

illuminating the Unseen: Investigating the Context-induced Harms in Behavioral Sensing

HAN ZHANG, University of Washington, USA

VEDANT DAS SWAIN, Northeastern University, USA

LEIJIE WANG, University of Washington, USA

NAN GAO, Tsinghua University, China

YILUN SHENG, University of Washington, USA

XUHAI XU, Massachusetts Institute of Technology, USA

FLORA D. SALIM, University of New South Wales, Australia

KOUSTUV SAHA, University of Illinois Urbana-Champaign, USA

ANIND K. DEY and JENNIFER MANKOFF, University of Washington, USA

Behavioral sensing technologies are rapidly evolving across a range of well-being applications. Despite its potential, concerns about the responsible use of such technology are escalating. In response, recent research within the sensing technology has started to address these issues. While promising, they primarily focus on broad demographic categories and overlook more nuanced, context-specific identities. These approaches lack grounding within domain-specific harms that arise from deploying sensing technology in diverse social, environmental, and technological settings. Additionally, existing frameworks for evaluating harms are designed for a generic ML life cycle, and fail to adapt to the dynamic and longitudinal considerations for behavioral sensing technology. To address these gaps, we introduce a framework specifically designed for evaluating behavioral sensing technologies. This framework emphasizes a comprehensive understanding of context, particularly the situated identities of users and the deployment settings of the sensing technology. It also highlights the necessity for iterative harm mitigation and continuous maintenance to adapt to the evolving nature of technology and its use. We demonstrate the feasibility and generalizability of our framework through post-hoc evaluations on two real-world behavioral sensing studies conducted in different international contexts, involving varied population demographics and machine learning tasks. Our evaluations provide empirical evidence of both situated identity-based harm and more domain-specific harms, and discuss the trade-offs introduced by implementing bias mitigation techniques.

CCS Concepts: • **Human-centered computing** → **Empirical studies in ubiquitous and mobile computing**.

Additional Key Words and Phrases: sensing technologies, AI harms, well-being, context-sensitivity, framework, responsibility

Authors' addresses: Han Zhang, micohan@cs.washington.edu, University of Washington, USA; Vedant Das Swain, v.dasswain@northeastern.edu, Northeastern University, USA; Leijie Wang, leijiew@cs.washington.edu, University of Washington, USA; Nan Gao, nangao@tsinghua.edu.cn, Tsinghua University, China; Yilun Sheng, ylsheng@cs.washington.edu, University of Washington, USA; Xuhai Xu, xoxu@mit.edu, Massachusetts Institute of Technology, USA; Flora D. Salim, flora.salim@unsw.edu.au, University of New South Wales, Australia; Koustuv Saha, ksaha2@illinois.edu, University of Illinois Urbana-Champaign, USA; Anind K. Dey, anind@uw.edu; Jennifer Mankoff, jmankoff@cs.washington.edu, University of Washington, USA.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Association for Computing Machinery.

XXXX-XXXX/2024/5-ART \$15.00

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

ACM Reference Format:

Han Zhang, Vedant Das Swain, Leijie Wang, Nan Gao, Yilun Sheng, Xuhai Xu, Flora D. Salim, Koustuv Saha, Anind K. Dey, and Jennifer Mankoff. 2024. Illuminating the Unseen: Investigating the Context-induced Harms in Behavioral Sensing. 1, 1 (May 2024), 34 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

The rapid evolution of sensing technologies has unlocked new possibilities for tracking and understanding human activities. Behavioral sensing technology, which involves using sensing to capture, model, and predict human behaviors, offers a broad spectrum of applications. These include, but are not limited to, well-being monitoring (e.g., mental health prediction [4, 34, 96, 109]), human activity recognition (e.g., identifying activities like “running”, “sitting”, and “walking” [76, 99]), and personalized recommendations (e.g., personalized music and taxi charging recommendation systems [80, 122]). This technology, in contrast to the traditional manual approach of using questionnaire-collected data for the same tasks, facilitates continuous, automated, and unobtrusive gathering of *context* [32, 36]. Here, context refers to capturing all information related to the interactions among users, applications, and their environment [39, 40].

In recent years, concerns about the responsible use of behavioral sensing technology have been growing [27, 33, 72, 73]. These concerns arise from the prevalent top-down design approach, which typically involves technology builders—a collective of researchers, designers, developers, and engineers—developing tools based on their assumptions of users’ goals, needs, or preferences [73, 94]. This approach often leads to a lack of sufficient understanding of the users’ diverse backgrounds and the situation in which the technology is used and deployed, a concept described as *context sensitivity* [38, 39]. Consequently, important contextual factors that may not initially appear relevant to its primary purpose are often disregarded during the early design phases. Such neglect can lead to technologies that fail to adequately meet the diverse and real-world needs of users effectively and may even introduce potential harms [135].

Understanding the differences in model performance and the potential harms of sensing technology is critical, as also highlighted in prior research [5, 10, 91, 135, 136, 138]. Researchers have started exploring how models may perform differently based on various personal sensitive attributes, a contextual factor we refer to as *identity-based harm*. These attributes include demographics, socioeconomic status, country of residence, and health conditions [5, 91, 136]. Accordingly, we recognize three gaps in the existing literature.

Gap 1: There is a limited understanding of more *situated identities* relevant to specific contexts, such as first-generation college students, immigration status, or disability status—groups that also confront profound societal inequalities [14, 127, 137]. The under-representation of these identities fails to capture a more comprehensive spectrum of user experiences and needs within sensing technology design and application.

Gap 2: A distinct aspect of sensing technology lies in what we term as *situation-based harm*, which remains understudied. This type of harm emerges specifically from the inherent versatility of sensing technology, which finds applications across a diverse array of scenarios—from varied technological infrastructures and environmental conditions to different device types [1, 134]. This diversity in application settings can lead to unique contextual harms when the technology is designed or evaluated without thorough consideration of these situational differences. Such harms are particularly concerning because they often do not intuitively link directly to users and can go unanticipated, thereby increasing the risk of irresponsible technology deployment and exacerbating the potential for ultimately negative impacts on users.

Gap 3: Understanding the intricacies of bias mitigation, particularly the potential trade-offs introduced by adopting such techniques, remains underexplored. Moreover, the current evaluation frameworks do not adequately cater to the unique demands of sensing technology pipelines. These pipelines require iterative bias mitigation and continuous maintenance to adapt to the dynamic nature of sensing data, challenges that are not fully addressed by existing frameworks (e.g., [67, 83, 117]). Moreover, there is a lack of a systematic approach to effectively design,

build, and evaluate sensing technologies, leaving gaps in how these sensing technologies are developed and how their impacts are assessed across different user groups and usage scenarios.

To address the identified gaps in existing research, we propose a theory-driven, human-centered framework specifically designed for evaluating sensing technology. This framework builds upon evaluation frameworks in machine learning [67, 83, 117] by incorporating a thorough assessment of both broad and domain-specific context-induced harms to users during the evaluation process. Moreover, our proposed framework emphasizes the need for iterative harm mitigation and continuous maintenance, critical components that are essential to accommodate the dynamic nature of sensing technology.

We conduct post-hoc evaluations on two real-world behavioral sensing studies to demonstrate the feasibility of our proposed framework and to derive new insights from the evaluation results. To ensure the generalizability of our framework, we select two studies conducted in different countries, with different populations, and involving distinct ML tasks. Specifically, one study focused on college student mental health detection [132], a classification task; and another on teenage student learning engagement prediction [59], a regression task.

The first evaluation provides empirical evidence of situated identity-based harm, extending our understanding beyond demographic categories such as race and age. The second evaluation, showcases robustness to identity-based harm, but instead, highlights evidence for situation-based harm by showing significant performance differences in algorithms across environmental situations, such as varying temperature conditions. Furthermore, our experiment on a bias mitigation technique reveals a critical trade-off: efforts to mitigate harm for one attribute inadvertently introduced harm to other attributes. In summary, our contributions are as follows.

- In Section 3, we propose a framework that integrates consideration for both identity-based and situation-based harms, and emphasizes the need for iterative bias mitigation and continuous maintenance to effectively manage the dynamic challenges posed by sensing technologies.
- In Sections 4.1 and 4.2, We conduct evaluations on two real-world behavioral sensing studies. We provide empirical evidence on the nuanced forms of identity-based harms and highlight situational dependencies that affect algorithm performance. Furthermore, we make our analysis codebase openly accessible for reproducibility¹.
- In Section 5, we offer key insights from our evaluation studies, focusing on the potential harms linked to more situated identity-based factors and the newly observed situation-based harms in behavioral sensing. Additionally, we discuss the trade-offs introduced by implementing bias mitigation techniques.

We further offer reflections within and beyond our proposed framework. Our work aims to contribute both conceptually and practically to the field, focusing on more responsible behavioral sensing technologies.

2 BACKGROUND AND RELATED WORK

As behavioral sensing technologies increasingly become a tool for tracking and reasoning about human activities, they present a blend of promising opportunities and potential risks [138]. In this section, we first review behavioral sensing technology, exploring its evolution and current landscape (Section 2.1). We then review a promising application domain of these technologies, specifically in well-being prediction (Section 2.2). Following this, we discuss the potential harms associated with these technologies, including both identity-based and situation-based harms, particularly arising from a lack of context-sensitivity (Section 2.3). We conclude this section by reviewing the human-centered design approach, aimed at addressing these harms (Section 2.4). This background forms the basis for proposing our evaluation framework.

¹We will release our codebase at publication.

2.1 Evolution of Behavioral Sensing Technology

Since the late 1990s, researchers in sensing technology have increasingly recognized the importance of enabling computing devices to enhance application performance by incorporating knowledge of the context in which they are used [2, 39, 111]. “Context” in this sense refers to all information related to the interactions among users, applications, and their environment [40]. Alongside this realization, there was growing advocacy for the creation of sensing systems designed to offer information or services that are relevant to the specific tasks of users, a concept known as *context-awareness* [2, 39]. Building upon this foundational concept, research efforts have since concentrated on creating various toolkits and frameworks to capture, infer, and generate context through diverse sensors [40, 50, 82, 110]. These initiatives have evolved to focus on employing these toolkits for passive data collection, aiming to infer human behavior, such as their phone usage, location, sleep, and steps [42]. More recently, the integration of Machine Learning (ML) and Artificial Intelligence (AI) algorithms into these technologies has further transformed this field. Researchers have begun integrating these advanced techniques into sensing technologies, not only for modeling human behavior but also for making predictions [4, 11, 37].

2.2 Well-being Prediction in Behavioral Sensing

Well-being prediction is one of the promising and extensively studied application domains for behavioral sensing technologies. This domain includes various aspects such as predicting mental health [4, 25, 113] and forecasting performance, engagement, as well as productivity [59, 124]. Specifically, in the area of mental health prediction, considerable research has been dedicated to depression prediction. For instance, studies have utilized passively sensed data such as physical activities, phone usage, sleep patterns, and step counts of participants to predict depressive symptoms [25, 125, 129]. Additionally, there have been significant research efforts aimed at understanding mood-related health concerns among students and workers, employing various types of passively sensed data [79, 91, 96, 108].

In parallel, evaluating performance, engagement, and productivity as a facet of mental health-related well-being prediction has also attracted considerable attention. Many studies have focused on student populations. For example, Wang et al. conducted a study using passively collected data from smartphones and wearables of college students to predict their cumulative GPA. In another work, Ahuja et al. developed a classroom sensing system to capture student facial expressions and body gestures from audio and video data. Their approach allowed for the analysis of students’ engagement levels based on these sensory inputs. In a different study, Gao et al. [57, 59] employed indoor environmental data, such as temperature, humidity, CO2 levels, and sounds, alongside physical activity data, to predict three dimensions of student engagement levels. Beyond the academic setting, behavioral sensing technology is increasingly being used in the corporate sector to assess workplace productivity and employee well-being [102]. For example, Mirjafar [93] trained machine learning algorithms on sensing data to differentiate performance levels in workplaces, offering insights for workspace optimization and stress management [37].

2.3 Potential Harms in Behavioral Sensing

While these behavioral sensing applications offer significant opportunities to improve and support human well-being, they often employ a top-down design approach, which is predominantly driven by technology builders’ assumptions of users’ goals, needs, or preferences [19, 94]. However, developing technologies based solely on these assumptions, alongside what is easily possible to sense, without a thorough understanding of the users’ diverse backgrounds and the situation in which the technology is used and deployed, can lead these technologies to a lack of *context sensitivity*. In this paper, we expand context sensitivity, traditionally linked to context-awareness [38, 39], and redefine it to highlight the responsible aspects of these technologies regarding diverse user groups and situations.

*“A technology is **context-sensitive** when it accounts for diverse user backgrounds, needs, and situations of use to provide value to users.”*

As behavioral sensing technologies advance toward practical, real-world applications, it becomes increasingly important to ensure that these technologies are context-sensitive. This consideration is crucial to mitigate potential harms to users. Below, we review existing work on one type of harms, identity-based harm, in sensing technology and introduce a domain-specific harm, situation-based harm.

2.3.1 Identity-based Harms. As highlighted by Yfantidou et al. in their review of sensing technology research from 2018 to 2022, a mere 5% of studies investigated algorithmic harms to users with sensitive attributes. Alarmingly, 90% of these studies limited their focus to only gender and age, and primarily relied on accuracy or error metrics for evaluation. In recent years, Ubicomp researchers have begun to understand the potential harms of sensing technology [5, 10, 91, 135, 136, 138]. For example, evaluations conducted on three existing datasets, based on an established framework to identify and understand sources of identity-based harm throughout the ML life cycle [116], have provided empirical evidence of identity-based harms across data generation, model building, and implementation processes [136]. Adler et al. showed that sensed-behaviors that predict depression risk were inconsistent across demographic and socioeconomic subgroups. In addition, researchers investigated into how models perform across different countries or cultures—factors that are also tied to identity—have shed light on variations in applications such as activity recognition [10] and mood inference [91]. While these developments are promising, these works primarily focused on broad demographic categories and overlooks more nuanced, context-specific identities. This oversight is concerning, because people with these attributes have a history of facing societal inequalities, as extensively documented in psychology and social work research [26, 55, 71, 112]. For example, Hangartner et al. found in their study of online recruitment platforms that individuals from immigrant and minority ethnic groups received 4-19% fewer recruiter contacts. Similarly, Blaser et al. noted the significant absence of disability reporting in tech companies and their media coverage. Moreover, Erete et al. employed autoethnography [47] and testimonial authority [29] to share their experiences as Black women academics during a pandemic disproportionately impacting their communities and in the context of civil unrest due to racial injustice.

2.3.2 Situation-based Harm. Another potential harm that could emerge due to a lack of context sensitivity in behavioral sensing technology is what we identify as situation-based harm. This type of harm could occur when sensing technology is implemented in diverse situations or settings. As this aspect of potential harm is relatively under-explored, we provide an example to help readers conceptually understand it. Specifically, if a behavioral sensing algorithm is predominantly based on data from iOS-based smartphones, it may not be effective on Android-based smartphones due to representational bias [92]. This may potentially lead to a disproportionate impact on individuals of lower socioeconomic status or those in developing countries who commonly use more affordable Android devices. Reportedly, iOS-based smartphones tend to be more than twice as expensive as their Android-based counterparts [68, 97].

