Report from Dagstuhl Seminar 25261

Future of Human-Centered Privacy

Zinaida Benenson[*1], Simone Fischer-Hübner[*2], Heather Richter Lipford[*3], and William Seymour[*4]

1 **Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), DE.** zinaida.benenson@fau.de
2 **Karlstad University, Chalmers University of Technology & University of Gothenburg, SE.** simone.fischer-huebner@kau.se
3 **University of North Carolina at Charlotte, US.** heather.lipford@uncc.edu
4 **King's College London, UK.** william.seymour@kcl.ac.uk

Abstract

The Dagstuhl Seminar on The Future of Human-Centered Privacy (25261), held from June 22–27, 2025, brought together researchers from academia and industry to discuss key issues at the intersection of privacy and human-computer interaction (HCI) research. This article summarizes the main discussion topics, and presents the summary of the outputs of five working groups that discussed: i) Measurement, Methods, and Ethics; ii) Supporting Developers; iii) AI for Privacy/Privacy for AI; iv) Consent, Control, and Communication; and v) Collective Privacy. This seminar was a continuation of a previous seminar held at King's College London on June 5–7, 2023 which laid the groundwork for the present seminar through it's discussion on the topics of inclusive privacy, multiuser privacy, privacy and AI, and privacy communication.

Seminar June 22–27, 2025 – https://www.dagstuhl.de/25261
2012 ACM Subject Classification Security and privacy → Human and societal aspects of security and privacy
Keywords and phrases Privacy, Human-computer Interaction, AI
Digital Object Identifier 10.4230/DagRep.15.6.84

1 Executive Summary

Zinaida Benenson (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), DE)
Simone Fischer-Hübner (Karlstad University, Chalmers University of Technology & University of Gothenburg, SE)
Heather Richter Lipford (University of North Carolina at Charlotte, US)
William Seymour (King's College London, UK)

Human-centered privacy resides in the intersection of privacy and human-computer interaction (HCI) research. It investigates users' privacy perceptions, concerns, and awareness in various settings, and also the understanding, usefulness, and usage of various privacy-enhancing technologies. On the one hand, the advance of Internet of Things, smart spaces, and AI have raised new questions that need to be investigated, e.g., how to negotiate privacy settings in the presence of different users of the same system, or how to improve the transparency of AI systems. On the other hand, there are many questions that have been explored for decades, but need to be adapted to these new areas and domains, such as "What is a privacy decision?" and "What information do users need to make a privacy decision?". Moreover, the multitude

[*] Editor / Organizer

of users also includes at-risk and vulnerable populations that interact (sometimes unwillingly or unknowingly) with digital systems, and require additional research to understand their needs.

This seminar brought together academic and industry experts from a range of disciplines to discuss these issues, which had been seeded at a preceding workshop hosted at King's College London in 2023.

Seminar participants initially convened for an opening session with the project organisers that included a short round of introductions with sharing of current research interests by all participants. Following this, there was a talk from Sören Preibusch with a practical perspective on human-centred privacy and a panel debate on Privacy in the age of AI. During the short introductory presentation sessions, invited talk, and the panel sessions, seminar participants noted key human-centred privacy aspects of interest and in need of future research. These notes were then clustered by a small group of participants into ten different themes. After voting and discussion, working groups were formed to elaborate on the following five:

- Measurement, Methods, and Ethics
- Supporting Developers
- AI for Privacy, Privacy for AI
- Consent, Control, and Communication
- Collective Privacy

These working groups met throughout the remaining duration of the seminar, periodically reporting back to the main group. The outputs of the five groups form the main body of this seminar report, and include an overview of the main open problems in the area, an overview of current approaches, and promising directions for future work. These take the form of key research questions, potential solutions in need of study, and roadmaps for the development of research capacity.

To round up the seminar, all working groups presented their future work directions on the last day. Afterwards, we discussed future intergroup activities: closed workshops, open workshops (e.g., at SOUPS or CHI conferences), joint projects and publications, possible funding (e.g., EU Cost Actions, Erasmus+) and research exchange visits at various universities.

2 Table of Contents

3 Overview of Talks and Panels

3.1 Human-centered Privacy in Practice

Sören Preibusch (Bundesinstitut für Risikobewertung – Berlin, DE, soeren.preibusch@bfr.bund.de)

In my talk, I deliver a practical perspective on human-centred privacy, drawing on my experience in a large tech company, a data protection supervisory authority, and a federal agency where I'm currently the head of IT. Whereas the goals in each of the roles vary, there is often both an operational and a strategic aspect to the work and we must consider three questions when building for privacy: What are we optimising for (e.g., compliance, improved experiences)? Whom are we designing for (e.g., employees, citizens, users)? Who is the adversary (e.g., external attackers, insider threats)?

I conclude that the practical effort in achieving good privacy arises both at the product-level (to be solved through managerial decisions) and at the programme-level (to be solved through capacity building).

I encourage our seminar to embark on research that tackles two challenges: First, how might we focus human labour on the most value-added activities – and automate the rest? Second, how might we do privacy right when it's not the primary task?

3.2 Panel Discussion on "Privacy in the age of AI"

Bettina Berendt (Weizenbaum Institute, TU Berlin, DE & KU Leuven, BE, berendt@tu-berlin.de)
Marc Langheinrich (Università della Svizzera italiana – Lugano, CH, marc.langheinrich@usi.ch)
Cristiana Santos (School of Law, Utrecht University, NL, c.teixeirasantos@uu.nl)

3.2.1 Statements

Marc Langheinrich

There will be more AI-generated attacks and scams which we will need to address. There is also the potential for AI to be used as personalized privacy/security advisors to keep users informed of these risks. For example, AI could be used to give people a better idea of what the consequences and inferences of their data disclosures.

Cristiana Santos

Developers need to make privacy-related decisions without sufficient legal guidance. They may now use LLMs, whose answers can be generic, lack legal relevance, and be unreliable. Thus, the community needs to help developers to not overly rely on AI.

Bettina Berendt

Thesis 1: We need to consider both "users" and "privacy" in a larger sense if we want to explore "the future": (a) from users to stakeholders or "affected persons" (as in the AI Act); (b) from privacy as confidentiality or control to "the freedom from unreasonable constraints on the construction of one's identity" – and maybe move beyond this focus on identity.

Thesis 2: Right now, we (whether as communities, countries, or humankind) arguably have bigger fish to fry than – to be polemic – avoiding unwanted ads: 1. at least putting restrictions on the ever-increasing concentration of economic (and with it political) power, a concentration much furthered by AI; 2. digital sovereignty, including cybersecurity for critical infrastructures; 3. the survival of the planet; 4. (re)instating the rule of law, internally, in international relations, and with a view to human rights. The bad news is that threats to all four appear to be spurred on by AI.

Thesis 3: The good news is that the protection of personal data, in particular via the core principle of data minimisation, can serve as a starting point in the fight for 1., 2. and 3. As regards 4., this probably exceeds the scope of our discussion here ...

3.2.2 Discussion

The seminar participants discussed the challenge of defining privacy when an AI model is the potential adversary. This may lead to different or new interpretations of privacy. Who and how we address these challenges will depend on how stakeholders define privacy.

4 Working Group Reports

4.1 Measurement, Methods, and Ethics

Cori Faklaris (University of North Carolina – Charlotte, US, cfaklari@charlotte.edu)
Yixin Zou (Max Planck Inst. for Security and Privacy – Bochum, DE, yixin.zou@mpi-sp.org)
Adam Jenkins (King's College – London, UK, adam.jenkins@kcl.ac.uk)
Anastasia Sergeeva (University of Luxembourg, LU, anastasia.sergeeva@uni.lu)
Apu Kapadia (Indiana University Bloomington, US, kapadia@iu.edu)
Daricia Wilkinson (Arizona State University – Tempe, US, daricia.wilkinson@asu.edu)
Sameer Patil (University of Utah – Salt Lake City, US, sameer.patil@utah.edu)
Simran Munot (Max Planck Institute for Informatics, DE, smunot@mpi-inf.mpg.de)

4.1.1 Introduction

We explored the challenges related to methodological rigor, participant consent, the validity of new data sources, and the ethical boundaries of AI-generated insights. Our discussions also covered foundational issues in research, including the reporting of demographics, the evolution of privacy education, and the persistent gaps in how we measure and evaluate privacy itself. This summary outlines the key problems, the current state of the art, open research questions, and a potential roadmap for the HCP research community.

4.1.2 Problems and Challenges

AI-Mediated Research. The use of AI tools, such as large language models (LLMs) for thematic analysis or AI companions for transcription, introduces questions of rigor, bias, and consent. There is significant scepticism about whether participants can genuinely consent to AI-mediated processes they may not fully understand and whether these methods can preserve the empathetic, nuanced core of qualitative inquiry. Furthermore, publishing findings based on AI-generated or synthetic data carries risks to authenticity and may undermine trust in research outcomes.

Synthetic Data. While industry interest in using synthetic data for UX research is growing, this approach risks undermining the authenticity and empathy that are central to qualitative research. Synthetic data may flatten or misrepresent the nuanced experiences of real users, and its use as a privacy-enhancing technology is being questioned, as it may not be immune to linkage attacks or attribute inference.

Demographics Reporting. The HCP field lacks common guidelines for collecting and reporting demographic data. This is complicated by GDPR constraints that mandate data minimization, which can conflict with the goal of comprehensive demographic reporting for transparency and bias detection. Moreover, first-order attributes such as age and gender are often coarse proxies for the lived experiences (e.g., discrimination) that are truly relevant.

Measurement and Evaluation Gaps. Many established scales and measurements for privacy concepts (e.g., privacy concern) are being challenged. Recent work shows that participants often interpret survey items differently than researchers intend, questioning the validity of these instruments. The rapid evolution of AI is likely further changing privacy expectations and behaviors, making our current measurement tools potentially obsolete.

Education and Curriculum. The integration of AI into education fosters an over-reliance on tools that can diminish critical thinking and problem-solving skills. It also introduces new avenues for academic dishonesty and erodes the essential human interaction in learning. For privacy education specifically, there is an urgent need to update curricula to address the complexities and interplay of AI, data transparency, and the often-flawed user mental models of how these systems operate.

4.1.3 State of the Art

The current state of HCP is similar to that of computing as a whole. It is characterized by rapid technological advancement running ahead of methodological and ethical consensus. Notably, the **use of AI in research** is seen as inevitable and is already widespread, with researchers using it for faster data processing and code development. However, studies comparing human and AI coders on qualitative data to date show significant discrepancies and only moderate overlap, highlighting reliability issues. Similarly, the **use of synthetic data** has been proposed for decades and applied in fields like medicine and political science. However, recent studies in HCI show that LLM-generated data, while plausible, is less diverse, prone to factual errors, and can contain biases (e.g., related to age) compared with human-generated data.

While the HCP field lacks its own standards for **demographics reporting**, other fields such as computing education and clinical research have established guidelines that standardize such publications and make them easily comparable. Specific guidelines are also available for asking about sensitive dimensions such as gender and ethnicity. However, broader societal shifts in understanding identity may require these to be updated.

A significant body of work exists on **measuring privacy concern and related constructs**, originating from law, public policy, and information systems. However, recent critiques using techniques like corpus linguistics have challenged the construct validity of many widely used scales, revealing a disconnect between researcher intent and participant interpretation.

AI in Education (AIED) is increasingly prevalent for its potential to enhance teaching and learning. However, the focus has often been on the technological implementation, with a growing body of work now highlighting the negative impacts on human cognition, critical thinking, and academic integrity.

4.1.4 Key Research Questions

To move forward, our community must address several key research questions:

1. **AI-Mediated Research & Synthetic Data:**
   - Under what circumstances is it feasible and ethical to use synthetic data or AI-assisted analysis in HCP research?
   - How can we develop benchmarks to validate the quality, representativeness, and authenticity of synthetic data?
   - How can we design and implement practices that embed meaningful transparency and trust-building into AI-mediated research, ensuring genuine participant consent?
   - How can we best collect, de-identify, and disseminate inclusive datasets for training AI models to offset dominant Western-focused cultural biases and amplify marginalized voices?
2. **Demographics and Measurement:**
   - How can we, as a community, develop guidelines to balance the goals of scientific transparency and rigor with the ethical needs for participant privacy and data minimization, especially under constraints like GDPR?
   - Moving beyond coarse proxies such as age and race, which are second- and third-order attributes (e.g., experienced discrimination, technical literacy) should we focus on to better understand our participants and their contexts?
   - How can we develop and validate new measurement instruments for privacy constructs that are robust to misinterpretation and reflect contemporary understandings of privacy in an AI-driven world?
3. **Education:**
   - How should we redesign HCP curricula to teach students not just how to *use* AI tools, but how to *critically reflect* on their societal, ethical, and personal impact?
   - What competency frameworks are needed to define what students should know (awareness) and be able to do (proficiency) regarding AI and privacy?

4.1.5 Solution Ideas and Directions

Addressing these questions requires a multi-faceted approach combining technical, methodological, and community-driven efforts.

- *Develop Community Guidelines and Standards:* The HCP community should work towards shared guidelines for the ethical use of AI in research, the reporting of demographics, and the validation of synthetic data. This could involve requiring explicit discussions of these considerations in paper submissions and review forms.

- *Adopt a "Human-in-the-Loop" Ethos:* For both research and education, we must emphasize critical human oversight. In research, this means not blindly accepting AI-generated analysis but using it as a tool to augment human intellect. In education, this means designing assignments where students must reflect on and critique the AI's output, with the student's critical judgment being the final word.
- *Create and Validate New Instruments and Taxonomies:* We need to develop and empirically validate new survey instruments for privacy. Additionally, creating a taxonomy of first-, second-, and third-order demographic attributes could help researchers make more intentional choices about data collection.
- *Establish Educational Resources:* There is an urgent need for repositories of syllabi, course resources, and hands-on learning modules for HCP education. These resources should be modelled on successful initiatives in related fields (e.g., usable security) and should explicitly address AI ethics and critical reflection.
- *Promote Transparency:* Researchers should be transparent about their methods, including documenting when and how AI was used. [Ex: The first author used Gemini Pro and Copilot for Education for helping to condense and aggregate the working group's documents.] For demographics, this includes reflecting on why certain attributes were collected and others were not. For education, this may mean making it a rule that students must disclose their use of AI.

