamin, C.S. Hathorn, S. Ibrahim, P.N. Gichuhi, and E.M. Bromfield. 2009. Healthy food choices and physical activity opportunities in two contrasting Alabama cities. *Health & Place* 15, 2 (June 2009), 429–438. <https://doi.org/10.1016/j.healthplace.2008.08.001>
- [27] Geoffrey C. Bowker and Susan Leigh Star. 2008. *Sorting things out: classification and its consequences* (1. paperback ed., 8. print ed.). MIT Press, Cambridge, Mass.
- [28] Karen 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* CSCW (2023).
- [29] Daniel E. Brown and Lynnette Leidy Sievert. 2016. Making Visible the Invisible. In *Biological Measures of Human Experience across the Lifespan*, Lynnette Leidy Sievert and Daniel E. Brown (Eds.). Springer International Publishing, Cham, 1–10. https://doi.org/10.1007/978-3-319-44103-0_1
- [30] Phil Brown. 1995. Naming and Framing: The Social Construction of Diagnosis and Illness. *Journal of Health and Social Behavior* 35 (1995), 34. <https://doi.org/10.2307/2626956>
- [31] Sandra Bucci, Matthias Schwannauer, and Natalie Berry. 2019. The digital revolution and its impact on mental health care. *Psychology and Psychotherapy: Theory, Research and Practice* 92, 2 (2019), 277–297. <https://doi.org/10.1111/papt.12222> _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/papt.12222>.
- [32] Taina Bucher. 2020. The right-time web: Theorizing the kairologic of algorithmic media. *New Media & Society* 22, 9 (Sept. 2020), 1699–1714. <https://doi.org/10.1177/1461444820913560> Place: London Publisher: Sage Publications Ltd WOS:000566711000011.
- [33] Dan L. Burk and Jessica Reyman. 2014. Patents as Genre: A Prospectus. *Law & Literature* 26, 2 (May 2014), 163–190. <https://doi.org/10.1080/1535685X.2014.888193>
- [34] Kelly Caine. 2016. Privacy Is Healthy. *IEEE Pervasive Computing* 15, 4 (Oct. 2016), 14–19. <https://doi.org/10.1109/MPRV.2016.61>
- [35] Tina Caliendo and Olga Hilas. 2019. The Promise and Pitfalls of Digital Medication. *US Pharmacist* 44, 7 (2019), 22–24. <https://www.uspharmacist.com/article/the-promise-and-pitfalls-of-digital-medication>
- [36] Rafael A Calvo and Sidney D’Mello. 2010. Affect Detection: An Interdisciplinary Review of Models, Methods, and Their Applications. *IEEE Transactions on Affective Computing* 1, 1 (Jan. 2010), 18–37. <https://doi.org/10.1109/T-AFFC.2010.1>
- [37] Yuanyuan Cao, Junjun Li, Xinghong Qin, and Baoliang Hu. 2020. Examining the Effect of Overload on the MHealth Application Resistance Behavior of Elderly Users: An SOR Perspective. *International Journal of Environmental Research and Public Health* 17, 18 (Sept. 2020), 6658. <https://doi.org/10.3390/ijerph17186658>
- [38] Stefan Carmien, Melissa Dawe, Gerhard Fischer, Andrew Gorman, Anja Kintsch, and James F. Sullivan. 2005. Socio-technical environments supporting people with cognitive disabilities using public transportation. *ACM Transactions on Computer-Human Interaction* 12, 2 (June 2005), 233–262. <https://doi.org/10.1145/1067860.1067865>
- [39] Noel Carroll, Ita Richardson, and Raja Abbas. 2020. “The Algorithm Will See You Now”: Exploring the Implications of Algorithmic Decision-making in Connected Health. In *Proceedings of the 13th International Joint Conference on Biomedical Engineering Systems and Technologies*. SCITEPRESS - Science and Technology Publications, Valletta, Malta, 758–765. <https://doi.org/10.5220/0009178507580765>
- [40] Stacy M. Carter, Wendy Rogers, Khin Than Win, Helen Frazer, Bernadette Richards, and Nehmat Houssami. 2020. The ethical, legal and social implications of using artificial intelligence systems in breast cancer care. *The Breast* 49 (Feb. 2020), 25–32. <https://doi.org/10.1016/j.breast.2019.10.001>
- [41] Pao-Long Chang, Chao-Chan Wu, and Hoang-Jyh Leu. 2010. Using patent analyses to monitor the technological trends in an emerging field of technology: a case of carbon nanotube field emission display. *Scientometrics* 82, 1 (Jan. 2010), 5–19. <https://doi.org/10.1007/s11192-009-0033-y>
- [42] Pao-Long Chang, Chao-Chan Wu, and Hoang-Jyh Leu. 2012. Investigation of technological trends in flexible display fabrication through patent analysis. *Displays* 33, 2 (April 2012), 68–73. <https://doi.org/10.1016/j.displa.2012.03.003>
- [43] Jin Chen, Cheng Chen, Joseph B. Walther, and S. Shyam Sundar. 2021. Do You Feel Special When an AI Doctor Remembers You? Individuation Effects of AI vs. Human Doctors on User Experience. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–7. <https://doi.org/10.1145/>

3411763.3451735

- [44] Yang Cheng and Hua Jiang. 2020. AI-Powered mental health chatbots: Examining users' motivations, active communicative action and engagement after mass-shooting disasters. *Journal of Contingencies and Crisis Management* 28, 3 (2020), 339–354. <https://doi.org/10.1111/1468-5973.12319> _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-5973.12319>.
- [45] D. J. Chinn, M. White, J. Harland, C. Drinkwater, and S. Raybould. 1999. Barriers to physical activity and socioeconomic position: implications for health promotion. *Journal of Epidemiology & Community Health* 53, 3 (March 1999), 191–192. <https://doi.org/10.1136/jech.53.3.191>
- [46] Munmun De Choudhury and Sushovan De. 2014. Mental Health Discourse on reddit: Self-Disclosure, Social Support, and Anonymity. In *Eighth International AAAI Conference on Weblogs and Social Media*. <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM14/paper/view/8075>
- [47] Chia-Fang Chung, Kristin Dew, Allison Cole, Jasmine Zia, James Fogarty, Julie A. Kientz, and Sean A. Munson. 2016. Boundary Negotiating Artifacts in Personal Informatics: Patient-Provider Collaboration with Patient-Generated Data. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. Association for Computing Machinery, New York, NY, USA, 770–786. <https://doi.org/10.1145/2818048.2819926>
- [48] Danielle Keats Citron and Daniel J. Solove. 2021. Privacy Harms. <https://doi.org/10.2139/ssrn.3782222>
- [49] Franck Cochoy and Bastien Soutjis. 2020. Back to the future of digital price display: Analyzing patents and other archives to understand contemporary market innovations. *Social Studies of Science* 50, 1 (Feb. 2020), 3–29. <https://doi.org/10.1177/0306312719884643>
- [50] Office of the Commissioner. 2020. FDA approves pill with sensor that digitally tracks if patients have ingested their medication. <https://www.fda.gov/news-events/press-announcements/fda-approves-pill-sensor-digitally-tracks-if-patients-have-ingested-their-medication> Publisher: FDA.
- [51] Juliet M. Corbin and Anselm L. Strauss. 2015. *Basics of qualitative research: techniques and procedures for developing grounded theory* (fourth edition ed.). SAGE, Los Angeles.