2.4 Human-Centered Approach for Investigating Harms

The broader ML community has developed various frameworks to evaluate bias [67, 83, 103, 117]. However, these frameworks often fall short in prioritizing users' needs. To identify and mitigate potential harms to humans, there is an increasing call within the HCI, CSCW, and Ubicomp communities for a focus on *human-centered AI* (HCAI) design [3, 9, 30, 46, 78, 118, 123]. While definitions of HCAI vary [23], the central theme revolves around designing AI technologies that are deeply attuned to the needs, values, and agency of human users, partners, and operators [23]. This approach focuses on ensuring that AI systems are technically efficient and align with human-centric values and ethical standards. Research efforts have been made to provide guidelines for designers to

create more ethical AI systems. For example, Amershi et al. distilled over 150 AI-related design recommendations into 18 broadly applicable guidelines, including ensuring AI systems are cautious about social biases and enabling users to provide feedback during regular interaction with the AI systems. Adopting this approach, researchers have actively engaged with stakeholders to gain an in-depth understanding of their experiences and perceptions regarding AI in general and sensing technology in particular. This engagement has focused on various aspects, including stakeholders' trust in AI systems [12, 75, 86, 119], their privacy concerns related to the use of sensing technologies [6, 104, 106], and the specific impacts of these technologies in different settings [27, 33, 35, 36, 73].

Value-sensitive design is another widely adopted approach that can address potential harms to humans [52]. This approach is grounded in the principle of integrating human values into the design process thoroughly and systematically. It utilizes a tripartite methodology that is both integrative and iterative, involving conceptual, empirical, and technical investigations [53, 54]. To explicate this approach, Friedman et al. presented three case studies in their work. Leveraging this concept, Zhu et al. applied value-sensitive design to algorithm development. In their approach, they actively engaged with stakeholders during the early stages of algorithm creation, incorporating their values, knowledge, and insights to ensure that the technology aligns with user needs and ethical standards.

3 AN EVALUATION FRAMEWORK FOR INVESTIGATING CONTEXT-INDUCED HARMS IN BEHAVIORAL SENSING

3.1 Our Evaluation Perspective: Human-Centered and Context-Sensitive AI

Our evaluation approach is grounded in the principles of *human-centered AI* (HCAI) [70, 114] and *value-sensitive* design [54, 139], which advocates for a shift away from traditional practices where behavioral sensing technology builders rely on predetermined assumptions about user needs and preferences [94]. Rather, it emphasizes a thorough consideration of potential harms to users throughout the technology design process. This aims to ensure that the behavior of the developed technologies does not reinforce negative stereotypes or biases against users, in line with the guidelines proposed by Amershi et al. [9]. Moreover, we emphasize the importance of iteratively mitigating potential harms to users and maintaining continuous oversight to address the dynamic nature of sensing technology [138].

3.2 Overview of the Framework

In this section, we first present our developed framework for evaluating and mitigating the context-induced harms in behavioral sensing. We then compare our framework with the prevailing evaluation approach in behavioral sensing technology. Our framework, shown in Figure 1, has six steps.

- **Step 1: Comprehensively understand the context.** In this initial phase of developing behavioral sensing technologies, a comprehensive understanding of the *context* is necessary. This involves the awareness of users' diverse backgrounds and engaging with them to understand their specific needs, as well as considering the variety of situated settings (such as technology infrastructure and environmental conditions).
- **Step 2: Establish criteria for evaluating harms, and make sure the bias is not attributed to random choice.** After obtaining a detailed understanding of the context, technology builders should select metrics that can effectively discern algorithmic variances in different contexts. To ensure that these differences are not attributed to random chance, technology builders should employ a rigorous quantitative method, predominantly a statistical analysis [135], for their assessment. The choice of these methods demands careful deliberation to address issues such as Type I errors (false positives) and variations from different groups.
- **Step 3: Evaluate whether the collected data includes inclusive and comprehensive contextual information.** This step involves evaluating whether the collected data adequately reflects the comprehensive

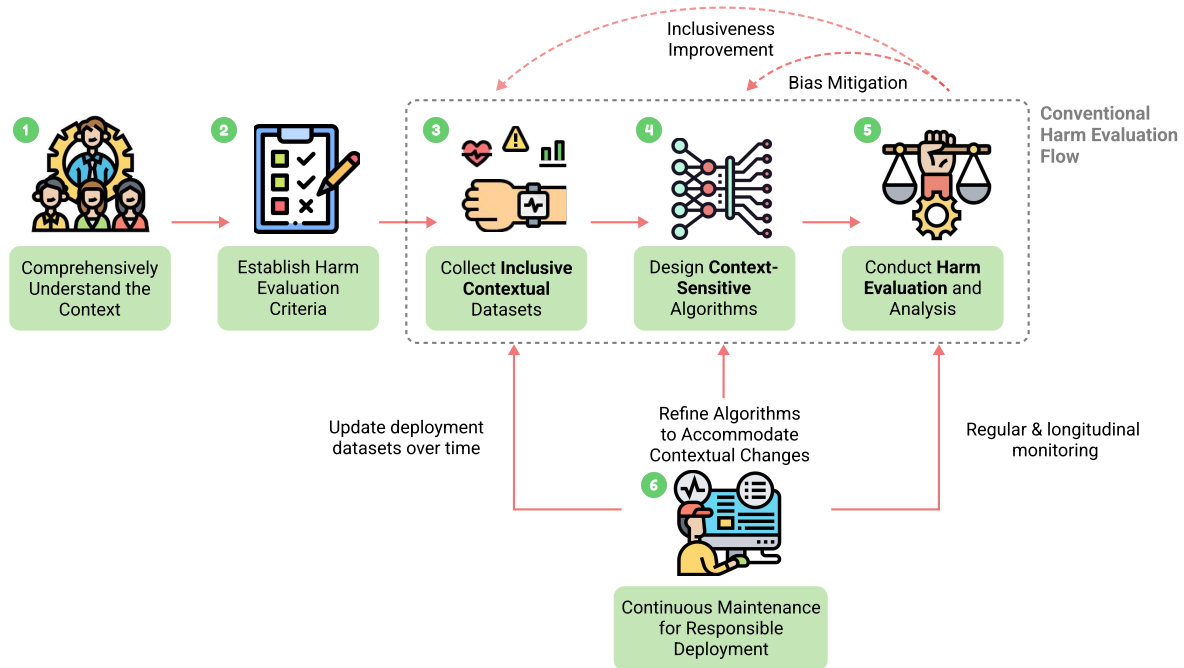


Fig. 1. Overview of the Framework for Evaluating and Mitigating Context-induced Harms in Behavioral Sensing. Steps 3, 4, and 5 cover the conventional evaluation flow for identifying harms.

context identified in Step 1. It is essential to ensure broad representation across a range of demographics and to be acutely aware of situational factors during data collection.

- **Step 4: Evaluate whether the technology uses context-sensitive algorithms.** This step focuses on assessing whether technology builders have engineered algorithms that are sensitive to potential harms and biases that may emerge from varying contexts. It often requires a careful approach to data selection for training and testing these algorithms. Moreover, technology builders should commit to a continuous cycle of refinement and improvement of these algorithms, especially when harm or bias is detected.
- **Step 5: Conduct harm evaluation on the behavioral sensing system, combining user feedback and techniques to mitigate harms.** This step involves assessing whether technology builders comprehensively evaluate the behavioral sensing system to identify and address potential harms or biases. A crucial element of this evaluation is the integration of user feedback, which offers valuable insights into the technology's performance in real-world scenarios and its effects on various user groups. If any biases are detected, technology builders need to revisit Step 4. There, they must refine the algorithms by leveraging the insights gained from user feedback, ensuring the technology not only performs optimally but also responsibly.
- **Step 6: Continuous maintenance of data and algorithms for responsible deployment.** This step of evaluation is critical in the lifecycle of behavioral sensing technology. Once the technology is deployed, technology builders should continuously monitor and update both the data and algorithms and evaluate the model performance to ensure that the technology remains up-to-date and adapts to various contexts.

Table 1. **Overview of Evaluation Criteria and Methods for Each Evaluation Study.** This table enumerates the various elements assessed in each evaluation study and lists the specific methods employed for evaluating the algorithms.

Framework Outline	Evaluation 1: Depression Detection	Evaluation 2: Engagement Prediction
Step 1: Comprehensively understand the context.	Identity-based harms: gender, sexual orientation, race, immigration status, first-generation college student status, and disability status. Situation-based bias: device types and data collection time. Engagement with users: understand their concerns and insights on depression detection sensing technologies.	Identity-based harms: gender, disability status, homeless youth, and religious minorities. Situation-based bias: Temperature condition, location, and class group. Engagement with users: understand their concerns and insights on student engagement prediction sensing technologies.
Step 2: Establish criteria for evaluating harms, and make sure the bias is not attributed to random choice.	Fairness metrics: disparities in accuracy, false negative rate, and false positive rate. Significance test: Mann-Whitney U test with Benjamini-Hochberg correction.	Fairness metrics: disparities in mean squared error. Significance test: Linear mixed model.
Step 3: Evaluate whether the collected data includes inclusive and comprehensive contextual information.	Consider collecting data that could introduce identity-based harms and/or Situation-based bias due to the contextual factors identified in Step 1.	
Step 4: Evaluate whether the technology uses context-sensitive algorithms.	Ensure that the algorithms are aware of harms to users and can adapt to contextual changes.	
Step 5: Conduct harm evaluation on the behavioral sensing system, combining user feedback and techniques to mitigate harms.	Assess algorithms for potential harms or biases and verify their mitigation with user feedback.	
Step 6: Continuous maintenance of data and algorithms for responsible deployment.	Implement strategies and actions to regularly update and maintain the data and algorithms.	

4 EVALUATING EXISTING BEHAVIORAL SENSING TECHNOLOGIES – TWO EVALUATION STUDIES

The objective of this section is two-fold. First, we aim to validate the feasibility of our proposed framework. Second, we strive to provide empirical evidence on the nuanced forms of identity-based harms and situation-based harms. To achieve these goals, we conduct post-hoc evaluations on two real-world behavioral well-being sensing technology studies, each within a different domain and involving a distinct ML task. The first study focuses on classifying students with depressive symptoms (Section 4.1). The second study aims to regress students' engagement levels (Section 4.2).

For each evaluation, we start with a background section, delineating the real-world problem, the datasets used, as well as the ML task and algorithms chosen for our evaluation. We then describe our evaluation process (summarized in Table 1). This is followed by the evaluation results. Our evaluations mainly focus on two aspects.

- Evaluate the extent to which the steps proposed in our framework have been considered in previous efforts in the design and implementation of well-being sensing technologies.
- Identify the potential harms and biases these technologies might introduce to users, by performing a quantitative evaluation of those algorithms.

We provide detailed descriptions of the evaluation process for steps 1 and 2 within the respective sections of each evaluation study, as these steps are customized to each specific case. Additionally, to gain a deeper understanding of algorithmic harms, we conduct an experiment focusing on bias mitigation in each evaluation study.

4.1 Evaluation Study 1: Depression Detection

Research has been conducted using longitudinal passive sensing data from smartphones and wearable devices to predict and detect depression (e.g., [120, 125, 130]). However, these studies often face challenges related to the limited access to datasets and algorithms, hindering reproducibility and transparency in the field. To address these issues, Xu *et al.* introduced GLOBEM [131], an open-sourced benchmark platform that includes implementations of nine depression detection algorithms and ten domain generalization algorithms. All depression detection algorithms focus on a common binary classification task: distinguishing whether users had at least mild depressive symptoms. They also released a four-year longitudinal passive sensing behavioral dataset from college students [132], aimed to highlight challenges in generalizing and reproducing longitudinal behavior models. In our evaluation study, we examine these depression detection algorithms through a lens focused on potential harms, employing the perspective provided by our proposed framework. Our goal is to assess whether the designs of these algorithms, or their implementations, have considered the steps outlined in our framework. If any of these steps were overlooked, what potential harms can we identify?

Datasets. We chose the four datasets from the GLOBEM study [132] for the evaluation of the depression detection algorithms. To facilitate analysis and comparison, these datasets were labeled chronologically according to the time of their collection (D1 to D4). These four datasets used in our evaluation consist of approximately 700 person-terms of data from around 500 unique participants who were enrolled in the same institution over 10 weeks during the Spring term between 2018 and 2021. These datasets include a wide range of passively sensed behavioral data, including sleep patterns, phone usage statistics, physical activity levels, and phone call records. The datasets also include a wide range of demographics, such as gender, race, first-generation status, immigration status, sexual orientation, and disability status. This data was continuously collected 24 hours per day from smartphones and Fitbits. Additionally, they also include self-reported depression data. We opted to use Beck Depression Inventory-II (BDI-II) scores [13], which were collected once per person at the end of each term in each dataset, as the ground truth.

Depression Detection Algorithms. We chose eight depression detection algorithms implemented by Xu *et al.* [131]. These algorithms consist of a combination of support vector machine [22, 49, 120], logistic regression [107, 125], random forest [120], Adaboost [129], multi-task learning [84], and collaborative-filtering-based model [130]. We excluded one algorithm in the implementation work developed by Chikersal *et al.* [25] from our evaluation study due to a significant disparity between our reproduced results and the reported results in the implementation work [131] (shown in Table 2).

4.1.1 Evaluation Methods and Results. In this subsection, we elaborate on the decision-making processes involved in each step of our framework, followed by presenting our evaluation results.

Step 1: Comprehensively understand the context. In the development of behavioral sensing technologies for depression detection, having a nuanced understanding of user diversity is crucial, as certain sub-populations exhibit higher depression rates. Studies indicate an increased prevalence of depression in specific demographics, such as women [8], first-generation college students [69], immigrants [56], non-heterosexual individuals [140], racial minorities [18], and disability status [137]. These findings underscore the importance for depression detection technology builders to be aware of this context – users’ sensitive attributes – to avoid societal biases and ensure equitable outcomes. Additionally, temporal factors and the type of devices used during data collection are crucial elements to consider. Research has indicated that depressive symptoms can fluctuate based on the time of day (e.g., morning vs. evening) [24, 100] and that user behaviors might be affected by the specific settings of their devices [28]. Recognizing these contexts – timing and device types during data collection – is crucial. Such considerations enable technology builders to accurately model and predict depressive symptoms under different conditions, thereby ensuring the technology’s adaptability and fairness in diverse data collection scenarios.

Table 2. **Reproduction Results.** Balanced accuracy of the nine depression prediction algorithms on four datasets: DS1 and DS2, which were used in prior work [131], and DS3 and DS4, which are newly reported in this study. The comparison of algorithm performance on DS1 and DS2 ensures the reliability of our fairness evaluation. The Δ column represents the difference between our reproduced results and the previously reported results.

Algorithms	DS1 (2018)			DS2 (2019)			DS3 (2020)	DS4 (2021)
	Prior results	Our results	Diff (Δ)	Prior results	Our results	Diff (Δ)	Our results	Our results
Wahle <i>et al.</i> [120]	0.526	0.538	0.012	0.527	0.518	-0.009	0.514	0.514
Saeb <i>et al.</i> [107]	0.539	0.539	0.000	0.508	0.513	0.005	0.588	0.500
Farhan <i>et al.</i> [49]	0.552	0.552	0.000	0.609	0.609	0.000	0.563	0.609
Canzian <i>et al.</i> [22]	0.559	0.538	-0.021	0.516	0.516	0.000	0.541	0.502
Wang <i>et al.</i> [125]	0.566	0.565	-0.001	0.500	0.500	0.000	0.577	0.516
Lu <i>et al.</i> [84]	0.574	0.574	0.000	0.558	0.558	0.000	0.611	0.553
Xu <i>et al.</i> - Interpretable [129]	0.722	0.688	-0.034	0.623	0.667	0.044	0.833	0.733
Xu <i>et al.</i> - Personalized [130]	0.723	0.753	0.030	0.699	0.690	-0.009	0.791	0.686
Chikersal <i>et al.</i> (removed) [25]	0.728	0.618	-0.110	0.776	0.670	-0.106	0.581	0.641

Furthermore, engaging users early in the technology development process and incorporating their values and feedback is vital for increasing user acceptance and engagement [54, 139]. This value-sensitive and human-centered approach ensures that the technology is not only technically sound but also resonates with the users' needs and preferences [9, 94].

