Article Type:

Grand Challenges in Human-Centered Privacy

Ruba Abu-Salma, *King's College London, UK*
Pauline Anthonysamy, *Google Zürich, Switzerland*
Zinaida Benenson, *Friedrich-Alexander-Universität, Erlangen, Germany*
Benjamin Berens, *Karlsruhe Institute of Technology, Karlsruhe, Germany*
Kovila P. L. Coopamootoo, *King's College London, UK*
Andreas Gutmann, *University College London, UK & Ofcom, UK*
Adam Jenkins, *King's College London, UK*
Sameer Patil, *University of Utah, Salt Lake City, UT, USA*
Sören Preibusch, *German Federal Institute for Risk Assessment, Berlin, Germany*
Florian Schaub, *University of Michigan, Ann Arbor, MI, USA*
William Seymour, *King's College London, UK*
Jose Such, *King's College London, UK & Universitat Politecnica de Valencia, Spain*
Mohammad Tahaei, *International Computer Science Institute & eBay, USA*
Aybars Tuncdogan, *King's College London, UK*
Max Van Kleek, *University of Oxford, Oxford, UK*
Daricia Wilkinson, *Arizona State University, Mesa, AZ, USA*

Abstract—We report the most salient themes and future directions for human-centered privacy that emerged from the discussions at the Future of Human-Centred Privacy *event, which brought together leading international experts from academia, industry, and government to discuss grand challenges in the field.*

Over the past two decades, the field of Human-Centered Privacy has emerged as a critical area at the intersection of Privacy and Human-Computer Interaction research. The importance of the area stems from the inherently *human* nature of privacy, the interplay between privacy and technology, and how humans interact with technology. The Human-Centered Privacy field aims to understand human perceptions, concerns, needs, and awareness regarding privacy issues in computers, smart devices, and online technologies. Additionally, it focuses on evaluating the user experience of existing Privacy-Enhancing Technologies (PETs) and designing novel, user-friendly PETs.

The proliferation of smart devices and advances in Artificial Intelligence (AI) present considerable challenges for human-centered privacy research and for the design and development of human-centered privacy solutions. These challenges are exacerbated by the growing complexity of devices and systems, the increasingly multi-user nature of technology, and the importance of supporting higher-risk populations during research and development.

In this article, we summarize the discussions and outcomes of *The Future of Human-Centred*[1] *Privacy* (FHCP) event[2], which took place at King's College London over three days in June 2023. This unique event, which blended elements of a workshop and a retreat, brought together 26 leading international experts

XXXX-XXX © 2025 IEEE
Digital Object Identifier 10.1109/XXX.0000.0000000

[1] While the article uses American spellings and punctuation, we use the British spelling to refer accurately to the title of the event that took place in the United Kingdom.

[2] This article represents the views and opinions of the author(s) and should not be taken as a statement of Ofcom policy/opinion.

from academia, industry, and government to engage in meaningful dialogue on emerging challenges in the field of Human-Centered Privacy. In this article, we report the most salient themes and future directions that emerged from the discussion among the FHCP attendees. These themes can inform privacy researchers as well as practitioners, such as UX designers, software developers, and privacy engineers, about emerging developments and challenges in the field.

Theme 1: Inclusive Privacy

Privacy harms are experienced differently by different societal groups; some encounter unique and severe privacy challenges, with risks further amplified for marginalized populations [1]. Identity characteristics (e.g., age, gender, sexual orientation, ethnicity, [dis]abilities, health conditions, religion), personal circumstances (e.g., immigration status, [un]availability of legal protections, risky professional responsibilities, digital literacy), and sociocultural factors (e.g., conceptualizations of privacy, interpersonal relations, societal structures, governance, business practices) often shape privacy risks, needs, and harms in distinct ways. This diversity is important to address in the design, engineering, and operation of digital systems.