To carry out this work will require resources for supporting the actual research and for developing and sustaining the types of people required. It will involve interdisciplinary collaboration between HCI researchers, ethicists, legal scholars (especially those specializing in data protection), data scientists, and educators. Skills in psychometrics, corpus linguistics, and qualitative methods are essential for developing and validating new measurements and analysis techniques. Furthermore, the HCP community needs funding to support the creation of inclusive datasets, the development of new measurement tools, and the organization of community-building events. We also need robust, shared infrastructure, such as secure repositories for data and educational materials.

Finally, we encourage recognizing the development of datasets, tools, and curricula as valuable academic contributions (on par with traditional publications). This will align incentives with the outlined agenda and with their importance for actually being able to create a more secure, private, and trustworthy future of computing.

4.1.6 Roadmap Development

Towards addressing the above concerns, we envision a progression of research and community action over the next decade.

Immediate (Next 1-2 Years).

- **Initiate Dialogue:** Form working groups and organize workshops at major HCI conferences (e.g., CHI, SOUPS, CSCW) dedicated to creating draft guidelines for AI in research and demographic reporting.
- **Curriculum Redesign:** Begin redesigning individual courses to incorporate critical AI reflection and hands-on learning modules. Start building an ad-hoc, shared repository for syllabi and course materials.
- **Incorporate into CFPs:** Program committees for major venues can begin to encourage or require reflection statements on the use of AI and demographic reporting choices in submissions.

Next Three Years.

- **Publish Community Guidelines:** Formalize and publish initial versions of community guidelines for AI use and demographic reporting.
- **Develop Competency Frameworks:** Establish a clear competency framework for AI and privacy education, defining core knowledge and skills.
- **Validate New Measures:** Conduct large-scale validation studies of new and existing privacy measurement instruments to establish a new set of reliable tools for the community.

Next Decade.

- **Establish Accreditation and Standards:** Work towards formal accreditation for educational programs that teach AI and privacy, similar to standards in cybersecurity.
- **Build Robust Infrastructure:** Develop and maintain community-wide infrastructure, including repositories for inclusive, de-identified datasets for training AI models, and platforms for sharing validated educational resources.
- **Longitudinal Studies:** Conduct long-term research tracking how technological affordances (like human-like chatbots) impact privacy behaviors and how societal understandings of privacy evolve over time.

References

1. Abdulrahman M. Al-Zahrani and Talal M. Alasmari. 2024. Exploring the impact of artificial intelligence on higher education: The dynamics of ethical, social, and educational implications. Humanities & social sciences communications 11, 1: 1–12. https://doi.org/10.1057/s41599-024-03432-4
2. Michael Benisch, Patrick Gage Kelley, Norman Sadeh, and Lorrie Faith Cranor. 2011. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. Personal and Ubiquitous Computing 15, 7: 679–694. https://doi.org/10.1007/s00779-010-0346-0
3. Jessica Colnago, Lorrie Faith Cranor, Alessandro Acquisti, and Kate Hazel Stanton. 2022. Is it a concern or a preference? An investigation into the ability of privacy scales to capture and distinguish granular privacy constructs. and Security (SOUPS 2022). Retrieved from https://www.usenix.org/conference/soups2022/presentation/colnago
4. Serge Egelman and Eyal Peer. 2015. Predicting Privacy and Security Attitudes. SIGCAS Comput. Soc. 45, 1: 22–28. https://doi.org/10.1145/2738210.2738215
5. Perttu Hämäläinen, Mikke Tavast, and Anton Kunnari. 2023. Evaluating Large Language Models in Generating Synthetic HCI Research Data: a Case Study. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 433, 1–19. https://doi.org/10.1145/3544548.3580688
6. Zhehui Liao, Maria Antoniak, Inyoung Cheong, Evie Yu-Yen Cheng, Ai-Heng Lee, Kyle Lo, Joseph Chee Chang, and Amy X. Zhang. 2024. LLMs as research tools: A large scale survey of researchers' usage and perceptions. arXiv [cs.CL]. Retrieved from http://arxiv.org/abs/2411.05025
7. Abdul Majeed. 2023. Attribute-centric and synthetic data based privacy preserving methods: A systematic review. Journal of Cybersecurity and Privacy 3, 3: 638–661. https://doi.org/10.3390/jcp3030030
8. Naresh K. Malhotra, Sung S. Kim, and James Agarwal. 2004. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. Information Systems Research 15, 4: 336–355. https://doi.org/10.1287/isre.1040.0032

9 Sameer Patil, Greg Norcie, Apu Kapadia, and Adam J. Lee. 2012. Reasons, rewards, regrets: privacy considerations in location sharing as an interactive practice. In Proceedings of the Eighth Symposium on Usable Privacy and Security. https://doi.org/10.1145/2335356.2335363

10 Sören Preibusch. 2013. Guide to Measuring Privacy Concern: Review of Survey and Observational Instruments. Int. J. Hum. -Comput. Stud. 71, 12: 1133–1143. https://doi.org/10.1016/j.ijhcs.2013.09.002

11 Anastasia Sergeeva, Björn Rohles, Verena Distler, and Vincent Koenig. 2023. "We Need a Big Revolution in Email Advertising": Users' Perception of Persuasion in Permission-based Advertising Emails. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 652, 1–21. https://doi.org/10.1145/3544548.3581163

12 Sarah Spiekermann and Lorrie Faith Cranor. 2009. Engineering Privacy. IEEE transactions on software engineering 35, 1: 67–82. https://doi.org/10.1109/tse.2008.88

13 Sarah Tabassum, Nishka Mathew, and Cori Faklaris. 2025. Privacy on the Move: Understanding Educational Migrants' Social Media Practices through the Lens of Communication Privacy Management Theory. In Proceedings of the ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS '25). Association for Computing Machinery, New York, NY, USA, 1–18. https://doi.org/10.1145/3715335.3735453

14 Mohammad Tahaei, Adam Jenkins, Kami Vaniea, and Maria Wolters. 2021. "I don't know too much about it": On the security mindsets of Computer Science students. arXiv [cs.CR], 27–46. https://doi.org/10.1007/978-3-030-55958-8_2

15 Kurt Thomas, Patrick Gage Kelley, David Tao, Sarah Meiklejohn, Owen Vallis, Shunwen Tan, Blaž Bratanič, Felipe Tiengo Ferreira, Vijay Kumar Eranti, and Elie Bursztein. 2025. Supporting Human Raters with the Detection of Harmful Content using Large Language Models. 2772–2789. Retrieved from https://ieeexplore.ieee.org/abstract/document/11023319/

16 Richard Van Noorden and Jeffrey M. Perkel. 2023. AI and science: what 1,600 researchers think. Nature 621, 7980: 672–675. https://doi.org/10.1038/d41586-023-02980-0

17 Daricia Wilkinson, Moses Namara, Karla Badillo-Urquiola, Pamela J. Wisniewski, Bart P. Knijnenburg, Xinru Page, Eran Toch, and Jen Romano-Bergstrom. 2018. Moving Beyond a "one-size fits all": Exploring Individual Differences in Privacy. In Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems, 1–8. Retrieved from https://dl.acm.org/doi/abs/10.1145/3170427.3170617

18 Pamela J. Wisniewski, Bart P. Knijnenburg, and Heather Richter Lipford. 2017. Making privacy personal: Profiling social network users to inform privacy education and nudging. International journal of human-computer studies 98: 95–108. https://doi.org/10.1016/j.ijhcs.2016.09.006

4.2 Supporting Developers

Sören Preibusch (BfR – Berlin, DE, soeren.preibusch@bfr.bund.de)
Nataliia Bielova (Inria – Rennes, FR, nataliia.bielova@inria.fr)
Dominik Herrmann (University of Bamberg, DE, dominik.herrmann@uni-bamberg.de)
Alena Naiakshina (University of Cologne, DE, alena.naiakshina@uni-koeln.de)
Cristiana Santos (School of Law, Utrecht University, NL, c.teixeirasantos@uu.nl)
Ha Dao (MPI-INF – Saarbrücken, DE, hadao@mpi-inf.mpg.de)
Bettina Berendt (Weizenbaum Institute, TU Berlin, DE & KU Leuven, BE, berendt@tu-berlin.de)

4.2.1 Introduction

Software products and services shape societies and peoples' lives. These software products don't magically appear – they are created by developers and then used by many more users, potentially by orders of magnitude more. Thus, the design and coding choices of developers determine the quality of software products – including the quality of privacy[1] implementations. This differentiates developers from consumers who usually only make privacy choices for themselves or a small number of others (see Chapter 4.5 on Collective Privacy). Unfortunately, many developers are ill-equipped to shoulder this broad impact: they lack both expertise and advice on privacy, and formal training in software engineering. In addition, being a developer can mean many different things, depending on the roles and the contexts: First, software development is never just about writing code. It encompasses the entire Software Development Lifecycle, from planning and design to deployment and maintenance, as well as all the internal processes, guidelines, and communication that enable and sustain that lifecycle. The resulting jobs to be done include writing code, testing, product and project management, requirements engineering and software architecting – each one associated with specific privacy questions and decisions [3]. In an ideal world, these jobs are distributed among experts who build specialized skills. In reality, governance structures differ and some jobs are combined, skipped, or executed by a "one-person band" [29].

Second, different contexts like the development of a new website or app, software maintenance, ad-hoc scripting to automate admin tasks, or database development all come with their own support needs. In each one of these contexts, developers routinely switch between producing code and consuming code (e.g., embedding third-party services or libraries, reusing code snippets).

4.2.2 Legal Background

This section defines the scope of relevant EU laws applicable to this report and offers conceptual legal clarifications, such as responsibility, accountability, data controller, data processor, liability.

General Data Protection Regulation (GDPR). The GDPR applies to the processing of personal data. Personal data refers to any information that renders a person identified or identifiable.

[1] We understand "privacy" as mediated via data protection.

Responsibility: Responsibility for compliance with GDPR obligations is determined based on the roles of the parties involved in the processing of personal data. The GDPR acknowledges the concept of "control" to assign responsibility for GDPR compliance (and administrative fines for non-compliance) to the data controller.

Controller: The controller is the entity that determines the "purposes" and "means" of processing personal data.

- Purposes: for example, developers determine the "purposes" of processing once they define the functionality needed for their own webpage/app service (either for compliance, site protection, user authentication, advertising, analytics, etc.).
- Means: concern more practical aspects of implementation, such as the choice of a particular type of hard- or software or the detailed security measures which may be left to the processor to decide on. Means are to be determined by the controller and processor. Developers determine the "means" when they embed that tool in its webpage service, thus triggering the start of the processing of personal data which would not be possible otherwise.
- Essential means: are traditionally and inherently reserved to the controller. "Essential means" are means that are closely linked to the purpose and the scope of the processing. Data controllers determine the "essential means" of processing by defining the type of data to be processed, the data recipients, data subjects and duration of processing. As such, a developer determines essential means (and it is thus a data controller) when they can terminate the processing, by simply removing the tool from its webpage.

The controller bears primary responsibility for compliance with GDPR principles. The controller is responsible for ensuring a legal basis for processing (e.g., consent), transparency and information provision, that data is processed for determined and specific purposes, data minimization, data protection by design and by default, responding to data subject rights, and adopting security measures. Case-law from the Court of Justice of the EU established that a given actor can be a controller if there is a purpose for data processing, even: (i) without having access to the data; (ii) without knowing that one processes personal data (not deliberately targeting personal data as such or wrongfully assessing that one does not process personal data); (iii) when choosing a third-party tool that allows the processing of personal data. These requirements for legal responsibility are important considering the challenges that several developers face, described in section 4.2.3, herein briefly referenced for clarity (in particular: privacy is not a priority for developers; developers do not feel responsible for privacy compliance; developers have difficulty understanding privacy requirements; developers depend on and trust third-party services; developers are manipulated by third-party tools).

Processor: The processor is the entity that processes data on behalf of the controller. The processor is responsible to follow only documented instructions from the controller, implement appropriate security, maintain records of processing, and notify data breaches to the controller.

Fines for GDPR infringements: When deciding on the amount of the administrative fine in each individual case, data protection authorities shall give due regard to the intentional or negligent character of the infringement; the nature, gravity and duration of the infringement taking into account the nature, scope or purpose of the processing concerned as well as the number of data subjects affected and the level of damage suffered by them. For especially severe violations, listed in Art. 83(5) GDPR, the fine can be up to 20 million euros, or in the case of an undertaking, up to 4% of their total global turnover of the preceding fiscal year, whichever is higher. But even the catalogue of less severe

violations in Art. 83(4) GDPR sets forth fines of up to 10 million euros, or, in the case of an undertaking, up to 2% of its entire global turnover of the preceding fiscal year, whichever is higher.

Liability for damages caused by a GDPR violation: A data subject has the right to compensation for material or non-material damage caused by a GDPR violation (Art. 82). A controller or processor is liable unless they can prove they were not responsible. Liability may be sole (if only one party was at fault), joint, or several (if multiple parties contributed).

Product Liability Directive (PLD). The PLD aims to protect consumers from defective products. It also covers digital products, such as software, digital manufacturing files, and AI systems placed on the market. Personal data is not a product under the PLD definition. The GDPR provides a dedicated liability framework for personal data.