Building upon the above analysis, our evaluation focuses on assessing whether the designs of these depression detection algorithms take into account three critical aspects: the potential for identity-based harm, situation-based harm, and the extent of technology builders' engagement with users to understand their concerns about mental health sensing technologies (summarized in Table 1).

Step 2: Establish criteria for evaluating harms, and make sure the bias is not attributed to random choice. In this step, we evaluate whether prior work established criteria for evaluating potential harms. To facilitate us to thoroughly examine the potential harms introduced by these depression detection algorithms, we set two key evaluation criteria: classification fairness metrics and thresholds for quantifying differences and biases. In the following subsections, we elaborate on the decision-making process that guided our choices of these criteria. Additionally, we detail the experimental implementation in Appendix A.2, for the sake of transparency and to facilitate reproducibility.

Criterion 1: Classification Fairness Metrics. We used three fairness metrics: disparity in accuracy, disparity in false negative rate, and disparity in false positive rate. These metrics were applied to assess algorithm performance across individuals with sensitive attributes and those without sensitive attributes. We intentionally chose not to adopt commonly used fairness metrics such as demographic parity (*e.g.*, [21]), which aim to ensure equal treatment across different groups. This decision was based on prior research findings indicating that individuals with sensitive attributes are more likely to experience depressive symptoms (*e.g.*, [61, 85, 90]). Using demographic parity, which aims for equal rates of predicted depressive symptoms across groups, could conflict with empirical evidence suggesting inherent disparities in depression prevalence. Our dataset analysis confirmed this, showing notably higher depression levels in certain sensitive groups (first-generation college students, immigrants, and non-male students) from 2018 and 2021² (see Figure 4 in Appendix). This highlights the critical need for selecting fairness metrics that reflect real-world disparities.

²We performed a Mann-Whitney U test with the Benjamini-Hochberg correction for significance testing (more details are in Section 4.1.1).

Criterion 2: Threshold for Quantifying Differences and Biases. We further added a criterion: a threshold quantifying differences in algorithmic performances across various groups. We implemented this to mitigate the impact of random variations. For this purpose, we chose established statistical tests, specifically opting for a non-parametric approach, considering the non-normal distribution of the chosen datasets. We utilized the Mann-Whitney U test, a widely recognized method for comparing means between two independent samples, irrespective of their distribution [87, 128]. We further employed the Benjamini-Hochberg (B-H) correction method to manage the Type I error rate associated with multiple comparisons within the same dataset [15]. We set a stringent False Discovery Rate (FDR) threshold at 0.05 [16, 62], ensuring that the rate of false positives is carefully controlled at 5%.

Steps 3 to 6: Collect inclusive datasets including comprehensive contextual information; develop context-sensitive algorithms; evaluate the behavioral sensing, combining user feedback and techniques to mitigate harms; and continuous maintenance of data and algorithms for responsible deployment. In evaluating steps 3 to 6, we assessed how the existing depression detection sensing technology builders handled several crucial aspects. First, we looked at whether they took into account the potential for identity-based and situation-based harms during data collection, in line with the contextual factors highlighted in Step 1. Second, we examined the design of the algorithms to determine if they were conscious of potential harms or biases and if they could adapt to contextual changes. Third, we evaluated how these technologies mitigated harms identified in their algorithms, particularly focusing on the use of user feedback. Finally, we assessed whether there were effective strategies and actions in place for the regular updating and maintenance of the data and algorithms.

Below, we present our evaluation findings of current depression detection algorithm designs, focusing on how well they align with our proposed framework. We then present the potential harms identified in these algorithms, based on the context identified in Step 1 and the criteria established in Step 2. This is followed by the experiment of bias mitigation and its results.

Table 3. Evaluation Results for Two Behavioral Sensing Technology Applications. This table summarizes the assessment of eight depression detection algorithms (third to tenth rows) based on their original publications, and a re-implementation study (eleventh row). The final line evaluates a student engagement prediction model. Symbols indicate consideration levels: ✓ for full consideration, ✖ for partial consideration, and ✗ for no consideration.

Algorithm Design / Implementation	Comprehensively Understand Context	Establish Harm Evaluation Criteria	Collect Inclusive Contextual Datasets	Design Context-Sensitive Algorithms	Harm Evaluation and Mitigation	Continuous Maintenance
Evaluation Study 1: Depression Detection						
Wahle <i>et al.</i>	✖	✗	✓	✗	✗	✗
Saeb <i>et al.</i>	✖	✗	✗	✗	✗	✗
Farhan <i>et al.</i>	✖	✗	✓	✗	✗	✗
Canzian <i>et al.</i>	✖	✗	✗	✗	✗	✗
Wang <i>et al.</i>	✖	✗	✓	✗	✗	✗
Lu <i>et al.</i>	✖	✗	✓	✗	✗	✗
Xu <i>et al.</i> - Interpretable	✖	✗	✗	✗	✗	✗
Xu <i>et al.</i> - Personalized	✖	✗	✓	✗	✗	✗
Xu <i>et al.</i> - Implementation	✖	✗	✓	✗	✗	✗
Evaluation Study 2: Engagement Prediction						
Gao <i>et al.</i> - En-gage	✖	✗	✓	✗	✗	✗

Evaluation Results. In our review of nine papers related to the design and implementation of eight depression detection algorithms, we observed that none of the prior work discussed potential harms to users, neither of them engaged with users to better understand their needs. All prior work considered identity-based context, *i.e.*, sensitive attributes. However, consistent with previous sensing technology research, most studies only focused on two sensitive attributes: age [49, 84, 120] and gender [49, 84, 125, 130, 131]. A few also considered

race [49, 84, 125, 131], but other sensitive attributes were largely overlooked. In terms of situated aspects, while most studies accounted for data collection time, consideration of device types was less common. Importantly, while these studies reported on this context information, many did not disclose the proportion of data pertaining to each, potentially leading to representative issues. More critically, none of the studies established criteria for evaluating potential harms, nor was there evidence of context-sensitive algorithm design or processes for harm evaluation and mitigation, particularly incorporating user feedback during the whole design process. Furthermore, there was a lack of strategies for the regular maintenance and updating of data and algorithms. These findings are summarized in Table 3, providing an overview of our evaluation results.

To assess the possibility of potential harms arising from a lack of context sensitivity in these depression detection algorithms, we carried out a quantitative analysis. Specifically, we leveraged the evaluation criteria defined in Step 2 (in Section 4.1.1) and evaluated the eight depression detection algorithms on the five sensitive attributes identified in Step 1 (gender, first-generation college student status, immigration status, race, and sexual orientation). Note that, disability status was removed due to the small sample size. The results, detailed in Table 4, revealed several insights.

Firstly, we observed biases in all algorithms towards certain sensitive attributes, *i.e.*, their disparities in accuracy, false negative rates, and false positive rates (highlighted in red in Table 3). Notably, algorithms with higher balanced accuracy [129, 133] tended to show fewer biases across these attributes when evaluated with the three fairness metrics. In particular, the algorithm **Xu_interpretable** [129] did not exhibit bias in terms of accuracy and false positive rate disparities.

Another interesting finding was the reduced bias in all algorithms on DS3, the dataset collected at the start of the COVID-19 outbreak in 2020. This suggests that the significant impact of COVID-19 might have overshadowed other sensitive attributes, leading to this pattern of decreased bias. Additionally, we did not see a consistent pattern indicating which algorithms consistently demonstrated fair performance regarding the sensitive attributes.

Additional Experiment on Bias Mitigation. Recognizing the presence of biases in the eight algorithms, we took steps to mitigate these algorithmic biases. Our approach involved an in-processing technique, where sensitive attributes were incorporated into both the training and testing phases [121, 135]. This method allows algorithms to understand and learn from the relationships between sensitive attributes and the target variable (*i.e.*, BDI-II scores). Previous research indicates that such an approach can help diminish discriminatory patterns present in the data, thereby enhancing the fairness of the models across diverse groups. [121, 135]. Our goal in this experiment was not to develop specific fair algorithms or mitigation techniques but rather to demonstrate a method for reducing bias and obtaining new insights.

As an example, we selected **Xu_interpretable** algorithm [129] due to its relatively high detection performance in depression detection (shown in Table 2) and its relatively low level of disparities across three fairness metrics in the four datasets and five sensitive attributes. We focused on mitigating bias related to first-generation college student status, a sensitive attribute where this algorithm showed bias in terms of disparity of false negative rate. We integrated this attribute into both the training and testing phases of the algorithm and re-evaluated the algorithm's performance.

Our evaluation result, as detailed in Table 5, demonstrates the effectiveness of including the status of being a first-generation college student in the training and testing phases to reduce algorithmic harms. This method led to a fair treatment of this particular sensitive attribute across all datasets, evaluated using three different fairness metrics. However, it is worth noting that while this approach improved fairness for first-generation college student status, it seemed to adversely impact fairness concerning other sensitive attributes such as sexual orientation and gender. A more comprehensive discussion of such trade-offs is described in Section 5.

Table 4. **Algorithmic Harm Evaluation Results.** Results of algorithmic harms through the disparity in accuracy, the disparity in false negative rate, and the disparity in false positive rate (without incorporating demographic data into the training and testing process). The results are adjusted p-values by Benjamini-Hochberg correction after the Mann-Whitney U test. Significance is highlighted in red. Acc, Fnr, and Fpr are the abbreviations of the disparity in accuracy, the disparity in false negative rate, and the disparity in false positive rate.

Algorithms	Sensitive Attributes	DS1 (2018)			DS2 (2019)			DS3 (2020)			DS4 (2021)		
		Acc	Fnr	Fpr	Acc	Fnr	Fpr	Acc	Fnr	Fpr	Acc	Fnr	Fpr
Wahle <i>et al.</i> [120]	First-gen College Student	0.020	0.030	0.030	0.010	0.050	0.010	0.020	0.030	0.040	0.010	0.050	0.020
	Gender	0.030	0.030	0.030	0.020	0.030	0.030	0.050	0.010	0.020	0.050	0.030	0.050
	Immigration Status	0.040	0.030	0.030	0.040	0.040	0.050	0.010	0.020	0.040	0.040	0.010	0.030
	Race	0.010	0.030	0.030	0.030	0.010	0.020	0.030	0.050	0.030	0.020	0.020	0.010
	Sexual Orientation	0.050	0.030	0.030	0.050	0.020	0.040	0.040	0.040	0.010	0.030	0.040	0.040
Saeb <i>et al.</i> [107]	First-gen College Student	0.010	0.030	0.010	0.020	0.030	0.030	0.010	0.050	0.050	0.020	0.030	0.050
	Gender	0.050	0.040	0.050	0.030	0.030	0.030	0.020	0.030	0.020	0.040	0.040	0.030
	Immigration Status	0.020	0.010	0.040	0.010	0.030	0.030	0.040	0.040	0.030	0.050	0.050	0.020
	Race	0.030	0.050	0.020	0.040	0.030	0.030	0.050	0.010	0.010	0.030	0.020	0.010
	Sexual Orientation	0.040	0.020	0.030	0.050	0.030	0.030	0.030	0.020	0.040	0.010	0.010	0.040
Farhan <i>et al.</i> [49]	First-gen College Student	0.030	0.020	0.040	0.020	0.030	0.040	0.030	0.040	0.040	0.030	0.010	0.010
	Gender	0.020	0.030	0.030	0.010	0.020	0.010	0.040	0.030	0.050	0.010	0.040	0.030
	Immigration Status	0.040	0.040	0.020	0.040	0.010	0.030	0.050	0.010	0.030	0.050	0.050	0.050
	Race	0.050	0.050	0.010	0.050	0.050	0.050	0.010	0.020	0.010	0.040	0.020	0.040
	Sexual Orientation	0.010	0.010	0.050	0.030	0.040	0.020	0.020	0.050	0.020	0.020	0.030	0.020
Canzian <i>et al.</i> [22]	First-gen College Student	0.020	0.030	0.030	0.020	0.010	0.020	0.020	0.020	0.030	0.010	0.030	0.020
	Gender	0.030	0.030	0.030	0.010	0.020	0.010	0.020	0.030	0.020	0.030	0.020	0.030
	Immigration Status	0.040	0.030	0.030	0.030	0.050	0.050	0.050	0.040	0.040	0.050	0.040	0.050
	Race	0.010	0.010	0.030	0.050	0.030	0.040	0.010	0.020	0.010	0.040	0.030	0.040
	Sexual Orientation	0.050	0.030	0.030	0.040	0.030	0.040	0.040	0.050	0.050	0.020	0.050	0.010
Wang <i>et al.</i> [125]	First-gen College Student	0.020	0.050	0.040	0.020	0.030	0.030	0.040	0.030	0.050	0.020	0.040	0.030
	Gender	0.040	0.040	0.030	0.030	0.030	0.030	0.050	0.040	0.020	0.040	0.020	0.040
	Immigration Status	0.010	0.020	0.010	0.010	0.030	0.030	0.010	0.010	0.030	0.050	0.010	0.050
	Race	0.030	0.010	0.020	0.040	0.030	0.030	0.020	0.020	0.040	0.030	0.030	0.020
	Sexual Orientation	0.050	0.030	0.050	0.050	0.030	0.030	0.030	0.050	0.010	0.010	0.050	0.010
Lu <i>et al.</i> [84]	First-gen College Student	0.010	0.020	0.030	0.010	0.010	0.050	0.040	0.040	0.040	0.050	0.040	0.030
	Gender	0.050	0.030	0.020	0.040	0.020	0.020	0.030	0.040	0.050	0.020	0.030	0.020
	Immigration Status	0.040	0.050	0.040	0.050	0.030	0.040	0.050	0.010	0.020	0.010	0.050	0.010
	Race	0.020	0.010	0.010	0.030	0.040	0.030	0.010	0.020	0.030	0.040	0.010	0.050
	Sexual Orientation	0.030	0.040	0.030	0.020	0.050	0.050	0.020	0.030	0.050	0.030	0.020	0.020
Xu_interpretable <i>et al.</i> [129]	First-gen College Student	0.040	0.020	0.050	0.010	0.030	0.010	0.010	0.010	0.020	0.030	0.010	0.030
	Gender	0.020	0.040	0.040	0.020	0.010	0.020	0.040	0.030	0.040	0.010	0.050	0.010
	Immigration Status	0.030	0.010	0.020	0.050	0.050	0.030	0.030	0.020	0.030	0.050	0.030	0.040
	Race	0.010	0.030	0.010	0.030	0.020	0.040	0.020	0.040	0.010	0.040	0.020	0.050
	Sexual Orientation	0.050	0.050	0.030	0.040	0.040	0.050	0.050	0.050	0.050	0.020	0.040	0.020
Xu_personalized <i>et al.</i> [130]	First-gen College Student	0.030	0.040	0.020	0.030	0.030	0.040	0.030	0.040	0.020	0.020	0.020	0.010
	Gender	0.010	0.050	0.040	0.040	0.050	0.030	0.050	0.010	0.010	0.050	0.030	0.020
	Immigration Status	0.050	0.020	0.050	0.020	0.040	0.020	0.010	0.030	0.020	0.040	0.050	0.030
	Race	0.020	0.010	0.030	0.010	0.020	0.010	0.020	0.020	0.050	0.030	0.040	0.040
	Sexual Orientation	0.040	0.030	0.010	0.050	0.010	0.050	0.040	0.050	0.040	0.010	0.010	0.050

4.2 Evaluation Study 2: Student Engagement Prediction

4.2.1 Background. In recent years, addressing the growing concerns of poor academic performance and student disinterest has led to a heightened interest in understanding student engagement, emotions, and daily behavior. This shift has coincided with significant advances in sensing technology, paving the way for novel methods to unobtrusively monitor and analyze student behavior and mental well-being in educational settings. A significant milestone in this domain is the introduction of the *En-Gage* dataset by Gao *et al.* [57]. This dataset, available at

Table 5. **Results After Implementing Bias Mitigation Techniques.** Example of algorithmic fairness evaluation results of Xu_Interpretable [129] through the three fairness metrics (with incorporating first-generation college student status into the training and testing process). The first row in each sub-table showcases the result of our evaluation pertaining to the first-generation college student status, subsequent to the incorporation of this sensitive attribute into both the training and testing phases. In the context of q-values obtained before and after the inclusion of first-generation college student status, ♠ represents the former, while ■ represents the latter. Acc, Fnr, and Fpr are the abbreviations of the disparity in accuracy, the disparity in false negative rate, and the disparity in false positive rate.