- [52] Adriana Cordova and Francesco Moschella. 2008. Algorithm for clinical evaluation and surgical treatment of gynaecomastia. *Journal of Plastic, Reconstructive & Aesthetic Surgery* 61, 1 (Jan. 2008), 41–49. <https://doi.org/10.1016/j.bjps.2007.09.033>
- [53] Patrick Corrigan. 2004. How stigma interferes with mental health care. *American Psychologist* 59, 7 (Oct. 2004), 614–625. <https://doi.org/10.1037/0003-066X.59.7.614>
- [54] Shanley Corvite, Kat Roemmich, Tillie 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* CSCW (2023).
- [55] Alan Cowen, Disa Sauter, Jessica L. Tracy, and Dacher Keltner. 2019. Mapping the Passions: Toward a High-Dimensional Taxonomy of Emotional Experience and Expression. *Psychological Science in the Public Interest* 20, 1 (July 2019), 69–90. <https://doi.org/10.1177/1529100619850176> Publisher: SAGE Publications Inc.
- [56] Kate Crawford. 2021. Artificial Intelligence Is Misreading Human Emotion. <https://www.theatlantic.com/technology/archive/2021/04/artificial-intelligence-misreading-human-emotion/618696/> Section: Technology.
- [57] Kate Crawford. 2021. *Atlas of AI: power, politics, and the planetary costs of artificial intelligence*. Yale University Press, New Haven. OCLC: on1111967630.
- [58] Jessica Crowe, Constance Lacy, and Yolanda Columbus. 2018. Barriers to Food Security and Community Stress in an Urban Food Desert. *Urban Science* 2, 2 (May 2018), 46. <https://doi.org/10.3390/urbansci2020046>
- [59] Shaundra B. Daily, Melva T. James, David Cherry, John J. Porter III, Shelby S. Darnell, Joseph Isaac, and Tania Roy. 2017. Affective computing: historical foundations, current applications, and future trends. *Emotions and affect in human factors and human-computer interaction* (2017), 213–231. Publisher: Elsevier.
- [60] Tugrul U. Daim, Guillermo Rueda, Hilary Martin, and Pisek Gerdstri. 2006. Forecasting emerging technologies: Use of bibliometrics and patent analysis. *Technological Forecasting and Social Change* 73, 8 (Oct. 2006), 981–1012. <https://doi.org/10.1016/j.techfore.2006.04.004>
- [61] Simon Davis. 2002. Brief Report: Autonomy Versus Coercion: Reconciling Competing Perspectives in Community Mental Health. *Community Mental Health Journal* 38, 3 (June 2002), 239–250. <https://doi.org/10.1023/A:1015267707856>
- [62] Alessandro Delfanti and Bronwyn Frey. 2021. Humanly Extended Automation or the Future of Work Seen through Amazon Patents. *Science, Technology, & Human Values* 46, 3 (May 2021), 655–682. <https://doi.org/10.1177/0162243920943665> Publisher: SAGE Publications Inc.
- [63] Department of Health and Human Services. 2022. Mental Health Myths and Facts. <https://www.mentalhealth.gov/basics/mental-health-myths-facts>
- [64] Lindsay H. Dewa, Mary Lavelle, Katy Pickles, Caroline Kalorkoti, Jack Jaques, Sofia Pappa, and Paul Aylin. 2019. Young adults' perceptions of using wearables, social media and other technologies to detect worsening mental health: A qualitative study. *PLOS ONE* 14, 9 (Sept. 2019), e0222655. <https://doi.org/10.1371/journal.pone.0222655>

- [65] Tawanna R Dillahunt, Juan F. Maestre, Vaishnav Kameswaran, Erica Poon, John Osorio Torres, Mia Gallardo, Samantha E. Rasmussen, Patrick C. Shih, Alice Bagley, Samuel L. A. Young, and Tiffany C. Veinot. 2022. Trust, Reciprocity, and the Role of Timebanks as Intermediaries: Design Implications for Addressing Healthcare Transportation Barriers. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–22. <https://doi.org/10.1145/3491102.3502494>
- [66] Marlese Durr. 2015. What is the Difference between Slave Patrols and Modern Day Policing? Institutional Violence in a Community of Color. *Critical Sociology* 41, 6 (Sept. 2015), 873–879. <https://doi.org/10.1177/0896920515594766>
- [67] Andrius Dzedzickis, Artūras Kaklauskas, and Vytautas Bucinskas. 2020. Human Emotion Recognition: Review of Sensors and Methods. *Sensors* 20, 3 (Jan. 2020), 592. <https://doi.org/10.3390/s20030592>
- [68] Simon D'Alfonso. 2020. AI in mental health. *Current Opinion in Psychology* 36 (Dec. 2020), 112–117. <https://doi.org/10.1016/j.copsyc.2020.04.005>
- [69] Frank Edwards, Michael H. Esposito, and Hedwig Lee. 2018. Risk of Police-Involved Death by Race/Ethnicity and Place, United States, 2012–2018. *American Journal of Public Health* 108, 9 (Sept. 2018), 1241–1248. <https://doi.org/10.2105/AJPH.2018.304559> Publisher: American Public Health Association.
- [70] Claudia Egher. 2023. Epistemic Inroads from the Asylum to Digital Psychiatry. In *Digital Healthcare and Expertise: Mental Health and New Knowledge Practices*, Claudia Egher (Ed.). Springer Nature Singapore, Singapore, 37–70. https://doi.org/10.1007/978-981-16-9178-2_2
- [71] Paul Ekman. 1971. Universals and cultural differences in facial expressions of emotion.. In *Nebraska symposium on motivation*, Vol. 19. University of Nebraska Press.
- [72] Darren Ellis. 2020. Techno-Securitisation of Everyday Life and Cultures of Surveillance-Apatheia. *Science as Culture* 29, 1 (Jan. 2020), 11–29. <https://doi.org/10.1080/09505431.2018.1561660>
- [73] Frank Engster and Phoebe V Moore. 2020. The search for (artificial) intelligence, in capitalism. *Capital & Class* 44, 2 (June 2020), 201–218. <https://doi.org/10.1177/0309816820902055>
- [74] Burches Enrique and Burches Marta. 2020. Efficacy, Effectiveness and Efficiency in the Health Care: The Need for an Agreement to Clarify its Meaning. *International Archives of Public Health and Community Medicine* 4, 1 (Jan. 2020). <https://doi.org/10.23937/2643-4512/1710035>
- [75] Alessandro Evangelista, Lorenzo Ardito, Antonio Boccaccio, Michele Fiorentino, Antonio Messeni Petruzzelli, and Antonio E. Uva. 2020. Unveiling the technological trends of augmented reality: A patent analysis. *Computers in Industry* 118 (June 2020), 103221. <https://doi.org/10.1016/j.compind.2020.103221>
- [76] Paul Festor, Ibrahim Habli, Yan Jia, Anthony Gordon, A. Aldo Faisal, and Matthieu Komorowski. 2021. Levels of Autonomy and Safety Assurance for AI-Based Clinical Decision Systems. In *Computer Safety, Reliability, and Security. SAFECOMP 2021 Workshops*, Ibrahim Habli, Mark Suján, Simos Gerasimou, Erwin Schoitsch, and Friedemann Bitsch (Eds.). Vol. 12853. Springer International Publishing, Cham, 291–296. https://doi.org/10.1007/978-3-030-83906-2_24 Series Title: Lecture Notes in Computer Science.