Manufacturer: Manufacturers include both individuals and legal entities who develop, manufacture, or produce products, even if they create products solely for their own use. Additionally, anyone who designs a product or has it manufactured and then attaches their name to it ("quasi-manufacturer"), is also considered a manufacturer.

Liability for defective products: Manufacturers can be held liable if their products are defective. A product is deemed defective if it does not meet the expected or legally required level of safety. When assessing defectiveness, factors such as the product's presentation, characteristics, labelling, foreseeable use, and the impact of other products used in conjunction with it are taken into account (Art. 7). These requirements can stem from other EU laws, including the AI Act (AIA) or Cyber Resilience Act (CRA). The directive excludes pure violations of privacy.

4.2.3 Challenges Description

Basic setting: Developers' economic setting, attitudes, and knowledge. In this section, we examine how developers' economic context, organizational constraints, and knowledge limitations shape their attitudes and practices toward privacy.

- **Privacy is not a priority for developers.** Developers often do not prioritize privacy in their work. Prior studies have found that organizational dynamics can actively discourage developers from engaging with privacy. In particular, negative privacy climates – where privacy is viewed as a barrier or secondary concern – can shape developer behavior and expectations, disincentivizing attention to privacy concerns [10]. Beyond organizational culture, practical constraints such as limited budgets, tight deadlines, and competing feature demands further reduce the incentive to prioritize privacy during development [4]. Even among website owners, who may be considered a distinct subgroup of developers, privacy is not a consistent consideration. Their decisions are largely driven by business goals, personal motivations, and available resources, which often override concerns about data protection [29]. Notably, privacy rarely factors into their selection of third-party services, despite the privacy risks these integrations may pose [37].
- **Developers do not feel responsible for privacy compliance.** Developers often do not perceive themselves as responsible for ensuring the protection of privacy. Instead, privacy is frequently regarded as a legal matter, leading many developers to defer responsibility to legal or compliance teams. This perception creates a sense of limited control and can result in frustration when developers are required to engage with legal frameworks or implement regulatory requirements they do not feel empowered to influence [12]. Moreover, responsibility for privacy is often diffuse within organizations, so developers

may not know whom to turn to. Responsibility shifts between stakeholders – developers, product managers, legal counsel, and third-party service providers – without a stable point of accountability. This ambiguity complicates efforts to embed privacy into the development process and impedes the establishment of effective compliance mechanisms [31].

- **Developers have difficulty understanding privacy requirements.** Developers often face significant challenges in understanding and interpreting privacy requirements. A key obstacle to GDPR onboarding is the general lack of familiarity with its principles among developers [1]. Rather than engaging directly with privacy concepts, developers frequently adopt the vocabulary and mindset of data security to approach privacy-related tasks, which can lead to misaligned implementations [10]. In practice, many developers lack a comprehensive understanding of the behavior and implications of third-party SDKs they integrate into their applications [2]. Core principles such as data minimization, fairness, and data protection by design and by default are perceived as abstract and difficult to operationalize in everyday development practices. This lack of clarity is especially problematic as developers have legal responsibilities under data protection law if they become a data controller, see Section 4.2.2 and [8].

Relationships created by the software development life cycle and the state of the software development market. Developers have implicit relationships with various third-party software providers since today's developers (here, as code users) depend on third-party services and in most contexts, it has become almost impossible to build software without third-party components. However, developers also rely on third-party privacy and compliance solutions that are not always effective. Finally, third-party services are shipped with privacy-unfriendly default settings and sometimes manipulate developers towards non-compliance.

- **Developers depend on and trust third-party services.**
  Web ecosystem: Websites trust and heavily depend on various types of third-party tools and services. Research surveying top-500 popular websites in 50 countries together covering all inhabited continents found out that dependencies on a third-party Domain Name System (DNS) and on Content Distribution Networks (CDNs) or Certificate Authorities (CAs) provider vary widely around the world, ranging from 19% to as much as 76% of websites, across all countries [16]. However, there is a highly concentrated market of third-party providers: three providers across all countries serve an average of 92% of websites and Google, by itself, serves an average of 70% of the surveyed websites. Websites rely on third parties for useful and visible content: Ikram et al. found that 50% of first-party websites render content that they did not directly load [14]. Additionally, multiple studies have measured the presence of third-party tracking on thousands of websites even if it is not always clear whether such trackers were included directly by the website developers or included in functional third-party content. Moreover, new forms of such tracking is being included: recent research [20, 39] has shown that "nearly 90% of all websites use at least one tracking first-party cookie, 96% of which are in fact set by third-party scripts running in a first-party context" [38]. Recent work shows that even if website owners include very popular third-party services, the services often provide privacy policies and technical documentation that do not match each other, and moreover do not match the actual data that will be collected from end users by these third-party services. This was studied in the context of Google Tag Manager [18].

Mobile app developers: Mobile app stores require app developers to include a privacy policy. Yet, they do not provide information about how to write a privacy policy. Developers therefore often need to comply with regulations and requirements without knowing what needs to go into the privacy policies [34]. There is minimal research on helping developers craft privacy policies. The lack of support can be seen in the wild, where there are still numerous apps without privacy policies as well as privacy policies that contain misleading and contradictory statements [35]. Ad networks do not always provide guidance about what a developer should include in a privacy policy [33].

- **Developers delegate privacy and compliance to third-party solutions that are often not effective**.

  Compliance solutions: Website owners use "GDPR compliance solutions" that would scan their websites and identify privacy violations, such as presence of trackers, believing that these solutions would ensure compliance. In practice, such compliance solutions contain both false positives – deceiving website publishers into believing that a consent banner is needed on an empty website without trackers [36] – and false negatives – compliance solutions only scan cookies, but miss other Web tracking technologies, such as browser fingerprinting and hence data processed without legal basis [25]. In addition, only few tools help developers configure privacy solutions and disclosures, like user-friendly consent popups and accurate privacy policies, and even fewer tools take into account both regulations and the current behaviours of common third-party APIs [30].

  Privacy solutions: Research shows that developers believe that privacy-oriented libraries would provide an effective privacy solution without truly understanding functionality of such libraries. Song et al. [28] used qualitative methods and mental models approaches to analyze the differences between conceptual models used to design open-source Differential privacy (DP) libraries and mental models of DP held by users. They found that comparing developers' conceptual models with users' mental models elucidates crucial gaps between theory and practice of DP libraries.

- **Rather than privacy by default, there is manipulation against privacy settings and towards non-compliance.**

  No privacy by default in third-party libraries: Developers tend to follow the guidelines and requirements provided by the platforms [27, 35]. However, online services are often built with no compliance features embedded. Third-party services (libraries, SDKs) that Web/app developers rely on are very often not privacy-preserving and not privacy compliant by default [23]. Moreover, developers are reluctant to change the default settings: for example, in the context of mobile app SDKs, developers largely keep ad networks' SDK default configuration [19].

  Manipulation of developers by third-party tools: Developers (here, those using the code) are shown to be manipulated by the user interfaces of the compliance tools, such as Consent Management Platforms (CMPs) provided by third-parties, who nudge developers towards installing non-compliant cookie consent banners [36].

Expertise and developers' relationship to it. The following challenges regard the impact of the relationships in which developers seek and use advice.

- **Communication gaps between developers and privacy experts abound.** Communication gaps between developers and (legal) privacy experts (arising already from basics such as mismatches between legal and technological terminology and conceptual systems [26]) hinder the effective implementation of current GDPR compliant software [12]. In addition, developers tend to omit contacting privacy experts for support to

avoid burdening them [13]. The creation of knowledge resources such as repositories is often viewed as a remedy. However, such repositories, which – in spite of creating initial enthusiasm on all sides – tend to turn out costly and laborious to maintain, and thus tend to give rise to projects that are discontinued [26].

- **Developers use questionable online sources for code/third-party components and legal information.** As noted above, developers have difficulty understanding privacy requirements. They often turn to online sources, both for legal information and for code meant to implement legal requirements, in order to be compliant with data protection law. Traditionally, developers have turned to communities and their sites such as StackOverflow, and also Google or YouTube, for guidance [31]. These platforms have become informal spaces where developers discuss privacy, ethics, and values, highlighting a gap in accessible, authoritative resources [9].
  Currently, people are increasingly turning to chatbots for advice. A steep decrease in StackOverflow usage was noted already in 2023 [5]; this trend continues [21]. As a consequence, community knowledge becomes centralized and subject to non-transparent processing. More specifically, question and advice texts are processed as (re)training data of privately owned and difficult-to-scrutinize LLMs. In addition, by virtue of the personalized and ephemeral nature of chatbot dialogues, this knowledge becomes inaccessible for oversight by experts.
  Various research has examined how the release of ChatGPT-3.5 has affected user engagement on StackOverflow and other StackExchange platforms. A recent empirical study [11] arrived at a cautiously optimistic conclusion about interactions on community platforms and the likely quality of advice: While overall activity on StackOverflow declined sharply, dropping to less than 30% of its pre-ChatGPT levels, the effects varied across different types of users. Low-reputation users, who often ask basic questions, significantly reduced their participation, likely turning to Generative AI (GenAI) tools such as ChatGPT for quick answers. In contrast, high-reputation users maintained or slightly shifted their activity, engaging more with each other and focusing on more advanced, peer-level discussions. This suggests that GenAI may play a dual role: automating basic Q&A for novices while complementing expert-level knowledge sharing, which may help sustain valuable online communities and high-quality training data for AI.
  The quality of advice given by chatbots may be affected by various factors. For example, provider bias has been observed: GenAI models show systematic preferences for services from specific providers in their recommendations (e.g., favoring Google Cloud over Microsoft Azure) [40].
- **Further research challenges arise on a meta-level.** In addition to the problems for products and processes that are created by developers drawing on questionable advice, this last point also presents a methodological research challenge. In a world where developers are using GenAI tools instead of Stack Overflow for advice seeking, it is unclear how researchers can observe and measure their behavior. In addition, it is getting even more difficult to research artifacts/objects they are creating as was done for StackOverflow because these resources are not accessible any more and – due to different contexts – not reproducible independently. Also, it is not clear any longer which resources were produced by humans and which by GenAI. Finally, chatbot interactions allow for richer advice (tailored to personal pre-knowledge and scenario at hand) than generic advice, so it is unclear what the benchmarks are for measuring "quality": is it "correct or wrong answers" or "efficiency of interaction" or "specificity of advice"? Which metrics are appropriate and which ones can be measured?

- **Who should provide guidance for which types of developers?** Currently, there are several actors that provide privacy guidance for developers, such as regulatory bodies who provide recommendations and guidelines for best practices to implement the data protection law. In addition, companies, and standardization bodies (such as NIST) also provide guidance.

4.2.4 Key Research Questions

1. Responsibility: Who should be responsible for privacy compliance according to their role?
   - How to identify different user groups we consider "developers" in a specific context?
   - Should third-party providers be responsible since they both control the tools and data collection, even if they are not the ones controlling the primary application and therefore may argue that they are not responsible?
   - What organizational or cultural factors influence how developers understand and engage with privacy requirements – or don't?
2. Knowledge: What knowledge gaps and misconceptions do developers have about privacy, and how do these shape their practices?
   - What kinds of legal information do developers need but often lack?
   - What are the common misunderstandings or contradictions developers hold about privacy and user expectations?
   - How do developers' assumptions about user attitudes influence their privacy-related decisions?
3. Delivery of Guidance: What forms of guidance and support mechanisms are most effective in helping developers implement privacy practices?
   - What forms of guidance (lightweight or embedded) are most useful in practice?
   - What prevents developers from consulting privacy experts or using available resources?
   - Is it necessary to build closer relationships between developers and privacy experts?
   - Can AI tools meaningfully support developers in privacy tasks, and how reliable are they?
   - What form of guidance from third-party services would be useful to help developers understand the privacy implications in their decision-making?
4. Evaluation: How can we evaluate the effectiveness of privacy guidance and support for developers?
   - What criteria or metrics should be used to measure success (e.g., compliance, usability, adoption)?
   - What kinds of experimental or observational studies can assess whether developers actually use and benefit from guidance?
   - How can we simulate or test real-world developer behavior in privacy-critical scenarios?
5. Meta-analysis: In an age where advice seeking is happening in private spaces (chatbots), what are the appropriate ways to study developer behavior and traces they left in the public space (e.g., artifacts/objects – benchmark on correct or wrong)?
   - How to research what chatbot interaction looks like?
   - Can we expect developers to donate their data (e.g., GenAI-prompting history)? We need to consider research ethics in this context.
   - How about intellectual property issues and authorship?

4.2.5 Solution ideas and directions

Solutions range from foundational education to increasingly in-the-moment support.

1. Build capacity and educate programmers with a privacy-aware mindset:
   - The next generation of computer scientists needs holistic education and effective training, in contrast to generic annual privacy training. These efforts should hone a certain mindset rather than focus on teaching yet another modular skill.
   - This education and training faces uncertainty around the value of skill sets amid the increasing use of generative AI. A large-scale deskilling regarding critical thinking [17] may also endanger democratic agency. The risk of deskilling should be kept in mind, researched, and potentially countered.
2. Equip developers with safe software engineering practices:
   - Support mechanisms should deliver clarity of privacy goals through requirements engineering and its translation into software architecture, considering the mental models that developers have of end-users [7].
3. Provide contextual support:
   - IDEs should offer easily accessible contextual programming aids for implementing privacy – potentially leveraging AI-based code-generation and AI-based advice (akin to secure development lifecycle and existing tools like Copilot). Barriers to use need to be lowered, and quality assurance of the advice given to developers strengthened – regardless of how and where in the messy software development process people seek advice.
   - Part of the solution might be a repository of privacy-information-seeking behavior of developers that can be studied by researchers (especially important in a present and future in which developers seek advice from chatbots, i.e., not in the open, and every developer potentially getting personalized different kinds of advice).
   - Good privacy must be the easy choice: Dark patterns must not be perpetuated and there should be privacy-friendly options when consuming and updating third-party code [32], such as coarse-grained defaults and code snippets for location APIs [15]. Limited monetisation options can be a struggle towards better privacy [6].
   - Higher transparency regarding data practices (collection, processing, storage etc.) in the market of software products and for third-party libraries is needed, so that developers are empowered to choose wisely which libraries or services they link into their software. Appropriate metrics can support end-to-end accountability aided by transparency along the software supply chain. Similar consumer-oriented portals already exist, such as Mozilla's "Privacy not included" (which is a "naming and shaming portal for IoT devices").

