



A Framework for Designing Fair Ubiquitous Computing Systems

Han Zhang
micohan@cs.washington.edu
University of Washington
Seattle, USA

Leijie Wang
leijiew@cs.washington.edu
University of Washington
Seattle, USA

Yilun Sheng
ylsheng@cs.washington.edu
University of Washington
Seattle, USA

Xuhai Xu
xuhaixu@uw.edu
University of Washington
Seattle, USA

Jennifer Mankoff
jmankoff@cs.washington.edu
University of Washington
Seattle, USA

Anind K. Dey
anind@uw.edu
University of Washington
Seattle, USA

ABSTRACT

Over the past few decades, ubiquitous sensors and systems have been an integral part of humans' everyday life. They augment human capabilities and provide personalized experiences across diverse contexts such as healthcare, education, and transportation. However, the widespread adoption of ubiquitous computing has also brought forth concerns regarding fairness and equitable treatment. As these systems can make automated decisions that impact individuals, it is essential to ensure that they do not perpetuate biases or discriminate against specific groups. While fairness in ubiquitous computing has been an acknowledged concern since the 1990s, it remains understudied within the field. To bridge this gap, we propose a framework that incorporates fairness considerations into system design, including prioritizing stakeholder perspectives, inclusive data collection, fairness-aware algorithms, appropriate evaluation criteria, enhancing human engagement while addressing privacy concerns, and interactive improvement and regular monitoring. Our framework aims to guide the development of fair and unbiased ubiquitous computing systems, ensuring equal treatment and positive societal impact.

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing; Ubiquitous and mobile computing theory, concepts and paradigms; Ubiquitous and mobile computing design and evaluation methods.**

KEYWORDS

Fairness; Ubiquitous Computing Systems; Framework

ACM Reference Format:

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 (UbiComp/ISWC '23 Adjunct)*, October 08–12,

2023, Cancun, Quintana Roo, Mexico. ACM, New York, NY, USA, 8 pages.
<https://doi.org/10.1145/3594739.3610677>

1 INTRODUCTION

The proliferation of smart devices and the growth of the Internet of Things (IoT) have revolutionized the way we interact with technology and the world around us. These devices have made it increasingly feasible to capture human behavior passively and continuously, opening up new avenues for research and applications. Data of various types collected by these smart devices (e.g., smartphones, smartwatches, and fitness trackers) have been used to design machine learning (ML) algorithms and systems that assess mental health [41, 48, 55, 58], monitor cognitive load and work performance [13, 42], recognize human activities [29, 43], and track academic performance [8, 27, 56]. Ultimately, these systems play a pivotal role in aiding the decision-making process.

However, despite the potential for these algorithms and systems to improve people's lives, they may also introduce harm due to a lack of fairness considerations. The design of algorithmic decision-making systems can incorporate discriminatory biases, leading to unjust outcomes, especially for marginalized communities already facing societal inequalities [6, 44]. Therefore, it is crucial to ensure that the benefits of these systems are available to all, and that no group is unfairly disadvantaged. By addressing issues of fairness, we can deploy these technologies with better fairness, transparency, and accountability.

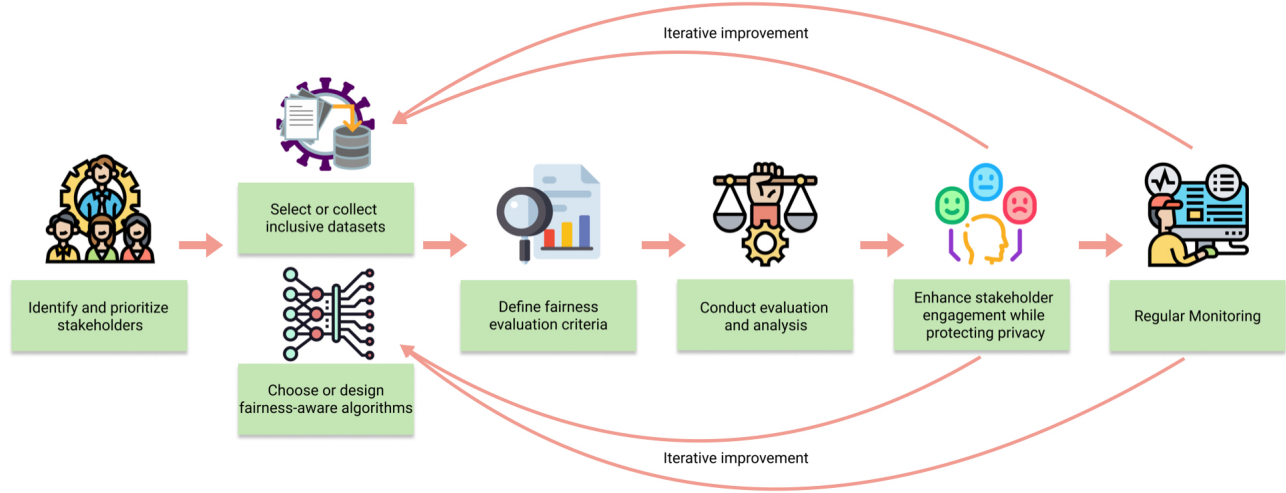
While fairness in ubiquitous computing has been acknowledged as a pressing concern since the 1990s [57], research on fairness in the field of mobile, wearable, and ubiquitous computing has received comparatively less attention than in the field of machine learning. Recent work shows that only 5% of papers published in the Proceedings of the ACM Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT) journal from 2018 to 2022 included assessments of fairness, while most other studies focused on accuracy or error metrics [61]. The lack of attention to fairness in ubiquitous computing may be caused by several factors. First, unlike researchers in the machine learning fairness community, who typically use large-scale and publicly available datasets, ubiquitous computing researchers often collect small-scale and private author-collected datasets, which poses challenges to the design of common metrics and benchmarks for assessing fairness [59]. Second, the sequential and dynamic nature of the data collected in ubiquitous computing (e.g., steps and sleep patterns) makes it more difficult to identify biases from the surface compared to the



