
Do Different Groups Have Comparable Privacy Tradeoffs?

Rezvan Joshghani

Department of Computer Science
College of Engineering
Boise State University
Boise, ID 83702, USA
RezvanJoshghani@boisestate.edu

Michael D. Ekstrand

Department of Computer Science
College of Engineering
Boise State University
Boise, ID 83702, USA
MichaelEkstrand@boisestate.edu

Bart Knijnenburg

School of Computing
Clemson University
215 McAdams Hall
Clemson, SC, 29634, USA
BartK@clemson.edu

Hoda Mehrpouyan

Department of Computer Science
College of Engineering
Boise State University
Boise, ID 83702, USA
HodaMehrpouyan@boisestate.edu

Abstract

Personalized systems increasingly employ Privacy Enhancing Technologies (PETs) to protect the identity of their users. In this paper, we are interested in whether the cost-benefit tradeoff — the underlying economics of the privacy calculus — is fairly distributed, or whether some groups of people experience a lower return on investment for their privacy decisions.

Author Keywords

User-Centric Privacy; Differential Privacy; K-anonymity; Fairness.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous; See <http://acm.org/about/class/1998> for the full list of ACM classifiers. This section is required.

Introduction

Using techniques such as data deletion or obfuscation, these systems provide privacy guarantees that can be expressed as ‘k-anonymity’ [19, 20] or ‘differential privacy’ [1, 17, 18]. The application of these techniques often comes with a certain reduction in personalization accuracy, and an active area of research tries to maximize privacy protections while minimizing the impact on accuracy.

Paste the appropriate copyright/license statement here. ACM now supports three different publication options:

- **ACM copyright:** ACM holds the copyright on the work. This is the historical approach.
- **License:** The author(s) retain copyright, but ACM receives an exclusive publication license.
- **Open Access:** The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single-spaced in Verdana 7 point font. Please do not change the size of this text box.

Each submission will be assigned a unique DOI string to be included here.

The privacy guarantees provided by PETs typically apply to a dataset as a whole, and the resulting reduction in accuracy is usually measured at this level as well. However, this does not necessarily mean that all users enjoy the same level of privacy protection, nor are they necessarily equally affected by the reduction in personalization accuracy. This reality raises concerns about the ways in which the impact of these PETs on privacy and accuracy may correlate with user membership in particular groups, such as demographic classes. This is particularly a concern if there is an adverse relationship between a subject's membership in a class subject to historical discrimination, such as a racial minority, and the impact of PETs on their privacy risks and benefits. We have treated the general case of this topic previously [6], identifying ways in which at-risk users may be additionally subject to greater privacy risks or greater cost of achieving privacy benefits.

Importantly, though, not every user desires the same level of privacy. Considering privacy as a tradeoff between costs and benefits, research shows that not all users prefer the same balance between these aspects [2-4]. Considering individuals' differences in privacy preferences and behaviors, instead of attempting to cover all users with a single global policy, is a valuable move in making privacy protections responsive to the needs and desires of individual users [5]. However, the jump from the high-level need for privacy to the individual's interaction with privacy in sociotechnical systems may miss important meso-level considerations. Particularly, if we employ a user-tailored and individualized approach to privacy, will members of different classes *obtain* or *perceive* a comparable return-on-investment for their privacy

behaviors? How does the privacy calculus itself, rather than the absolute difference in risk or reward, change for different subsets of the population, and how does this local tradeoff relate to the balance of accuracy and protection provided by the PETs they encounter in their daily lives?

Background

The fundamental question driving our inquiry into fairness and privacy is "who pays, and who benefits?" [7] in a sociotechnical system, particularly as it applies to the costs and benefits of privacy protections and the broader system in which they are implemented [6].

The question of fairness - or of its opposite for our purposes, bias or discrimination - in computing systems is not a new issue [8]. In the last decade, significant work on fairness such as machine learning tools has resulted in several concepts and operationalizations of fairness, such as individual fairness [9] that requires similar subjects to receive similar judgements, statistical parity that requires different groups to experience similar rates of judgements, and disparate mistreatment [10]. These (often mutually-exclusive [11,12]) approaches share a similar high-level goal: a person's experience with a computer system should not depend in irrelevant ways on their social or demographic characteristics.