4.2.6 Roadmap Development

Directions for for academia, policy-makers, industry should proceed along the following questions.

1. What immediate research questions need to be answered?
   - What motivates developers to care for privacy in different contexts?
   - How can we integrate privacy support in developers's day-to-day work?
   - What kind of guidance is effective for different kinds of developers?

- What kinds of privacy-related development tasks/requirements can be supported with existing off-the-shelf GenAI tools or coding aids that are used today? Where do they make mistakes?
- What could motivate third-party service providers to provide services with privacy-preserving defaults?

2. What are challenges and questions to be answered in the next three years?
   - How can GenAI tools be adapted/improved to provide helpful guidance for developers?
   - What role does the current and foreseeable market power of a small number of commercial Big GenAI tools play in the answer to the previous question?
   - How to motivate developers in public bodies (administration/governmental institutions) to care for privacy?
   - How to make privacy a priority task for developers?
   - How to motivate third-party service providers to provide compliant services?
3. What are challenges and questions to be answered within the next decade?
   - What are the effects of the quickly evolving GenAI tools, which changes how developers work and who develops software?
4. What skills and collaborations are necessary to address them?
   - Collaboration between scholars from computer science and law, psychology, ethics, economics
   - Collaboration with practitioners (controllers, processors) and data protection authorities
   - Collaboration with legal experts that consult companies on data protection requirements
   - Collaboration with industry bodies (like VDI/VDE or IHK/Handwerkskammer in DE) that advise organizations
5. What tools, infrastructure, funding, or incentives are needed?
   - Tools that assist developers in making more privacy-friendly decisions
   - Actual testbeds to try some new interventions. (technical support and funding for support staff)
6. What roadblocks may slow down or discourage research progress?
   - Perceived poor enforcement of GDPR violations may disincentivize developers to care for privacy.
   - Changes to GDPR that increase or reduce requirements for controllers.

Methodologically, this may work best by working with case studies. Candidate areas for developers are consent management products and tag managers.

References

1 Abdulrahman Alhazmi and Nalin Asanka Gamagedara Arachchilage. I'm all ears! Listening to software developers on putting gdpr principles into software development practice. *Personal and Ubiquitous Computing*, 25(5):879–892, 2021.

2 Noura Alomar and Serge Egelman. Developers say the darnedest things: Privacy compliance processes followed by developers of child-directed apps. *Proceedings on Privacy Enhancing Technologies*, 2022.

3 Rebecca Balebako and Lorrie Cranor. Improving app privacy: Nudging app developers to protect user privacy. *IEEE Security & Privacy*, 12(4):55–58, 2014.

4 Partha Das Chowdhury, Joseph Hallett, Nikhil Patnaik, Mohammad Tahaei, and Awais Rashid. Developers are neither enemies nor users: they are collaborators. In *2021 IEEE Secure Development Conference (SecDev)*, pages 47–55. IEEE, 2021.

5 R Maria del Rio-Chanona, Nadzeya Laurentsyeva, and Johannes Wachs. Large language models reduce public knowledge sharing on online Q&A platforms. *PNAS nexus*, 3(9):pgae400, 2024.

6 Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. "Money makes the world go around": Identifying barriers to better privacy in children's apps from developers' perspectives. In Yoshifumi Kitamura, Aaron Quigley, Katherine Isbister, Takeo Igarashi, Pernille Bjørn, and Steven Mark Drucker, editors, *CHI '21: CHI Conference on Human Factors in Computing Systems, Virtual Event / Yokohama, Japan, May 8-13, 2021*, pages 46:1–46:15. ACM, 2021.

7 Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. How can we design privacy-friendly apps for children? Using a research through design process to understand developers' needs and challenges. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW2):1–29, 2023.

8 Michèle Finck. Cobwebs of control: the two imaginations of the data controller in EU law. *International Data Privacy Law*, 11(4):333–347, 2021.

9 Rohan Grover. Encoding privacy: Sociotechnical dynamics of data protection compliance work. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2024.

10 Irit Hadar, Tomer Hasson, Oshrat Ayalon, Eran Toch, Michael Birnhack, Sofia Sherman, and Arod Balissa. Privacy by designers: software developers' privacy mindset. *Empirical Software Engineering*, 23(1):259–289, 2018.

11 Babak Heydari and Negin Maddah. The shifting dynamics of online knowledge platforms and the implications for generative ai sustainability, 2025. Available at SSRN: `https://ssrn.com/abstract=5117087`.

12 Stefan Albert Horstmann, Samuel Domiks, Marco Gutfleisch, Mindy Tran, Yasemin Acar, Veelasha Moonsamy, and Alena Naiakshina. "Those things are written by lawyers, and programmers are reading that." Mapping the communication gap between software developers and privacy experts. *Proceedings on Privacy Enhancing Technologies*, 2024.

13 Stefan Albert Horstmann, Sandy Hong, David Klein, Raphael Serafini, Martin Degeling, Martin Johns, Veelasha Moonsamy, and Alena Naiakshina. "Sorry for bugging you so much." Exploring developers' behavior towards privacy-compliant implementation. In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 1215–1233. IEEE, 2025.

14 Muhammad Ikram, Rahat Masood, Gareth Tyson, Mohamed Ali Kaafar, Noha Loizon, and Roya Ensafi. The chain of implicit trust: An analysis of the web third-party resources loading. In *The World Wide Web Conference*, pages 2851–2857, 2019.

15 Shubham Jain, Janne Lindqvist, et al. Should I protect you? Understanding developers' behavior to privacy-preserving APIs. In *Workshop on Usable Security (USEC'14)*, 2014.

16 Rashna Kumar, Sana Asif, Elise Lee, and Fabian E Bustamante. Each at its own pace: Third-party dependency and centralization around the world. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 7(1):1–29, 2023.

17 Hao-Ping Lee, Advait Sarkar, Lev Tankelevitch, Ian Drosos, Sean Rintel, Richard Banks, and Nicholas Wilson. The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI conference on human factors in computing systems*, pages 1–22, 2025.

18 Gilles Mertens, Nataliia Bielova, Vincent Roca, and Cristiana Santos. You can't trust your tag neither: Privacy leaks and potential legal violations within the google tag manager. In *EuroS&P 2025-10th IEEE European Symposium on Security and Privacy*, 2025.

19 Abraham H Mhaidli, Yixin Zou, and Florian Schaub. "We can't live without them!" App developers' adoption of ad networks and their considerations of consumer risks. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, pages 225–244, 2019.

20 Shaoor Munir, Sandra Siby, Umar Iqbal, Steven Englehardt, Zubair Shafiq, and Carmela Troncoso. Cookiegraph: Understanding and detecting first-party tracking cookies. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 3490–3504, 2023.

21 Seokran Park and Dongyeon Kim. Impacts of AI-based search on user engagement: Evidence from stack overflow's overflowai. In Michael D. Myers, Rose Alinda Alias, Wai Fong Boh, Robert M. Davisin, Barney Tan, and Nor Zairah Ab Rahim, editors, *29th Pacific Asia Conference on Information Systems, PACIS 2025, Kuala Lumpur, Malaysia, July 6-9, 2025*, 2025.

22 Article 29 Data Protection Working Party. Guidelines on transparency under regulation 2016/679, 17/en wp260 rev.01, April 2018.

23 David Rodriguez, Joseph A Calandrino, Jose M Del Alamo, and Norman Sadeh. Privacy settings of third-party libraries in Android apps: A study of Facebook SDKs. *Proceedings on Privacy Enhancing Technologies*, 2025.

24 Arianna Rossi, Rossana Ducato, Helena Haapio, and Stefania Passera. When design met law: Design patterns for information transparency. *Droit de la Consommation = Consumenterecht: DCCR*, (122–123):79–121, 2019.

25 Cristiana Santos, Midas Nouwens, Michael Toth, Nataliia Bielova, and Vincent Roca. Consent management platforms under the GDPR: processors and/or controllers? In *Annual Privacy Forum*, pages 47–69. Springer, 2021.

26 Stefan Schiffner, Bettina Berendt, Triin Siil, Martin Degeling, Robert Riemann, Florian Schaub, Kim Wuyts, Massimo Attoresi, Seda Gürses, Achim Klabunde, et al. Towards a roadmap for privacy technologies and the general data protection regulation: a transatlantic initiative. In *Annual privacy forum*, pages 24–42. Springer, 2018.

27 Katie Shilton and Daniel Greene. Linking platforms, practices, and developer ethics: Levers for privacy discourse in mobile application development. *Journal of Business Ethics*, 155(1):131–146, 2019.

28 Patrick Song, Jayshree Sarathy, Michael Shoemate, and Salil Vadhan. "I inherently just trust that it works"': Investigating mental models of open-source libraries for differential privacy. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW2):1–39, 2024.

29 Alina Stöver, Nina Gerber, Henning Pridöhl, Max Maass, Sebastian Bretthauer, Matthias Hollick, Dominik Herrmann, et al. How website owners face privacy issues: Thematic analysis of responses from a covert notification study reveals diverse circumstances and challenges. *Proceedings on Privacy Enhancing Technologies*, 2023.

30 Ruoxi Sun and Minhui Xue. Quality assessment of online automated privacy policy generators: an empirical study. In *Proceedings of the 24th International Conference on Evaluation and Assessment in Software Engineering*, pages 270–275, 2020.

31 Mohammad Tahaei, Marvin Ramokapane, Tianshi Li, Jason I Hong, and Awais Rashid. Charting app developers' journey through privacy regulation features in ad networks. In *The 22nd Privacy Enhancing Technologies Symposium*, pages 33–56. De Gruyter Open Ltd., 2022.

32 Mohammad Tahaei and Kami Vaniea. A survey on developer-centred security. In *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 129–138. IEEE, 2019.

33 Mohammad Tahaei and Kami Vaniea. "Developers are responsible": What ad networks tell developers about privacy. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–11, 2021.

34 Mohammad Tahaei, Kami Vaniea, and Awais Rashid. Embedding privacy into design through software developers: Challenges and solutions. *IEEE Security & Privacy*, 21(1):49–57, 2022.

35 Mohammad Tahaei, Kami Vaniea, and Naomi Saphra. Understanding privacy-related questions on stack overflow. In *Proceedings of the 2020 CHI conference on human factors in computing systems*, pages 1–14, 2020.

36 Michael Toth, Nataliia Bielova, and Vincent Roca. On dark patterns and manipulation of website publishers by CMPs. *Proceedings on Privacy Enhancing Technologies*, 2022(3):478–497, 2022.

37 Christine Utz, Sabrina Amft, Martin Degeling, Thorsten Holz, Sascha Fahl, and Florian Schaub. Privacy rarely considered: Exploring considerations in the adoption of third-party services by websites. *Proceedings on Privacy Enhancing Technologies*, 2023.

38 Yash Vekaria, Yohan Beugin, Shaoor Munir, Gunes Acar, Nataliia Bielova, Steven Englehardt, Umar Iqbal, Alexandros Kapravelos, Pierre Laperdrix, Nick Nikiforakis, Jason Polakis, Franziska Roesner, Zubair Shafiq, and Sebastian Zimmeck. SoK: Advances and open problems in web tracking. *CoRR*, abs/2506.14057, 2025.

39 Yash Vekaria, Benjamin Standaert, Max Ostapenko, Abdul Haddi Amjad, Yana Dimova, Shaoor Munir, Chris Böttger, and Umar Iqbal. Chapter 10: Privacy. In *The 2024 Web Almanac*. HTTP Archive, 2024. [Online]. Available: `https://almanac.httparchive.org/en/2024/privacy`.

40 Xiaoyu Zhang, Juan Zhai, Shiqing Ma, Qingshuang Bao, Weipeng Jiang, Qian Wang, Chao Shen, and Yang Liu. The invisible hand: Unveiling provider bias in large language models for code generation. In Wanxiang Che, Joyce Nabende, Ekaterina Shutova, and Mohammad Taher Pilehvar, editors, *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 21376–21403, Vienna, Austria, July 2025. Association for Computational Linguistics.

4.3 AI for Privacy, Privacy for AI

Benjamin Berens (Karlsruhe Institute of Technology, DE, benjamin.berens@kit.edu)
Simone Fischer-Hübner (Karlstad University, Chalmers University of Technology & University of Gothenburg, SE, simone.fischer-Hübner@kau.se)
Andreas Gutmann (University College London, UK)
Bailey Kacsmar (University of Alberta – Edmonton, CA, kacsmar@ualberta.ca)
Agnieszka Kitkowska (Jönköping University, SE, agnieszka.kitkowska@ju.se)
Marc Langheinrich (Università della Svizzera italiana – Lugano, CH, marc.langheinrich@usi.ch)
Mainack Mondal (Indian Institute of Technology Kharagpur, IN, mainack@cse.iitkgp.ac.in)
Elissa Redmiles (Georgetown University – Washington DC, US, elissa.redmiles@georgetown.edu)

4.3.1 Introduction

The current human-AI interaction paradigm is making computational machines more pervasively trained and deployed on much more data (than previous generations of AI). This raises both the opportunity to use AI tools to enhance societal outcomes, including privacy, and threats from the exacerbation and introduction of risks and harms. We posit that going

forward, we must address the question "how does human-centered privacy change in a world where AI-based automation (e.g., AI-agents) is incorporated into the many facets of digital life, for better or worse?"