Fairness Metric	Sensitive Attribute	DS1 (2018)		DS2 (2019)		DS3 (2020)		DS4 (2021)	
		q value ♠	q value ■	q value ♠	q value ■	q value ♠	q value ■	q value ♠	q value ■
Disparity in Acc	First-gen College Student	0.040	0.030	0.010	0.020	0.010	0.040	0.030	0.030
	Gender	0.020	0.010	0.020	0.010	0.040	0.020	0.010	*0.010
	Immigration Status	0.030	0.050	0.050	0.050	0.030	0.030	0.050	0.040
	Race	0.010	0.040	0.030	0.040	0.020	0.010	0.040	0.020
	Sexual Orientation	0.050	0.020	0.040	0.030	0.050	0.050	0.020	0.050
Disparity in Fnr	First-gen College Student	0.020	0.040	0.030	0.020	*0.010	0.030	0.010	0.020
	Gender	0.040	0.010	0.010	0.030	0.030	0.050	0.050	0.050
	Immigration Status	0.010	0.030	0.050	0.050	0.020	0.040	0.030	0.040
	Race	0.030	0.020	0.020	0.040	0.040	0.020	0.020	0.030
	Sexual Orientation	0.050	0.050	0.040	0.010	0.050	0.010	0.040	0.010
Disparity in Fpr	First-gen College Student	0.050	0.030	0.010	0.020	0.020	0.040	0.030	0.030
	Gender	0.040	0.010	0.020	0.010	0.040	0.020	0.010	*0.010
	Immigration Status	0.020	0.020	0.030	0.040	0.030	0.050	0.040	0.050
	Race	0.010	0.050	0.040	0.030	0.010	0.010	0.050	0.020
	Sexual Orientation	0.030	0.035	0.040	0.038	0.050	0.027	0.030	0.040

PhysioNet ³, is distinguished as the largest and most diverse dataset in environmental and affect sensing within the educational field, offering unparalleled insights into student engagement patterns and classroom dynamics through a diverse array of sensing technologies.

Dataset. The En-Gage dataset includes a four-week cross-sectional study involving 23 Year-10 students (15–17 years old, 13 female and 10 male) and 6 teachers (33–62 years old, four female and two male) in a mixed-gender K12 private school. It utilizes wearable sensors to collect physiological data and daily surveys to gather information on the participants' thermal comfort (the comfort level of students regarding the perceived temperature at the time), learning engagement, seating locations, and emotions during school hours. An initial online survey was conducted to obtain participants' background information, including age, gender, general thermal comfort, and class groups. The dataset reflects the students' organization into different groups (Form group, Math group, and Language group), aiding in tracking their classroom locations. To clarify, students are typically enrolled in courses based on their form group division, except for math courses which are determined by their math group division, and language courses which are determined by their language group division.

Throughout the study, the participants were asked to wear *Empatica E4* wristbands [89] during school time, which capture 3-axis accelerometer readings, electrodermal activity, photoplethysmography (PPG), and skin temperature. They were also asked to complete online surveys three times a day, posted after certain classes. These surveys capture detailed insights into participants' behavioral, emotional, and cognitive engagement, as well as their emotions, thermal comfort and seating locations [58]. In total, the dataset comprises 291 survey responses and 1415.56 hours of physiological data from all participants.

Engagement Prediction Models. We chose the engagement regression model, LightGBM Regressors [98], developed by Gao *et al.* [59]. The regression model is designed to predict student engagement across three dimensions: *emotional*, *cognitive*, and *behavioral* engagement. Emotional engagement evaluates their feelings of belonging

³The dataset download link: <https://physionet.org/content/in-gauge-and-en-gage/1.0.0/>

and emotional reaction to the educational environment, cognitive engagement assesses their effort to understand complex ideas and skills, and behavioral engagement looks at students' participation in academic and extracurricular activities. The 1 to 5 Likert scale was used for scoring engagement levels, where 1 represents low and 5 high engagement. To predict these multidimensional scores, a variety of features were extracted, including data from wearable devices and weather stations. It is worth noting that, data such as gender, thermal comfort, and class groups, were not used for the engagement prediction.

4.2.2 Evaluation Methods and Results. In this subsection, following the approach used in Subsection 4.1.1, we first explain the decision-making process for each step of our framework, followed by the results of our evaluation.

Step 1: Comprehensively understand the context. Previous research has highlighted that marginalized groups, including women of color, students with disabilities, homeless youth, and religious minority students, often face feelings of alienation and isolation, which can significantly affect their learning engagement [101, 115]. Studies have also pointed out that variables like the perceived temperature and the timing of data collection can impact student engagement [60, 95]. Moreover, the role of social learning spaces, derived from engaging with student participants, has been recognized as a factor contributing to enhanced engagement [88].

Given these insights, our evaluation of student engagement prediction technology focused on whether its design took these contextual factors into account and aligned with the values and experiences of users. An overview of contextual factors we evaluated can be found in Table 1.

Step 2: Establish criteria for evaluating harms, and make sure the bias is not attributed to random choice. Given the regression-based prediction task, we used the disparity in Mean Squared Error (MSE) as our primary fairness metric to identify biases in model performance. MSE, the average of squared discrepancies between predicted and actual values, is widely recognized for assessing regression model accuracy [126]. Additionally, to discern if biases were systematic or due to random variation, and considering repeated measurements from individuals, we adopted a linear mixed model method [77]. This approach involved calculating residuals (differences between actual values and predictions) across various engagement prediction tasks. Subsequently, we utilized a linear mixed model, executed in Python, to examine whether these residuals significantly varied among different groups (e.g., gender and thermal comfort). This statistical method is beneficial for its ability to account for both within-group and between-group variations in the data, thereby offering a deeper insight into the biases present in model performance.

Steps 3 to 6. Our approach to decision-making and evaluation for Steps 3 to 6 in this evaluation study mirrors the method we employed in the first evaluation study, detailed in Section 4.1.1. A comprehensive overview of the criteria and methods used in these steps can also be found in Table 1.

Evaluation Results. In line with the results from our first evaluation study, our examination of relevant papers in this study [57, 59] indicates that researchers did not engage with the users to understand their needs or considered potential harms to users, and only very limited contextual factors were considered. These factors included gender, thermal comfort at the time of data collection, and information about the courses and classrooms that participants were involved in prior to data collection. Additionally, a key observation is that while this contextual data was considered during the data collection phase, it was not actively incorporated into the training and testing phases of their algorithms. Moreover, the researchers did not address the potential harms of their algorithms. They did not establish criteria for evaluating such harms or implement techniques, including student feedback, to mitigate potential biases. Additionally, there was no evidence of strategies for regular maintenance and updates of the data and algorithms.

We carried out a quantitative analysis to assess the potential negative impacts derived from neglecting certain contextual factors. The findings, detailed in Table 6, indicate that specific situated contexts – such as thermal comfort, the group division (e.g., language and math groups), and the courses students were engaged in prior to

Table 6. Results of Linear Mixed Models Analysis. This table displays the results from linear mixed models, focusing on identifying the significance of differences in regression models across diverse contexts within different engagement prediction tasks. Levels of significance are denoted as follows: * for $p < 0.05$, ** for $p < 0.01$, and *** for $p < 0.001$. For each contextual factor, one group is designated as the reference (or baseline) category, for example, the Female group in Gender. The “Interpret” represents the average effect for the reference group when all other predictors are held at their reference level (for categorical variables).

Contextual Factors	Model Variables	Emotional Engagement			Cognitive Engagement			Behavioral Engagement		
		Coef.	Std. Error	P> z	Coef.	Std. Error	P> z	Coef.	Std. Error	P> z
Gender	Intercept	0.006	0.119	0.962	-0.055	0.121	0.651	-0.033	0.105	0.753
	groups [T.Male]	-0.025	0.179	0.891	0.080	0.182	0.659	-0.038	0.157	0.810
	Group Var	0.121	0.072		0.120	0.070		0.076	0.044	
Thermal Comfort	Intercept	-0.140	0.116	0.225	-0.113	0.118	0.338	-0.136	0.113	0.232
	groups [T.No change]	0.286	0.112	*0.011	0.181	0.119	0.130	0.215	0.117	0.067
	groups [T.Warmer]	-0.117	0.150	0.436	-0.013	0.160	0.937	-0.190	0.156	0.225
	Group Var	0.112	0.067		0.098	0.054		0.081	0.052	
Language Group	Intercept	0.001	0.094	0.991	0.101	0.114	0.377	0.021	0.102	0.836
	groups [T.Room 41]	0.721	0.244	**0.003	-0.492	0.299	0.100	0.327	0.264	0.215
	groups [T.Room 43]	-0.094	0.184	0.611	-0.213	0.219	0.331	-0.248	0.198	0.210
	groups [T.Room 68]	-0.351	0.198	0.076	-0.181	0.241	0.452	-0.324	0.214	0.129
	Group Var	0.057	0.045		0.097	0.070		0.068	0.052	
Math Group	Intercept	0.194	0.166	0.242	-0.314	0.155	*0.043	0.125	0.153	0.414
	groups [T.Room 41]	-0.292	0.215	0.175	0.330	0.201	0.101	-0.335	0.198	0.091
	groups [T.Room 43]	-0.251	0.223	0.261	0.500	0.209	0.017	-0.131	0.205	0.524
	Group Var	0.111	0.070		0.086	0.05		0.082	0.055	
Course	Intercept	-0.285	0.237	0.228	-0.376	0.241	0.119	-0.022	0.237	0.926
	groups [T.English]	0.488	0.241	*0.043	0.396	0.246	0.107	0.290	0.246	0.239
	groups [T.Health]	0.400	0.353	0.257	0.395	0.360	0.273	-0.086	0.360	0.811
	groups [T.Language]	0.075	0.255	0.769	-0.148	0.260	0.569	-0.400	0.260	0.124
	groups [T.Maths]	0.274	0.240	0.253	0.530	0.245	*0.030	0.035	0.245	0.885
	groups [T.PE]	0.485	0.356	0.173	0.558	0.363	0.125	0.226	0.363	0.532
	groups [T.Politics]	0.174	0.251	0.489	0.212	0.256	0.407	-0.343	0.257	0.182
	groups [T.Science]	0.240	0.262	0.360	0.634	0.267	*0.018	-0.085	0.267	0.749
	Group Var	0.128	0.085		0.126	0.073		0.084	0.054	

data collection – significantly influence the performance of the prediction algorithm. For example, as illustrated in Table 6, the algorithm’s ability to accurately assess emotional engagement was statistically different between students who were comfortable with the room temperature and those who were not (feeling either too cold or too warm). To delve deeper into this observation, we analyzed the mean squared error (MSE) of the regression algorithm across different levels of thermal comfort. As reported in Table 7, the algorithm showed a notably lower error rate ($MSE = 0.631$, $p = 0.011$) when predicting the emotional engagement of students who were comfortable with the temperature, compared to those who were not ($MSE = 0.822$ for students feeling the temperature should be cooler and $MSE = 0.742$ for students feeling the temperature should be warmer). Similarly, our analysis indicated a significantly higher error rate ($MSE = 0.849$, $p = 0.043$) in predicting the cognitive engagement of students in Room 40 for their math class, as opposed to those in other math groups ($MSE = 0.715$ for Room 41 and $MSE = 0.753$ for Room 43). Interestingly, our analysis revealed no evidence of algorithmic bias or harm, both with gender and in predicting student behavioral engagement.

Table 7. **Overview of Basic Statistics.** MSE refers to the Mean Squared Error, indicating the average of the squares of the errors. 'Residual' denotes the difference between the ground truth and prediction. MR represents the Mean Residual, which is the average of residuals within each group. "Ind" and "Obs" stand for individuals and observations, respectively.

Context Factors	Groups	Counts (Ind/Obs)	Emotional Engagement		Cognitive Engagement		Behavioral Engagement	
			MSE	MR	MSE	MR	MSE	MR
Gender	Female	13/149	0.708	-0.014	0.711	-0.005	0.800	-0.023
	Male	10/142	0.693	0.033	0.822	-0.004	0.674	0.002
Thermal Comfort	No Change	22/163	0.631	0.158	0.774	0.069	0.687	0.129
	Cooler	20/77	0.822	-0.140	0.755	-0.135	0.730	-0.115
	Warmer	14/51	0.742	-0.242	0.751	-0.045	0.915	-0.300
Language Group	Room 40	13/155	0.618	0.008	0.690	0.156	0.681	0.051
	Room 41	2/53	0.886	0.703	1.019	-0.563	0.799	0.427
	Room 43	5/52	0.526	-0.075	0.779	-0.082	0.592	-0.186
	Room 68	3/53	1.007	-0.312	0.823	-0.072	1.016	-0.276
Math Group	Room 40	7/80	0.671	0.247	0.849	-0.327	0.640	0.178
	Room 41	9/110	0.763	-0.114	0.715	0.044	0.783	-0.185
	Room 43	7/101	0.657	-0.046	0.753	0.197	0.768	0.030
Course	Chapel	11/12	0.938	-0.268	1.100	-0.297	0.693	0.021
	English	18/71	0.599	0.255	0.484	0.057	0.639	0.340
	Health	8/8	0.991	0.155	1.538	0.149	0.776	-0.035
	Language	20/38	0.986	-0.206	0.970	-0.512	0.975	-0.372
	Maths	20/79	0.551	-0.033	0.816	0.150	0.779	0.0177
	PE	8/8	1.044	0.235	1.314	0.103	0.845	0.224
	Politics	19/43	0.754	-0.119	0.784	-0.101	0.643	-0.341
Science	19/32	0.641	0.006	0.537	0.252	0.687	-0.050	

Table 8. **Results After Implementing Bias Mitigation Techniques.** The outcomes following the inclusion of language group assignment data in the training and testing phases of the emotional engagement prediction algorithm.