This work is licensed under a Creative Commons Attribution International 4.0 License.

UbiComp/ISWC '23 Adjunct, October 08–12, 2023, Cancun, Quintana Roo, Mexico
© 2023 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0200-6/23/10.
<https://doi.org/10.1145/3594739.3610677>

Figure 1: Overview of the framework for designing fair ubiquitous computing systems.



more structured and static tabular format of data (e.g., language corpus from social media) used by the machine learning fairness community [61]. Finally, the limited body of research on fairness in ubiquitous computing is also compounded by the lack of guidance for designing fair ubiquitous computing systems. Without such guidance, it is difficult to ensure that technologies designed for ubiquitous computing are being used in a fair and ethical manner.

In this work, we propose a framework for designing fair systems in the context of ubiquitous computing. As illustrated in Figure 1, our framework starts by identifying relevant stakeholders in different contexts, determining who will use the algorithmic tool and who will be impacted by the algorithmic outcomes. Next, researchers gather appropriate and inclusive datasets and select or develop algorithms for evaluation. Evaluation criteria are then defined, incorporating fairness metrics and other performance measures, with a clear rationale for their selection (e.g., why it is more reasonable to use the disparity of a traditional performance metric, such as false negative rate, as a fairness metric in a specific context). Algorithms are evaluated using the defined criteria and fairness metrics, and the results are analyzed within the context of stakeholder priorities. Transparency is ensured by communicating the findings and soliciting feedback from stakeholders. Finally, to adapt to the dynamic nature of ubiquitous computing, the systems are iteratively improved, and regularly monitored.

The contributions of this work are as follows. Firstly, we emphasize the significance of integrating fairness considerations into ubiquitous computing. Secondly, we present a framework that guides the design of fair ubiquitous computing systems. Lastly, we provide a detailed rationale for each component of the proposed framework. Our intention for this work is to serve as a valuable resource for future endeavors aiming to incorporate fairness into the design of ubiquitous computing systems.

2 BACKGROUND AND RELATED WORK

In this section, we first provide a broader review of prior work on fairness (Section 2.1). We then review the existing fairness literature in ubiquitous computing (Section 2.2).

2.1 A Broader Review of Existing Fairness Literature

As ML/AI is now being used in many decision-making systems, concerns have been raised about the fairness of these systems [6, 62]. In response to these concerns, Mehrabi *et al.* conducted a comprehensive systematic review of prior research, examining various sources of biases that can impact AI applications [40]. Their study identified two key sources of unfairness in machine learning outcomes: biases originating from the data and biases arising from the algorithms themselves. To mitigate biases stemming from the data, researchers have proposed the adoption of inclusive benchmark datasets [9]. These datasets aim to enhance the representation and diversity of the training data, thereby reducing the potential for biased and discriminatory outcomes in machine learning models.

Meanwhile, researchers have made substantial contributions to laying the foundation for understanding and mitigating algorithmic biases [16, 24, 62]. This work has resulted in the proposal of influential notions and frameworks aimed at promoting fairness in algorithmic decision-making systems.