Risks and harms may come as a result of decisions made by individuals (who use an AI-employing system), by others (collateral privacy through inferences made about you even as a non-user), and by society (to prioritize an AI capability over a threat to some individuals). Therefore, to understand these different sources of harm, we first need to surmise if existing privacy and behavioral decision-making theories, and decision-making consequences, explain and predict such decisions as well as potential privacy issues that consequently emerge. If not, we need to update these theories and frameworks to better explain and predict potential privacy issues that can emerge from individual human-AI interaction and societal decisions about AI deployment. From there we can develop a better understanding of decision making, risks and threats specific to this domain, as well as how AI can be used to benefit society helping rather than hindering privacy in our digital spaces.

Finally, we emphasize that while AI poses new challenges in the space of human-centered privacy, it also offers positive potential for providing usable end-user guidance or even for automated or semi-automated privacy decisions.

4.3.2 Problems/Challenges Description

Transferability of Existing Theories and Methods. The existing models of decision-making and behavioral outcomes in the context of privacy assume that the same factors that enable individual information disclosures for more traditional technologies also apply to the context of AI-based technologies, e.g., when individuals interact with genAI [2] embedded in chatbots or social robots. However, we have little data or theory investigating or explaining how AI-based technologies might actually be influencing the behaviors of individuals or groups, as well as how such influences can change over time and across different contexts (e.g., usage scenarios).

Creating interaction-grounded frameworks that could provide us with an opportunity to better understand how individual and societal privacy decisions and behaviors change because of AI used in the systems requires handling additional challenges, as today AI transforms privacy decision-making by automating decisions that influence human choices through design, personalization, or implementation of anthropomorphic features. Overall, these features have the potential to influence the trust users put in AI-driven systems. It remains unknown whether the effects of such influences on privacy decisions and behaviors are positive or, on the contrary, negatively affect user privacy and control over their data.

Threats, Risks, and Harms. A fundamental question faces privacy for AI: what threats are essentially exacerbations of prior known privacy concerns, through AI's mechanism of scale, and what threats are novel or considerably altered as a result of capabilities and scalability introduced by human-AI interaction. Prior work on privacy of AI has centred heavily on model-focused attacks on data confidentiality, integrity, correctness, memorization [17], extraction [15], and inference attacks [16], leaving less explored the impact of human-AI interaction and decision-making. Among the threats that exist, or may come to be, we need to assess the risk (impact $\times$ likelihood) each poses, build robust threat models that map the harms that may materialise from these risks (considering differences in individuals and circumstances), and design sociotechnical mitigations into products / processes / society (laws, regulators, etc).

[2] genAI throughout the rest of the chapter refers to generative AI models.

AI exacerbates existing privacy risks. First, AI anthropomorphism may exacerbate people's greater willingness to share information with machines. Second, AI exacerbates existing surveillance risk, including collateral privacy – the impact of others' data/interactions on an individual's risk. Third, AI exacerbates data risks of both extraction (intentional) and leakage (intentional or unintentional). Examples include leaks through information sharing during benign use of benign tools and through interaction with malicious tools.

AI also presents at least two novel privacy risks. First, agentic AI (action-taking agents) may be used maliciously to execute and scale attacks, e.g., social engineering, or may leak private information during benign action. Second, AI capabilities allow for increased plausibility in the generation of recognizable digital replicas (i.e., deepfakes).

A cross-cutting theme is that both existing and novel risks will be amplified by a desire for greater utility: more capabilities, with higher quality. Given the probabilistic nature of AI, some threats cannot be defended against in the near term. Thus, mitigation may include not only technical solutions but societal-level trade-offs between tool capabilities and corresponding indefensible risks and/or identification of social-norm shaping interventions to bound the use of capabilities.

Using AI for usable personalized privacy. In recent years, different AI tools, including personalized privacy assistants, have been developed for assisting and guiding users with privacy decision making (e.g., [8]). (Semi-)automation of privacy decisions founded on AI-based personalised guidance can enhance usable privacy management for users who are often overwhelmed with privacy decision requests, by proposing or enforcing privacy decisions meeting these users personal preferences. AI tools offering personalised privacy assistance can also consider different preferences that users with different demographic backgrounds (e.g. culture, gender, age) typically have, and can enhance accessibility by adapting privacy information to specific user needs.

On the other hand, automation of privacy decisions as well as "privacy nudges" also raise ethical and legal questions, especially regarding the users' autonomy and control over their data, which is an essential privacy principle highlighted in Recital 7 GDPR [3]. Some privacy decisions, such as consent, legally require human actions, i.e., cannot be fully automated. Another problem is that users may overly trust AI-proposed decisions (automation bias). While AI-based tools have been developed that propose decisions matching users' expectations with a very high accuracy, many AI techniques are probabilistic and hence cannot provide any guarantees that the proposed actions reflect the users' wishes. In addition, LLM-based approaches may provide information that is inaccurate or plain wrong. Hence, finding the right balance between AI-based personalised decision guidance and enforcement, and user involvement – keeping users in the loop – constitutes a challenge.

Moreover, ML-based personalized privacy assistants typically analyze the users' attitudinal or behavioral privacy preferences, metadata or content, or other data types. Hence, they in turn need to process personal data and user profiles. Therefore, it is essential that personalised privacy assistants themselves are developed following a privacy-by-design approach [28], especially if the models are not trained locally on the users' devices and under the users' control. Most research work on personalised privacy assistants does not (or not adequately) address such crucial privacy-by-design aspects.

AI-based tools can also be used for personalizing how privacy policies are communicated (e.g., format, language, style). This can improve accessibility for diverse users, including vulnerable groups. However, simplified or voice-based communication may not meet GDPR

[3] https://gdpr-info.eu/recitals/no-7/

requirements (Articles 12, 13) [29]. And if compromised, AI-based privacy assistants could manipulate users into disclosing sensitive information or disclose more information than they would have disclosed otherwise, which constitutes a deceptive design (dark pattern) that is illegal in some jurisdictions (e.g., in the EU Digital Services Act, which is reflected also in AI Act, Digital Markets Act, and Consumer Rights Directive).

4.3.3 State of the Art

Different fields (HCI, AI, SWE, law, and more) have each developed disparate definitions of safety as well as classification taxonomies and evaluation frameworks for privacy risks and harms. Some of these build upon prior theories of privacy, such as using Solove's framework [18], while others introduce their own structure. However, while these provide organization to the domain, they each use different conceptual terminology, different categorisations, and largely remain at an overview level, aiming to capture the overall field. This creates challenges and limitations in terms of applicability for those creating, deploying, or overseeing these systems as: (i) By trying to cover the breadth of the space, nuance cannot be captured, (ii) these frameworks are often rooted in academic/theoretical/philosophical perspectives such that those wanting to apply these insights need to translate and transfer them to match their requirements. Unfortunately, these can often not transfer well or easily to such industry processes and priorities (e.g. in regards to when or how harms must be mitigated).

We know from usability foundations that adoption rises when "things" are designed with end-user requirements and use cases in mind from the start, not added at the end. Regulators and governments (for instance EU enacted, Canada tabled, UK's AI Bill tabled) have similarly tried to capture risk as either:

- All AI (which has contentious definitional bounds)
- AI for specific tasks, with some using domain-specific or industry-specific existing guidelines or laws as baselines and others trying to create new ones that specifically target "AI in X"

While we can articulate notions of privacy risk conceptually in this space, which concrete harms will follow from them is less clear, difficult to measure with existing tooling and models, which correspondingly leads to contention as to next steps for all stakeholders and relevant parties.

Following the human-centric approach, to ensure a comprehensive understanding of both the positive and negative effects of AI on people, one must first gain a thorough understanding of how individuals and groups make decisions regarding their privacy and security when using technologies that incorporate AI. Traditional psychological models like the Theory of Reasoned Action (TRA) and Theory of Planned Behavior (TPB) fall short in explaining how people make privacy decisions under uncertainty or cognitive overload [19, 20, 21, 22]. Dual-process theories – System 1 (S1) and System 2 (S2) – provide a more comprehensive framework [22, 21]. S1 is fast, intuitive, and driven by heuristics (e.g., affect heuristic), while S2 is slow, analytical, and effortful. Most everyday decisions rely on S1, which is sensitive to various peripheral cues (e.g., visuals), heuristics and biases, and other factors that might affect cognitive processing of information (e.g., time pressure). Privacy decision-making models build on psychological theories, such as TPB, TRA, or dual process theories, at times combining them into more holistic but theoretical models. One such model, APCO (Antecedents–Privacy Concerns–Outcomes), centers privacy concerns while acknowledging background factors like personality, culture, trust, and skills [23]. This conceptual macro-model incorporates System 1 thinking, recognizing the role of fast, intuitive judgments in privacy choices. Following the model, privacy concerns seem to play a central role to

decision-making, yet, they might be hard to assess and various scales to measure them were developed (e.g., Internet Users's Information Privacy Concerns (IUIPC) [24], Mobile Users Privacy Concerns [25], and even Social Robots Privacy Concerns [26]). AI development at a fast pace resulted in the first attempts to measure concerns in this context with the Privacy Concerns related to AI Misuse (PC-AIM) scale [27], which builds on existing models like the IUIPC but adds dimensions like data permanence, profiling, reduced judgment, and algorithmic bias. However, more studies are needed to validate its accuracy and applicability to various contexts of interactions with AI-based technologies. Considering the complexity of decision-making processes, other measurement instruments/methods for latent factors affecting decisions and behavior in the context of privacy and AI might be needed.

In AI interactions, anthropomorphization can influence privacy decisions, sometimes causing discomfort (the "uncanny valley" effect). However, this can be offset by empathetic responses and immediate implicit feedback (e.g., through multi-turn interactions), which increase user comfort and data disclosure. These effects, though impactful, remain under explored, raising concerns about potential privacy and security risks. AI agents often trigger affective heuristics through emotional cues like personalization and human-like behavior, which can override rational risk assessment. As a result, users may underestimate privacy risks and overshare data, especially under cognitive load or time constraints.

Studies show users are often more willing to share sensitive data with AI than with humans – particularly if AI is perceived as non-judgmental [30, 31]. However, if AI seems too powerful or invasive, privacy concerns intensify. Cultural factors also shape responses: U.S. users tend to trust AI more unless data sensitivity is high, whereas Chinese users are more [32]. It is possible that older adults may accept AI, similarly to how they accept social robots [33], if its utility outweighs privacy risks, but still demand control and transparency; however, research is needed to confirm such assumptions.

Regarding the beneficial use of AI in the context of privacy, i.e., using AI to support stakeholders with privacy-related decision-making, a rapid review of the existing literature sees three main research areas emerging:

(A) Using AI to assess privacy issues in to-be-shared social media content (e.g., [2, 3, 4, 5]): The key challenges are the lack of real-world deployments, the issue of cold-start, and the social complexity inherent in such sharing decisions. Multi-party privacy conflicts are also difficult to solve, with or without AI.

(B) Using GenAI/LLMs for privacy policy understanding (e.g., [34, 7] [Tongue and Caragea 2020]): The key challenges here are unreliability of LLMs, the legal validity of consenting to AI-summarized policies, and, of course, the inherent complexity of policies written not for humans but lawyers

(C) AI-based privacy management (e.g., decision making) (e.g., [10, 6][Bhattacharya, 2024]).

Area (A) has traditionally seen much work due to the importance of social media. Since 2023, area (B) is growing quickly. Area (C) is just emerging.

4.3.4 Key Research Questions

We identified four key research questions in this domain. This report focuses on the first three. However, we include the fourth as we also need to communicate amongst stakeholders all the aspects of the other three research questions.

- RQ1: How can we use AI to enhance usable privacy?
- RQ2: How does AI affect (positively/negatively) human privacy decision-making and behavior?

- RQ3: How does AI pose novel or exacerbate existing risks to people's privacy?
- RQ4: How can developers, companies, researchers, or governments (etc.) communicate privacy (cf RQ1-3), both reactively and proactively, (e.g., risks, potential manipulation) in the context of ubiquitous AI-based systems?

4.3.5 Solution ideas and directions

Guided by our research questions and the literature, we synthesize the path forward as Table 1.

4.3.6 Resources required

Interdisciplinary collaboration: the role of social and psychological factors necessitates collaborations with at minimum Psychology and Sociology. Moreover, cooperation with legal experts is needed for eliciting and enforcing legal requirements.

Hardware (GPU, robotics infra, etc.). Assessments may require direct system interaction Secure funding to support longitudinal research and collection of large datasets (perhaps on an ongoing basis)

Data. Evaluation and benchmarking datasets, including multi-turn interaction data Incident tracking for privacy issues/harms/consequences & regulation for data access to commercial data (cf. research access provisions in EU and UK regulation of online safety)

4.3.7 Roadblocks

Speed of innovation. High volatility of everything with AI (possibly relating to views on "outdatedness" as well as changes in applications and mitigations of issues)

Disconnection For collaborative efforts between technical, legal, social sciences, and design knowledge bases, there are correspondingly interdisciplinary challenges in the collaboration between people with deep technical AI knowledge and other relevant skill sets (law, HCI, psychology, sociology).

