Understanding the long-term impact and perceptions of privacy-enhancing technologies for bystander obscuration in AR

Brendan David-John*
Computer Science, Virginia Tech

Bo Ji†
Computer Science, Virginia Tech

Evan Selinger‡
Philosophy, Rochester Institute of Technology

ABSTRACT

Wearable Augmented Reality (AR) devices are poised to provide a bright future of immersive ubiquitous experiences and interactions. While the powerful suite of sensors on modern AR devices are necessary for enabling the future of spatial computing, they can create unease in bystanders due to privacy concerns. A primary source of concern is related to the risk of identification and surveillance from recent advancements in facial recognition. Our position paper outlines a vision of wearable AR displays and sensors in which privacy-enhancing technologies (PETs) protect bystanders' identities and behaviors. Within this vision, we identify several long-term research questions related to perception, usability, and deployment of bystander PETs within the AR ecosystem that impact the normalization of wearable AR, privacy concerns for diverse users and use cases, and the resulting ethical considerations.

Index Terms: bystander privacy, augmented reality, obscuration.

1 INTRODUCTION

Wearable Augmented Reality (AR) devices, such as the Microsoft HoloLens or Magic Leap, are set apart from other mobile devices by the immersive experience they offer. While the powerful suite of sensors on modern AR devices are necessary for enabling such an immersive experience, they can create unease in bystanders due to privacy concerns [9, 4]. A primary source of concern is related to the risk of identification and surveillance from recent advancements in facial recognition [24]. For example, an AR user may download apps that capture always-on camera data to provide assistive benefits such as heads-up navigation [28], translating text [3], and relaying facial expressions to neuro-divergent users within social settings [10]. Without the user's knowledge, the app could be storing captured camera data remotely for later use. A user walking by a Black Lives Matter protest while wearing the glasses may unintentionally identify the event organizer or attendees in the data by capturing their faces, which could be purchased by agencies targeting specific groups and organizations [14]. Taken together with public backlash for earlier AR displays such as the Google Glass [19, 18], addressing the perception and privacy risks of AR sensors is a necessary step to their deployment and developing an ethical understanding of their impact.

The described scenario highlights several critical issues of future AR devices: AR sensor data erodes privacy in public spaces, the privacy risks introduced by such devices can negatively impact public perception, and a negative public perception could result in the harm or ostracization of users and a lack of adoption or access to AR assistive technologies that otherwise provide a practical trade-off of privacy for utility.

Our future vision is the use of privacy-enhancing technologies (PETs) to protect bystander privacy while enabling the future of AR. State-of-the-art bystander PETs automatically detect and remove personally identifying information, with a focus on faces within camera data [5, 21]. Several long-term research questions must be answered before our vision can be achieved:

- RQ_1 : how does the public perceive AR PETs and how will they be deployed in practice?
- RQ_2 : How well should PETs perform in different settings to be considered effective?
- RQ_3 : how can we avoid jeopardizing privacy in the long run, particularly if PETs normalize wearable AR in public but are limited to opt-in or difficult to configure systems?

*e-mail: bmdj@vt.edu
†e-mail: boji@vt.edu
‡e-mail: evan.selinger@rit.edu

2 RELATED WORK

2.1 Bystander Privacy

The technical perspective on bystander privacy and AR headsets has focused on balancing the level of privacy provided by PETs (e.g., face removal or blurring) and the impact on user experience or utility of AR applications [6]. The systems for identifying bystanders and managing consent for data capture can operate in a passive or active manner, spanning typical images captured and shared through social media [8, 13], videos captured to support lifelogging [15], and those that focus on wearable AR displays [5]. The recurring theme across these methods is how they deal with consent from bystanders, as active systems rely on a known gesture or additional technology integration to signal whether their face or identifying features must be removed from the sensor data stream [16]. Passive systems automatically monitor AR data streams to identify and redact bystanders and their private information [21]. The trade-off between active and passive bystander systems depends on the public availability and understanding of the technology, indicating that passive systems running automatically without any prior knowledge are ideal for public AR use cases [1, 6]. The discussed bystander protections and PETs vary in how they are implemented, how successful they are in wiping bystander information from the data stream, and their social perceptions. Thus, we identify a gap in comparative studies and deployments that jointly understand the impact of wearable AR on privacy and the different types of PETs on public perceptions.