Moving from Qualitative to Systematic Understanding

Researchers have investigated and documented the impact of the identity characteristics, personal circumstances, and sociocultural factors mentioned above on the privacy needs, experiences, and harms of various at-risk or exposure-sensitive populations and communities. Such research efforts require addressing several structural impediments, such as a lack of access to participants, inadequate trust toward researchers, and language barriers. While online platforms have made it easier to recruit participants, these issues continue to hinder large-scale studies with appropriate regional and demographic representation. As a result, research with such populations tends to be qualitative (i.e., interview-based), with a relatively small number of participants. Therefore, assessing whether the findings generalize beyond the sample and tracing and comparing the issues across populations is often difficult. Consolidating and synthesizing privacy-related findings across populations and contexts to formulate general theories and models of privacy vulnerability and needs remains a key challenge. While attempts to address the challenge have started to emerge (see, for example, Warford et al.'s [2] framework for unifying research with at-risk users), more work is needed to advance from mostly exploratory, qualitative research to large-scale confirmatory studies. Such studies can solidify characterizations of privacy risks and needs for specific populations into generalizable theories. In addition, the outcomes of the studies can directly inform the design and development of solutions to enhance privacy for all types of users.

Designing Privacy Solutions for Specific and General Audiences

Privacy solutions ought to avoid overly broad definitions of privacy that do not support the nuanced needs of diverse groups. The design of privacy-related solutions should involve assessing whether the privacy needs of any groups are neglected and identifying harms that may arise as a consequence. Design encodes *values*. Therefore, it is essential to consider the trade-offs for different user groups in terms of the underlying values. While an inclusive approach to design can empower individuals to manage privacy more effectively, the evolving definitions of privacy and safety require careful consideration of sociotechnical realities when balancing privacy risks with safety measures. Effectively designing privacy solutions for the diverse needs of individuals and communities is challenging because of the intersectionality of identities and experiences. In inclusive privacy design, it is crucial to consider the balance between individual and collective privacy needs, acknowledging situations where one may supersede the other due to prevailing power dynamics. Focusing on a particular context by identifying relevant stakeholders and eliciting their contextual privacy needs at the design stage can help reconcile personally specific needs with those of the larger population. Such an approach can help design human-centered privacy solutions that achieve a reasonable balance in dealing with interpersonal, corporate, and governmental demands that may intrude upon individual privacy.

Being Attentive to Responsible Research and Practice with At-Risk Populations

When researching or designing for at-risk populations, it is crucial to be extra attentive to ethical and respectful conduct to avoid causing inadvertent harm through the effort meant to help them. For instance, additional privacy and security safeguards in data collection, analysis, and reporting may be required when studying vulnerable populations, such as children or survivors of abuse. Specialized training in ethical and community-specific best practices is essential when working with such groups. In addition, researchers and practitioners studying vulnerable populations need to be mindful of their *positionality*, i.e., the potential influence of their identity and experiences in relation to the populations under study. Such careful attention can help researchers and practitioners make positive contributions to support the privacy needs of the studied communities and serve as their advocates [3].

Theme 2: Privacy-related Decisions and Communication

Users are often asked to make many privacy decisions, such as accepting privacy policies, consenting to cookies, or selecting app permissions. During decision-making, users require clear choices and information about the consequences of their decisions. Therefore, it is important to communicate the privacy benefits and risks of each choice to empower users to make informed and meaningful decisions. However, existing consent mechanisms, permission dialogs, cookie banners, and warning messages may not align with user needs and may not balance stakeholder interests (e.g., an app may request more permissions than it needs to function properly). Assisting users in managing privacy requires understanding the context of the decision, bridging the gap between the mental models of users and the business models of the product or service, and creating usable and accessible interventions that allow users to make, reconsider, or revoke a decision.

Working with Habituation

Being constantly asked to make privacy-related decisions can lead to *habituation*, wherein users tend to respond quickly without much thinking. Habituation is a core aspect of human information processing and often serves vital functions, such as preventing cognitive overload and stress. However, habituation can be counterproductive in privacy-related situations because it reduces attention to information over time, potentially impairing people's ability to make appropriate decisions (e.g., by inadvertently accepting invasive cookies or privacy practices). Efforts to counter habituation when making privacy-related decisions typically focus on breaking habits that can cause privacy invasion. To understand how habituation can be leveraged positively at a fundamental research level, researchers may employ methods used in psychology and neuroscience, such as Magnetic Resonance Imaging (MRI), Electroencephalogram (EEG) readings, cortisol tests, heart rate monitoring, and galvanic skin resistance measurements [4]. In practice, those who build the products and services could explore user experiences to detect and avoid habituation. For instance, one strategy to counter habituation could be to measure how long users view messages before taking action and reflect the behavior back to users at appropriate times.