- [77] Casey Fiesler. 2019. Ethical Considerations for Research Involving (Speculative) Public Data. *Proceedings of the ACM on Human-Computer Interaction* 3, GROUP (Dec. 2019), 249:1–249:13. <https://doi.org/10.1145/3370271>
- [78] Casey Fiesler. 2021. *Innovating Like an Optimist, Preparing Like a Pessimist: Ethical Speculation and the Legal Imagination*. SSRN Scholarly Paper 3779036. Social Science Research Network, Rochester, NY. <https://papers.ssrn.com/abstract=3779036>
- [79] Max Fisher and Josh Keller. 2017. Why Does the U.S. Have So Many Mass Shootings? Research Is Clear: Guns. *The New York Times* (Nov. 2017). <https://www.nytimes.com/2017/11/07/world/americas/mass-shootings-us-international.html>
- [80] Jessica L. Ford and Emily N. Scheinfeld. 2016. Exploring the Effects of Workplace Health Promotions: A Critical Examination of a Familiar Organizational Practice. *Annals of the International Communication Association* 40, 1 (Jan. 2016), 277–305. <https://doi.org/10.1080/23808985.2015.11735263>
- [81] Annie B. Fox, Brian N. Smith, and Dawne Vogt. 2018. How and when does mental illness stigma impact treatment seeking? Longitudinal examination of relationships between anticipated and internalized stigma, symptom severity, and mental health service use. *Psychiatry Research* 268 (Oct. 2018), 15–20. <https://doi.org/10.1016/j.psychres.2018.06.036>
- [82] Fight for the Future. 2022. Letter to Zoom on Emotion Analysis Software. <https://www.fightforthefuture.org/news/2022-05-10-letter-to-zoom>
- [83] Ryan Gabrielson, Eric Sagara, and Ryan Grochowski Jones. 2014. Deadly Force, in Black and White. <https://www.propublica.org/article/deadly-force-in-black-and-white>
- [84] Maria Gendron and Lisa Feldman Barrett. 2009. Reconstructing the Past: A Century of Ideas About Emotion in Psychology. *Emotion Review* 1, 4 (Oct. 2009), 316–339. <https://doi.org/10.1177/1754073909338877>
- [85] Erving Goffman. 1986. *Stigma: notes on the management of spoiled identity* (26th pr ed.). Simon & Schuster, New York. OCLC: 299745941.

- [86] Cristina Gorrostieta, Reza Lotfian, Kye Taylor, Richard Brutti, and John Kane. 2019. Gender De-Biasing in Speech Emotion Recognition. In *Interspeech 2019*. ISCA, 2823–2827. <https://doi.org/10.21437/Interspeech.2019-1708>
- [87] Ben Green. 2021. Data Science as Political Action: Grounding Data Science in a Politics of Justice. *Journal of Social Computing* 2, 3 (Sept. 2021), 249–265. <https://doi.org/10.23919/JSC.2021.0029>
- [88] Gretchen Greene. 2022. *The Ethics of AI and Emotional Intelligence: Data sources, applications, and questions for evaluating ethics risk*. Technical Report. Partnership on AI. <https://partnershiponai.org/paper/the-ethics-of-ai-and-emotional-intelligence/>
- [89] Thomas Grote and Philipp Berens. 2020. On the ethics of algorithmic decision-making in healthcare. *Journal of Medical Ethics* 46, 3 (March 2020), 205–211. <https://doi.org/10.1136/medethics-2019-105586>
- [90] Alana Gunn. 2022. Stigma, surveillance, and wounded healing: Promoting a critical ethics of care in research with formerly incarcerated Black women. *Journal of Community Psychology* (March 2022), jcop.22845. <https://doi.org/10.1002/jcop.22845>
- [91] Katarina Hamberg. 2008. Gender Bias in Medicine. *Women's Health* 4, 3 (May 2008), 237–243. <https://doi.org/10.2217/17455057.4.3.237> Publisher: SAGE Publications Ltd STM.
- [92] Jay Hancock. 2015. Work wellness programs put employee privacy at risk. <http://www.cnn.com/2015/09/28/health/workplace-wellness-privacy-risk-exclusive/> Publication Title: CNN Health.
- [93] Douglas Heaven. 2020. Why faces don't always tell the truth about feelings. *Nature* 578, 7796 (Feb. 2020), 502–504. <https://doi.org/10.1038/d41586-020-00507-5>
- [94] Gregory M. Herek, John P. Capitanio, and Keith F. Widaman. 2003. Stigma, social risk, and health policy: Public attitudes toward HIV surveillance policies and the social construction of illness. *Health Psychology* 22, 5 (2003), 533–540. <https://doi.org/10.1037/0278-6133.22.5.533>
- [95] Stephen Hilgartner. 2015. Capturing the imaginary. *Science and Democracy: Making Knowledge and Making Power in the Biosciences and Beyond*. Abingdon: Routledge (2015).
- [96] L K. Hironaka and M. K. Paasche-Orlow. 2008. The implications of health literacy on patient-provider communication. *Archives of Disease in Childhood* 93, 5 (May 2008), 428–432. <https://doi.org/10.1136/adc.2007.131516>
- [97] Jill R. Horwitz, Brenna D. Kelly, and John E. DiNardo. 2013. Wellness Incentives In The Workplace: Cost Savings Through Cost Shifting To Unhealthy Workers. *Health Affairs* 32, 3 (March 2013), 468–476. <https://doi.org/10.1377/hlthaff.2012.0683>
- [98] Jane Hu. 2022. Abilify MyCite tells your doctor if you don't take your pill, raising ethical questions. *Slate* (April 2022). <https://slate.com/technology/2022/04/abilify-mycite-pill-tracking-mental-health.html>
- [99] Andrew Iliadis and Amelia Acker. 2022. The seer and the seen: Surveying Palantir's surveillance platform. *The Information Society* 0, 0 (Aug. 2022), 1–30. <https://doi.org/10.1080/01972243.2022.2100851> Publisher: Routledge _eprint: <https://doi.org/10.1080/01972243.2022.2100851>.
- [100] Rachael E. Jack, Caroline Blais, Christoph Scheepers, Philippe G. Schyns, and Roberto Caldara. 2009. Cultural Confusions Show that Facial Expressions Are Not Universal. *Current Biology* 19, 18 (Sept. 2009), 1543–1548. <https://doi.org/10.1016/j.cub.2009.07.051>
- [101] Maia Jacobs, Melanie F. Pradier, Thomas H. McCoy, Roy H. Perlis, Finale Doshi-Velez, and Krzysztof Z. Gajos. 2021. How machine-learning recommendations influence clinician treatment selections: the example of antidepressant selection. *Translational Psychiatry* 11, 1 (Feb. 2021), 1–9. <https://doi.org/10.1038/s41398-021-01224-x> Number: 1 Publisher: Nature Publishing Group.