One influential notion is **fairness through awareness**, proposed by Dwork *et al.* [16]. This notion emphasizes the consideration of **individual fairness**, which suggests treating similar individuals similarly. The authors highlighted the importance of ensuring ML models make consistent decisions and avoid discrimination based on protected attributes (such as race, gender, or age). Dwork *et al.* [16] also discussed **group fairness**, which focuses on fairness at the group level. They emphasized that the demographics of those receiving positive or negative classifications should align with the overall population. The concepts of individual fairness and

group fairness contribute to the broader understanding and pursuit of fairness in machine learning and decision-making systems.

To enforce algorithmic fairness, researchers have proposed various mathematical formulations and frameworks. Zafar *et al.* [62] explored the notion of **disparate mistreatment**, where ground truth is available for historical decisions used during the training phase. They provided a mathematical formulation for incorporating fairness criteria into the training process of machine learning models. This work allows practitioners to define and optimize fairness goals during modeling training, aiming to reduce bias in the resulting predictions. Another widely adopted framework, introduced by Hardt *et al.* [24], is **equality of opportunity**. This framework highlights the importance of equalizing the true positive rates across different demographic groups to ensure fairness. The authors provided theoretical analysis and practical algorithms for achieving equality of opportunity, and they demonstrated the effectiveness of their approach through empirical evaluations.

2.2 Existing Ubiquitous Computing Fairness Literature

While fairness research has made significant progress in addressing bias and discrimination in the machine learning fairness community, it is also crucial to consider these issues within the context of ubiquitous computing. Ubiquitous computing, characterized by the integration of computing power into everyday environments [1, 2, 14], presents unique challenges and opportunities for ensuring fairness. For example, one specific challenge arises from the dynamic nature of datasets collected in ubiquitous computing, which often makes it difficult to obtain an accurate understanding of the environment and capture every individual's needs. Consequently, biases and discrimination can be perpetuated and even amplified, leading to unfair outcomes for individuals. For instance, consider a health and fitness app that aims to provide personalized health recommendations based on user data collected from various wearable sensors. If the app's models are trained on outdated or inconsistent data due to irregular updates from users, it might provide inaccurate recommendations, potentially disadvantaging certain individuals.

To address these concerns, researchers have recently begun to study fairness specifically in the field of ubiquitous computing. For instance, researchers conducted a systematic review of papers published in the IMWUT journal between 2018 and 2022 on algorithmic fairness [61]. They found that only 5% of the published papers included fairness reports, indicating a need for more attention to fairness in the field of ubiquitous computing. Recently, researchers have started designing fairness-aware ubiquitous computing systems. One such work proposed a method that combines personalized federated learning with hierarchical clustering techniques to enhance the accuracy, robustness, and fairness of activity recognition systems [33]. By leveraging hierarchical clustering based on both activity similarity and user similarity, personalized models are created within smaller user groups, ensuring equitable treatment of individuals within the system. In another study, researchers developed a data-driven fairness-aware charging recommendation system specifically designed for large-scale electric taxi fleets [54]. The authors employed data-driven techniques and incorporated various self-defined fairness metrics (*e.g.*, reduction of traveling

time and charging station occupation rates) to recommend charging stations to electric taxi drivers, aiming to ensure fairness in the allocation of resources.

3 TOWARDS FAIR UBIQUITOUS COMPUTING

In this section, we provide a summary of the challenges and limitations that must be addressed to achieve fair ubiquitous computing (Section 3.1). This is followed by the objective of this work (Section 3.2).

3.1 Challenges and Limitations towards Fair Ubiquitous Computing

Despite the emerging body of research that has started to investigate the fairness of systems in ubiquitous computing, as the example discussed in Section 2.2, there are still several challenges and inherent limitations due to the unique nature of ubiquitous computing. Below, we identify six challenges and limitations and emphasize the need for addressing them.

Limited consideration of sensitive attributes in specific contexts. In contrast to the machine learning fairness community, which has explored a broad spectrum of sensitive attributes such as race, sexuality, disability, and nationality [4, 5, 10, 25, 53], recent studies have highlighted a disparity in the focus of fairness work within the field of ubiquitous computing. Over the past five years, the research in ubiquitous computing has predominantly concentrated on gender and age attributes [21, 22, 64], with these attributes being mentioned in almost 90% of the included papers [61]. On the other hand, many sensitive attributes, which are often associated with discrimination, have received limited attention in the ubiquitous computing literature [31, 38, 63]. Only a few of these papers discuss sensitive attributes through the lens of fairness.

For instance, one group of stakeholders that has been overlooked in the literature is individuals with disabilities [18, 52]. Their unique needs require specific attention to ensure fair and inclusive ubiquitous computing experiences. Sexual orientation [23] is another important attribute that should be considered in fairness work for ubiquitous computing when modeling mental health.