Defining What Is Communicated to Whom

When supporting users in making privacy-related decisions, it is crucial to distinguish between necessary and optional information. Overload from optional information can hinder effective decision-making, especially in time-sensitive circumstances. Additionally, clarity is needed on whether the information is relevant for lay users or only for experts, such as developers or IT administrators. Information that is irrelevant, or perceived as such, is disregarded more easily. Another option to reduce decision fatigue among users is to design trust-based mechanisms for delegating privacy-related decisions to others, including to a group of people (e.g., to a collective or crowdworkers) or AI (see Theme 3).

Communicating Appropriately During Each Phase of Decision-Making

Greater emphasis should be placed on understanding the entire process of making privacy-related decisions. Instead of viewing the process as *uniform*, it is important to recognize its distinct *pre-decisional*, *post-decisional*, *actional*, and *post-actional* phases [5]. Consequently, there is a need to develop distinct communication strategies for each phase of the process to guide users in making effective privacy-related decisions and managing corresponding consequences. In addition, it is vital to design the strategies by recognizing that users may operate with inaccurate or incomplete mental models stemming from inaccurate or outdated information. For example, supporting users during each phase of the privacy-related decision-making process could employ the following communication strategies: i) inform users about privacy issues (*pre-decisional*); ii) confirm whether users would like to be more aware of privacy issues (*post-decisional*); iii) alert users when their privacy is at risk (*actional*); and iv) provide users with feedback on how their decision affected their privacy (*post-actional*).

Supporting Reversal of Past Decisions

It is essential to acknowledge and address that users can make incorrect decisions. Yet, current systems often make it difficult to reverse prior privacy-related decisions, even though restrictions on reversal are typically unnecessary. The value and potential of reversible decisions are highlighted by recent advances, such as the functionality allowing users to withdraw or delete sent messages in messaging apps. Research on the reversal of past privacy-related decisions could mimic such functionality. Making privacy-related decisions easier to change in practice could alleviate the stress associated with making privacy-related decisions while helping address habituation at the same time.

Theme 3: Privacy and AI

As advances in AI continue rapidly, the role of AI in human-centered privacy is increasingly a part of the public debate. In the public's perspective on AI, privacy is a significant concern [6]. The interplay between privacy and AI has a dual dynamic. On the one hand, the inherent tendency of AI to magnify the scope and intricacies of data-related issues can lead to notable reconfiguration and amplification of known and established privacy issues and raise novel privacy concerns. On the other hand, AI could help address existing privacy issues. For instance, AI can be applied to learn user privacy preferences and help users manage privacy accordingly. This dual dynamic makes it complex to investigate privacy issues in practical systems incorporating AI and, in turn, engineer solutions that aid effective privacy management in such systems. Privacy-related challenges in the AI context could be categorized into the following six high-level themes.

Providing User Control

The lack of control over AI-driven decisions that impact privacy poses a prominent challenge. It can sometimes be difficult to determine where the control lies. AI can use personal data in unexpected ways that exacerbate existing privacy concerns and the consequent potential harm. Striking a balance between protecting privacy by minimizing data collection and collecting sufficient data to enable personalized AI is a key consideration for future explorations in research and practice.

Supporting Transparency and Explainability

AI systems that can impact privacy need greater transparency and explainability to comprehend AI decision-making and attribute responsibility for errors. The opaque nature of AI algorithms poses challenges in elucidating inherent biases in the data and algorithms and their potential impact on privacy. The negative impact of AI algorithms on the privacy of those outside the status quo is particularly concerning because of the potential to amplify and further entrench existing privacy harms. AI systems must foster trust and accountability by including clear explanations, avoiding exaggerated claims, and providing the rationale behind decisions. Practitioners need to prioritize explainable AI techniques in high-stakes privacy scenarios in particular.

Empowering Lay Users

Addressing the power and knowledge imbalances between experts and non-experts is necessary to support human-centered privacy in the AI domain. In practice, this requires clear, usable, and understandable privacy controls for lay users that align with their privacy needs. Further, the inclusion of diverse user groups in the design and development of AI systems is essential to ensure equitable privacy outcomes. Underrepresented communities may experience distinct barriers in accessing and benefiting from privacy protections within AI systems (see Theme 1 for inclusive privacy). An inclusive approach must be coupled with the obligation to conduct research and pursue business objectives without causing harm or (mis)using AI for deception, such as spreading disinformation, generating deceptive user interfaces, etc.