- [102] Maia L. Jacobs, James Clawson, and Elizabeth D. Mynatt. 2015. Comparing Health Information Sharing Preferences of Cancer Patients, Doctors, and Navigators. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. Association for Computing Machinery, New York, NY, USA, 808–818. <https://doi.org/10.1145/2675133.2675252>
- [103] Sheila Jasanoff and Sang-Hyun Kim. 2009. Containing the Atom: Sociotechnical Imaginaries and Nuclear Power in the United States and South Korea. *Minerva* 47, 2 (June 2009), 119–146. <https://doi.org/10.1007/s11024-009-9124-4>
- [104] Suneel Jethani. 2020. Doing Time in the Home-Space: Ankle Monitors, Script Analysis, and Anticipatory Methodology. *Embodied Computing: Wearables, Implantables, Embeddables, Ingestibles* (2020), 161–186. ISBN: 0262357801 Publisher: MIT Press.
- [105] David K. Johnson. 2006. *The lavender scare: the Cold War persecution of gays and lesbians in the federal government* (paperb. ed ed.). The Univ. of Chicago Press, Chicago, Ill.
- [106] Edward B Kang. 2022. Biometric imaginaries: Formatting voice, body, identity to data. *Social Studies of Science* 52, 4 (Aug. 2022), 581–602. <https://doi.org/10.1177/03063127221079599> Publisher: SAGE Publications Ltd.
- [107] Andreas Kaplan and Michael Haenlein. 2019. Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons* 62, 1 (Jan. 2019), 15–25. <https://doi.org/10.1016/j.bushor.2018.08.004>

- [108] Matthew M Kavanagh, Stefan D Baral, Maureen Milanga, and Jeremy Sugarman. 2019. Biometrics and public health surveillance in criminalised and key populations: policy, ethics, and human rights considerations. *The Lancet HIV* 6, 1 (Jan. 2019), e51–e59. [https://doi.org/10.1016/S2352-3018\(18\)30243-1](https://doi.org/10.1016/S2352-3018(18)30243-1)
- [109] Christopher J. Kelly, Alan Karthikesalingam, Mustafa Suleyman, Greg Corrado, and Dominic King. 2019. Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine* 17, 1 (Oct. 2019), 195. <https://doi.org/10.1186/s12916-019-1426-2>
- [110] Eugenia Kim, De'Aira Bryant, Deepak Srikanth, and Ayanna Howard. 2021. Age Bias in Emotion Detection: An Analysis of Facial Emotion Recognition Performance on Young, Middle-Aged, and Older Adults. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, Virtual Event USA, 638–644. <https://doi.org/10.1145/3461702.3462609>
- [111] Young Gil Kim, Jong Hwan Suh, and Sang Chan Park. 2008. Visualization of patent analysis for emerging technology. *Expert Systems with Applications* 34, 3 (April 2008), 1804–1812. <https://doi.org/10.1016/j.eswa.2007.01.033>
- [112] Victoria Knight. 2019. America to face a shortage of primary care physicians within a decade or so. *Washington Post* (2019).
- [113] Redi Koobak and Suruchi Thapar-Björkert. 2014. Writing the place from which one speaks. *Writing academic texts differently: Intersectional feminist methodologies and the playful art of writing* (2014), 47–61. Publisher: Routledge 605 Third Avenue, New York, NY 10017.
- [114] Kira Kretzschmar, Holly Tyroll, Gabriela Pavarini, Arianna Manzini, Ilina Singh, and NeurOx Young People's Advisory Group. 2019. Can Your Phone Be Your Therapist? Young People's Ethical Perspectives on the Use of Fully Automated Conversational Agents (Chatbots) in Mental Health Support. *Biomedical Informatics Insights* 11 (2019), 1178222619829083. <https://doi.org/10.1177/1178222619829083>
- [115] Shailesh Kumar. 2016. Burnout and Doctors: Prevalence, Prevention and Intervention. *Healthcare* 4, 3 (June 2016), 37. <https://doi.org/10.3390/healthcare4030037>
- [116] Alfons Labisch. 1992. The Social Construction of Health. From Early Modern Times to the Beginnings of the Industrialization. *The Social Construction of Illness. Illness and Medical Knowledge in Past and Present*, Stuttgart, Franz Steiner (1992), 85–101.
- [117] Sarah Esther Lageson. 2022. Criminal Record Stigma and Surveillance in the Digital Age. *Annual Review of Criminology* 5, 1 (Jan. 2022), 67–90. <https://doi.org/10.1146/annurev-criminol-030920-092833>
- [118] K. K. Rebecca Lai, Bill Marsh, and Anjali Singhvi. 2020. Here Are the 100 U.S. Cities Where Protesters Were Tear-Gassed. *The New York Times* (June 2020). <https://www.nytimes.com/interactive/2020/06/16/us/george-floyd-protests-police-tear-gas.html>
- [119] Emily LaRosa and David Danks. 2018. Impacts on Trust of Healthcare AI. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. ACM, New Orleans LA USA, 210–215. <https://doi.org/10.1145/3278721.3278771>
- [120] Min Kyung Lee and Katherine Rich. 2021. Who Is Included in Human Perceptions of AI?: Trust and Perceived Fairness around Healthcare AI and Cultural Mistrust. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–14. <https://doi.org/10.1145/3411764.3445570>
- [121] Helen Lewis, Block. 1971. Shame and guilt in neurosis. *Psychoanalytic review* 58, 3 (1971), 419–438. Publisher: National Psychological Association for Psychoanalysis.
- [122] ShiPian Li, LinRu Hou, KeWen Chen, ZhengLiang Zhe, Ce Ying, Yi Guo, and HaiYan Ren. 2022. On the present and future of medical humanistic care in the era of artificial intelligence. *History and Philosophy of Medicine* 4, 2 (2022), 5. <https://doi.org/10.53388/HPM20220401013>
- [123] Andrej Lunekski, Panagiotis Bamidis, and Madga Hitoglou-Antoniadou. 2008. Affective computing and medical informatics: State of the art in emotion-aware medical applications. *Studies in health technology and informatics* 136 (Feb. 2008), 517–22. <https://doi.org/10.3233/978-1-58603-864-9-517>
- [124] Deborah Lupton. 2000. The social construction of medicine and the body. *Handbook of social studies in health and medicine* (2000), 50–63. Publisher: SAGE Publications Ltd UK.
- [125] Kristin M. Madison. 2016. The Risks Of Using Workplace Wellness Programs To Foster A Culture Of Health. *Health Affairs* 35, 11 (Nov. 2016), 2068–2074. <https://doi.org/10.1377/hlthaff.2016.0729>
- [126] Pauline Katharina Mantell, Annika Baumeister, Stephan Ruhrmann, Anna Janhsen, and Christiane Wooten. 2021. Attitudes towards Risk Prediction in a Help Seeking Population of Early Detection Centers for Mental Disorders—A Qualitative Approach. *International Journal of Environmental Research and Public Health* 18, 3 (Jan. 2021), 1036. <https://doi.org/10.3390/ijerph18031036>
- [127] Marco Marabelli and Sue Newell. 2019. Reflections on Algorithmic Decision-making in the Healthcare Industry. In *Organizational Learning, Knowledge and Capabilities*. Brighton, UK.
- [128] Aaron K Martin, Rosamunde E Van Brakel, and Daniel J Bernhard. 2009. Understanding resistance to digital surveillance: Towards a multi-disciplinary, multi-actor framework. *Surveillance & Society* 6, 3 (April 2009), 213–232. <https://doi.org/10.24908/ss.v6i3.3282>

- [129] Paul Martin. 2015. Commercialising neurofutures: Promissory economies, value creation and the making of a new industry. *BioSocieties* 10, 4 (Dec. 2015), 422–443. <https://doi.org/10.1057/biosoc.2014.40>
- [130] Cynthia K. Matthews and Nancy Grant Harrington. 2000. Invisible disability. In *Handbook of communication and people with disabilities: Research and application*. Lawrence Erlbaum Associates Publishers, Mahwah, NJ, US, 405–421.