Contextual Factors	Model Variables	Emotional Engagement		
		Coef.	Std. Error	P> z
Gender	Intercept	-0.026	0.096	0.785
	groups [T.Male]	0.030	0.143	0.834
	Group Var	0.060	0.045	
Thermal Comfort	Intercept	-0.154	0.103	0.137
	groups [T.No change]	0.288	0.110	**0.009
	groups [T.Warmer]	-0.094	0.147	0.523
Group Var	0.055	0.042		
Language Group	Intercept	0.022	0.093	0.811
	groups [T.Room 41]	0.236	0.240	0.325
	groups [T.Room 43]	-0.060	0.181	0.740
	groups [T.Room 68]	-0.282	0.195	0.147
	Group Var	0.053	0.043	

Contextual Factors	Model Variables	Emotional Engagement		
		Coef.	Std. Error	P> z
Math Room	Intercept	0.164	0.131	0.210
	groups [T.Room 41]	-0.299	0.169	0.077
	groups [T.Room 43]	-0.172	0.175	0.327
	Group Var	0.050	0.042	
Course	Intercept	-0.345	0.228	0.131
	groups [T.English]	0.579	0.240	*0.016
	groups [T.Health]	0.554	0.351	0.115
	groups [T.Language]	0.111	0.254	0.663
	groups [T.Maths]	0.304	0.239	0.203
	groups [T.PE]	0.490	0.354	0.166
	groups [T.Politics]	0.202	0.251	0.420
groups [T.Science]	0.320	0.261	0.219	
Group Var	0.051	0.040		

Additional Experiment on Bias Mitigation. To determine if the findings from our first evaluation study can be replicated using the same bias mitigation technique – incorporate context data into both the training and testing phase of the algorithm – we conducted an additional experiment in this evaluation study.

As an example, we aimed to mitigate the algorithmic bias caused by the lack of detailed information about students' assignments to different language groups, specifically in the context of predicting student emotional engagement. As indicated in Table 8, incorporating this information into both the training and testing phases of the algorithm proved effective in reducing algorithmic harm. Compared to Table 3, this method resulted in a

more equitable prediction performance across students assigned to various language groups. However, it was less effective in addressing biases related to different levels of thermal comfort and the variety of courses students were taking.

5 DISCUSSION

In this section, we begin by summarizing the key insights we derived from our evaluation studies (Section 5.1). This summary covers the various findings, their implications, and how they contribute to our understanding of designing behavioral sensing technologies. Following this, we delve into a reflection on our framework, examining its strengths, limitations, and considering perspectives that extend beyond its current scope (Section 5.2).

5.1 Key Insights on Evaluation Studies

Our evaluations of two real-world behavioral well-being sensing technology studies demonstrated the practicality and effectiveness of our proposed framework. Throughout both evaluation studies, we identified a range of commonalities as well as unique findings, which we detail below.

5.1.1 Potential Harms to Marginalized Groups Due to Context-insensitivity. In both of our evaluation studies, we uncovered a critical and consistent issue with existing behavioral sensing technology designs: a widespread disregard for potential harms to users. Our evaluation, as detailed in Table 3, revealed that none of the designs thoroughly considered steps 2, 4, 5, and 6 proposed in our framework during their design processes. Furthermore, while a few designs did consider the collection of more diverse contextual datasets (e.g., [129, 130]), this type of data was not utilized effectively during the algorithm training and testing phases. Our quantitative analysis of algorithm performance substantiates the concern of potential harms due to this oversight. Both studies identified significant issues, either identity-based harm or situation-based harm. Identity-based harm, which is more straightforward, can directly impact marginalized groups. In contrast, the concept of situation-based harm is more nuanced and can be subtle in its impact on these groups. To illustrate this further, in addition to the example discussed in Section 2.3.2, our second evaluation study provides another insightful instance. Specifically, we found that the algorithm for predicting emotional engagement was less effective for students who felt uncomfortably cold or warm compared to those who were comfortable with the temperature (as detailed in Tables 6 and 7). This finding may imply a potential indirect harm to individuals of lower socioeconomic status, who may have restricted access to air conditioning and thus are more likely to experience algorithmic harms [64].

5.1.2 Balance Trade-offs in Achieving Algorithmic Fairness. In our first evaluation study, as detailed in Section 4.1.1, we encountered a trade-off when attempting to mitigate harms. We observed that while using an in-processing mechanism (i.e., incorporating the sensitive attribute into the training and testing process) helped reduce bias for that particular attribute, it unexpectedly introduced a bias towards other sensitive attributes. This outcome highlights the complex and nuanced nature of mitigating algorithmic harms. It suggests that while certain mitigation strategies might address specific biases, they can also unintentionally create new harms.

Understanding these intricate trade-offs, technology builders should explicitly ask the question when developing context-sensitive algorithms: which group of users should be prioritized and to what extent? Our framework emphasized the importance of comprehensively understanding users' backgrounds and specific needs to answer the first part of this question. Our emphasis on engaging more with users and involving them throughout the design of the behavioral sensing technologies answers the latter part of this question.

The complex field of algorithmic fairness often presents trade-offs not only between different user groups but also among various fairness metrics [66]. Each metric provides a unique lens on bias, focusing on different aspects of equity. However, optimizing for one metric may lead to unintended negative outcomes in another, creating challenging scenarios [31]. For example, in our first evaluation study, when sensitive attributes were incorporated

into Saeb *et al.*'s algorithm training and testing, it reduced bias in accuracy disparity for most sensitive attributes, as shown in Table 9. Yet, a detailed examination of other key fairness metrics like disparity in false negative rates and positive rates (Tables 10 and 11) reveals significant variations. This highlights the complex dynamics involved in fairness optimization, where achieving fairness in one dimension might inadvertently lead to imbalances in others, emphasizing the multifaceted nature of achieving algorithmic fairness. This dynamic becomes even more crucial in behavioral sensing technologies, where data collection remains continuous, and system behavior is deeply adaptive to changing contexts. This recognition sets the stage for our subsequent discussion (Section 5.2), delving into the essential requirement for regular and systematic monitoring.

5.1.3 A Need for Engage Users Throughout the Design Process. Another key finding from both of our evaluation studies is the complete absence of user involvement throughout the design process of existing behavioral sensing technologies. Given the widespread use of behavioral sensing technologies, particularly in the mental health domain, this is concerning. As argued by Zhu *et al.* [139], engaging with users in the early stage of the design process can ensure that technologies are designed with a deep understanding of users' needs and values, which can significantly enhance user acceptance and satisfaction. Furthermore, the engagement of users extends beyond the initial design phase to include ongoing feedback loops. Regular interactions with users allow for iterative improvements and adjustments based on evolving needs, emerging challenges, and changing social contexts [9]. However, it is important to recognize the balance between involving users to mitigate technology harms and minimizing demands on their time and resources. This is especially necessary for people with different needs [137]. Striking this balance ensures that users' contributions are meaningful and sustainable, and that their valuable input genuinely shapes the direction of the technology while respecting their availability and capacity.

5.2 Towards More Responsible Behavioral Sensing

In this section, we discuss various aspects both within and beyond our current framework. These include the necessity of continuous maintenance for longitudinal behavioral sensing deployment while minimizing human labor, considerations of harms in other components of behavioral sensing technology, as well as the incorporation of other responsible considerations. Our intention is to inspire researchers and designers towards the conception and realization of more responsible behavioral sensing technologies.

5.2.1 Need for Regular Maintenance while Alleviating Excessive Human Labor. During our evaluation studies, we discovered that the reason behind the lack of continuous maintenance for responsible deployment (as outlined in step 6 of our framework) in these behavioral sensing technology designs was that the technologies were not truly deployed in real-world settings. This limitation arises from the nature of the limited datasets and the absence of deployable algorithmic systems [131, 138].

Behavioral sensing technologies in real-world applications operate in dynamic environments and depend on continuously evolving data streams. This dynamic nature increases the risk of situation-based harms, demanding continuous vigilance to guarantee that the system's accuracy and fair alignment persist over time. Regular maintenance is a key step in achieving this goal. By continuously updating deployment datasets, refining algorithms to accommodate temporal dynamics, and regularly monitoring the system's performance, the system can uphold its reliability and fulfill its ethical responsibility towards users and users. However, it is also important to avoid overburdening human resources with excessive maintenance demands. High human labor requirements can lead to operational inefficiencies, increased costs, and hindered scalability [51]. Striking a balance between rigorous maintenance and an approach that minimizes the burden on human resources is pivotal. Leveraging automation, intelligent monitoring, and adaptive algorithms can potentially alleviate this issue.

5.2.2 Other Aspects of Harms in Behavioral Sensing Technology. In this work, our emphasis centers on addressing harms within the algorithmic aspects of behavioral sensing technology. Nonetheless, it is important to acknowledge that considerations of harms should extend beyond the algorithm design process and include other critical components of the system. One dimension, for example, to focus on is the user interface and user interaction.

A user interface that is designed without incorporating the consideration of harms to users can inadvertently influence users to make certain choices or take specific actions. When these nudges disproportionately benefit particular groups, it can result in disparate outcomes that perpetuate inequality. Additionally, if user interfaces are not designed with accessibility in mind, individuals with disabilities might face barriers in accessing the interacting with the system. Future endeavors should take this aspect into account when designing their user interface. Finally, approaches to transparently informing users about potential fairness concerns, similar to transparent information about accuracy concerns, should be incorporated into a deployed fair behavioral sensing technology.

5.2.3 Expanding Responsible Considerations to Address Additional Needs. While our study primarily concentrates on algorithmic harms in the context of behavioral sensing technology, it is essential to recognize that responsible considerations encompass a broader spectrum of dimensions, such as transparency, privacy, and accountability (e.g., [45, 81]). As behavioral sensing technology becomes more widely used, ensuring transparency becomes a pressing concern. A lack of transparency can lead to opacity and a lack of user trust [12, 20]. Interpretability and explanation techniques are crucial in addressing this issue, allowing users to understand algorithmic decisions and aiding system developers in identifying potential harms [43]. Furthermore, continuous data collection in behavioral sensing raises significant privacy challenges [104]. The risk of unauthorized and unintended data sharing is ever-present. Researchers can develop privacy-preserving algorithms and techniques tailored specifically to behavioral sensing environments and delve into privacy-enhancing technologies, such as secure multi-party computation [63, 74], federated learning [105], and differential privacy [44], and incorporate them into the framework. In addition, accountability is about establishing mechanisms to hold responsible parties accountable for the outcomes of their algorithms and systems [41]. In behavioral sensing, accountability can be challenging due to complex decision-making processes and interactions between various components. Technology builders of these technologies must be held answerable for their impact on users. Expanding our proposed framework to include all the above-discussed aspects can create a more comprehensive foundation for responsible behavioral sensing technology design and deployment.

6 LIMITATIONS AND FUTURE WORK

While our research included two comprehensive evaluations of real-world behavioral sensing technologies across various domains and machine ML tasks, aiming to derive broader conclusions, we acknowledge that both evaluation studies are situated within the overarching theme of well-being prediction. This specific focus may limit the generalizability of our findings to other applications outside of well-being prediction. Furthermore, in both of our evaluation studies, we identified various instances of identity-based and situation-based harms. However, we note that some aspects might still be overlooked. Moreover, the datasets used in our evaluations presented their own set of constraints. For instance, specific instances of harms, such as disability status, were either unrepresented or underrepresented due to limited data collection or small sample sizes. This data limitation restricts our ability to make conclusive statements about these groups. Future work should continue to explore the potential harms in a broader array of behavioral sensing technology applications, identify additional instances both within the two potential harms we discussed and beyond, as well as collect more inclusive datasets for a more comprehensive analysis. Additionally, while our study demonstrates the feasibility of our proposed framework for evaluating and mitigating harms, its practical usability and broader application still require validation. Comparing

our framework with others, although beyond the scope of this study, will be crucial in future meta-evaluations to benchmark its effectiveness and utility in the field.

7 CONCLUSION

In conclusion, this paper identified three gaps in the existing literature. We introduce a specific framework for evaluating and mitigating the context-induced harms in behavioral sensing. Our framework highlights a comprehensive consideration of potential broad and domain-specific harms due to a lack of context sensitivity and the need for iterative harm mitigation and continuous maintenance for responsible technology deployment. Through our two evaluation studies, we showcase the feasibility of our proposed framework. By conducting quantitative analyses, we uncover empirical evidence of identity-based and situation-based harms in existing behavioral sensing technologies and validate the framework's capability in identifying and mitigating these harms. We discuss the insights learned from the evaluation studies and other aspects within and beyond the scope of our proposed framework. We hope our work inspires technology builders in our field to amplify their attention to the significance of incorporating harm considerations and other responsible considerations in behavioral sensing technologies.

REFERENCES

- [1] Gregory D Abowd. 2012. What next, ubicomp? Celebrating an intellectual disappearing act. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*. 31–40.
- [2] Gregory D Abowd, Anind K Dey, Peter J Brown, Nigel Davies, Mark Smith, and Pete Steggle. 1999. Towards a better understanding of context and context-awareness. In *Handheld and Ubiquitous Computing: First International Symposium, HUC'99 Karlsruhe, Germany, September 27–29, 1999 Proceedings 1*. Springer, 304–307.
- [3] Daron Acemoglu. 2021. *Harms of AI*. Technical Report. National Bureau of Economic Research.
- [4] Daniel A Adler, Dror Ben-Zeev, Vincent WS Tseng, John M Kane, Rachel Brian, Andrew T Campbell, Marta Hauser, Emily A Scherer, and Tanzeem Choudhury. 2020. Predicting early warning signs of psychotic relapse from passive sensing data: an approach using encoder-decoder neural networks. *JMIR mHealth and uHealth* 8, 8 (2020), e19962.
- [5] Daniel A Adler, Caitlin A Stamatis, Jonah Meyerhoff, David C Mohr, Fei Wang, Gabriel J Aranovich, Srijan Sen, and Tanzeem Choudhury. 2024. Measuring algorithmic bias to analyze the reliability of AI tools that predict depression risk using smartphone sensed-behavioral data. *npj Mental Health Research* 3, 1 (2024), 17.
- [6] Daniel A Adler, Emily Tseng, Khatiya C Moon, John Q Young, John M Kane, Emanuel Moss, David C Mohr, and Tanzeem Choudhury. 2022. Burnout and the Quantified Workplace: Tensions around Personal Sensing Interventions for Stress in Resident Physicians. *Proceedings of the ACM on Human-computer Interaction* 6, CSCW2 (2022), 1–48.
- [7] Karan Ahuja, Dohyun Kim, Franceska Xhakaj, Virag Varga, Anne Xie, Stanley Zhang, Jay Eric Townsend, Chris Harrison, Amy Ogan, and Yuvraj Agarwal. 2019. EduSense: Practical classroom sensing at Scale. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 3 (2019), 1–26.
- [8] Paul R Albert. 2015. Why is depression more prevalent in women? *Journal of psychiatry & neuroscience: JPN* 40, 4 (2015), 219.
- [9] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- [10] Karim Assi, Lakmal Meegahapola, William Droz, Peter Kun, Amalia De Götzen, Miriam Bidoglia, Sally Stares, George Gaskell, Altangerel Chagnaa, Amarsanaa Ganbold, et al. 2023. Complex daily activities, country-level diversity, and smartphone sensing: A study in denmark, italy, mongolia, paraguay, and uk. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–23.
- [11] Nikola Banovic, Tofi Buzali, Fanny Chevalier, Jennifer Mankoff, and Anind K Dey. 2016. Modeling and understanding human routine behavior. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 248–260.
- [12] Nikola Banovic, Zhuoran Yang, Aditya Ramesh, and Alice Liu. 2023. Being Trustworthy is Not Enough: How Untrustworthy Artificial Intelligence (AI) Can Deceive the End-Users and Gain Their Trust. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–17.
- [13] Aaron T Beck, Robert A Steer, Roberta Ball, and William F Ranieri. 1996. Comparison of Beck Depression Inventories-IA and-II in psychiatric outpatients. *Journal of personality assessment* 67, 3 (1996), 588–597.
- [14] Alireza Behtoui and Anders Neergaard. 2010. Social capital and wage disadvantages among immigrant workers. *Work, Employment and Society* 24, 4 (2010), 761–779.