References

1 Morel, V., Iwaya, L. & Fischer-Hübner, S. AI-driven Personalized Privacy Assistants: a Systematic Literature Review. *IEEE Access*. (2025)
2 Freiberger, V., Fleig, A. & Buchmann, E. "You don't need a university degree to comprehend data protection this way": LLM-Powered Interactive Privacy Policy Assessment. *Proceedings Of The Extended Abstracts Of The CHI Conference On Human Factors In Computing Systems*. pp. 1-12 (2025)
3 Freiberger, V., Fleig, A. & Buchmann, E. Explainable AI in Usable Privacy and Security: Challenges and Opportunities. *ArXiv Preprint ArXiv:2504.12931*. (2025)
4 Hamid, A., Samidi, H., Finin, T., Pappachan, P. & Yus, R. GenAIPABench: A benchmark for generative AI-based privacy assistants. *ArXiv Preprint ArXiv:2309.05138*. (2023)
5 Aydin, I., Diebel-Fischer, H., Freiberger, V., Möller-Klapperich, J., Buchmann, E., Färber, M., Lauber-Rönsberg, A. & Platow, B. Assessing Privacy Policies with AI: Ethical, Legal, and Technical Challenges. *ArXiv Preprint ArXiv:2410.08381*. (2024)
6 Najana, M., Balakrishnan, A. & Bhattacharya, S. Conceptualizing Copilots as Privacy Assistants: A Theoretical Framework. (2024)
7 Ayci, G., Sensoy, M., özgür, A. & Yolum, P. Uncertainty-aware personal assistant for making personalized privacy decisions. *ACM Transactions On Internet Technology*. **23**, 1-24 (2023)

Table 1 The roadmap of strategies towards advancement in the AI for Privacy and Privacy for AI domain per research question.

Horizon	Strategy	RQ1	RQ2	RQ3
Immediate	• Establish threat models that map risks to harms	X		X
	• Validate existing theoretical and empirical models of human-centered privacy in the context of emerging AI capabilities, e.g., re-evaluate the impact of novel surveillance capabilities, future power dynamics, and advanced interaction dynamics on privacy expectations, decisions, and achievability	X	X	X
	• Standardize methodologies for considering tradeoffs between PETs and other AI considerations (fairness, representation, etc.)			X
	• Develop privacy-by-design guidelines for personalized AI-based privacy assistants (i.e., where the AI assistant's capabilities require significant access to the user's data)	X		
	• Elicit legal, regulatory, and/or end user requirements	X	X	X
3+ years	• Operationalize risk assessment measurement tools (mixed methods, need qualitative and quantitative, would need both model-centric, and human-interaction style ones) for a range of AI capabilities including agentic AI	X		X
	• Human-centered design templates for semi-automated tools that action privacy on user's behalf that address issues such as accuracy, semantic understanding, cold start,–along with providing the corresponding validation mechanism that supports subsequent deployments	X		
	• Clarifying validity of AI-Supported (legal) Consent and Liability	X		
10+ years	• Longitudinal studies of Human-AI interaction & changing modalities (esp. Anthropomorphic effects) on human behavior, and of the effect of AI on general privacy decisions, concerns, and behavior	X	X	X
	• Extending semi-automated tool design templates to address interdependent privacy & multi-user collaboration or conflict	X		
	• Regulation and schemes to enable third-party audits		X	X

8 Zhan, N., Sarkadi, S. & Such, J. Privacy-enhanced personal assistants based on dialogues and case similarity. *European Conference On Artificial Intelligence*. (2023)

9 Ayci, G., özgür, A., Sensoy, M. & Yolum, P. Explain to me: Towards understanding privacy decisions. *Proceedings Of The 2023 International Conference On Autonomous Agents And Multiagent Systems*. pp. 2790-2791 (2023)

10 Morel, V. & Fischer-Hübner, S. Automating privacy decisions-where to draw the line?. *2023 IEEE European Symposium On Security And Privacy Workshops (EuroS&PW)*. pp. 108-116 (2023)

11 Carter, S., D'Aquin, M., Spagnuelo, D., Tiddi, I., Cormican, K. & Felzmann, H. The Privacy-Value-App Relationship and the Value-Centered Privacy Assistant. *ArXiv Preprint ArXiv:2308.05700*. (2023)

12 Ischen, C., Araujo, T., Voorveld, H., Noort, G. & Smit, E. Privacy concerns in chatbot interactions. *Chatbot Research And Design: Third International Workshop, CONVERSATIONS 2019, Amsterdam, The Netherlands, November 19–20, 2019, Revised Selected Papers 3*. pp. 34-48 (2020)

13 Belen Saglam, R., Nurse, J. & Hodges, D. Privacy concerns in chatbot interactions: When to trust and when to worry. *HCI International 2021-Posters: 23rd HCI International Conference, HCII 2021, Virtual Event, July 24–29, 2021, Proceedings, Part II 23*. pp. 391-399 (2021)

14 Zhan, X., Carrillo, J., Seymour, W. & Such, J. Malicious LLM-Based Conversational AI Makes Users Reveal Personal Information. *USENIX Security*. (2025)

15 Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U. & Others Extracting training data from large language models. *30th USENIX Security Symposium (USENIX Security 21)*. pp. 2633-2650 (2021)

16 Staab, R., Vero, M., Balunović, M. & Vechev, M. Beyond memorization: Violating privacy via inference with large language models. *ArXiv Preprint ArXiv:2310.07298*. (2023)

17 Kim, S., Yun, S., Lee, H., Gubri, M., Yoon, S. & Oh, S. Propile: Probing privacy leakage in large language models. *Advances In Neural Information Processing Systems*. **36** pp. 20750-20762 (2023)

18 Solove, D. A taxonomy of privacy. *University Of Pennsylvania Law Review.*, 477-560 (2006)

19 Madden, T., Ellen, P. & Ajzen, I. A Comparison of the Theory of Planned Behavior and the Theory of Reasoned Action. *PSPB*. **18** pp. 3-9 (1992)

20 Ajzen, I. Nature and Operation of Attitudes. *Annual Reviews Of Psychology*. **52** pp. 27-58 (2001)

21 Kahneman, D. A Perspective on Judgment and Choice. *American Psychologist*. **3**, 7-18 (2003)

22 Evans, J. & Stanovich, K. Dual-Process Theories of Higher Cognition: Advancing the Debate. *Perspectives On Psychological Science*. **8**, 223-241 (2013)

23 Dinev, T., Mcconnell, A., Smith, H., Dinev, T., Raton, B. & Smith, H. Economics : Thinking Outside the "APCO" Box Systems, Psychology, and Behavioral Economics : Thinking Outside the "APCO" Box. *Information Systems Research*. **26**, 639-655 (2015)

24 Malhotra, N., Kim, S. & Agarwal, J. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*. **15**, 336-355 (2004)

25 Xu, H., Gupta, S., Rosson, M. & Carroll, J. Measuring mobile users' concerns for information privacy. (2012)

26 Jia, S., Chi, O. & Lu, L. Social Robot Privacy Concern (SRPC): Rethinking privacy concerns within the hospitality domain. *International Journal Of Hospitality Management*. **122** (2024,9)

27 Menard, P. & Bott, G. Artificial intelligence misuse and concern for information privacy: New construct validation and future directions. *Information Systems Journal*. **35**, 322-367 (2025)

28 Cavoukian, A. Understanding How to Implement Privacy by Design, One Step at a Time. *IEEE Consumer Electronics Magazine*. **9**, 78-82 (2020,3)

29 Commission, E. Regulation (EU) 2016/679 Of The European Parliament And Of The Council of 27 April 2016. *Official Journal Of The European Union*. (2016)

30 Sohn, S., Labrecque, L., Siemon, D. & Morana, S. Artificial intelligence versus human service agents: How their presence shapes consumer information privacy concerns. *Journal Of Retailing*. (2025)

31 Kim, T., Jiang, L., Duhachek, A., Lee, H. & Garvey, A. Do you mind if I ask you a personal question? How AI service agents alter consumer self-disclosure. *Journal Of Service Research*. **25**, 649-666 (2022)

32 Liu, Y., Yan, W., Hu, B., Lin, Z. & Song, Y. Chatbots or humans? Effects of agent identity and information sensitivity on users' privacy management and behavioral intentions: A comparative experimental study between China and the United States. *International Journal Of Human–Computer Interaction*. **40**, 5632-5647 (2024)

33 Reinhardt, D., Khurana, M. & Acosta, L. "I still need my privacy": Exploring the level of comfort and privacy preferences of German-speaking older adults in the case of mobile assistant robots. *Pervasive And Mobile Computing*. **74** pp. 101397 (2021)

34 Kqiku, L. Privacy-Decision Assisting Techniques. (2024)

4.4 Consent, Control, and Communication

Arianna Rossi (Sant'Anna School of Advanced Studies – Pisa, IT, arianna.rossi@santannapisa.it)
Farzaneh Karegar (Karlstad University, SE, farzaneh.karegar@kau.se)
Florian Alt (LMU München, DE, florian.alt@ifi.lmu.de)
Maija Poikela (Charité – Berlin, DE, maija.poikela@bih-charite.de)
Mark Warner (University College London, UK, mark.warner@ucl.ac.uk)
Sophie Grimme (OFFIS – Oldenburg, DE, sophie.grimme@offis.de)
William Seymour (King's College London, UK, william.seymour@kcl.ac.uk)
Zinaida Benenson (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), DE, zinaida.benenson@fau.de)

4.4.1 Introduction

We are bombarded with cookie banners which often violate informed consent requirements [1], privacy policies are not written to be understood [2], and real time bidding – the economic basis of the contemporary web – is "structurally difficult to reconcile with European data protection law" [3]. While it is clear that the transparency and consent model for privacy decisions is broken [4], its adoption as the cornerstone of consumer and data protection regulation means that it is here to stay. This calls for the need to envision and design solutions that address the continuing problems of transparency and consent. At the same time, the continued use and development of artificial intelligence has created a data-oriented world where it seems inevitable that data flows are destined to grow. This development can

be beneficial for individuals and societies (e.g., for data-driven medical research), providing new opportunities. At the same time, emerging technologies and novel interaction modalities also pose new threats which the research community must anticipate.

In this context, it is essential that we work together to improve the mechanisms that we have to control how our data is used and how we are asked for consent. In our discussions, we strive to create effective, usable, lawful, and accessible tools for consent, control, and communication around the use of personal data.

We make a distinction between the idea of *consent* as an explicit decision point allowing for positive and negative choices, and *control* as the broader category of mechanisms that people have to influence how their data is used. Notably, whilst consent is always something that the data subject negotiates with a data controller, control as used here encompasses resistance, activism, and any other actions taken by an individual, or some other entity on their behalf, that change how their data is used.

4.4.2 State of the Art

Despite decades of research and regulatory attention, privacy notices and consent forms continue to fall short of their intended purposes: attracting user attention, informing individuals about data practices, enabling comprehension, and ensuring meaningful control over personal data. Although during the years privacy policies of service providers have been improved, they are still often compliance-driven rather than user-centred. For example, privacy policy forms, which are one form of privacy notices, have to fulfil requirements from different stakeholders [5, 6]. While users expect to receive clear, simple information about data practices and privacy controls, service providers employ privacy policies to demonstrate compliance with legal requirements.

These shortcomings are further compounded by usability issues such as cluttered layouts, inconsistent placement of privacy controls, or technical barriers that make opting out difficult [7]. In addition, consent mechanisms on websites and digital platforms often rely on design choices nudging users toward acceptance instead of fostering genuine understanding or control. Empirical studies show that interface manipulations such as default buttons, pre-selected options, or the removal and obscuring of rejection choices can raise acceptance [8]. A (perceived) lack of choice and decoupled notices are other reasons why users do not pay attention to privacy notices [5, 6].

This results in notice fatigue, habituation, post-decision regret, and superficial consent, where users agree without fully understanding the implications [9, 10]. Consequently, although many implementations formally satisfy regulatory requirements, they fail to meet the spirit of informed and freely given consent, thereby undermining the user's ability to exercise meaningful control over personal data.

To partly address the usability issues, researchers have explored ways to improve the usability of privacy notices and consent mechanisms. Proposed solutions include privacy nutrition labels [11], layered and short-form summaries [12, 22], personalised notices [13, 14], icons [15, 16], and even comic-based interfaces [17], among other design patterns [24]. These approaches attempt to make complex data practices more transparent and digestible, but each comes with trade-offs: condensed formats risk oversimplification, while visual representations, such as the meaning behind icons, often require additional user learning. Beyond presentation, other efforts focus on interactive consent, using mechanisms like drag-and-drop, swiping, or checkboxes, to actively engage users with the content, improving both attention and retention [18].

Emerging technologies intensify the challenges in these contexts. Artificial intelligence systems often operate opaquely, making it difficult to communicate how data is processed or how decisions are made, although they have the potential for improving usable personalised privacy and decision-making. In AR/VR and metaverse environments, the traditional notice and choice paradigm is even less effective and sometimes impossible. Immersive systems collect highly sensitive biometric and behavioural data, much of which is invisible to the user. The immersive nature of these platforms also makes interruptions for lengthy consent impractical, raising the risk that consent becomes implicit or illusory.

4.4.3 Case Studies

The working group considered a range of different use cases in order to map out the problem space. Each case study was chosen because it surfaced relevant issues and solutions related to consent, control, and/or communication of data practices.

- **Medical data donation for research purposes.** This included discussion around the values and reasons behind the donation of medical data, the potential positive and negative consequences of doing so, and how this could be effectively, lawfully, and ethically communicated.
- **Consent and control for virtual/mixed reality.** We discussed if and when we should design consent decisions to avoid pulling users out of immersive experiences, as well as alternatives such as making choices before/after immersive moments or through other interaction modalities.
- **Conversational Agents.** We talked about how conversation, particularly when spoken, is a relatively low bandwidth means of interaction, but does present opportunities for more engaging conversations about privacy. It foregrounds a question from multi-layer privacy policies about what should be in the first layer.
- **Data collection by smart cars.** The discussion covered how people don't expect this data collection, leading to cases where consent isn't freely given. There are also tangible consequences that arise from data collected by cars, particularly related to car insurance.

When considering these case studies, two main areas of discussion emerged. The first of these was the lack of communication around consequences for data sharing, making it difficult to answer the question of why people should care about privacy. The second centred on the development of AI agents that (amongst other things) would be tasked with making privacy decisions on our behalf.