The initial understanding of social perceptions for wearable AR began with incidents resulting from the initial roll out of Google Glass [18, 19] and carry through to the most popular mixed-reality devices today, the Quest 3 and Apple Vision Pro which has seen increased use in public spaces such as coffee shops and airplanes [17]. Prior works from research groups include user studies to characterize AR bystander perceptions [20, 9]. Their findings suggest that PETs should be integrated into AR systems and consider contextual information such as location and the relationship to a bystander, however, there is a lack of understanding in how bystanders or users feel when presented with the protections resulting from different types of PETs or their quantitative performance when attempting to protect the bystander within every frame of data.

2.2 Obscurity

We focus on the privacy concept of obscurity as developed by law and philosophy scholars [12, 25]. Obscurity theory proposes that

when it is costly (e.g., through time, effort, or money) to find or correctly interpret information, people will be disinclined to try to locate and correctly interpret it. Obscurity is thus a probabilistic account of protecting information that views cost as a deterrent, such as how long one must observe AR camera data to successfully identify a bystander and with high confidence link them to some private information. An example of obscurity in public from the original paper includes eavesdropping on a stranger's conversation. If you overhear somebody revealing that a family member has been checked into a rehab facility and hear just their first name it is theoretically possible to link this information with an identity. However, the work needed to identify the speaker, which family members share this first name, and whether this information has any external relevance (i.e., they are a politician or celebrity) would dissuade an adversary and provide privacy protections.

In the protest example above, a PET that reduces the number of frames containing the organizer's face from many to just one could make confident identification and surveillance at scale significantly more difficult. However, obscurity theory *cautions* that current and future surveillance technologies can significantly reduce obscurity protections that society has taken for granted when they are not legally protected. Care must be taken to ensure practical privacy risks are not taken for granted in case relevant transaction costs are greatly reduced through technological advancements or new use cases. The obscurity landscape for wearable AR and bystander PETs is still emerging as the technology is not yet normalized at scale.

3 Methodology

3.1 What is a key security or privacy harm that you consider critical and/or challenging to address for a future AR ecosystem?

As outlined above, addressing bystander concerns over the presence of wearable AR displays is critical in the wide-spread usage of these technologies. Furthermore, a broader understanding of the perceptions across diverse users and the specific impacts of PETs are necessary. The challenge of understanding user and bystander perceptions across cultural background, economic status, ethnicity, and technical literacy for modeling future real-world deployments led us to pose RQ_1. The goal of understanding the specific settings and key performance metrics that lead to PETs as an ideal solution to bystander privacy led us to RQ_2. Finally, we revisited obscurity theory to pose RQ_3 on whether PETs could result in a net loss to privacy, in cases where the public is normalized to their presence but external factors jeopardize bystander privacy.

3.2 How do you envision improving AR design and development practices to address this harm?

Current design work to support privacy and security efforts establish a ground-up approach to address AR challenges [23, 22], including bystander privacy [2, 26]. These research methods develop tools to collect, measure, and implement privacy-focused design elements into the AR authoring process. We see new future research directions as contributing to the design of more effective bystander PETs by profiling what elements support public acceptance and where they will be deployed (RQ_1), benchmarking PETs to inform practical deployments (RQ_2), and finally informing the ethical considerations of future AR designs (RQ_3).

To address RQ_1, we envision the use of longitudinal studies that explore the deployment of representative PETs and measure user perceptions across a diverse set of user populations. The goal of longitudinal studies are to validate the impact of prior findings and provide evidence that can inform design principles for future AR deployments. For example, studies may identify that specific types of visual indicators for informing bystanders may be effective for abled users but alternatives are necessary for certain populations. Likewise, default privacy configurations and the necessity of training or educational materials for new users may vary widely across economic status and depend on the policies and legal frameworks in certain countries (e.g., GDPR [27]). This research question also considers modeling *how* PETs will be deployed in practice, and could make use of measurement studies to profile current AR application use and how users would configure PETs if provided to them in the form of prototypes or plug-ins within existing applications.