- [15] Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)* 57, 1 (1995), 289–300.
- [16] Yoav Benjamini and Daniel Yekutieli. 2005. Quantitative trait loci analysis using the false discovery rate. *Genetics* 171, 2 (2005), 783–790.
- [17] Brianna Blaser, Cynthia Bennett, Richard E Ladner, Sheryl E Burgstahler, Jennifer Mankoff, C Frieze, and JL Quesenberry. 2019. *Perspectives of women with disabilities in computing*. Cambridge, UK: Cambridge Univ. Press.
- [18] Joshua Breslau, Sergio Aguilar-Gaxiola, Kenneth S Kendler, Maxwell Su, David Williams, and Ronald C Kessler. 2006. Specifying race-ethnic differences in risk for psychiatric disorder in a USA national sample. *Psychological medicine* 36, 1 (2006), 57–68.
- [19] MA Brodie, EM Pliner, A Ho, Kalina Li, Z Chen, SC Gandevia, and SR Lord. 2018. Big data vs accurate data in health research: large-scale physical activity monitoring, smartphones, wearable devices and risk of unconscious bias. *Medical hypotheses* 119 (2018), 32–36.
- [20] Zana Buçinca, Maja Barbara Malaya, and Krzysztof Z Gajos. 2021. To trust or to think: cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–21.
- [21] Francois Buet-Golfouse and Islam Utyagulov. 2022. Towards fair unsupervised learning. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 1399–1409.
- [22] Luca Canzian and Mirco Musolesi. 2015. Trajectories of depression: unobtrusive monitoring of depressive states by means of smartphone mobility traces analysis. In *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*. 1293–1304.
- [23] Tara Capel and Margot Brereton. 2023. What is Human-Centered about Human-Centered AI? A Map of the Research Landscape. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–23.
- [24] Iwona Chelminski, F Richard Ferraro, Thomas V Petros, and Joseph J Plaud. 1999. An analysis of the “eveningness–morningness” dimension in “depressive” college students. *Journal of affective disorders* 52, 1-3 (1999), 19–29.
- [25] Prerna Chikersal, Afsaneh Doryab, Michael Tumminia, Daniella K Villalba, Janine M Dutcher, Xinwen Liu, Sheldon Cohen, Kasey G Creswell, Jennifer Mankoff, J David Creswell, et al. 2021. Detecting depression and predicting its onset using longitudinal symptoms captured by passive sensing: a machine learning approach with robust feature selection. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 1 (2021), 1–41.
- [26] Tina Chou, Anu Asnaani, and Stefan G Hofmann. 2012. Perception of racial discrimination and psychopathology across three US ethnic minority groups. *Cultural Diversity and Ethnic Minority Psychology* 18, 1 (2012), 74.
- [27] Shreya Chowdhary, Anna Kawakami, Mary L Gray, Jina Suh, Alexandra Olteanu, and Koustuv Saha. 2023. Can Workers Meaningfully Consent to Workplace Wellbeing Technologies?. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 569–582.
- [28] Karen Church, Denzil Ferreira, Nikola Banovic, and Kent Lyons. 2015. Understanding the challenges of mobile phone usage data. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services*. 504–514.
- [29] Patricia Hill Collins. 2019. *Intersectionality as critical social theory*. Duke University Press.
- [30] European Commission. 2019. Ethics guidelines for trustworthy AI. <https://www.aepd.es/sites/default/files/2019-12/ai-ethics-guidelines.pdf>.
- [31] Sam Corbett-Davies and Sharad Goel. 2018. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023* (2018).
- [32] Victor P Cornet and Richard J Holden. 2018. Systematic review of smartphone-based passive sensing for health and wellbeing. *Journal of biomedical informatics* 77 (2018), 120–132.
- [33] Shanley Corvite, Kat Roemmich, Tillie Ilana Rosenberg, and Nazanin Andalibi. 2023. Data Subjects’ Perspectives on Emotion Artificial Intelligence Use in the Workplace: A Relational Ethics Lens. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–38.
- [34] Vedant Das Swain, Victor Chen, Shrija Mishra, Stephen M Mattingly, Gregory D Abowd, and Munmun De Choudhury. 2022. Semantic gap in predicting mental wellbeing through passive sensing. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [35] Vedant Das Swain, Lan Gao, Abhirup Mondal, , Gregory D Abowd, and Munmun De Choudhury. 2024. Sensible and Sensitive AI for Worker Wellbeing: Factors that Inform Adoption and Resistance for Information Workers. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*.
- [36] Vedant Das Swain, Lan Gao, William A Wood, Srikruthi C Matli, Gregory D Abowd, and Munmun De Choudhury. 2023. Algorithmic Power or Punishment: Information Worker Perspectives on Passive Sensing Enabled AI Phenotyping of Performance and Wellbeing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [37] Vedant Das Swain, Koustuv Saha, Manikanta D Reddy, Hemang Rajvanshy, Gregory D Abowd, and Munmun De Choudhury. 2020. Modeling organizational culture with workplace experiences shared on glassdoor. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–15.
- [38] Nigel Davies, Keith Mitchell, Keith Cheverst, and Gordon Blair. 1998. Developing a context sensitive tourist guide. In *1st workshop on human computer interaction with Mobile devices, GIST technical report G98-1*, Vol. 1.
- [39] Anind K Dey. 2001. Understanding and using context. *Personal and ubiquitous computing* 5 (2001), 4–7.

- [40] Anind K Dey, Gregory D Abowd, and Daniel Salber. 2001. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction* 16, 2-4 (2001), 97–166.
- [41] Nicholas Diakopoulos. 2016. Accountability in algorithmic decision making. *Commun. ACM* 59, 2 (2016), 56–62.
- [42] Afsaneh Doryab, Prerna Chikarsel, Xinwen Liu, and Anind K Dey. 2018. Extraction of behavioral features from smartphone and wearable data. *arXiv preprint arXiv:1812.10394* (2018).
- [43] Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608* (2017).
- [44] Cynthia Dwork. 2006. Differential privacy. In *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II* 33. Springer, 1–12.
- [45] Upol Ehsan, Q Vera Liao, Michael Muller, Mark O Riedl, and Justin D Weisz. 2021. Expanding explainability: Towards social transparency in ai systems. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [46] Upol Ehsan, Koustuv Saha, Munmun De Choudhury, and Mark O Riedl. 2023. Charting the Sociotechnical Gap in Explainable AI: A Framework to Address the Gap in XAI. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–32.
- [47] Carolyn Ellis, Tony E Adams, and Arthur P Bochner. 2011. Autoethnography: an overview. *Historical social research/Historische sozialforschung* (2011), 273–290.
- [48] Sheena Erete, Yolanda A Rankin, and Jakita O Thomas. 2021. I can’t breathe: Reflections from Black women in CSCW and HCI. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW3 (2021), 1–23.
- [49] Asma Ahmad Farhan, Chaoqun Yue, Reynaldo Morillo, Shweta Ware, Jin Lu, Jinbo Bi, Jayesh Kamath, Alexander Russell, Athanasios Bamis, and Bing Wang. 2016. Behavior vs. introspection: refining prediction of clinical depression via smartphone sensing data. In *2016 IEEE wireless health (WH)*. IEEE, 1–8.
- [50] Denzil Ferreira, Vassilis Kostakos, and Anind K Dey. 2015. AWARE: mobile context instrumentation framework. *Frontiers in ICT* 2 (2015), 6.
- [51] Carl Benedikt Frey and Michael A Osborne. 2017. The future of employment: How susceptible are jobs to computerisation? *Technological forecasting and social change* 114 (2017), 254–280.
- [52] Batya Friedman. 1996. Value-sensitive design. *interactions* 3, 6 (1996), 16–23.
- [53] Batya Friedman, David G Hendry, Alan Borning, et al. 2017. A survey of value sensitive design methods. *Foundations and Trends® in Human-Computer Interaction* 11, 2 (2017), 63–125.
- [54] Batya Friedman, Peter H Kahn, Alan Borning, and Alina Hultgren. 2013. Value sensitive design and information systems. *Early engagement and new technologies: Opening up the laboratory* (2013), 55–95.
- [55] David M Frost. 2011. Social stigma and its consequences for the socially stigmatized. *Social and Personality Psychology Compass* 5, 11 (2011), 824–839.
- [56] Kenneth Fung and Cindy-Lee Dennis. 2010. Postpartum depression among immigrant women. *Current opinion in Psychiatry* 23, 4 (2010), 342–348.
- [57] Nan Gao, Max Marschall, Jane Burry, Simon Watkins, and Flora D Salim. 2022. Understanding occupants’ behaviour, engagement, emotion, and comfort indoors with heterogeneous sensors and wearables. *Scientific Data* 9, 1 (2022), 261.
- [58] Nan Gao, Mohammad Saiedur Rahaman, Wei Shao, Kaixin Ji, and Flora D Salim. 2022. Individual and group-wise classroom seating experience: Effects on student engagement in different courses. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–23.
- [59] Nan Gao, Wei Shao, Mohammad Saiedur Rahaman, and Flora D Salim. 2020. n-gage: Predicting in-class emotional, behavioural and cognitive engagement in the wild. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020), 1–26.
- [60] Eva Garrosa, Luis Manuel Blanco-Donoso, Isabel Carmona-Cobo, and Bernardo Moreno-Jiménez. 2017. How do curiosity, meaning in life, and search for meaning predict college students’ daily emotional exhaustion and engagement? *Journal of Happiness Studies* 18 (2017), 17–40.
- [61] Jane L Givens, Thomas K Houston, Benjamin W Van Voorhees, Daniel E Ford, and Lisa A Cooper. 2007. Ethnicity and preferences for depression treatment. *General hospital psychiatry* 29, 3 (2007), 182–191.
- [62] Mark E Glickman, Sowmya R Rao, and Mark R Schultz. 2014. False discovery rate control is a recommended alternative to Bonferroni-type adjustments in health studies. *Journal of clinical epidemiology* 67, 8 (2014), 850–857.
- [63] Oded Goldreich. 1998. Secure multi-party computation. *Manuscript. Preliminary version* 78, 110 (1998).
- [64] Carina J Gronlund. 2014. Racial and socioeconomic disparities in heat-related health effects and their mechanisms: a review. *Current epidemiology reports* 1 (2014), 165–173.
- [65] Dominik Hangartner, Daniel Kopp, and Michael Siegenthaler. 2021. Monitoring hiring discrimination through online recruitment platforms. *Nature* 589, 7843 (2021), 572–576.
- [66] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems* 29 (2016).

- [67] Wiebke Toussaint Hutiri and Aaron Yi Ding. 2022. Bias in automated speaker recognition. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. 230–247.
- [68] Maral Jamalova and M Constantinovits. 2019. The comparative study of the relationship between smartphone choice and socio-economic indicators. *International Journal of Marketing Studies* 11, 3 (2019), 11.
- [69] Sharon Rae Jenkins, Aimee Belanger, Melissa Londoño Connally, Adriel Boals, and Kelly M Durón. 2013. First-generation undergraduate students' social support, depression, and life satisfaction. *Journal of College Counseling* 16, 2 (2013), 129–142.
- [70] Michael I Jordan. 2019. Artificial intelligence—the revolution hasn't happened yet. *Harvard Data Science Review* 1, 1 (2019), 1–9.
- [71] Saffron Karlsen and James Y Nazroo. 2002. Relation between racial discrimination, social class, and health among ethnic minority groups. *American journal of public health* 92, 4 (2002), 624–631.
- [72] Harmanpreet Kaur, Daniel McDuff, Alex C Williams, Jaime Teevan, and Shamsi T Iqbal. 2022. "I didn't know I looked angry": Characterizing observed emotion and reported affect at work. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [73] Anna Kawakami, Shreya Chowdhary, Shamsi T Iqbal, Q Vera Liao, Alexandra Olteanu, Jina Suh, and Koustuv Saha. 2023. Sensing Wellbeing in the Workplace, Why and For Whom? Envisioning Impacts with Organizational Stakeholders. *arXiv preprint arXiv:2303.06794* (2023).
- [74] Marcel Keller. 2020. MP-SPDZ: A versatile framework for multi-party computation. In *Proceedings of the 2020 ACM SIGSAC conference on computer and communications security*. 1575–1590.
- [75] Sunnie SY Kim, Elizabeth Anne Watkins, Olga Russakovsky, Ruth Fong, and Andrés Monroy-Hernández. 2023. Humans, AI, and Context: Understanding End-Users' Trust in a Real-World Computer Vision Application. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 77–88.
- [76] Hyeokhyen Kwon, Catherine Tong, Harish Haresamudram, Yan Gao, Gregory D Abowd, Nicholas D Lane, and Thomas Ploetz. 2020. IMUTube: Automatic extraction of virtual on-body accelerometry from video for human activity recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020), 1–29.
- [77] Nan M Laird and James H Ware. 1982. Random-effects models for longitudinal data. *Biometrics* (1982), 963–974.
- [78] Min Kyung Lee, Nina Grgić-Hlača, Michael Carl Tschantz, Reuben Binns, Adrian Weller, Michelle Carney, and Kori Inkpen. 2020. Human-centered approaches to fair and responsible AI. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–8.
- [79] Boning Li and Akane Sano. 2020. Extraction and interpretation of deep autoencoder-based temporal features from wearables for forecasting personalized mood, health, and stress. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 2 (2020), 1–26.
- [80] Jiayu Li, Zhiyu He, Yumeng Cui, Chenyang Wang, Chong Chen, Chun Yu, Min Zhang, Yiqun Liu, and Shaoping Ma. 2022. Towards Ubiquitous Personalized Music Recommendation with Smart Bracelets. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (2022), 1–34.
- [81] Q Vera Liao and Jennifer Wortman Vaughan. 2023. AI Transparency in the Age of LLMs: A Human-Centered Research Roadmap. *arXiv preprint arXiv:2306.01941* (2023).
- [82] Brian Y Lim and Anind K Dey. 2010. Toolkit to support intelligibility in context-aware applications. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*. 13–22.
- [83] Lydia T Liu, Serena Wang, Tolani Britton, and Rediet Abebe. 2023. Reimagining the machine learning life cycle to improve educational outcomes of students. *Proceedings of the National Academy of Sciences* 120, 9 (2023), e2204781120.
- [84] Jin Lu, Chao Shang, Chaoqun Yue, Reynaldo Morillo, Shweta Ware, Jayesh Kamath, Athanasios Bamis, Alexander Russell, Bing Wang, and Jinbo Bi. 2018. Joint modeling of heterogeneous sensing data for depression assessment via multi-task learning. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (2018), 1–21.
- [85] Nissa B Lucero, Renea L Beckstrand, Lynn Clark Callister, and Ana C Sanchez Birkhead. 2012. Prevalence of postpartum depression among Hispanic immigrant women. *Journal of the American Academy of Nurse Practitioners* 24, 12 (2012), 726–734.
- [86] Shuai Ma, Ying Lei, Xinru Wang, Chengbo Zheng, Chuhan Shi, Ming Yin, and Xiaojuan Ma. 2023. Who Should I Trust: AI or Myself? Leveraging Human and AI Correctness Likelihood to Promote Appropriate Trust in AI-Assisted Decision-Making. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [87] Henry B Mann and Donald R Whitney. 1947. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics* (1947), 50–60.
- [88] Kelly E Matthews, Victoria Andrews, and Peter Adams. 2011. Social learning spaces and student engagement. *Higher Education Research & Development* 30, 2 (2011), 105–120.
- [89] Cameron McCarthy, Nikhilesh Pradhan, Calum Redpath, and Andy Adler. 2016. Validation of the Empatica E4 wristband. In *2016 IEEE EMBS international student conference (ISC)*. IEEE, 1–4.
- [90] Deanna LH McFadden. 2016. Health and academic success: A look at the challenges of first-generation community college students. *Journal of the American Association of Nurse Practitioners* 28, 4 (2016), 227–232.