4.4.4 Focus I: Consequences of Privacy Decisions

This area of discussion was sparked by an observation in the main session that, as privacy experts, we often struggle to effectively communicate the importance of digital privacy. An obvious response would be to convey the tangible consequences (good and bad) and risks of making a privacy decision, but this is something that we don't currently have a good overview of. To this end, the group considered the specification of a consequences-based and cross-platform privacy decision manager that would help people to make privacy decisions that are aligned with their individual expectations and values. We preliminary call this system 3CM: *Consent, Control and Communication Manager* and imagine it to be similar to existing password managers. This idea is based on discussions of the group "Consequence-based Privacy Decisions" in Dagstuhl Seminar "My Life, Shared" from 2013[4] [19].

[4] https://www.dagstuhl.de/13312

A core component of 3CM would be a database of events and outcomes of sharing data, both positive (e.g., personalisation) and negative (e.g., data breaches). Given appropriate context, an LLM (Large Language Model) would be able to provide explanations of these consequences of sharing or not sharing before a decision is made. These decisions would be stored by the manager for future use and could be applied automatically where directed. This would add functionalities to existing tools such as Consent-O-Matic[5]. The privacy decision manager would learn from the decisions that the user makes, taking into account contextual elements that might impact sharing decisions.

Such a concept is not without challenges. In 2013, one of the main challenges was how the system would present the consequences and learn from past decisions. These challenges remain, but seem to be better manageable with the proliferation of LLMs. There are also implementation difficulties around curating data sources and creating experiences that function across devices. A consequences-based decision manager would not address underlying problems with the notice and consent model, but it would present an opportunity to push back and change the incentives for users (and thus organisations) around the choices they make in relation to their data. Ultimately, it would also need to be *easier to use than not use*, increasing user acceptance. This kind of system also brings with it the question of when or if to ask users if they want to reconsider their decisions, potentially leveraging so-called 'teachable moments' where people have the time and energy to make decisions outside of the task that led to the initial prompt to share data.

4.4.5 Focus II: Agents

Artificial agents are increasingly embedded in our day-to-day activities, whether professional or personal tasks. Especially since the widespread availability of conversational interfaces that enable anyone to interact with Large Language Models (LLMs), such as GPT, it is evident that humans find it useful and beneficial to resort to those instruments, even if their usage is often accompanied by a lack of awareness of the entailed risks. As in other contexts, in this kind of interactions, privacy and security are not the primary task of users. Thus, certain old-standing issues are being re-proposed with emerging technologies, in addition to the new challenges these technologies pose.

The group also discussed possible solutions offered by artificial agents to the so-called "privacy self-management dilemma". A trained artificial agent could take privacy decisions on behalf of the user while it performs other tasks (e.g., refuse profiling cookies while booking a flight). But this scenario raises additional questions, such as whether and to what extent an agent can make decisions on behalf of the user. Even if this could sometimes be determined contextually (e.g., not in high-risk scenarios), still, to make informed decisions, the agent should understand the contextual norms of using personal information. In this regard, it would be similar to having a human assistant. The participants discussed whether it is possible that the agent learns through a process of personalization or if it should be based on the profiles of others, as is already the case for some privacy assistants. Another open question concerns whether a threshold for the acceptable amount of errors that can be made should be set. Such an agent calls into question the very nature of "informed" decisions, a concept which should probably be reformulated. What is interesting is that taking privacy-friendly decisions not only concerns the type and amount of data that is shared or withheld, but also the capacity to make inferences based on such data.

[5] `https://consentomatic.au.dk`

In addition, the group discussed the types of requirements that should be formalized for the development of such an agent. These should concern the user interface (e.g., in terms of timing and modality of communication) and should envision the situations where a decision needs to be deferred to a human. This setting also emphasizes the central role that trust would play: requirements should concern the establishment and maintenance of user trust in the agent and cover cases where a breakdown of trust may occur. As a conclusion, the participants acknowledged the need to also include experts with AI and legal backgrounds to address all these questions meaningfully.

Bringing together the two focuses, the group was keen to explore the creation of a privacy decision agent that collated and explained the potential consequences of different privacy decisions.

4.4.6 Key Research Questions

Based on the discussions during the week, the group compiled a list of key research questions for the research community. These are sorted by topic area and timeline:

4.4.6.1 Human factors issues around consent and control

- Immediate research questions
  1. What are researchers' conceptualisations, and users' mental models of providing, revoking and negotiating consent? How do researchers' and users' views differ?
  2. What is the users' understanding of the current state of privacy policies compared to 10-15 years ago?
  3. Does enhanced user-centred transparency exist, and if so, has it been effective at all? What would more "effective" mean (i.e., noticeable, memorizable, understandable, etc)?
  4. Are users aware of the enhanced control promised in privacy rules and regulations, and to what extent do they exercise their rights? Does "enhanced control" (i.e., more actionable rights that are now available to users) make people feel more in control?
  5. If people feel more in control, do they feel more satisfied with their privacy?
  6. What would a systematic taxonomy of control mechanisms look like (e.g., contacting a DPA; changing your data; right to access, etc)?
- Challenges and questions for the next 3 years
  7. How can informed consent be conceptualised differently from a human factors perspective?
  8. Under which circumstances can the necessity of active decision-making be lifted from a legal, ethical, and human point of view?
  9. Do existing implementations of dynamic consent "work"? For example, do people revise their decisions? If not, why not? Do they feel information fatigue?
  10. How is the right to lodge a complaint implemented across different national DPAs? Who has invoked that right? Who hasn't? Why not? What recommendations can we give?
- Challenges and questions for the next 10 years
  11. How can we develop or support legal mechanisms that do not require people to constantly read and make decisions?
  12. Could trust models be utilised to help remove individual consent responsibility? For example, trust in AI agents, data cooperatives, and experts.
  13. Would data intermediaries help or exacerbate the situation?
  14. How can we modify the incentive structure to encourage service providers to adopt a different consent model?

4.4.6.2 Requirements for consent/control in different contexts

- Immediate research questions
  1. What do we understand by "context" of privacy decisions? What is all the relevant information to a decision?
  2. How might we gather and organize the consequences of data disclosures that we know? How might we leverage DPIAs to this end? Which existing taxonomy can be adapted for which purpose?
  3. Beyond consequences, how might we incorporate user-centred factors into our understanding of context, such as values, personality, motivation, health status, etc.?
  4. How might we design a *Consent, Control and Communication Manager* 3CM based on these data?
  5. Why did broadly similar concepts such as P3P fail?
  6. What kind of affirmative actions ("unambiguous indication of one's will") might be imagined and implemented in novel modalities such as VR/AR/voice?
- Challenges and questions for the next 3 years
  7. How does the threshold for being informed change depending on context when making privacy decisions? (e.g. do low stakes decisions require less information to be provided?)
  8. How might a privacy decision agent consider context beyond consequences? What kind of resources would be needed to build it and maintain it?
  9. How might we reimagine consent and control for different interaction modalities, such as mixed reality or voice interfaces?
  10. At which moment(s) should consent be asked in different contexts? Should it be disruptive or not?
  11. How should the consequences be communicated, and how can their (long-term and short-term) effects on users' decisions be measured? How might we counteract the potential biases of positive and negative consequences in the database?
  12. How can we responsibly facilitate reflection on privacy decisions?
- Challenges and questions for the next 10 years
  13. How could novel consent models be developed to help positively shape/influence data use and data sharing practices of organisations?
  14. How might we re-write our trained habits that are rooted in the current notice and consent model (e.g., to click accept without thinking)?
  15. How should guidelines and best practices be formulated to meaningfully guide developers/designers to select the best timing, modality, channel etc based on the context?

4.4.6.3 AI-driven (personalized) privacy assistants

- Immediate research questions
  1. What are the knowns and the unknowns of developing such an AI agent?
  2. Which legal standing does a browser extension making cookie decisions on behalf of the user have? Does that apply to other types of technologies?
  3. How do changing ways of interacting with digital technology change the way that people engage with consent and control? (e.g. smartphone, messaging, AI)
  4. What would happen if only data that is needed for the basic functionality of the service (if data minimization was really enforced) was collected, both in terms of benefits and limitations? How would a world without personalization look like?

- Challenges and questions next 3 years:
  5. When/if to ask users if they want to reconsider their decisions, potentially leveraging the teachable moments?
  6. How might an AI agent learn from your decisions as you make them, as well as considering the wider context (for example, new laws or regulations) that might impact sharing decisions?
  7. To what extent can an AI agent make decisions on your behalf? How does it decide whether it can or where it cannot?
  8. What would be different from a human assistant? Should it learn through a process of personalization, and if so, what kind of safeguards should be in place to protect user privacy? What kind of errors would we be willing to tolerate?
  9. Should it learn based on other people's preferences or decisions? If so, what should it learn from others' preferences?
  10. What do we mean by "privacy-friendly' agent? How would we implement that?
  11. Would privacy-enhancing technologies make it entirely unnecessary for us to manage our data and our privacy?
- Challenges and questions for the next 10 years
  12. Does the AI agent always need to make inferences and communicate them to the user? How should a system decide on whether those decisions should be taken by humans or by the agent?
  13. How does consent and control look in imagined futures? For example, if we are not using web portals anymore, but just using AI agents who use web portals on our behalf.
  14. How would the tradeoffs between robustness, accuracy, fairness and privacy be decided and implemented in such a system [20]?
  15. Who would / should develop such systems?
  16. How will different agents interact?

References

1. Polona Car and Filippo Cassetti. 2025. Regulating dark patterns in the EU: Towards digital fairness. *Prepared by the European Parliamentary Research Service for the European Parliament.*
2. Ewa Luger, Stuart Moran, and Tom Rodden. 2013. Consent for all: revealing the hidden complexity of terms and conditions. *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13). Association for Computing Machinery, New York, NY, USA, 2687–2696. https://doi.org/10.1145/2470654.2481371*
3. Michael Veale and Frederik Zuiderveen Borgesius. 2021. Adtech and Real-Time Bidding under European Data Protection Law. *In German Law Journal, Available at SSRN: https://ssrn.com/abstract=3896855*
4. Acquisti, A., Adjerid, I., & Brandimarte, L. (2013). Gone in 15 seconds: The limits of privacy transparency and control. IEEE Security & Privacy, 11(4), 72-74.
5. Schaub, F., Balebako, R. & Cranor, L. Designing effective privacy notices and controls. *IEEE Internet Computing*. **21**, 70-77 (2017)
6. Schaub, F., Balebako, R., Durity, A. & Cranor, L. A design space for effective privacy notices. *Eleventh Symposium On Usable Privacy And Security (SOUPS 2015)*. pp. 1-17 (2015)
7. Nouwens, M., Liccardi, I., Veale, M., Karger, D. & Kagal, L. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. *Proceedings Of The 2020 CHI Conference On Human Factors In Computing Systems*. pp. 1-13 (2020)

8 Bauer, J., Bergstrøm, R. & Foss-Madsen, R. Are you sure, you want a cookie?–The effects of choice architecture on users' decisions about sharing private online data. *Computers In Human Behavior*. **120** pp. 106729 (2021)

9 Böhme, R. & Köpsell, S. Trained to accept? A field experiment on consent dialogs. *Proceedings Of The SIGCHI Conference On Human Factors In Computing Systems*. pp. 2403-2406 (2010)

10 Turow, J., Hennessy, M., & Draper, N. (2015). The tradeoff fallacy: How marketers are misrepresenting American consumers and opening them up to exploitation. Available at SSRN 2820060.

11 Kelley, P., Cesca, L., Bresee, J. & Cranor, L. Standardizing privacy notices: an online study of the nutrition label approach. *Proceedings Of The SIGCHI Conference On Human Factors In Computing Systems*. pp. 1573-1582 (2010)

12 Gluck, J., Schaub, F., Friedman, A., Habib, H., Sadeh, N., Cranor, L. & Agarwal, Y. How short is too short? Implications of length and framing on the effectiveness of privacy notices. *Twelfth Symposium On Usable Privacy And Security (SOUPS 2016)*. pp. 321-340 (2016)

13 Harbach, M., Hettig, M., Weber, S. & Smith, M. Using personal examples to improve risk communication for security & privacy decisions. *Proceedings Of The SIGCHI Conference On Human Factors In Computing Systems*. pp. 2647-2656 (2014)

14 Wogalter, M., Conzola, V. & Smith-Jackson, T. Research-based guidelines for warning design and evaluation. *Applied Ergonomics*. **33**, 219-230 (2002)

15 Cranor, L., Guduru, P. & Arjula, M. User interfaces for privacy agents. *ACM Transactions On Computer-Human Interaction (TOCHI)*. **13**, 135-178 (2006)

16 Holtz, L., Zwingelberg, H. & Hansen, M. Privacy policy icons. *Privacy And Identity Management For Life*. pp. 279-285 (2011)

17 Tabassum, M., Alqhatani, A., Aldossari, M. & Richter Lipford, H. Increasing user attention with a comic-based policy. *Proceedings Of The 2018 CHI Conference On Human Factors In Computing Systems*. pp. 1-6 (2018)

18 Karegar, F., Pettersson, J. & Fischer-Hübner, S. The dilemma of user engagement in privacy notices: Effects of interaction modes and habituation on user attention. *ACM Transactions On Privacy And Security (TOPS)*. **23**, 1-38 (2020)

19 Zinaida Benenson, Delphine Christin, Alexander De Luca, Simone Fischer-Hübner, Thomas Heimann, Joachim Meyer. Consequence-based Privacy Decisions: a New Way to Better Privacy Management. *In: Alessandro Acquisti, Ioannis Krontiris, Marc Langheinrich, and Martina Angela Sasse. 'My Life, Shared' - Trust and Privacy in the Age of Ubiquitous Experience Sharing (Dagstuhl Seminar 13312). In Dagstuhl Reports, Volume 3, Issue 7, pp. 74-107, Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2013) https://doi.org/10.4230/DagRep.3.7.74*

20 Alex Gittens and Bülent Yener and Moti Yung. An Adversarial Perspective on Accuracy, Robustness, Fairness, and Privacy: Multilateral-Tradeoffs in Trustworthy ML. *In IEEE Access, Volume 10, 2022. https://doi.org/10.1109/ACCESS.2022.3218715*

4.5 Collective Privacy

Heather Richter Lipford (University of North Carolina at Charlotte, US, heather.lipford@uncc.edu)
Nina Gerber (Technical University of Darmstadt, DE, n.gerber@psychologie.tu-darmstadt.de)
Karola Marky (Ruhr University Bochum, DE, karola.marky@rub.de)
Jessica Vitak (University of Maryland – College Park, US, jvitak@umd.edu)
Camille Cobb (University of Illinois – Urbana-Champaign, US)

4.5.1 Introduction

Privacy has long been conceptualized at the individual level (e.g., [2, 36]). However, with the rise of social and mobile media, as well as the widespread data collection of digital trace data at scale, it is critical to move beyond individual considerations to focus on the privacy needs and risks at the group, organizational, and societal level. For example, in groups or dyads, individuals need to account for and navigate varying privacy preferences and norms; significant research has explored this as it relates to social media (e.g., [15]). Privacy in organizations is likely influenced by the more formalized and hierarchical structures that restrict information flows and create power imbalances between employers and employees. Likewise, technology developers yield significant power in determining the extent to which privacy can–or cannot–be enacted. Finally, privacy evaluations at the societal level run into challenges due to different regulatory landscapes as well as different values and norms.