Empirical Results

Addressing the 'fairness' of PETs in relation to users' personal risk-benefit tradeoff is a formidable task that involves several complex research questions:

- How do we measure the impact of PETs on individual users' level of privacy and accuracy?

- How do we quantify users' (contextualized) risk-benefit tradeoff?
- How do we assess the alignment between the privacy-accuracy balance provided by the PET and the user's desired balance between these factors?
- Can we map these factors for different demographic groups? Or, ideally, per individual user?
- Can we develop a 'user-tailored' version of the PET that improves the alignment with the privacy-accuracy balance the user desires?

In this paper, we make two small contributions to this research agenda. First, we map differences in privacy preferences by gender and ethnicity as found in publicly accessible research datasets. This will highlight demographic differences in users' risk-benefit tradeoff (albeit not contextualized to a specific system or scenario). Then, we analyze the impact of various differential privacy procedures on men versus women. This gives us an idea of how "generic" (non-user-tailored) PETs create a balance between privacy and characteristics.

Survey Data

Privacy research shows that individual people differ substantially in the kind of data they would prefer to protect [13], and the kind of privacy-related behaviors they choose to engage in [14]. Privacy perceptions, preferences, and behaviors differ at the cultural level as well [15, 16]. From a fairness perspective, we are interested to find out whether there are demographic differences in privacy attitudes and behaviors, especially when comparing underrepresented and marginalized groups (e.g. African-Americans, women) to their dominant, majority counterpart (e.g., Whites,

men). Moreover, when dealing with marginalized groups, it is also important to take into account the privacy-related socio-economic realities that may influence individuals' level of behavioral control over their privacy. To investigate the existence of such differences, we identified one datasets related to privacy, and report significant ($p < .01$) differences by gender and ethnicity. People's privacy can be violated by others spreading lies about them online.

Spreading lies

People's privacy can be violated by others spreading lies about them online. The Pew American Trends Panel Wave 24 (administered January 9 to 23, 2017) dataset shows the following results regarding gender:

- Men are more likely to have had someone post untrue information about them online than women.
- For men, such untrue information had more likely to do with their sexual identity, religion or political views, while for women it had more likely to do with their relationship or sexual history.
- For women, such untrue information had more far-reaching consequences, as it was more likely to impact romantic and personal/family relationships, or cause mental or emotional stress or problems at school than for men.
- Women were more likely to try to get the untrue information corrected or removed than men.

The same Pew dataset shows the following results regarding race:

- Whites are less likely to have had someone post untrue information about them online than non-whites.

- For non-whites, such untrue information had more likely to do with their sexual identity, gender identity, or health/medical history than for whites.
- For non-whites, such untrue information had more far-reaching consequences, as it was more likely to lead to trouble finding a job, or impact romantic and personal/family relationships than for whites.
- Non-whites were more likely to try to get the untrue information corrected or removed than whites.

Privacy

The Pew Internet Survey #4 (administered January 2015) dataset shows the following results regarding gender:

- Women find it more important to be in control of who can get information about them ($p < .001$), to not have someone watch them or listen to them without their permission ($p < .001$), to control what information is collected about them ($p = .007$), to not have individuals in social and work situations ask them things that are highly personal ($p = .011$), and to be able to go around in public without always being identified ($p = .013$).

The same Pew dataset shows the following results regarding race:

- Non-whites find it more important to control what information is collected about them ($p = .004$), to not have individuals in social and work situations ask them things that are highly personal ($p = .003$), and to be able to go around in public without always being identified ($p = .011$).

Conclusion and Future Work

In this paper, we call for an investigation of fairness in privacy, focusing on the alignment between users'

privacy calculus and the privacy-accuracy balance provided by PETs. We demonstrate that users' risk-benefit tradeoff likely differs by race and gender. We also demonstrate that existing PETs create a balance between privacy and accuracy that may differ by demographic characteristics.

Given the obvious limitations in our dataset, it is difficult for us to argue about the alignment between users' privacy calculus and PETs' privacy-accuracy balance: we only have very generic data about users' preferences, and one very specific recommender dataset with limited demographic details about the users. We encourage researchers to investigate these aspects in more detail, preferably covering both aspects within the same system.

For example, a movie subscription service employing a PET to protect its users' privacy could survey its users about their privacy preferences and risk-benefit tradeoff, and then test whether the distribution of these preferences aligns with the privacy enhancements (and accuracy reductions) provided by the PET across different demographics. Potentially, a mapping could even be established for each individual user!