- [91] Lakmal Meegahapola, William Droz, Peter Kun, Amalia De Götzen, Chaitanya Nutakki, Shyam Diwakar, Salvador Ruiz Correa, Donglei Song, Hao Xu, Miriam Bidoglia, et al. 2023. Generalization and Personalization of Mobile Sensing-Based Mood Inference Models: An Analysis of College Students in Eight Countries. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (2023), 1–32.
- [92] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)* 54, 6 (2021), 1–35.
- [93] Shayan Mirjafari, Kizito Masaba, Ted Grover, Weichen Wang, Pino Audia, Andrew T Campbell, Nitesh V Chawla, Vedant Das Swain, Munmun De Choudhury, Anind K Dey, et al. 2019. Differentiating higher and lower job performers in the workplace using mobile sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 2 (2019), 1–24.
- [94] David C Mohr, Ken R Weingardt, Madhu Reddy, and Stephen M Schueller. 2017. Three problems with current digital mental health research... and three things we can do about them. *Psychiatric services* 68, 5 (2017), 427–429.
- [95] AL Morrison, S Rozak, AU Gold, and JE Kay. 2020. Quantifying student engagement in learning about climate change using galvanic hand sensors in a controlled educational setting. *Climatic Change* 159 (2020), 17–36.
- [96] Mehrab Bin Morshed, Koustuv Saha, Richard Li, Sidney K D’Mello, Munmun De Choudhury, Gregory D Abowd, and Thomas Plötz. 2019. Prediction of mood instability with passive sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 3 (2019), 1–21.
- [97] Jema David Ndibwile, Edith Talina Luhanga, Doudou Fall, Daisuke Miyamoto, and Youki Kadobayashi. 2018. Smart4gap: Factors that influence smartphone security decisions in developing and developed countries. In *Proceedings of the 2018 10th International Conference on Information Management and Engineering*. 5–15.
- [98] Martin Nemeth, Dmitrii Borkin, and German Michalconok. 2019. The comparison of machine-learning methods XGBoost and LightGBM to predict energy development. In *Computational Statistics and Mathematical Modeling Methods in Intelligent Systems: Proceedings of 3rd Computational Methods in Systems and Software 2019, Vol. 2* 3. Springer, 208–215.
- [99] Liangying Peng, Ling Chen, Zhenan Ye, and Yi Zhang. 2018. Aroma: A deep multi-task learning based simple and complex human activity recognition method using wearable sensors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (2018), 1–16.
- [100] Arcady A Putilov. 2018. Associations of depression and seasonality with morning-evening preference: comparison of contributions of its morning and evening components. *Psychiatry research* 262 (2018), 609–617.
- [101] Stephen John Quaye, Shaun R Harper, and Sumun L Pendakur. 2019. *Student engagement in higher education: Theoretical perspectives and practical approaches for diverse populations*. Routledge.
- [102] Mohammad Saiedur Rahaman, Jonathan Liono, Yongli Ren, Jeffrey Chan, Shaw Kudo, Tim Rawling, and Flora D Salim. 2020. An ambient-physical system to infer concentration in open-plan workplace. *IEEE Internet of Things Journal* 7, 12 (2020), 11576–11586.
- [103] Inioluwa Deborah Raji, Andrew Smart, Rebecca N White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 33–44.
- [104] Carson Reynolds and Rosalind Picard. 2004. Affective sensors, privacy, and ethical contracts. In *CHI’04 extended abstracts on Human factors in computing systems*. 1103–1106.
- [105] Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N Galtier, Bennett A Landman, Klaus Maier-Hein, et al. 2020. The future of digital health with federated learning. *NPJ digital medicine* 3, 1 (2020), 119.
- [106] John Rooksby, Alistair Morrison, and Dave Murray-Rust. 2019. Student perspectives on digital phenotyping: The acceptability of using smartphone data to assess mental health. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [107] Sohrab Saeb, Mi Zhang, Christopher J Karr, Stephen M Schueller, Marya E Corden, Konrad P Kording, David C Mohr, et al. 2015. Mobile phone sensor correlates of depressive symptom severity in daily-life behavior: an exploratory study. *Journal of medical Internet research* 17, 7 (2015), e4273.
- [108] Koustuv Saha, Larry Chan, Kaya De Barbaro, Gregory D Abowd, and Munmun De Choudhury. 2017. Inferring mood instability on social media by leveraging ecological momentary assessments. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 3 (2017), 1–27.
- [109] Koustuv Saha, Ted Grover, Stephen M Mattingly, Vedant Das Swain, Pranshu Gupta, Gonzalo J Martinez, Pablo Robles-Granda, Gloria Mark, Aaron Striegel, and Munmun De Choudhury. 2021. Person-centered predictions of psychological constructs with social media contextualized by multimodal sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 1 (2021), 1–32.
- [110] Daniel Salber, Anind K Dey, and Gregory D Abowd. 1999. The context toolkit: Aiding the development of context-enabled applications. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 434–441.
- [111] Bill Schilit, Norman Adams, and Roy Want. 1994. Context-aware computing applications. In *1994 first workshop on mobile computing systems and applications*. IEEE, 85–90.


- [112] Michael T Schmitt and Nyla R Branscombe. 2002. The meaning and consequences of perceived discrimination in disadvantaged and privileged social groups. *European review of social psychology* 12, 1 (2002), 167–199.
- [113] Yasaman S Sefidgar, Woosuk Seo, Kevin S Kuehn, Tim Althoff, Anne Browning, Eve Riskin, Paula S Nurius, Anind K Dey, and Jennifer Mankoff. 2019. Passively-sensed behavioral correlates of discrimination events in college students. *Proceedings of the ACM on Human-computer Interaction* 3, CSCW (2019), 1–29.
- [114] Ben Shneiderman. 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction* 36, 6 (2020), 495–504.
- [115] Ashlee L Sjogren and Theresa N Melton. 2021. The complexities of student engagement for historically marginalized youth in an after-school program. *Journal of Youth Development* 16, 5 (2021), 105–121.
- [116] Harini Suresh and John Guttag. 2021. A framework for understanding sources of harm throughout the machine learning life cycle. In *Proceedings of the 1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*. 1–9.
- [117] Harini Suresh and John V Guttag. 2019. A framework for understanding unintended consequences of machine learning. *arXiv preprint arXiv:1901.10002* 2, 8 (2019).
- [118] Mohammad Tahaei, Marios Constantinides, Daniele Quercia, Sean Kennedy, Michael Muller, Simone Stumpf, Q Vera Liao, Ricardo Baeza-Yates, Lora Aroyo, Jess Holbrook, et al. 2023. Human-Centered Responsible Artificial Intelligence: Current & Future Trends. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–4.
- [119] Kailas Vodrahalli, Roxana Daneshjou, Tobias Gerstenberg, and James Zou. 2022. Do humans trust advice more if it comes from ai? an analysis of human-ai interactions. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society*. 763–777.
- [120] Fabian Wahle, Tobias Kowatsch, Elgar Fleisch, Michael Rufer, Steffi Weidt, et al. 2016. Mobile sensing and support for people with depression: a pilot trial in the wild. *JMIR mHealth and uHealth* 4, 3 (2016), e5960.
- [121] Mingyang Wan, Daochen Zha, Ninghao Liu, and Na Zou. 2023. In-processing modeling techniques for machine learning fairness: A survey. *ACM Transactions on Knowledge Discovery from Data* 17, 3 (2023), 1–27.
- [122] Guang Wang, Yongfeng Zhang, Zhihan Fang, Shuai Wang, Fan Zhang, and Desheng Zhang. 2020. FairCharge: A data-driven fairness-aware charging recommendation system for large-scale electric taxi fleets. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 1 (2020), 1–25.
- [123] Qiaosi Wang, Michael Madaio, Shaun Kane, Shivani Kapania, Michael Terry, and Lauren Wilcox. 2023. Designing Responsible AI: Adaptations of UX Practice to Meet Responsible AI Challenges. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [124] Rui Wang, Gabriella Harari, Peilin Hao, Xia Zhou, and Andrew T Campbell. 2015. SmartGPA: how smartphones can assess and predict academic performance of college students. In *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*. 295–306.
- [125] Rui Wang, Weichen Wang, Alex DaSilva, Jeremy F Huckins, William M Kelley, Todd F Heatherton, and Andrew T Campbell. 2018. Tracking depression dynamics in college students using mobile phone and wearable sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 1 (2018), 1–26.
- [126] Zhou Wang and Alan C Bovik. 2009. Mean squared error: Love it or leave it? A new look at signal fidelity measures. *IEEE signal processing magazine* 26, 1 (2009), 98–117.
- [127] Tabitha G Wilbur and Vincent J Roscigno. 2016. First-generation disadvantage and college enrollment/completion. *Socius* 2 (2016), 2378023116664351.
- [128] Frank Wilcoxon. 1992. *Individual comparisons by ranking methods*. Springer.
- [129] Xuhai Xu, Prerna Chikersal, Afsaneh Doryab, Daniella K Villalba, Janine M Dutcher, Michael J Tuminia, Tim Althoff, Sheldon Cohen, Kasey G Creswell, J David Creswell, et al. 2019. Leveraging routine behavior and contextually-filtered features for depression detection among college students. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 3 (2019), 1–33.
- [130] Xuhai Xu, Prerna Chikersal, Janine M Dutcher, Yasaman S Sefidgar, Woosuk Seo, Michael J Tuminia, Daniella K Villalba, Sheldon Cohen, Kasey G Creswell, J David Creswell, et al. 2021. Leveraging collaborative-filtering for personalized behavior modeling: a case study of depression detection among college students. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 1 (2021), 1–27.
- [131] Xuhai Xu, Xin Liu, Han Zhang, Weichen Wang, Subigy Nepal, Yasaman Sefidgar, Woosuk Seo, Kevin S Kuehn, Jeremy F Huckins, Margaret E Morris, et al. 2023. GLOBEM: Cross-Dataset Generalization of Longitudinal Human Behavior Modeling. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (2023), 1–34.
- [132] Xuhai Xu, Han Zhang, Yasaman Sefidgar, Yiyi Ren, Xin Liu, Woosuk Seo, Jennifer Brown, Kevin Kuehn, Mike Merrill, Paula Nurius, et al. 2022. GLOBEM Dataset: Multi-Year Datasets for Longitudinal Human Behavior Modeling Generalization. *arXiv preprint arXiv:2211.02733* (2022).
- [133] Zikang Xu, Jun Li, Qingsong Yao, Han Li, Xin Shi, and S Kevin Zhou. 2022. A Survey of Fairness in Medical Image Analysis: Concepts, Algorithms, Evaluations, and Challenges. *arXiv preprint arXiv:2209.13177* (2022).

- [134] Stephen S Yau and Fariaz Karim. 2004. An adaptive middleware for context-sensitive communications for real-time applications in ubiquitous computing environments. *Real-Time Systems* 26 (2004), 29–61.
- [135] Sofia Yfantidou, Marios Constantinides, Dimitris Spathis, Athena Vakali, Daniele Quercia, and Fahim Kawsar. 2023. Beyond Accuracy: A Critical Review of Fairness in Machine Learning for Mobile and Wearable Computing. *arXiv preprint arXiv:2303.15585* (2023).
- [136] Sofia Yfantidou, Pavlos Sermpezis, Athena Vakali, and Ricardo Baeza-Yates. 2023. Uncovering Bias in Personal Informatics. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 3 (2023), 1–30.
- [137] Han Zhang, Margaret Morris, Paula Nurius, Kelly Mack, Jennifer Brown, Kevin Kuehn, Yasaman Sefidgar, Xuhai Xu, Eve Riskin, Anind Dey, et al. 2022. Impact of Online Learning in the Context of COVID-19 on Undergraduates with Disabilities and Mental Health Concerns. *ACM Transactions on Accessible Computing* 15, 4 (2022), 1–27.
- [138] Han Zhang, Leijie Wang, Yilun Sheng, Xuhai Xu, Jennifer Mankoff, and Anind K Dey. 2023. A Framework for Designing Fair Ubiquitous Computing Systems. In *Adjunct Proceedings of the 2023 ACM International Joint Conference on Pervasive and Ubiquitous Computing & the 2023 ACM International Symposium on Wearable Computing*. 366–373.
- [139] Haiyi Zhu, Bowen Yu, Aaron Halfaker, and Loren Terveen. 2018. Value-sensitive algorithm design: Method, case study, and lessons. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–23.
- [140] Brendan P Zietsch, Karin JH Verweij, Andrew C Heath, Pamela AF Madden, Nicholas G Martin, Elliot C Nelson, and Michael T Lynskey. 2012. Do shared etiological factors contribute to the relationship between sexual orientation and depression? *Psychological Medicine* 42, 3 (2012), 521–532.

A APPENDIX

A.1 Percentage of Each Group within Each Sensitive Attribute

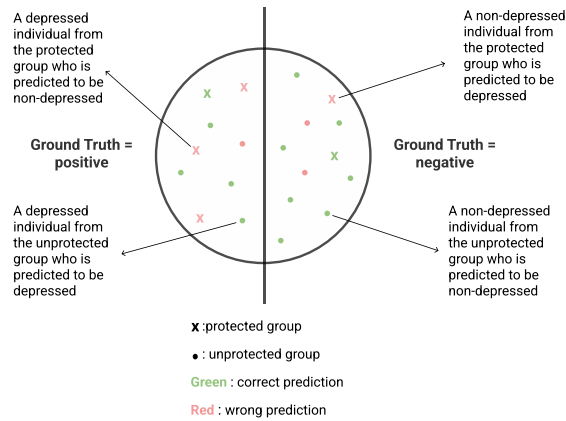
Fig. 2. Percentage of each group within each sensitive attribute. The protected group for each sensitive attribute (e.g., first-gen) is shaded in dark colors, while the unprotected group is shaded in light colors (e.g., non-first-gen). Non-male includes women, transgender individuals, and genderqueer individuals, non-heterosexual includes homosexual, bisexual, and asexual individuals, and non-white includes black, asian, latinx, and biracial.



A.2 Case Study 1: Example of the Statistical Evaluation and Experimental Implementation

Fig. 3. Example of fairness evaluation based on the disparities in accuracy, false negative rate, and false positive rate. (a) shows the synthetic data for 20 individuals, with 6 belonging to the protected group (represented by “x” marks) and 14 belonging to the unprotected group (represented by “.” marks). (b) visualizes the distribution and disparities of predictions for both groups, where correction predictions are depicted in green and incorrect predictions in red.

PID	True Label	Prediction	Protected Group	Assigned Value
1	1	0	Yes	0
2	1	1	Yes	1
3	1	1	No	1
4	0	0	No	1
5	0	0	No	1
6	1	1	No	1
7	1	1	No	1
8	0	1	Yes	0
9	0	0	No	1
10	1	0	Yes	0
11	0	1	No	0
12	0	0	Yes	1
13	1	0	Yes	0
14	1	1	No	1
15	0	0	No	1
16	0	0	No	1
17	1	0	No	0
18	0	0	No	1
19	0	1	No	0
20	0	0	No	1



(a) Synthetic data for 20 individuals.