Given the wide implementation of ubiquitous computing across diverse contexts, the limited focus on different groups of stakeholders who may experience disproportionate effects from automated decision-making systems raises concerns and emphasizes the need for a broader consideration of these marginalized groups. To advance fairness and inclusively in ubiquitous computing, it is essential to expand the research scope and give careful consideration to identifying and prioritizing stakeholders in the specific context.

Issue of Data Bias. Similar to the machine learning fairness community, bias in datasets used for training models is a significant concern within ubiquitous computing community [40, 61]. These biases can stem from various factors, such as the data collection process, sampling techniques, or the inherent societal biases present in the data. Additionally, as ubiquitous computing systems extensively use data gathered from diverse sources, including sensors, mobile devices, and online platforms [1, 28], these data sources often suffer from sparsity and uneven distribution, which can introduce biases and perpetuate existing inequities.

The scarcity or imbalance of data can result in underrepresentation of certain groups, leading to biased algorithms and discriminatory outcomes. Therefore, it is crucial to address these challenges in data collection and ensure that data used in ubiquitous computing systems is representative, diverse, and less inherently biased.

Limited Fairness-aware Algorithms in Ubiquitous Computing Systems. Despite an emerging body of research that has started to investigate algorithmic fairness and mitigate bias in algorithms (e.g., [24, 62]), the development of fair algorithms for ubiquitous computing systems lags significantly behind. Based on the comprehensive review of fairness papers published in IMWUT between 2018 and 2022, it is evident that only a limited number of studies (three papers [50, 51, 65]) have explored the incorporation of fairness considerations into machine learning algorithms during the training process [61].

This can be attributed to the distinct challenges presented by ubiquitous computing, setting it apart from other domains. For instance, ubiquitous computing systems rely heavily on contextual information, including user location, preferences, and social interactions. Contextual factors can introduce additional complexities when determining fairness, as fairness considerations might vary based on the specific context and user characteristics. Moreover, ubiquitous computing systems often make real-time or near-real-time decisions based on continuous time-series data. Ensuring algorithmic fairness in these dynamic, time-sensitive scenarios adds a layer of complexity.

Lack of Context-aware Evaluation Criteria. In contrast to the machine learning fairness community, which commonly relies on standard fairness metrics such as demographic parity [7], equalized odds [62], and equal opportunity [24], the ubiquitous computing community often employs performance metrics such as accuracy and error rate [61] without providing explicit justifications. Given the datasets used in the machine learning fairness community are static, while datasets used in ubiquitous computing are context-specific and sequential, it is crucial for ubiquitous computing researchers to carefully select appropriate fairness metrics that align with their specific contexts.

As an illustration, consider the case of modeling depressive behavior. In this instance, employing the disparity of false negative rates between prioritized stakeholders and others as a fairness metric proves more suitable than utilizing commonly employed metrics such as demographic parity and equalized odds commonly used in the machine learning fairness community. This selection is substantiated by existing research that demonstrates higher levels of mental health concerns among prioritized stakeholders, such as females with depressive symptoms [19, 35, 39]. By focusing on the disparity of false negative rates, which quantifies the variations in misclassification rates for individuals with depressive symptoms, researchers can more accurately capture the specific challenges and disparities faced by the prioritized stakeholders. In contrast, adopting standard machine learning fairness metrics, which primarily strive for equal treatment across groups by assuming comparable levels of depressive symptoms across genders in this example, may fail to adequately address the nuanced needs and disparities inherent to this particular context.

Additionally, another challenge in both the ubiquitous computing and the machine learning communities, *i.e.*, how to define a

threshold that determines the point at which models are classified as unfair across different groups, should be addressed [61]. This absence of a well-defined threshold for determining fairness in ubiquitous computing models makes it challenging to assess whether a model is truly fair and not simply due to chance.

Lack of Transparency and Explainability to stakeholders. Fairness is closely intertwined with the principles of transparency and explainability. As AI-powered systems play an increasingly significant role in consequential decision-making, their explainability becomes essential for end-users to make informed and accountable decisions [17, 34]. Over the past few years, significant advancements have been made in addressing this aspect within the machine learning fairness and human-computer interaction communities.

For instance, machine learning researchers have deliberately opted for certain machine learning models, such as decision trees and linear models, which possess transparent structures and inherently provide interpretable explanations for their problems (e.g., [11, 30]). Additionally, there has been substantial progress in developing techniques for model interpretability and explainability such as LIME [46] and SHAP [36, 37].