To address RQ_2, we are interested in solidifying the definition of success in the application of PETs by providing standard metrics and frameworks for benchmarking that leads to deployment recommendations. For example, new PETs are typically evaluated relative to past algorithms with new or existing datasets and measure how many frames are protected in the output datastream [7, 13, 5]. A standard protocol for characterizing target metrics and values with the public perceptions collected within RQ_1 would provide a viable path towards transferring research prototypes to practical deployments. For example, a common question posed after researching a new PET is whether obscuring bystander information 98% of the time is considered successful, or if 100% is necessary. We expect to link the answer to the concept of obscurity and difficulty in correctly interpreting information about bystanders, while considering the level of current facial recognition and privacy-invasive technologies. The answer also depends on the context of the evaluated scenario, and establishing standard targets based on crowd-sourced perceptions and representative inputs would allow researchers and platforms to make informed decisions when evaluating privacy-utility trade-offs that could influence real-world deployments.

RQ_3 is a more complex question to address, as it treats the obscurity gained by PETs as a double-edged sword. While the protection from PETs establishes obscurity by design, we are concerned with the impact normalization without proper enforcement would have on bystander privacy. The goal of this research thrust would support iterative ethical considerations in AR design as researchers and industry giants continuously influence what the future integration of AR within daily life will look like, and not just the technical limitations of future systems. For example, the marketed use cases for the Meta Ray Bans integrated glasses includes documenting travel in a way that provides enhanced functionality with less distractions or burdens compared to a cell phone in a convincing manner [11]. If these visions and marketing materials strongly influence future use cases and AR designs, it could be paired with accompanying messaging on the availability of PETs that may lead to a mass normalization of the technology. Specifically, since passive bystander PETs are more likely to be adopted for wearable AR it is easy to both misinterpret or neglect default settings and configurations which does not empower bystanders to ensure PETs are providing obscurity as expected. The result could produce a net loss in privacy without enforcement of PETs best practices and standards. Investigating this research question would provide an ethical understanding that could influence AR design standards and policy discussions.

4 Conclusion

This position paper outlines three critical research questions for providing practical bystander privacy for wearable AR systems. The goal of this work is to establish long-term challenges and directions that aid in identifying the collaborations, methods, experiments, and tools needed to achieve a privacy-preserving vision for the future of AR.

Acknowledgments

The authors acknowledge support from the Commonwealth Cyber Initiative and the Virginia Tech Center for Human-Computer Interaction.

REFERENCES

[1] M. Abraham, P. Saeghe, M. Mcgill, and M. Khamis. Implications of xr on privacy, security and behaviour: Insights from experts. In *Nordic Human-Computer Interaction Conference*, pp. 1–12, 2022. 1

[2] I. Ahmad, R. Farzan, A. Kapadia, and A. J. Lee. Tangible privacy: Towards user-centric sensor designs for bystander privacy. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2):1–28, 2020. 2

[3] T. Ahmed, A. Kapadia, V. Potluri, and M. Swaminathan. Up to a limit? privacy concerns of bystanders and their willingness to share additional information with visually impaired users of assistive technologies. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3):1–27, 2018. 1

[4] D. Bhardwaj, A. Ponticello, S. Tomar, A. Dabrowski, and K. Krombholz. In focus, out of privacy: The wearer's perspective on the privacy dilemma of camera glasses. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pp. 1–18, 2024. 1

[5] M. Corbett, B. David-John, J. Shang, Y. C. Hu, and B. Ji. Bystandar: Protecting bystander visual data in augmented reality systems. In *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*, pp. 370–382, 2023. 1, 2

[6] M. Corbett, B. David-John, J. Shang, Y. C. Hu, and B. Ji. Securing bystander privacy in mixed reality while protecting the user experience. *IEEE Security & Privacy*, 2023. 1

[7] D. Darling, A. Li, and Q. Li. Identification of subjects and bystanders in photos with feature-based machine learning. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pp. 1–6. IEEE, 2019. 2