If the alignment turns out to be skewed, the company could then attempt to *tailor* the protection of the PET to various demographic characteristics (or, ambitiously, to each individual user) in an attempt to improve the alignment, while keeping overall privacy guarantees intact. This user-tailored approach would not only guarantee its users a certain level of privacy, but also ascertain this this privacy is distributed fairly among its users.

Acknowledgements

1. Dwork, Cynthia, and Aaron Roth. "The algorithmic foundations of differential privacy." *Foundations and Trends® in Theoretical Computer Science* 9, no. 3-4 (2014): 211-407.
2. Spiekermann, Sarah, Jens Grossklags, and Bettina Berendt. "E-privacy in 2nd generation E-commerce: privacy preferences versus actual behavior." In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pp. 38-47. ACM, 2001.
3. Knijnenburg, Bart P., Alfred Kobsa, and Hongxia Jin. "Dimensionality of information disclosure behavior." *International Journal of Human-Computer Studies* 71, no. 12 (2013): 1144-1162.
4. Mehrpouyan, Hoda, Ion Madrazo Azpiazu, and Maria Soledad Pera. "Measuring Personality for Automatic Elicitation of Privacy Preferences." In *Privacy-Aware Computing (PAC), 2017 IEEE Symposium on*, pp. 84-95. IEEE, 2017.
5. Knijnenburg, Bart P. "Privacy? I Can't Even! Making a Case for User-Tailored Privacy." *IEEE Security & Privacy* 15, no. 4 (2017): 62-67.
6. Ekstrand, Michael D., Rezvan Joshaghani, and Hoda Mehrpouyan. "Privacy for All: Ensuring Fair and Equitable Privacy Protections."
7. Franklin, Ursula. *The real world of technology*. House of Anansi, 1999.
8. Friedman, Batya, and Helen Nissenbaum. "Bias in computer systems." *ACM Transactions on Information Systems (TOIS)* 14, no. 3 (1996): 330-347.
9. Dwork, Cynthia, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. "Fairness through awareness." In *Proceedings of the 3rd innovations in theoretical computer science conference*, pp. 214-226. ACM, 2012.
10. Zafar, Muhammad Bilal, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. "Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment." In *Proceedings of the 26th International Conference on World Wide Web*, pp. 1171-1180. International World Wide Web Conferences Steering Committee, 2017.
11. Chouldechova, Alexandra. "Fair prediction with disparate impact: A study of bias in recidivism prediction instruments." *Big data* 5, no. 2 (2017): 153-163.
12. Friedler, Sorelle A., Carlos Scheidegger, and Suresh Venkatasubramanian. "On the (im) possibility of fairness." arXiv preprint arXiv:1609.07236 (2016).
13. Knijnenburg, Bart P., Alfred Kobsa, and Hongxia Jin. "Dimensionality of information disclosure behavior." *International Journal of Human-Computer Studies* 71, no. 12 (2013): 1144-1162.
14. Wisniewski, Pamela J., Bart P. Knijnenburg, and Heather Richter Lipford. "Making privacy personal: Profiling social network users to inform privacy education and nudging." *International Journal of Human-Computer Studies* 98 (2017): 95-108.
15. Li, Yao, Alfred Kobsa, Bart P. Knijnenburg, and MH Carolyn Nguyen. "Cross-Cultural Privacy Prediction." *Proceedings on Privacy Enhancing Technologies* 2017, no. 2 (2017): 113-132.
16. Cho, H, Li, Y., Knijnenburg, B.P., Kobsa, A. (in press). *Collective Privacy Management in Social Media: A Cross-cultural Validation*. ACM Transaction on Computer-Human Interaction.
17. McSherry, Frank, and Kunal Talwar. "Mechanism design via differential privacy." In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pp. 94-103. IEEE, 2007.
18. Ji, Zhanglong, Zachary C. Lipton, and Charles Elkan. "Differential privacy and machine learning: a

survey and review." *arXiv preprint arXiv:1412.7584* (2014).

19. Aggarwal, Charu C. "On k-anonymity and the curse of dimensionality." In *Proceedings of the 31st international conference on Very large data bases*, pp. 901-909. VLDB Endowment, 2005.
20. Sweeney, Latanya. "k-anonymity: A model for protecting privacy." *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10, no. 05 (2002): 557-570.