HCI researchers have been focusing on designing machine learning models' outputs in a transparent and understandable manner for end-users to bridge the gap between complex machine learning models and human comprehension (e.g., [3, 66]). Additionally, researchers have been investigating ways to incorporate user feedback and control mechanisms in machine learning models to enhance transparency and trust [45].

In contrast to other research communities, the transparency and explainability of existing work in the field of ubiquitous computing pose unique challenges. One primary reason is that ubiquitous computing systems often operate on vast and diverse datasets comprising heterogeneous and time-series data from multiple sources. The processing and analysis of such intricate data necessitate the utilization of sophisticated algorithms and models, rendering it arduous to elucidate their underlying rationale in a clear and interpretable manner. Furthermore, these systems commonly handle sensitive user data, including personal health information [32], location data [58], or behavioral patterns [59]. Maintaining privacy and security is of utmost importance in such contexts, often entailing practices such as data anonymization, access restrictions, and the careful selection of interviewees. However, this pursuit of privacy introduces a trade-off between fairness and privacy [12], further exacerbating the complexity and challenges of transparency and explainability.

Need for Regular Monitoring. In contrast to other domains where data may be relatively static, ubiquitous computing systems require regular monitoring for fairness due to their dynamic and adaptive nature [14, 49]. These systems collect data dynamically in real-time, which introduces the need for ongoing monitoring to detect any biases that may emerge as the system adapts and learns from new information. Contextual factors such as location, time, and social surroundings play a significant role in ubiquitous computing systems, influencing their behavior. Regular fairness monitoring becomes essential to ensure fair treatment across different contexts and prevent biases from impacting users. Additionally, ubiquitous computing systems often make real-time decisions and

interact directly with users. Monitoring helps identify any biases or discriminatory patterns in these decisions and interactions, enabling corrective actions to be taken promptly. Furthermore, the adaptability of ubiquitous computing systems and their potential to amplify biases emphasize the necessity of regular monitoring to ensure that biases are not perpetuated or amplified, safeguarding fairness and equitable outcomes for users.

3.2 Objective of This Work

Our work is motivated by the challenges and limitations pertaining to fairness in ubiquitous computing. Our objective of this work is to adapt and integrate the existing frameworks and concepts (reviewed in Section 2.1) into a specialized framework tailored for ubiquitous computing. Through this work, we aim to advance the development and deployment of ubiquitous computing systems that prioritize fairness and effectively cater to the diverse needs of stakeholders. By advocating for the integration of fairness considerations into system design, we seek to pave the way for future research towards fair ubiquitous computing.

4 FRAMEWORK FOR DESIGNING FAIR UBIQUITOUS COMPUTING SYSTEMS

In this section, we present an overview of our proposed framework (shown in Figure 1), designed specifically to address the challenges and limitations discussed in the preceding section (Section 4.1). Additionally, we delve into potential avenues for future research and development (Section 4.2).

4.1 Overview of Proposed Framework

Our framework has six components in total, each representing one important stage in ubiquitous computing system implementation:

- (1) **Identify and prioritize stakeholders.** Identify relevant stakeholders, such as those who will use the systems, as well as those who will be affected and biased by the algorithmic outcomes. For example, in the context of education, students with minoritized identities may face biased outcomes, making them important stakeholders. Transparency is crucial for both students and instructors who will use the systems. In contrast to the existing prevalent ubiquitous computing fairness literature, which often concentrates on a restricted range of sensitive attributes, our framework highlights the crucial importance of meticulously considering diverse sensitive attributes within varying contexts, *e.g.*, taking sexual minority status into account when modeling mental health.
- (2a) **Select/collect inclusive datasets.** Collect or select representative datasets that include prioritized stakeholders for evaluation, based on the specific context. In comparison to the prevailing ubiquitous computing fairness literature that frequently relies on unrepresentative datasets, our framework underscores the importance of gathering more inclusive datasets. One example is the recently published GLOBEM dataset [60], where researchers intentionally oversampled diverse subpopulation groups (*e.g.*, gender, race, and immigration).
- (2b) **Choose/design fairness-aware algorithms.** Carefully select or design explainable fairness-aware algorithms that are

relevant to the identified context and align with the goals of the evaluation. In contrast to the prevailing ubiquitous computing fairness literature, which often neglects the consideration of fairness during the algorithm design process, our framework emphasizes the need for developing fairness-aware algorithms. Note that, step 2a and step 2b can be interchanged.