- [131] Andrew McStay. 2018. *Emotional AI: the rise of empathic media*. SAGE, Los Angeles.
- [132] Andrew McStay. 2020. Emotional AI, soft biometrics and the surveillance of emotional life: An unusual consensus on privacy. *Big Data & Society* 7, 1 (Jan. 2020), 205395172090438. <https://doi.org/10.1177/2053951720904386>
- [133] Cilia Mejia-Lancheros, James Lachaud, Julia Woodhall-Melnik, Patricia O’Campo, Stephen W. Hwang, and Vicky Stergiopoulos. 2021. Longitudinal interrelationships of mental health discrimination and stigma with housing and well-being outcomes in adults with mental illness and recent experience of homelessness. *Social Science & Medicine* 268 (Jan. 2021), 113463. <https://doi.org/10.1016/j.socscimed.2020.113463>
- [134] Evgeny Morozov. 2013. *To save everything, click here: the folly of technological solutionism*. PublicAffairs, New York.
- [135] Claire Moses. 2022. The Spread of Book Banning. *The New York Times* (July 2022). <https://www.nytimes.com/2022/07/31/briefing/book-banning-debate.html>
- [136] Arvind Narayanan. 2019. How to recognize AI snake oil. *Arthur Miller Lecture on Science and Ethics* (2019). Publisher: Massachusetts Institute of Technology.
- [137] National Institute of Health. 2015. New Paradigm Will Help Identify Leads for Drug Discovery. <https://www.nih.gov/news-events/news-releases/new-paradigm-will-help-identify-leads-drug-discovery>
- [138] National Institute of Health. 2015. NIH Clinical Center awarded prestigious certification for electronic medical record system. <https://www.nih.gov/news-events/news-releases/nih-clinical-center-awarded-prestigious-certification-electronic-medical-record-system>
- [139] National Institute of Health. 2017. Mirror image: Researchers create higher-quality pictures of biospecimens. <https://www.nih.gov/news-events/news-releases/mirror-image-researchers-create-higher-quality-pictures-biospecimens>
- [140] National Institute of Health. 2020. AI dual-stain approach improved accuracy, efficiency of cervical cancer screening in NIH study. <https://www.nih.gov/news-events/news-releases/ai-dual-stain-approach-improved-accuracy-efficiency-cervical-cancer-screening-nih-study>
- [141] Guisi Ni, Weisong Shi, and Prashant Mahajan. 2014. Appurtenant: enhancing completeness and efficiency of bidirectional patient-physician communication using automatic speech recognition. In *Proceedings of the 2014 workshop on Mobile augmented reality and robotic technology-based systems - MARS '14*. ACM Press, Bretton Woods, New Hampshire, USA, 35–40. <https://doi.org/10.1145/2609829.2609830>
- [142] Paige Nong, Minakshi Raj, and Jodyn Platt. 2022. Integrating predictive models into care: facilitating informed decision-making and communicating equity issues. *The American Journal of Managed Care* 28, 1 (Jan. 2022), 18–24. <https://doi.org/10.37765/ajmc.2022.88812>
- [143] Anne Nordberg, Mary K. Twis, Mark A. Stevens, and Schnavia Smith Hatcher. 2018. Precarity and structural racism in Black youth encounters with police. *Child and Adolescent Social Work Journal* 35, 5 (Oct. 2018), 511–518. <https://doi.org/10.1007/s10560-018-0540-x>
- [144] Office of Science and Technology Policy. 2021. Notice of Request for Information (RFI) on Public and Private Sector Uses of Biometric Technologies. *Federal Register* 86, 193 (Oct. 2021), 56300–56302. <https://www.govinfo.gov/content/pkg/FR-2021-10-08/pdf/2021-21975.pdf>
- [145] Julius Ohrnberger, Eleonora Fichera, and Matt Sutton. 2017. The relationship between physical and mental health: A mediation analysis. *Social Science & Medicine* 195 (Dec. 2017), 42–49. <https://doi.org/10.1016/j.socscimed.2017.11.008>
- [146] Fabian Okeke, Emily Tseng, Benedetta Piantella, Mikaela Brown, Harveen Kaur, Madeline R. Sterling, and Nicola Dell. 2019. Technology, home health care, and heart failure: a qualitative analysis with multiple stakeholders. In *Proceedings of the 2nd ACM SIGCAS Conference on Computing and Sustainable Societies*. ACM, Accra Ghana, 122–133. <https://doi.org/10.1145/3314344.3332487>
- [147] Andrew Ortony. 2022. Are All “Basic Emotions” Emotions? A Problem for the (Basic) Emotions Construct. *Perspectives on Psychological Science* 17, 1 (Jan. 2022), 41–61. <https://doi.org/10.1177/1745691620985415>
- [148] Thomas Owen. 2018. Twenty one years of HIV/AIDS medicines in the newspaper: patents, protest, and philanthropy. *Media, Culture & Society* 40, 1 (Jan. 2018), 75–93. <https://doi.org/10.1177/0163443717703795>
- [149] Stephen Palmer and David J Torgerson. 1999. Definitions of efficiency. *BMJ: British Medical Journal* 318, 7191 (April 1999), 1136. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1115526/>
- [150] Trishan Panch, Tom J. Pollard, Heather Mattie, Emily Lindemer, Pearse A. Keane, and Leo Anthony Celi. 2020. “Yes, but will it work for my patients?” Driving clinically relevant research with benchmark datasets. *npj Digital Medicine* 3, 1 (June 2020), 1–4. <https://doi.org/10.1038/s41746-020-0295-6> Number: 1 Publisher: Nature Publishing Group.
- [151] David Parisi. 2019. Rumble/control: Toward a critical history of touch feedback in video games. *ROMchip* 1, 2 (2019).
- [152] Sun Young Park, Pei-Yi Kuo, Andrea Barbarin, Elizabeth Kazianas, Astrid Chow, Karandeep Singh, Lauren Wilcox, and Walter S. Lasecki. 2019. Identifying Challenges and Opportunities in Human-AI Collaboration in Healthcare. In

- Conference Companion Publication of the 2019 on Computer Supported Cooperative Work and Social Computing*. ACM, Austin TX USA, 506–510. <https://doi.org/10.1145/3311957.3359433>
- [153] Shobita Parthasarathy. 2012. *Building genetic medicine: breast cancer, technology, and the comparative politics of health care*. MIT press.
- [154] Shobita Parthasarathy. 2017. *Patent Politics: Life Forms, Markets, and the Public Interest in the United States and Europe*. University of Chicago Press.