[8] D. Darling, A. Li, and Q. Li. Automated bystander detection and anonymization in mobile photography. In *Security and Privacy in Communication Networks: 16th EAI International Conference, SecureComm 2020, Washington, DC, USA, October 21-23, 2020, Proceedings, Part I 16*, pp. 402–424. Springer, 2020. 1

[9] T. Denning, Z. Dehlawi, and T. Kohno. In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2377–2386, 2014. 1

[10] C. Engel, J. Schmalfuß-Schwarz, D. Gollasch, M. Branig, S. Dirks, and G. Weber. Workshop on designing accessible extended reality: An opportunity for people with disabilities and disorders. 2023. 1

[11] I. Hamilton. Meta ray-ban glasses pair to quest 3 for a preview of our wearable future. 2

[12] W. Hartzog and E. Selinger. Surveillance as loss of obscurity. *Wash. & Lee L. Rev.*, 72:1343, 2015. 1

[13] R. Hasan, D. Crandall, M. Fritz, and A. Kapadia. Automatically detecting bystanders in photos to reduce privacy risks. In *2020 IEEE Symposium on Security and Privacy (SP)*, pp. 318–335. IEEE, 2020. 1, 2

[14] L. Hecht-Felella. Federal agencies are secretly buying consumer data. *Brennan Center for Justice, https://www. brennancenter. org*, 2021. 1

[15] R. Hoyle, R. Templeman, S. Armes, D. Anthony, D. Crandall, and A. Kapadia. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pp. 571–582, 2014. 1

[16] M. Koelle, S. Ananthanarayan, S. Czupalla, W. Heuten, and S. Boll. Your smart glasses' camera bothers me! exploring opt-in and opt-out gestures for privacy mediation. In *Proceedings of the 10th Nordic Conference on Human-Computer Interaction*, NordiCHI '18, p. 473–481. Association for Computing Machinery, New York, NY, USA, 2018. doi: 10.1145/3240167.3240174 1

[17] K. Kozuch. I flew 8,000 miles wearing apple vision pro — here's what it's really like, March 2024. 1

[18] K. Levy. A san francisco woman says she was attacked for wearing google glass in a bar, February 2014. 1

[19] K. Levy. A surprising number of places have banned google glass in san francisco, March 2014. 1

[20] J. O'Hagan, P. Saeghe, J. Gugenheimer, D. Medeiros, K. Marky, M. Khamis, and M. McGill. Privacy-enhancing technology and everyday augmented reality: Understanding bystanders' varying needs for awareness and consent. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 6(4):1–35, 2023. 1

[21] N. Raina, G. Somasundaram, K. Zheng, S. Saarinen, J. Messiner, M. Schwesinger, L. Pesqueira, I. Prasad, E. Miller, P. Gupta, et al. Egoblur: Responsible innovation in aria. *arXiv preprint arXiv:2308.13093*, 2023. 1

[22] S. Rajaram, C. Chen, F. Roesner, and M. Nebeling. Eliciting security & privacy-informed sharing techniques for multi-user augmented reality. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–17, 2023. 2

[23] S. Rajaram, F. Roesner, and M. Nebeling. Reframe: An augmented reality storyboarding tool for character-driven analysis of security & privacy concerns. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–15, 2023. 2

[24] F. Roesner, T. Kohno, and D. Molnar. Augmented reality: challenges & opportunities for security and privacy. *J Comput Secur Neurosci—Part*, 1(2), 2021. 1

[25] E. Selinger and W. Hartzog. Obscurity and privacy. In *Spaces for the Future*, pp. 119–129. Routledge, 2017. 1

[26] Y. Yao, J. R. Basdeo, O. R. Mcdonough, and Y. Wang. Privacy perceptions and designs of bystanders in smart homes. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–24, 2019. 2

[27] R. N. Zaeem and K. S. Barber. The effect of the gdpr on privacy policies: Recent progress and future promise. *ACM Transactions on Management Information Systems (TMIS)*, 12(1):1–20, 2020. 2

[28] Y. Zhao, J. Stefanucci, S. Creem-Regehr, and B. Bodenheimer. Evaluating augmented reality landmark cues and frame of reference displays with virtual reality. *IEEE Transactions on Visualization and Computer Graphics*, 29(5):2710–2720, 2023. 1