- (3) **Define evaluation criteria.** Determine appropriate fairness metrics and thresholds to quantify differences and biases across stakeholder groups. Provide explicit justifications for the chosen metrics and thresholds. For example, in the evaluation of a depression detection algorithm, the disparity in false negative rates across different groups serves as a critical fairness metric. This metric reflects the algorithm's failure to identify depression in certain populations, making it essential to address. To ensure that observed disparities are not merely due to chance occurrences, a statistical test can be conducted. In contrast to the existing literature on fairness in ubiquitous computing, which often lacks justification for the selected fairness criteria, our framework emphasizes the necessity of designing context-aware evaluation criteria.
- (4) **Conduct evaluation and analysis.** Use the selected fairness metrics to evaluate algorithms on selected datasets, and analyze the results based on the predetermined evaluation criteria within the specific context. Additionally, thoroughly discuss the potential harm to stakeholders that may arise from the algorithmic decisions. In contrast to the existing literature on fairness in ubiquitous computing, which often neglects the discourse on the harm to stakeholders, our proposed framework seeks to bridge this gap.
- (5) **Enhance stakeholder engagement while protecting privacy.** Communicate evaluation findings, recommendations, and potential limitations to stakeholders, fostering transparency, accountability, and stakeholder involvement in algorithmic decision-making. Safeguard stakeholder privacy throughout the process. In comparison to the prevailing literature on fairness in ubiquitous computing, which often overlooks the involvement of humans in the system design, our framework emphasizes the importance of including human perspectives in the design process. Additionally, our framework recognizes the significance of striking a balance between fairness and privacy considerations.
- (6) **Iterative improvement and regular monitoring.** Refine algorithms to address potential biases and unfairness. Iterate on algorithmic design, data collection, and preprocessing to enhance fairness. Continuously monitor real-world performance and update algorithms and evaluation processes to align with evolving fairness standards and best practices. In comparison to the prevailing literature on fairness in ubiquitous computing, which often lacks this step, our framework emphasizes the importance of iterative improvement and regular monitoring.

4.2 Potential Directions for Future Work

In this section, we outline potential directions for future work and extensions of the proposed framework for designing fair ubiquitous

computing systems. These avenues of research offer opportunities to advance the framework and enhance its applicability in real-world contexts.

4.2.1 Validation and Case Studies. To further validate and demonstrate the effectiveness of the proposed framework, future work can focus on conducting validation studies and case studies in real-world scenarios. For example, researchers can apply the framework to specific contexts and assess its practicality. By selecting representative use cases, researchers can demonstrate how the framework can be implemented and tailored to address fairness challenges in different scenarios. Case studies can involve evaluating or deploying fair algorithms in healthcare settings and educational environments. The findings from these case studies will provide valuable insights into the framework's feasibility, efficacy, and adaptability across diverse ubiquitous computing application domains. Moreover, to enrich the validation process, future work should include interviewing identified stakeholders to gather their feedback on the framework and the results of fairness testing in various scenarios. This qualitative feedback can provide additional context and perspectives, further refining and validating the proposed framework. By conducting these studies, it can also contribute to the identification of potential challenges and limitations of the framework.

4.2.2 Balancing Privacy and Fairness. As ubiquitous computing systems rely on collecting and analyzing vast amounts of personal data to make algorithmic decisions, addressing the tradeoff between privacy and fairness stands as a significant future direction within the context of the proposed framework for designing fair ubiquitous computing systems. One possible avenue for future research lies in the exploration of strategies and methodologies to reconcile the inherent tension between privacy and fairness. Researchers can develop privacy-preserving algorithms and techniques tailored specifically to ubiquitous computing environments and delve into privacy-enhancing technologies, such as secure multi-party computation [20, 26], federated learning [47], and differential privacy [15], and incorporate them into the framework. Another essential avenue for future investigation lies in understanding the impact of various privacy-preserving mechanisms on the fairness of algorithmic decision-making. Researchers can undertake a comprehensive examination of the tradeoff between privacy and fairness, critically analyzing how privacy-enhancing measures may influence the accuracy, reliability, and equity of algorithmic outcomes.

5 CONCLUSION

In conclusion, this work introduces a novel framework for designing fair ubiquitous computing systems, addressing the existing gap in the ubiquitous computing literature on fairness challenges. By presenting this framework, we contribute to the advancement of ubiquitous computing research, with the hope of providing a valuable resource for researchers and practitioners striving to develop more equitable and inclusive systems in the future. We envision that the proposed framework can serve as an initial stepping stone towards fostering fairness and ensuring that ubiquitous computing systems align with ethical principles and societal values.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. EDA-2009977 and the University of Washington College of Engineering, Department of Electrical and Computer Engineering, and the Paul G. Allen School of Computer Science and Engineering.