Considering Ethics and Responsibility

Ethical considerations surrounding the development and deployment of AI are essential when addressing privacy issues in AI systems. Currently, it is often unclear who should receive reports regarding privacy-related harms observed or experienced when using AI systems. In addition, it is essential to be proactive about safeguarding privacy rights by anticipating and mitigating potential harms and long-term consequences when designing AI systems and deploying them in practice. It is imperative to recognize that ethics and privacy are deeply entangled because privacy is heavily linked to *values*.

Developing AI-informed PETs

AI may simultaneously contribute to privacy harms *and* solutions. The unique privacy risks posed by AI, such as model inversion and membership inference attacks, necessitate the development of new PETs. For instance, protecting privacy in the age of AI may require new approaches designed to safeguard against these AI-specific vulnerabilities. In this regard, it is also necessary to pay attention to the risks associated with AI systems providing deceptive information or exploiting human weaknesses. In contrast, AI could be applied to improve and advance privacy protection by developing AI-enhanced PETs. For example, smart privacy assistants could help people manage their privacy more effectively [7].

Enhancing Data Governance

Balancing the tension between collecting data to train and operate AI systems and protecting privacy is a key challenge for research and practice. Currently, there is a significant opportunity to develop business models centered on sourcing *privacy-respecting, high-quality* data. In this context, the role of public, open data and responsible data stewardship warrants further exploration. Examining the implications of data governance and challenging the existing power dynamics in data-centric AI ecosystems is essential for human-centered privacy.

Theme 4: Multi-user Privacy

Multi-user and interdependent privacy issues arise when an individual's privacy is affected by information sharing by others [8]. In addition, such privacy concerns can arise because of voluntary or involuntary sharing of devices, online accounts, or physical spaces among multiple individuals.

Accommodating Shared Devices, Accounts, and Spaces

It is well known that people share devices and accounts. For instance, conceptualizations of a smart home inherently assume the sharing of devices in a communal living space. Yet, in practice, the one-user-one-account assumption permeates through all stages of system design, even for smart home devices. Similarly, design personas typically describe individual users, ignoring people's natural interconnectedness. In fact, there is a significant mismatch between the architecture of most devices and services and the realities of everyday life within families, communities, and societies. Bridging the gap between system mechanisms designed with the single-user assumption and the need to support interconnectedness and sharing is a significant challenge, with cascading effects on the privacy of individuals and collectives. Although solutions to address the challenge have started to emerge in research, these typically lack maturity and struggle to gain user acceptance. More attention from practitioners is needed to promote the mainstream adoption of such solutions.

Considering the context of the smart home as an example, situations such as caregiving and illness may necessitate sharing data and access with partners, families, friends, and community members. Conversely, previously shared data may need to be disentangled in certain circumstances, such as when children reach adulthood or romantic relationships end. The question of how users or providers should manage data and devices through continually evolving human relationships is not yet settled. Moreover, the presence of third parties, such as guests or domestic workers, introduces additional privacy concerns in smart homes. Designing smart home systems that protect the privacy of inhabitants, guests, and bystanders *from each other* is an ongoing challenge for researchers and practitioners.

Considering the Scale and Context of Data Collection by Others

Issues of multi-user privacy extend beyond individual devices or smart homes, arising at multiple scales, from specific social network platforms [9] to smart cities. Due to significant differences in nuance and complexity in these contexts, solutions that work well at one scale may not be effective at another. Responsibly designed AI and machine learning systems hold promise for supporting individuals and groups in making collective decisions about privacy across domains and scales. Yet, the promise comes with associated privacy challenges, as stated in Theme 3.

While privacy protection in public spaces has been an ongoing debate in research and practice, it has gained new significance with the advent of facial recognition capabilities in smartphones, smart glasses, and surveillance cameras. Further, widespread data collection can make it challenging to enforce privacy-related laws and regulations and to prevent unauthorized data collection and reuse. Public policy should incentivize innovative practical solutions, making the design goal of protecting privacy not only necessary but also profitable.

Protecting Against Inferences

Last but not least, *inferences* can be made about one user based on data from another. For example, it is nearly impossible to avoid making inferences about a user based on the information shared by the person's connections on social network services. Such privacy risks can be experienced even by those who do not use social network services because inferences can be made about them based on the information disclosed by others they know who use the services. Similarly, the decision of one person to sequence and share DNA (Deoxyribonucleic Acid) with a genetic database can inadvertently implicate all relatives of the person, even distant ones, since they share a substantial portion of the DNA [10]. Multi-user privacy concerns connected with inferences derived from the behavior of others have persisted for decades because they are difficult to avoid. Yet, protections against these inferences are essentially non-existent in practice.