- [155] Shobita Parthasarathy. 2018. Use the patent system to regulate gene editing. *Nature* 562, 7728 (Oct. 2018), 486–488. <https://doi.org/10.1038/d41586-018-07108-3>
- [156] Shobita Parthasarathy. 2023. Can Innovation Serve the Public Good? *Boston Review* (July 2023). <https://www.bostonreview.net/articles/can-innovation-serve-the-public-good/>
- [157] Patent and Trademark Office. 2020. Manual of Patent Examining Procedure, Ninth Edition, Revision of June 2020. *Federal Register* 85, 133 (July 2020). <https://www.federalregister.gov/documents/2020/07/10/2020-14931/manual-of-patent-examining-procedure-ninth-edition-revision-of-june-2020>
- [158] Jessica K. Paulus and David M. Kent. 2020. Predictably unequal: understanding and addressing concerns that algorithmic clinical prediction may increase health disparities. *npj Digital Medicine* 3, 1 (Dec. 2020), 99. <https://doi.org/10.1038/s41746-020-0304-9>
- [159] Luigi Pellizzoni. 2017. Intensifying embroilments: Technosciences, imaginaries and publics. *Public Understanding of Science* 26, 2 (Feb. 2017), 212–219. <https://doi.org/10.1177/0963662516663563>
- [160] Sachin R Pendse, Daniel Nkemelu, Nicola J Bidwell, Sushrut Jadhav, Soumitra Pathare, Munmun De Choudhury, and Neha Kumar. 2022. From Treatment to Healing: Envisioning a Decolonial Digital Mental Health. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–23. <https://doi.org/10.1145/3491102.3501982>
- [161] Rosalind W. Picard. 2003. Affective computing: challenges. *International Journal of Human-Computer Studies* 59, 1 (July 2003), 55–64. [https://doi.org/10.1016/S1071-5819\(03\)00052-1](https://doi.org/10.1016/S1071-5819(03)00052-1)
- [162] John D. Piette, Dean Schillinger, Michael B. Potter, and Michele Heisler. 2003. Dimensions of patient-provider communication and diabetes self-care in an ethnically diverse population. *Journal of General Internal Medicine* 18, 8 (Aug. 2003), 624–633. <https://doi.org/10.1046/j.1525-1497.2003.31968.x>
- [163] Martin Prince, Vikram Patel, Shekhar Saxena, Mario Maj, Joanna Maselko, Michael R Phillips, and Atif Rahman. 2007. No health without mental health. *The Lancet* 370, 9590 (Sept. 2007), 859–877. [https://doi.org/10.1016/S0140-6736\(07\)61238-0](https://doi.org/10.1016/S0140-6736(07)61238-0)
- [164] Emily Mower Provost, Melvin Mcinnis, John Henry Gideon, Katherine Anne Matton, and Soheil Khorram. 2020. Automatic speech-based longitudinal emotion and mood recognition for mental health treatment.
- [165] Heather L. Ramey, Mary-Ellen Rayner, Sharif S. Mahdy, Heather L. Lawford, Jordi Lancot, Miranda Campbell, Eileen Valenzuela, Joshua Miller, and Valerie Hazlett. 2019. The Young Canadians Roundtable on Health: promising practices for youth and adults working in partnership. *Canadian Journal of Public Health* 110, 5 (Oct. 2019), 626–632. <https://doi.org/10.17269/s41997-019-00254-9>
- [166] Thara Ravindran, Alton Chua Yeow Kuan, and Dion Goh Hoe Lian. 2014. Antecedents and effects of social network fatigue: Antecedents and Effects of Social Network Fatigue. *Journal of the Association for Information Science and Technology* 65, 11 (Nov. 2014), 2306–2320. <https://doi.org/10.1002/asi.23122>
- [167] Nicola J. Reavley, Anthony F. Jorm, and Amy J. Morgan. 2017. Discrimination and positive treatment toward people with mental health problems in workplace and education settings: Findings from an Australian National Survey. *Stigma and Health* 2, 4 (Nov. 2017), 254–265. <https://doi.org/10.1037/sah0000059>
- [168] Thomas P Reith. 2018. Burnout in United States Healthcare Professionals: A Narrative Review. *Cureus* (Dec. 2018). <https://doi.org/10.7759/cureus.3681>
- [169] Lauren Rhue. 2018. Racial Influence on Automated Perceptions of Emotions. <https://doi.org/10.2139/ssrn.3281765>
- [170] Lauren Rhue. 2019. Understanding the Hidden Bias in Emotion-Reading AIs. <https://singularityhub.com/2019/01/11/understanding-the-hidden-bias-in-emotion-reading-a-is/> Publication Title: Singularity Hub.
- [171] Jessica L. Roberts, Elizabeth Weeks, and Elizabeth Weeks Leonard. 2018. *Healthism: health-status discrimination and the law*. Cambridge University Press.
- [172] Kjetil Rommetveit and Brian Wynne. 2017. Technoscience, imagined publics and public imaginations. *Public Understanding of Science* 26, 2 (Feb. 2017), 133–147. <https://doi.org/10.1177/0963662516663057>
- [173] Marcus D. Rushing, Andre G. Montoya-Barthelemy, Fozia A. Abrar, Eduardo M. Medina, Helen A. O. Popoola-Samuel, and Zeke J. McKinney. 2022. Law Enforcement Violence in the Black Community: A Catalyst for Clinician Engagement in Social Justice. *American Journal of Preventive Medicine* 62, 1 (Jan. 2022), 122–127. <https://doi.org/10.1016/j.amepre.2021.07.002> Publisher: Elsevier.
- [174] James A. Russell. 1994. Is there universal recognition of emotion from facial expression? A review of the cross-cultural studies. *Psychological Bulletin* 115, 1 (1994), 102–141. <https://doi.org/10.1037/0033-2909.115.1.102>

- [175] James A. Russell (Ed.). 1995. *Everyday conceptions of emotion: an introduction to the psychology, anthropology, and linguistics of emotion*. Number no. 81 in NATO ASI series. Kluwer Academic Publishers, Dordrecht, Netherlands ; Boston. Meeting Name: NATO Advanced Research Workshop on Everyday Conceptions of Emotion.
- [176] James A. Russell and Lisa Feldman Barrett. 1999. Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology* 76, 5 (1999), 805–819. <https://doi.org/10.1037/0022-3514.76.5.805>
- [177] Wulf Rössler. 2016. The stigma of mental disorders: A millennia-long history of social exclusion and prejudices. *EMBO reports* 17, 9 (Sept. 2016), 1250–1253. <https://doi.org/10.15252/embr.201643041>
- [178] Jitender Sareen, Amit Jagdeo, Brian J. Cox, Ian Clara, Margreet ten Have, Shay-Lee Belik, Ron de Graaf, and Murray B. Stein. 2007. Perceived Barriers to Mental Health Service Utilization in the United States, Ontario, and the Netherlands. *Psychiatric Services* 58, 3 (March 2007), 357–364. <https://doi.org/10.1176/ps.2007.58.3.357>
- [179] Anvita Saxena, Ashish Khanna, and Deepak Gupta. 2020. Emotion Recognition and Detection Methods: A Comprehensive Survey. *Journal of Artificial Intelligence and Systems* 2, 1 (2020), 53–79. <https://doi.org/10.33969/AIS.2020.21005>
- [180] Harald Schmidt. 2012. Wellness Incentives, Equity, and the 5 Groups Problem. *American Journal of Public Health* 102, 1 (Jan. 2012), 49–54. <https://doi.org/10.2105/AJPH.2011.300348>
- [181] Evan Selinger and Woodrow Hartzog. 2016. Facebook’s emotional contagion study and the ethical problem of co-opted identity in mediated environments where users lack control. *Research Ethics* 12, 1 (Jan. 2016), 35–43. <https://doi.org/10.1177/1747016115579531> Publisher: SAGE Publications Ltd.