(b) Visualization of prediction distributions and disparities.

A.2.1 Example of the Statistical Evaluation. In Figure 3, we present an illustrative example to demonstrate our approach to fairness evaluation. In this example, we generated synthesized ground-truth labels and predictions from an algorithm for a sample of 20 individuals. Among these individuals, 6 are part of the protected group, while

14 belong to the unprotected group (as shown in Figure 3a). We assigned a value of “1” for accurate predictions and “0” for inaccurate predictions based on the correctness of the predictions.

Figure 3b visualizes the distribution of predictions for both the protected group (represented by “x”) and the unprotected group (represented by “.”). The circle in the figure represents the distribution of predictions, with the left side indicating cases where all ground truth values are positive (representing individuals with depression in our case study), and the right side representing cases where all ground truth values are negative (representing individuals without depression in our case study). The accuracy of the predictions is indicated by the color, with green representing correct predictions by the algorithm and red representing incorrect predictions.

In this example, when considering the disparity in accuracy, the algorithm made incorrect predictions for 4 out of 6 individuals from the protected group (represented by the red “x” marks among all both red and green “x” marks). Conversely, for the unprotected group, the algorithm made incorrect predictions for 3 out of 14 individuals (depicted by the red “.” marks among all red and green “.” marks). To assess the statistical significance of these disparities, we conducted the Mann-Whitney U test in combination with the Benjamini-Hochberg correction. Specifically, we applied this test to the 2 “1” values and 4 “0” values corresponding to the protected group, as well as the 11 “1” values and 3 “0” values corresponding to the unprotected group.

When examining the difference in false negative rates, the relevant information for statistical analysis is contained in the left portion of the circle depicted in Figure 3b. Specifically, we conducted a statistical test on the 1 “1” value and 3 “0” values in the protected group, as well as the 4 “1” values and 1 “0” value in the unprotected group. Similarly, an evaluation of the disparity in false positive rates was conducted on the marks on the right side of the circle in Figure 3b.

A.2.2 Experimental Implementation. We applied the two evaluation criteria as defined in Section 4.1.1 to evaluate the fairness of the eight depression detection algorithms. We provide a detailed explanation of our statistical analyses to capture disparities in accuracy, false negative rate, and false positive rate below (an example of this approach can be found in above). We provide open access to our evaluation codebase to enable reference and reproducibility for future research.

To perform the Benjamini-Hochberg correction, we first calculated the p values for all attributes using the Mann-Whitney U test. Then, we arranged the p values in ascending order and assign ranks to them, with the smallest p value receiving rank 1, the second smallest receiving rank 2, and so on. Next, we calculated the adjusted q values for each individual p value using the formula: $(i/m) \times Q$, where i is the rank of the individual p value, m is the total number of tests, and Q is the false discovery rate, 0.05. Finally, we compared the original p values to the calculated q values. Attributes with a p value smaller than the corresponding q value and less than 0.05 were considered to have significant differences (which we highlighted in red in Tables 4 and 5).

To examine potential disparities across various groups of one algorithm, we employed a systematic approach. Initially, we categorized algorithm predictions based on their correctness, assigning a value of “1” to instances where the algorithm accurately predicted the ground truth and a value of “0” to instances where the algorithm falsely predicted the ground truth. Subsequently, we applied the Mann-Whitney test in conjunction with the Benjamini-Hochberg correction to different subsets of the “0” and “1” values to evaluate the following three hypotheses. First, we conducted a thorough analysis to determine whether the algorithms exhibited comparable accuracy in predicting the ground truth for both the protected and unprotected groups, aiming to evaluate potential disparities in accuracy. To achieve this, we performed the Mann-Whitney test with the Benjamini-Hochberg correction on the complete set of “0” and “1” values. Second, our assessment focused on whether the algorithms demonstrated similar false negative rates in predicting the ground truth for both the protected and unprotected groups, to identify potential disparities in false negative rates. To accomplish this, we conducted the same test on the subset of “0” and “1” values where the ground truth labels were positive. Similarly, we proceeded to investigate whether the algorithms displayed comparable false positive rates for both the protected

and unprotected groups. This was achieved by applying the same test on the subset of “0” and “1” values where the ground truth labels were negative.

A.3 Comparisons of Depression Scores for Different Groups of Four Datasets.

Fig. 4. Comparisons of depression (BDI-II) scores for different groups of four datasets. The red dotted line indicates the cutoff point (i.e., 13) for BDI-II scores, which is used to distinguish between students with at least mild depressive symptoms (BDI-II score ≥ 13) and those without (BDI-II < 13). Significance levels after Benjamini-Hochberg (B-H) correction are marked with an asterisk ($*p < 0.05$) in red on the subplot. First-gen, BA, and HET represent first-generation college students, bachelor, and heterosexual, respectively.

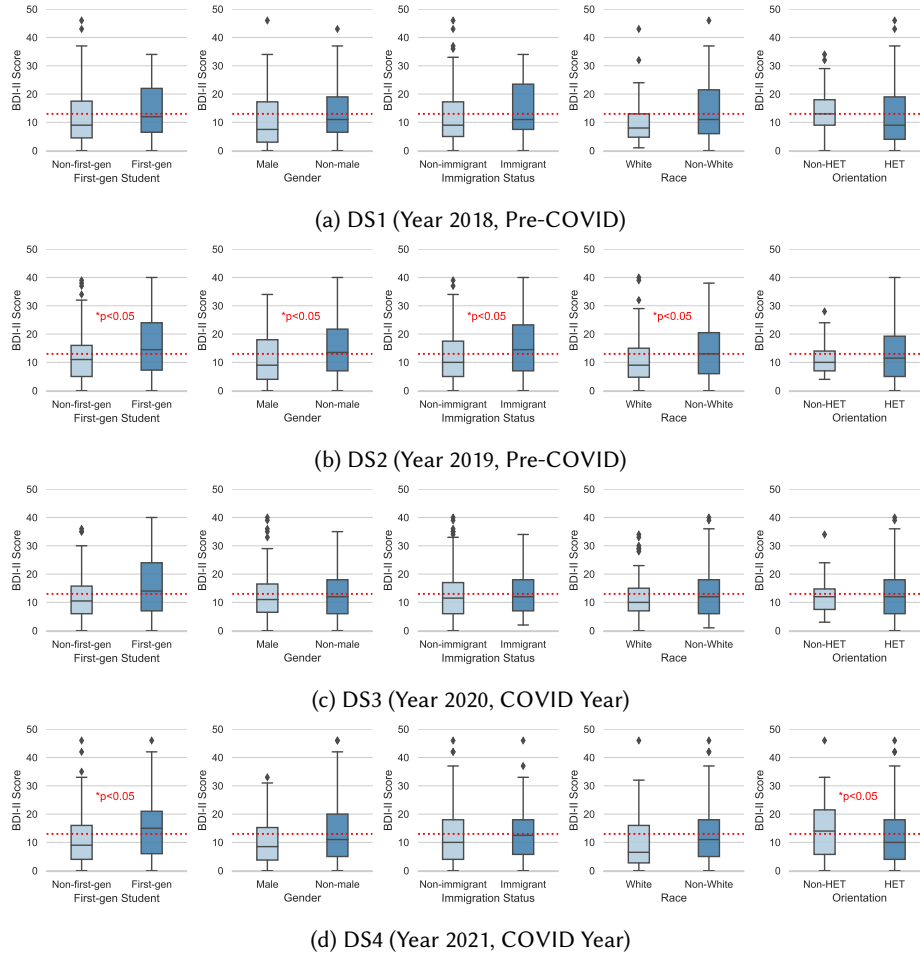


Table 9. Summary of bias changes with the addition of sensitive attributes in the training and testing process in terms of **disparity in accuracy**. This table provides an overview of bias alterations resulting from the inclusion of sensitive attributes during the training and testing processes, using disparity in accuracy as the fairness metrics. It encompasses bias amplification and reduction for each sensitive attribute across the four datasets. The comparison highlights the consequences of adding or excluding sensitive attributes in training and testing. Extra bias is denoted in red, while reduced bias is highlighted in green. For instance, considering the Xu_interpretable algorithm, Tables 4 and 5 present fairness evaluation outcomes before and after incorporating data related to first-generation college student status. When this sensitive attribute is introduced, an additional bias towards gender emerges in DS4, indicated by the label “1” in this table.

Algorithm	Added Attributes	Sensitive Attributes				
		First-gen college student	Gender	Immigration Status	Race	Sexual Orientation
Wahle <i>et al.</i>	First-gen college student	1	-1	0	0	0
	Gender	1	1	0	0	0
	Immigration Status	1	-1	0	0	0
	Race	0	-1	0	0	0
	Sexual Orientation	1	0	0	0	0
Saeb <i>et al.</i>	First-gen college student	-1	-1	-1	-1	-1
	Gender	-2	1	-1	-1	-1
	Immigration Status	0	0	-1	-1	0
	Race	0	0	0	0	0
	Sexual Orientation	-1	0	0	-1	-1
Farhan <i>et al.</i>	First-gen college student	0	0	0	0	0
	Gender	0	1	0	0	0
	Immigration Status	0	0	1	0	1
	Race	0	-1	0	1	1
	Sexual Orientation	0	1	0	0	0
Canzian <i>et al.</i>	First-gen college student	1	0	0	0	-1
	Gender	0	0	0	0	0
	Immigration Status	-1	-1	0	0	0
	Race	0	0	0	0	0
	Sexual Orientation	0	-1	2	0	0
Wang <i>et al.</i>	First-gen college student	1	0	-1	1	-1
	Gender	0	1	0	1	0
	Immigration Status	0	0	-1	1	0
	Race	0	1	-1	0	0
	Sexual Orientation	1	1	-1	0	0
Lu <i>et al.</i>	First-gen college student	0	1	0	2	1
	Gender	-1	0	0	1	0
	Immigration Status	0	0	0	1	0
	Race	0	0	0	1	0
	Sexual Orientation	-1	0	0	0	1
Xu_interpretable <i>et al.</i>	First-gen college student	0	1	0	0	0
	Gender	0	0	0	0	0
	Immigration Status	0	1	0	0	0
	Race	0	1	0	0	0
	Sexual Orientation	0	1	0	0	0
Xu_personalized <i>et al.</i>	First-gen college student	0	0	0	0	0
	Gender	0	0	0	0	0
	Immigration Status	0	0	0	0	0
	Race	0	0	0	0	0
	Sexual Orientation	0	0	0	0	0

Table 10. Summary of bias changes with the addition of sensitive attributes in the training and testing process in terms of **disparity in false negative rates**. This table provides an overview of bias alterations resulting from the inclusion of sensitive attributes during the training and testing processes, using disparity in false negative rates as the fairness metrics. It encompasses bias amplification and reduction for each sensitive attribute across the four datasets. The comparison highlights the consequences of adding or excluding sensitive attributes in training and testing. Extra bias is denoted in red, while reduced bias is highlighted in green.

Algorithm	Added Attribute	Sensitive Attribute				
		First-gen college student	Gender	Immigration Status	Race	Sexual Orientation
Wahle <i>et al.</i>	First-gen college student	2	0	0	1	0
	Gender	0	0	0	0	0
	Immigration Status	0	0	1	1	0
	Race	0	0	0	0	0
	Sexual Orientation	0	0	0	0	1
Saeb <i>et al.</i>	First-gen college student	3	2	0	2	-2
	Gender	1	2	0	1	-1
	Immigration Status	0	1	0	0	-1
	Race	0	0	0	2	-1
	Sexual Orientation	0	0	0	0	0
Farhan <i>et al.</i>	First-gen college student	1	0	-1	0	1
	Gender	0	0	-1	0	0
	Immigration Status	0	1	2	0	1
	Race	0	0	-1	0	0
	Sexual Orientation	0	0	0	0	1
Canzian <i>et al.</i>	First-gen college student	2	1	0	0	0
	Gender	0	0	0	0	0
	Immigration Status	0	0	3	1	1
	Race	1	0	0	1	0
	Sexual Orientation	0	0	0	0	1
Wang <i>et al.</i>	First-gen college student	1	0	0	1	0
	Gender	0	0	0	0	0
	Immigration Status	0	0	0	0	0
	Race	0	0	0	1	0
	Sexual Orientation	0	0	0	0	0
Lu <i>et al.</i>	First-gen college student	1	1	0	1	0
	Gender	0	2	0	0	0
	Immigration Status	-1	1	1	0	0
	Race	-1	0	0	3	0
	Sexual Orientation	-1	0	0	0	0
Xu_interpretable <i>et al.</i>	First-gen college student	-1	0	0	0	0
	Gender	0	1	0	0	0
	Immigration Status	-1	0	0	0	0
	Race	-1	0	0	0	0
	Sexual Orientation	-1	0	0	0	0
Xu_personalized <i>et al.</i>	First-gen college student	0	0	0	0	0
	Gender	0	0	0	0	0
	Immigration Status	0	0	0	0	0
	Race	0	0	0	0	0
	Sexual Orientation	0	0	0	0	0

Table 11. Summary of bias changes with the addition of sensitive attributes in the training and testing process in terms of **disparity in false positive rates**. This table provides an overview of bias alterations resulting from the inclusion of sensitive attributes during the training and testing processes, using disparity in false positive rates as the fairness metrics. It encompasses bias amplification and reduction for each sensitive attribute across the four datasets. The comparison highlights the consequences of adding or excluding sensitive attributes in training and testing. Extra bias is denoted in red, while reduced bias is highlighted in green.

Algorithm	Added Attribute	Sensitive Attribute				
		First-gen college student	Gender	Immigration Status	Race	Sexual Orientation
Wahle <i>et al.</i>	First-gen college student	3	0	1	1	0
	Gender	-1	1	0	0	0
	Immigration Status	0	0	1	0	0
	Race	1	0	0	2	0
	Sexual Orientation	0	0	0	1	1
Saeb <i>et al.</i>	First-gen college student	4	-1	0	1	0
	Gender	1	2	1	2	1
	Immigration Status	1	1	2	3	1
	Race	1	0	1	1	0
	Sexual Orientation	0	-1	0	0	1
Farhan <i>et al.</i>	First-gen college student	1	1	0	1	0
	Gender	0	2	0	1	-1
	Immigration Status	0	0	1	-1	-1
	Race	0	0	0	-1	-1
	Sexual Orientation	0	1	0	-1	-1
Canzian <i>et al.</i>	First-gen college student	1	-1	1	0	-1
	Gender	0	0	0	0	0
	Immigration Status	-1	-1	3	1	-1
	Race	0	0	0	-1	0
	Sexual Orientation	-1	-1	1	0	0
Wang <i>et al.</i>	First-gen college student	2	0	0	0	-1
	Gender	0	0	0	0	-1
	Immigration Status	0	0	1	0	0
	Race	1	1	0	1	-1
	Sexual Orientation	0	0	0	0	-1
Lu <i>et al.</i>	First-gen college student	3	1	-1	1	1
	Gender	0	3	-1	0	-1
	Immigration Status	1	0	1	1	-1
	Race	2	0	0	2	-1
	Sexual Orientation	0	0	0	0	0
Xu_interpretable <i>et al.</i>	First-gen college student	0	1	0	0	0
	Gender	0	0	0	0	0
	Immigration Status	1	1	0	0	0
	Race	1	1	0	0	0
	Sexual Orientation	1	1	0	0	0
Xu_personalized <i>et al.</i>	First-gen college student	0	0	0	-1	0
	Gender	0	0	0	-1	0
	Immigration Status	0	0	0	-1	0
	Race	0	0	0	-1	0
	Sexual Orientation	0	0	0	-1	0