Conclusion

In this article, we have delineated the most salient challenges in human-centered privacy organized around four themes: inclusive privacy, privacy-related decisions and communication, privacy and AI, and multi-user privacy. The description of these challenges is a call to action to motivate privacy researchers and practitioners to: (i) study and bolster the privacy of at-risk populations responsibly and ethically; (ii) design and develop better support and communication for (reversible) privacy-related decision-making; (iii) empower users to manage their privacy within AI systems by promoting transparent, explainable operation;

and (iv) embrace and support the multi-user nature of privacy. This article provides a starting point for those in academia, industry, or government interested in advancing the agenda of human-centered privacy and guiding human-focused research and practice to address privacy issues in modern technologies.

Acknowledgments

We are grateful to thank Emilee Rader for her feedback on draft versions of this article and for her participation in the discussions during the FHCP event. We thank Florian Alt, L. Jean Camp, Simone Fischer-Hübner, Elise Leray-Meyer, Alena Naiakshina, Elissa Redmiles, Heather Richter Lipford, and Paul Overbury for participating in the FHCP event. The article benefited from editorial input from the participants of the Human-Centered Computing seminar at the University of Utah. We are grateful to the Engineering and Physical Sciences Research Council (EPSRC) and the National Cyber Security Centre (NCSC) for supporting the FHCP event through EP/S018972/1.

References

[1] A. E. Marwick and d. boyd, “Understanding privacy at the margins — Introduction,” *International Journal of Communication*, vol. 12, pp. 1157–1165, 2018. [Online]. Available: https://ijoc.org/index.php/ijoc/article/view/7053.

[2] N. Warford, T. Matthews, K. Yang, *et al.*, “SoK: A framework for unifying at-risk user research,” in *2022 IEEE Symposium on Security and Privacy*, ser. IEEE S&P 2022, 2022, pp. 2344–2360. DOI: 10.1109/SP46214.2022.9833643.

[3] R. Bellini, E. Tseng, N. Warford, *et al.*, “SoK: Safer digital-safety research involving at-risk users,” in *2024 IEEE Symposium on Security and Privacy*, ser. IEEE S&P 2024, 2024, pp. 635–654. DOI: 10.1109/SP54263.2024.00071.

[4] A. Vance, B. Kirwan, D. Bjornn, J. Jenkins, and B. B. Anderson, “What do we really know about how habituation to warnings occurs over time? A longitudinal fMRI study of habituation and polymorphic warnings,” in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’17, Denver, Colorado, USA: Association for Computing Machinery, 2017, pp. 2215–2227, ISBN: 9781450346559. DOI: 10.1145/3025453.3025896.

[5] A. Achtziger and P. M. Gollwitzer, “Motivation und Volition im Handlungsverlauf,” in *Motivation und Handeln*, J. Heckhausen and H. Heckhausen, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2018, pp. 355–388, ISBN: 978-3-662-53927-9. DOI: 10.1007/978-3-662-53927-9_12.

[6] M. Tahaei, M. Constantinides, D. Quercia, and M. Muller, *A systematic literature review of human-centered, ethical, and responsible AI*, 2023. arXiv: 2302.05284 [cs.HC]. [Online]. Available: https://arxiv.org/abs/2302.05284.

[7] I. Krsek, A. Kabra, Y. Dou, *et al.*, *Measuring, modeling, and helping people account for privacy risks in online self-disclosures with AI*, 2024. arXiv: 2412.15047 [cs.HC]. [Online]. Available: https://arxiv.org/abs/2412.15047.

[8] M. Humbert, B. Trubert, and K. Huguenin, “A survey on interdependent privacy,” *ACM Comput. Surv.*, vol. 52, no. 6, Oct. 2019, ISSN: 0360-0300. DOI: 10.1145/3360498.

[9] J. M. Such and N. Criado, “Multiparty privacy in social media,” *Commun. ACM*, vol. 61, no. 8, pp. 74–81, Jul. 2018, ISSN: 0001-0782. DOI: 10.1145/3208039.

[10] M. Humbert, E. Ayday, J.-P. Hubaux, and A. Telenti, “Quantifying interdependent risks in genomic privacy,” *ACM Trans. Priv. Secur.*, vol. 20, no. 1, Feb. 2017, ISSN: 2471-2566. DOI: 10.1145/3035538.