REFERENCES

- [1] Gregory D Abowd, Anind K Dey, Peter J Brown, Nigel Davies, Mark Smith, and Pete Steggles. 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.
- [2] Gregory D Abowd and Elizabeth D Mynatt. 2000. Charting past, present, and future research in ubiquitous computing. *ACM Transactions on Computer-Human Interaction (TOCHI)* 7, 1 (2000), 29–58.
- [3] Chadia Abras, Diane Maloney-Krichmar, Jenny Preece, et al. 2004. User-centered design. *Bainbridge, W. Encyclopedia of Human-Computer Interaction*. Thousand Oaks: Sage Publications 37, 4 (2004), 445–456.
- [4] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. 2018. A reductions approach to fair classification. In *International Conference on Machine Learning*. PMLR, 60–69.
- [5] McKane Andrus and Sarah Villeneuve. 2022. Demographic-reliant algorithmic fairness: characterizing the risks of demographic data collection in the pursuit of fairness. In *2022 ACM Conference on Fairness, Accountability, and Transparency*. 1709–1721.
- [6] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2017. Fairness in machine learning. *Nips tutorial 1* (2017), 2017.
- [7] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *California law review* (2016), 671–732.
- [8] Mehrab Bin Morshed, Koustuv Saha, Munmun De Choudhury, Gregory D Abowd, and Thomas Plötz. 2020. Measuring self-esteem with passive sensing. In *Proceedings of the 14th EAI International Conference on Pervasive Computing Technologies for Healthcare*. 363–366.
- [9] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [10] Toon Calders and Sicco Verwer. 2010. Three naive bayes approaches for discrimination-free classification. *Data mining and knowledge discovery* 21 (2010), 277–292.
- [11] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*. 1721–1730.
- [12] Hongyan Chang and Reza Shokri. 2021. On the privacy risks of algorithmic fairness. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 292–303.
- [13] 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.
- [14] Anind K Dey. 2001. Understanding and using context. *Personal and ubiquitous computing* 5 (2001), 4–7.
- [15] 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.
- [16] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*. 214–226.
- [17] 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.
- [18] Pier Luigi Emiliani and Constantine Stephanidis. 2005. Universal access to ambient intelligence environments: Opportunities and challenges for people with disabilities. *IBM Systems Journal* 44, 3 (2005), 605–619.
- [19] 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.
- [20] Oded Goldreich. 1998. Secure multi-party computation. *Manuscript. Preliminary version* 78, 110 (1998).
- [21] Yanbin Gong, Qian Zhang, Bobby HP NG, and Wei Li. 2022. BreathMentor: Acoustic-based Diaphragmatic Breathing Monitor System. *Proceedings of the*

- ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 2 (2022), 1–28.
- [22] Bhanu Teja Gullapalli, Stephanie Carreiro, Brittany P Chapman, Deepak Ganesan, Jan Sjoquist, and Tauhidur Rahman. 2021. Optrack: a wearable-based clinical opioid use tracker with temporal convolutional attention networks. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 5, 3 (2021), 1–29.
 - [23] Samantha Guz, Shanna K Kattari, Brittanie Atteberry-Ash, Cary L Klemmer, Jarrod Call, and Leonardo Kattari. 2021. Depression and suicide risk at the cross-section of sexual orientation and gender identity for youth. *Journal of Adolescent Health* 68, 2 (2021), 317–323.
 - [24] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems* 29 (2016).
 - [25] Faisal Kamiran, Toon Calders, and Mykola Pechenizkiy. 2010. Discrimination aware decision tree learning. In *2010 IEEE international conference on data mining*. IEEE, 869–874.
 - [26] 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.
 - [27] Minhuyng Kim, Inyeop Kim, and Uichin Lee. 2021. Beneficial neglect: Instant message notification handling behaviors and academic performance. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 1 (2021), 1–26.
 - [28] John Krumm. 2018. *Ubiquitous computing fundamentals*. CRC Press.
 - [29] Hyeokhyen Kwon, Catherine Tong, Harish Haresamudram, Yan Gao, Gregory D Abowd, Nicholas D Lane, and Thomas Plötz. 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.
 - [30] Himabindu Lakkaraju, Stephen H Bach, and Jure Leskovec. 2016. Interpretable decision sets: A joint framework for description and prediction. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1675–1684.
 - [31] Nicol Turner Lee. 2018. Detecting racial bias in algorithms and machine learning. *Journal of Information, Communication and Ethics in Society* 16, 3 (2018), 252–260.
 - [32] 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.
 - [33] Youpeng Li, Xuyu Wang, and Lingling An. 2023. Hierarchical Clustering-based Personalized Federated Learning for Robust and Fair Human Activity Recognition. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 7, 1 (2023), 1–38.
 - [34] 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).
 - [35] 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.
 - [36] Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. 2020. From local explanations to global understanding with explainable AI for trees. *Nature machine intelligence* 2, 1 (2020), 56–67.
 - [37] Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30 (2017).
 - [38] Thomas Manzini, Yao Chong Lim, Yulia Tsvetkov, and Alan W Black. 2019. Black is to criminal as caucasian is to police: Detecting and removing multiclass bias in word embeddings. *arXiv preprint arXiv:1904.04047* (2019).
 - [39] 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.
 - [40] 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.
 - [41] 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.
 - [42] Subigya Nepal, Gonzalo J Martinez, Shayan Mirjafari, Koustuv Saha, Vedant Das Swain, Xuhai Xu, Pino G Audia, Munmun De Choudhury, Anind K Dey, Aaron Striegel, et al. 2022. A Survey of Passive Sensing in the Workplace. *arXiv preprint arXiv:2201.03074* (2022).
 - [43] 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.
 - [44] Dana Pessach and Erez Shmueli. 2022. A review on fairness in machine learning. *ACM Computing Surveys (CSUR)* 55, 3 (2022), 1–44.
 - [45] Jennifer Preece, Helen Sharp, and Yvonne Rogers. 2015. *Interaction design: beyond human-computer interaction*. John Wiley & Sons.
 - [46] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “Why should I trust you?” Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.
 - [47] 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.
 - [48] 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.
 - [49] Albrecht Schmidt, Michael Beigl, and Hans-W Gellersen. 1999. There is more to context than location. *Computers & Graphics* 23, 6 (1999), 893–901.
 - [50] Farhana Shahid, Wasifur Rahman, M Saifur Rahman, Sharmin Akther Purabi, Ayesha Seddiqua, Moin Mostakim, Farhan Feroz, Tanjir Rashid Soron, Fahmida Hossain, Nabila Khan, et al. 2020. Leveraging free-hand sketches for potential screening of PTSD. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 3 (2020), 1–22.
 - [51] Taoran Sheng and Manfred Huber. 2020. Weakly supervised multi-task representation learning for human activity analysis using wearables. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 2 (2020), 1–18.
 - [52] Constantine Stephanidis, Gavriel Salvendy, Margherita Antona, Jessie YC Chen, Jianming Dong, Vincent G Duffy, Xiaowen Fang, Cali Fidopiastis, Gino Fragoni, Limin Paul Fu, et al. 2019. Seven HCI grand challenges. *International Journal of Human-Computer Interaction* 35, 14 (2019), 1229–1269.
 - [53] Shari Trewin, Sara Basson, Michael Muller, Stacy Branham, Jutta Treviranus, Daniel Gruen, Daniel Hebert, Natalia Lyckowski, and Erich Manser. 2019. Considerations for AI fairness for people with disabilities. *AI Matters* 5, 3 (2019), 40–63.
 - [54] 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.
 - [55] Rui Wang, Fanglin Chen, Zhenyu Chen, Tianxing Li, Gabriella Harari, Stefanie Tignor, Xia Zhou, Dror Ben-Zeev, and Andrew T Campbell. 2014. StudentLife: assessing mental health, academic performance and behavioral trends of college students using smartphones. In *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*. 3–14.
 - [56] 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.
 - [57] Mark Weiser. 1993. Some computer science issues in ubiquitous computing. *Commun. ACM* 36, 7 (1993), 75–84.
 - [58] Xuhai Xu, Prerna Chikersal, Afsaneh Doryab, Daniella K Villalba, Janine M Dutcher, Michael J Tumminia, 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.
 - [59] Xuhai Xu, Xin Liu, Han Zhang, Weichen Wang, Subigya 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.
 - [60] 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).
 - [61] 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).
 - [62] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. 2017. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th international conference on world wide web*. 1171–1180.
 - [63] 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.

- [64] Hanbin Zhang, Chenhan Xu, Huining Li, Aditya Singh Rathore, Chen Song, Zhisheng Yan, Dongmei Li, Feng Lin, Kun Wang, and Wenyao Xu. 2019. Pdmov: Towards passive medication adherence monitoring of parkinson's disease using smartphone-based gait assessment. *Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies* 3, 3 (2019), 1–23.
- [65] Yunke Zhang, Fengli Xu, Tong Xia, and Yong Li. 2021. Quantifying the Causal Effect of Individual Mobility on Health Status in Urban Space. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 4 (2021), 1–30.
- [66] 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.