- [182] Aaron Shapiro. 2020. ‘Embodiments of the invention’: Patents and urban diagrammatics in the smart city. *Convergence: The International Journal of Research into New Media Technologies* 26, 4 (Aug. 2020), 751–774. <https://doi.org/10.1177/1354856520941801>
- [183] Gilly Sharpe. 2015. Precarious identities: ‘Young’ motherhood, desistance and stigma. *Criminology & Criminal Justice* 15, 4 (Sept. 2015), 407–422. <https://doi.org/10.1177/1748895815572163>
- [184] Lindsay Sheehan, Katherine Niewegłowski, and Patrick W. Corrigan. 2017. Structures and Types of Stigma. In *The Stigma of Mental Illness - End of the Story?*, Wolfgang Gaebel, Wulf Rössler, and Norman Sartorius (Eds.). Springer International Publishing, Cham, 43–66. https://doi.org/10.1007/978-3-319-27839-1_3
- [185] Richard N. Shiffman, Bryant T. Karras, Sujai Nath, Laura Engles-Horton, and Geoffrey J. Corb. 1999. Pen-based, mobile decision support in healthcare. *ACM SIGBIO Newsletter* 19, 2 (Aug. 1999), 5–7. <https://doi.org/10.1145/954507.954509>
- [186] Ajay Kumar Shukla, Rekha Mehani, and Balakrishnan Sadasivam. 2021. Abilify MyCite (Aripiprazole): A Critical Evaluation of the Novel Dosage Form. *Journal of Clinical Psychopharmacology* 41, 1 (Jan. 2021), 93–94. <https://doi.org/10.1097/JCP.0000000000001334>
- [187] Amy E Sichel, Jason D Seacat, and Nina A Nabors. 2014. Mental health stigma update: A review of consequences. *Advances in Mental Health* 12, 3 (Dec. 2014), 202–215. <https://doi.org/10.1080/18374905.2014.11081898>
- [188] Amy E Sichel, Jason D Seacat, and Nina A Nabors. 2019. Mental health stigma: Impact on mental health treatment attitudes and physical health. *Journal of Health Psychology* 24, 5 (April 2019), 586–599. <https://doi.org/10.1177/1359105316681430>
- [189] Piyarat Silapasuphakornwong and Kazutake Uehira. 2021. Smart Mirror for Elderly Emotion Monitoring. In *2021 IEEE 3rd Global Conference on Life Sciences and Technologies (LifeTech)*. 356–359. <https://doi.org/10.1109/LifeTech52111.2021.9391829>
- [190] Kalpana Srivastava, Suprakash Chaudhury, Sana Dhamija, Jyoti Prakash, and Kaushik Chatterjee. 2020. Digital technological interventions in mental health care. *Industrial Psychiatry Journal* 29, 2 (2020), 181–184. https://doi.org/10.4103/ipj.ipj_32_21
- [191] Luke Stark and Jesse Hoey. 2019. THE ETHICS OF EMOTION IN AI SYSTEMS. *AoIR Selected Papers of Internet Research* (Oct. 2019). <https://doi.org/10.5210/spir.v2019i0.11039>
- [192] Robert Steinbrook. 2006. Imposing Personal Responsibility for Health. *New England Journal of Medicine* 355, 8 (Aug. 2006), 753–756. <https://doi.org/10.1056/NEJMp068141>
- [193] Heather Stuart. 2006. Mental illness and employment discrimination. *Current Opinion in Psychiatry* 19, 5 (Sept. 2006), 522–526. <https://doi.org/10.1097/01.yco.0000238482.27270.5d>
- [194] Samina T. Syed, Ben S. Gerber, and Lisa K. Sharp. 2013. Traveling Towards Disease: Transportation Barriers to Health Care Access. *Journal of Community Health* 38, 5 (Oct. 2013), 976–993. <https://doi.org/10.1007/s10900-013-9681-1>
- [195] Rachael Tatman. 2017. Gender and Dialect Bias in YouTube’s Automatic Captions. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*. Association for Computational Linguistics, Valencia, Spain, 53–59. <https://doi.org/10.18653/v1/W17-1606>
- [196] Leimin Tian, Sharon Oviatt, Michal Muszynski, Brent C. Chamberlain, Jennifer Healey, and Akane Sano. 2022. *Applied affective computing* (first edition ed.). Number 41 in ACM books. ACM Books, New York, NY.
- [197] Kentaro Toyama. 2015. *Geek heresy: rescuing social change from the cult of technology*. PublicAffairs, New York.

- [198] Joseph Turow, author. 2021. *The voice catchers: how marketers listen in to exploit your feelings, your privacy, and your wallet*. Yale University Press. Publication Title: The voice catchers: how marketers listen in to exploit your feelings, your privacy, and your wallet.
- [199] Theo van Leeuwen. 2008. *Discourse and Practice: New Tools for Critical Analysis*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195323306.001.0001>
- [200] Anjali Vats. 2020. *The color of creatorship: intellectual property, race, and the making of Americans*. Stanford University Press.
- [201] Bruce Vielmetti. 2020. A Black man who was arrested in his new house after neighbor called police is suing Monona, officers. <https://www.jsonline.com/story/news/2020/09/17/wisconsin-black-man-arrested-his-own-new-home-suing-monona/3483421001/>
- [202] Carolina Villegas-Galaviz. 2022. Ethics of Care as Moral Grounding for AI. In *Ethics of data and analytics* (first edition ed.). CRC Press, Boca Raton.
- [203] Rachael C. Walker, Allison Tong, Kirsten Howard, and Suetonia C. Palmer. 2019. Patient expectations and experiences of remote monitoring for chronic diseases: Systematic review and thematic synthesis of qualitative studies. *International Journal of Medical Informatics* 124 (April 2019), 78–85. <https://doi.org/10.1016/j.ijmedinf.2019.01.013>
- [204] Dakuo Wang, Liuping Wang, Zhan Zhang, Ding Wang, Haiyi Zhu, Yvonne Gao, Xiangmin Fan, and Feng Tian. 2021. “Brilliant AI Doctor” in Rural Clinics: Challenges in AI-Powered Clinical Decision Support System Deployment. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. <https://doi.org/10.1145/3411764.3445432>
- [205] Yueheng Wang, Kun Qian, Jacob Nelson, Hiromichi Yagi, Akifumi Kishi, Kenji Morita, and Yoshiharu Yamamoto. 2020. Can Affective Computing Better the Mental Status of the Electronic Games Player? A Perspective. In *2020 IEEE 2nd Global Conference on Life Sciences and Technologies (LifeTech)*, 366–367. <https://doi.org/10.1109/LifeTech48969.2020.1570620404>
- [206] Samuel D. Warren and Louis D. Brandeis. 1890. The Right to Privacy. *Harvard Law Review* 4, 5 (Dec. 1890), 193. <https://doi.org/10.2307/1321160>
- [207] C. P. West, L. N. Dyrbye, and T. D. Shanafelt. 2018. Physician burnout: contributors, consequences and solutions. *Journal of Internal Medicine* 283, 6 (June 2018), 516–529. <https://doi.org/10.1111/joim.12752>
- [208] David R. Williams and Ronald Wyatt. 2015. Racial Bias in Health Care and Health: Challenges and Opportunities. *JAMA* 314, 6 (Aug. 2015), 555–556. <https://doi.org/10.1001/jama.2015.9260>
- [209] Andrew Wong, Erkin Otles, John P. Donnelly, Andrew Krumm, Jeffrey McCullough, Olivia DeTroyer-Cooley, Justin Pestrue, Marie Phillips, Judy Konye, Carleen Penozza, Muhammad Ghous, and Karandeep Singh. 2021. External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients. *JAMA Internal Medicine* 181, 8 (Aug. 2021), 1065–1070. <https://doi.org/10.1001/jamainternmed.2021.2626>
- [210] Laura Wood. 2020. Global Emotion Detection & Recognition Market Size is Projected to Grow from USD 21.6 Billion in 2019 to USD 56.0 Billion by 2024, at a CAGR of 21.0% - ResearchAndMarkets.com. <https://www.businesswire.com/news/home/20200213005614/en/Global-Emotion-Detection-Recognition-Market-Size-is-Projected-to-Grow-from-USD-21.6-Billion-in-2019-to-USD-56.0-Billion-by-2024-at-a-CAGR-of-21.0---ResearchAndMarkets.com>
- [211] Steve Woolgar. 1990. Configuring the User: The Case of Usability Trials. *The Sociological Review* 38, 1_suppl (May 1990), 58–99. <https://doi.org/10.1111/j.1467-954X.1990.tb03349.x> Publisher: SAGE Publications Ltd.
- [212] David Wright, Rowena Rodrigues, Charles Raab, Richard Jones, Iván Székely, Kirstie Ball, Rocco Bellanova, and Stine Bergersen. 2015. Questioning surveillance. *Computer Law & Security Review* 31, 2 (April 2015), 280–292. <https://doi.org/10.1016/j.clsr.2015.01.006>
- [213] S. Wyatt. 2008. Technological determinism is dead; Long live technological determinism. In *Handbook of Science and Technology Studies*, E. Hackett, O. Amsterdamska, M. Lynch, and J. Wajcman (Eds.). MIT Press, Cambridge, 165–180.
- [214] Jun Yang, Rui Wang, Xin Guan, Mohammad Mehedi Hassan, Ahmad Almogren, and Ahmed Alsanad. 2020. AI-enabled emotion-aware robot: The fusion of smart clothing, edge clouds and robotics. *Future Generation Computer Systems* 102 (Jan. 2020), 701–709. <https://doi.org/10.1016/j.future.2019.09.029>
- [215] Shoshana Zuboff. 2020. You Are Now Remotely Controlled. *New York Times* (2020), 8.
- [216] Ethan Zuckerman and Arvind Narayanan. 2022. See Through AI Hype with Arvind Narayanan. <https://publicinfrastructure.org/podcast/see-through-ai-hype-with-arvind-narayanan/>

Received January 2023; revised July 2023; accepted November 2023

Table 3. Patent Applications in Our Dataset

Patent Number	Patent Publication Number	Filing Year	Applicant and/or Affiliated Organization or Company
P1	US 2018_0032126 A1	2016	National University Of Defense Technology, Liu Yadong
P2	US 2018_0039745 A1	2017	Chevalier Timothy W, Atlas5D Inc
P3	US 10478111 B2	2015	Sri Intl Inc, Knoth Bruce
P4	US 2020_0194125 A1	2019	Fichner Rathus Lois
P5	US 2019_0252059 A1	2019	Gleason Brad
P6	US20200364445	2019	Suzuki Masato, Panasonic
P7	US 2020_0245918 A1	2019	Mindstrong Health, Dagum Paul
P8	US 2020_0075039 A1	2019	Sentio Solutions Inc, Eleftheriou Georgios
P9	US 9173567 B2	2011	Jain Jawahar, Fujitsu Ltd
P10	US 9028405 B2	2009	Koninkl Philips Nv
P11	US 2020_0139077 A1	2020	Biradar Poornima, Sony Corp
P12	0068994 A1	2016	Slomkowski, Robin S
P13	0285700 A1	2019	Narayanan Shrikanth Sambasivan, University Of Southern California
P14	US 2020_0275873 A1	2019	Boe Tech Group Co Ltd
P15	10410655 B2	2017	Fujitsu Ltd
P16	0050837 A1	2019	Nuralogix
P17	US 2020_0214626 A1	2019	Marco Tech Llc
P18	US20200286603	2018	Ajilore Olusola, Univ Illinois Urbana Champaign
P19	US 2020_0064986 A1	2019	Huang John, Caressa
P20	10417484 B2	2017	Wipro Ltd
P21	US 10736555 B2	2018	Carr-Jordan Erin Marie, Arizona State University
P22	US 9141604 B2	2013	Thirumalainambi Rajkumar, Riaex Inc
P23	US 2019_0110728 A1	2017	Sbodio Marco Luca, Intl Business Machines Corp
P24	US 2019_0206424 A1	2019	Joshua Feast, Cogito Corp
P25	US 2020_0250276 A1	2019	Intl Business Machines Corp
P26	US 2017_0311864 A1	2017	Manabe Seiichi, Omron Corp
P27	US 2017_0344713 A1	2015	Riistama Jarno Mikael, Koninkl Philips Nv
P28	US 2019_0052724 A1	2017	Dancel Ivan Tumbocon
P29	US 2019_0254581 A1	2016	Rutgers State Univ New Jersey
P30	US 2020_0225963 A1	2019	Noh Kyoung Ju, Electronics and Telecom Rsch Inst
P31	US 2018_0375809 A1	2018	Lo Kit Yi
P32	US 2015_0324634 A1	2013	Brosens-Kessels Angeliq Carin Johanna Maria, Koninkl Philips Nv
P33	US 2015_0079560 A1	2014	Cowan Jonathan Daniel
P34	US 2018_0285528 A1	2017	Healey Jennifer Anne, Intel
P35	US 10839201 B2	2019	Johnson Jason, Akili Interactive Labs Inc
P36	US 7236963 B1	2003	Lamuth John E, Chocolate Goddess Inc
P37	US 10289201 B2	2017	Cruz-Hernandez Juan Manuel, Immersion Corp
P38	US 2015_0026111 A1	2013	Greatcall Inc
P39	US 10580435 B2	2017	Ashoori Maryam, Ibm
P40	US 2017_0042713 A1	2015	Nurmikko Arto V, Brown Univ
P41	US2020075040A1	2018	Emily Mower, University Of Michigan
P42	US 9934363 B1	2016	Intl Business Machines Corp
P43	US 2012_0277594 A1	2012	Pryor Timothy R
P44	US20170262606	2017	Abdullah Saeed, Cornell Univ
P45	US 10325066 B2	2017	Yeh, Ta-Chuan
P46	US 2020_0234827 A1	2019	Internicola Charles J, Mira Therapeutics Inc
P47	US 2020_0020447 A1	2018	Intl Business Machines Corp
P48	9418390 B2	2012	Intel Corp
P49	0110727 A1	2018	Hitachi
P50	US 10251591 B2	2017	Soza; Ana Maria (Santiago, Cl)
P51	US 2009_0002178 A1	2007	Microsoft
P52	US 2019_0272466 A1	2019	Univ Southern California
P53	US 9514281 B2	2012	Winterlight Labs Inc
P54	US 2019_0272725 A1	2019	New Sun Technologies, Inc., Sunnyvale Ca
P55	US 10589087 B2	2007	Individual Investors
P56	US 10803992 B2	2017	Individual Investors
P57	US 2019_0074028 A1	2017	Howard; Newton
P58	US 2020_0188629 A1	2018	Trevor Ai