Beyond these, privacy discussions must now account for the interdependent nature of data creation, use, and ownership. However, accounting for the various entanglements of people and data in a networked world is incredibly challenging and inherently interdisciplinary. Researchers have begun addressing these complexities, with work spanning the computational (e.g., [26]) and social (e.g., [3]) sciences as well as legal scholarship (e.g., [33, 34]).

Researchers have proposed various terms to describe this process, including interdependent privacy [4], networked privacy (e.g., [6]), collective privacy [16], privacy as contextual integrity [24], and comparative privacy [17]; however, these definitions do not fully encompass the complexities involved in a world where both social and technical aspects of data disclosure, collection, and use are constantly evolving and vary based on cultural, community, and contextual factors. Technical solutions alone cannot account for these factors; thus, a more human-centered approach is necessary for imagining and implementing privacy-enhancing solutions for the design of, communication about, and regulation of technologies.

Thus, we propose the following privacy framework to account for these factors:

Socio-collective privacy refers to the idea that privacy is co-constructed and negotiated within social relationships, groups, and society at large. Socio-collective privacy involves the interplay between individual privacy choices and the influence of others – peers, communities, institutions, and socio-cultural norms – in shaping those choices. From this perspective, privacy is not just a personal issue but a shared, social process, embedded in group behavior and societal expectations.

In this working group report, we describe a research roadmap for six aspects of privacy that emerged from the group's conversations at the seminar. Over multiple days, we identified these six clusters by first reviewing and discussing ideas raised during a brainstorming session with all seminar attendees. We discussed intersections and overlaps across these notes, and then expanded into related areas. We decided to focus on three levels beyond the individual

– group, organizational, and societal – and began identifying the areas that had pressing research questions at one or more of these levels. From this discussion, we see this line of research as addressing two overarching goals:

1. understanding how the interplay between the dynamics of and across levels – individual, group, organizational, and societal – impacts key privacy outcomes (e.g., perceptions and mental models; awareness and knowledge; decisions and behaviors); and
2. supporting communities of people to interact and negotiate around privacy.

In the following sections, we describe each of the six research clusters, including their core challenges, a brief summary of the state of the art, and a set of research questions to address that topic at the group, organizational, and/or societal level. We then share a roadmap that outlines research priorities over the short-, medium-, and long-term.

4.5.2 Research Clusters

4.5.2.1 The Role of Collective Actions and Structures

In the discussion, we found that in digitization, the intersection of using digital systems, privacy, and societal values [9] should become a more critical area of study. As digital services increasingly function as societal infrastructure utilized by various kinds of collectives – such as private groups and organizations – we consider it essential to investigate how collective opinions and experiences influence the design of future digital services and their privacy aspects. This includes to what extent the services need to capture and process data, but also the way in which individuals interact with privacy mechanisms (e.g., to manage consent [5]). Based on that, we recognize the role of collectives and their input in shaping digital systems and propose the following overarching research question:

- How can collective opinions/societies shape the design of digital products (inc. privacy aspects)?

Furthermore, communication structures within collectives and societies can serve as vital infrastructures for diffusing essential privacy-related knowledge. This diffusion can occur on multiple levels. For instance, privacy adepts can support their peers in private settings [11]. In doing so, privacy-related knowledge, or at least better privacy settings, are improved in informal contexts. Such support may take the form of "digital housekeeping" [35], which involves individuals helping each other manage their digital privacy practices. Existing research indicates that while these informal support systems are emerging, several barriers exist on the social level [11] and the available technological infrastructure for support is limited [12]. Beyond the private context, it is crucial to examine the privacy practices within organizations. A deeper investigation into these practices can help identify win-win situations that benefit both the organization and its employees, fostering a culture of privacy that enhances trust and collaboration. Here, we propose the following research questions:

- How can privacy-related knowledge be diffused throughout a society/collective?
- How can we create win-win situations to promote privacy (e.g., in organizations)?

Several domains have successfully leveraged collectives and societal structures through existing infrastructures, such as public healthcare initiatives, for instance, aimed at anti-smoking efforts [37] and the automotive ecosystem, which includes manufacturers, repair shops, and official inspection centers. These domains have demonstrated effectiveness in promoting collective well-being and safety in the analog world. Their success prompts an important question: how can we translate these effective strategies into digital systems? We argue that understanding the mechanisms that have facilitated collective action in these

established domains could provide valuable insights for fostering similar initiatives in digital environments. By examining the principles and practices that underpin these successful analog efforts, we can explore ways to adapt and implement them in the context of digital privacy. Based on that, we propose:

- What can we learn from other successful domains (e.g., health) to create a societal infrastructure for privacy and security?

Finally, investigating the research questions detailed above goes together with further research challenges, as it remains unclear how to effectively identify collective measures, establish specific collective goals, and track their progression from a methodological perspective. We argue that addressing these challenges is needed for ultimately developing frameworks that can guide collective action in privacy initiatives as then, we can better understand how to "mobilize" communities and assess the impact of their efforts in promoting privacy in digital environments. Based on that, we propose the following research questions:

- What are collective measures and goals that we can aim for?
- How can we collectively track progress on privacy?

4.5.2.2 Social Norms and Influences

Many theories of privacy (e.g., contextual integrity and boundary regulation (cf. [24, 27]) integrate norm-based conceptualizations of privacy. These existing theories recognize that norms may differ between groups (or organizations, cultures, etc.), and researchers have sought to understand and describe existing norms. Law and policy also regularly reference norms (or "expectations"). Following somewhat directly from these existing ideas, we call out the importance of continuing to pursue a systematic understanding of privacy norms in various groups:

- What are existing privacy norms?
- How do social norms shape/influence privacy behaviors in different contexts?

Understanding the norms of certain groups may be especially relevant to informing conversations about privacy. For example, WEIRD groups have historically had disproportionate influence on the development and design of technologies and their (lack of) privacy affordances. As we return to below, these privacy norms are important to understand in part because of their (likely) influences on norms within other groups.

However, since these societal power structures have also led to knowledge production (i.e., research) disproportionately originating from similar groups, our understanding of their privacy norms already goes beyond our understanding of norms within more niche or marginalized groups. Thus, researchers should also prioritize exploring privacy norms in groups that have been less studied. Additionally, McDonald and Forte [20] argued that taking norm-based perspectives systematically excludes the privacy preferences and needs of individuals who do not fit norms – which is often connected to them being marginalized or vulnerable. Thus, researchers pursuing studies about norms should be sensitive to the impact of vulnerability and consider divergence from the norm, even within studies on non-WEIRD groups. Researchers focused on norms must, therefore, figure out:

- How can we integrate the perspectives of diverse users (e.g. fundamentalists, vulnerable people, children) in research on socio-collective privacy?

Beyond understanding what norms exist, we must also deepen our understanding of how these norms impact individuals' privacy-related behaviors:

- How is individual behavior influenced by groups', organization's and societal norms and behaviors?

- What are the mechanisms in group dynamics that influence individual user behaviors?
- What are the differing effects for different individuals, and from different types of groups?
- What aspects of privacy norms discourage conversations about privacy?
- To what extent is technology design consistent with norms?
- How can collective opinions/societies shape design? Or, how can design reflect collective norms and opinions?

Recognizing that norms can and do shift over time and differ between groups, it is important to understand how privacy-related norms are formed and, in particular, how anti-privacy (or privacy-indifferent) norms come to be. We propose asking:

- What blocks the formation of social privacy norms?
- How do individuals – including those who already exist within normative societal bounds as well as ones at the margins or who have more fundamentalist privacy perspectives – influence groups, organizations, or societies?
- Considering sub-groups whose privacy norms differ from broader societal norms, to what extent are these groups formed around shared beliefs about privacy vs. swayed by some aspect of group dynamics toward a new norm?

Social influences seem to be shifting us toward being less concerned with privacy, so influence can be both a barrier and a potential solution to improving privacy. Perhaps a better understanding of norm formation could empower more marginalized or vulnerable group members to have more meaningful influence on the norms in the groups they belong to:

- What efforts does it take for individuals to have a meaningful influence on groups, organizations, or societies?
- What and how can public influencers (e.g. creators on TikTok or YouTube) have an impact on privacy norms and outcomes?
- How can we design such systems to support meaningful social influences regarding privacy norms?

Finally, we acknowledge that our positionality as privacy researchers comes with an interest in improving privacy and shifting the norm within societies and groups to which we belong to be more privacy-positive. To do this, we must pursue research that informs us about how to achieve these goals. For example:

- How can we create/shape/shift social privacy norms?
- How can we foster societal conversations regarding privacy?
- Can we evaluate the efficacy of focused, short-term privacy campaigns (e.g., against privacy-invasive legislation) compared to sustained, long-term efforts?

4.5.2.3 Privacy Conversations

Social influences such as stories have been found to be an effective trigger for prompting secure and privacy-conscious behavior [29, 28]. In the security context, for example, anti-phishing training that relies on discussions and role-playing have been found to be more effective in increasing self-efficacy and support-seeking than traditional training approaches [7]. Still, privacy especially is a topic that rarely comes up in conversations among non-expert users [8], and is even considered a social taboo by some people [10]. Even experts are hesitant to bring the topic up towards non-experts due to fear of disinterest and negative reactions [11]. Interestingly, in a recent study with a representative U.S. sample most people reported to be interested in having privacy conversations, but felt that other people did not care enough

about the topic, and lacked natural conversation starters [10]. These findings indicate that to fully leverage the potential of social dynamics through privacy conversations, we first have to understand why people decide to (not) engage in conversations about privacy, and, based on the results, identify strategies to facilitate privacy discussions.

Based on these considerations, the following key research questions emerge:

- What are the barriers to privacy conversations?
- How can we trigger privacy conversations?
- How can we foster ongoing privacy conversations?
- What kinds of conversations are beneficial? How are they beneficial?
- What privacy-related topics do people want to talk about?

There has been some initial research on the reasons and effects of conversations mainly focusing on the cybersecurity field. For example, Das et al. [8] conducted an interview study and identified the intention to warn others, share protection strategies, and seek advice were the main motivators of starting security conversations. Their results further imply that people might be hesitant to bring up privacy and security topics since they fear being considered as paranoid, socially inappropriate, or preachy. These findings have later been confirmed for expert users, who also questioned their moral authority to comment upon other people's privacy decisions [11]. Rader et al. [29], replicated by Pfeffer et al. [28], investigated the effect of security stories, finding that stories told in a private context were more likely to drive behavior change, while stories told by security experts were more likely to be retold. In a subsequent analysis, Rader and Wash [30] found that when telling security stories, experts tend to focus on attack mechanisms and prevention measures, while non-experts are more interested in who executed the attack and for which reasons. Further, conversations among developers on platforms such as Stack Overflow have been found to increase knowledge and awareness [14].

Other fields have been studying conversation starters in general, showing that user interfaces can successfully trigger conversations and engagement. For example, public displays have been utilized as conversation triggers, and have been found particularly effective in waiting and non-time-critical situations [1, 23]. Other approaches created interesting nuggets of information from contextual factors, e.g., 'the current temperature is equal to the coldest temperature ever measured in Sao Paulo [21]. Recently, Murtezaj et al. [22] have highlighted the potential of using public security user interfaces for increasing awareness and promoting behavior change.

4.5.2.4 Digital Literacy

One of the goals of privacy interventions is to contribute to the awareness and knowledge of individuals. Digital literacy is also a precursor to making informed privacy decisions and adopting privacy tools. As the above section discussed, one of the ways that users learn about security and privacy is through informal stories shared among friends and family. Organizations and societies also spread knowledge through a variety of formal and informal channels, such as media, communications, influencers, and so on. Thus, understanding and improving digital literacy requires a focus on communities and how knowledge is spread within and through them. Research questions then arise regarding:

- What level of knowledge is sufficient (for individuals or the group)?
- What "digital survival skills" should groups of people have?
- How can that knowledge be diffused throughout a group/community/society?

- To what extent is it important for this knowledge to be spread throughout a community (e.g., all individuals have knowledge) vs. concentrated amongst experts within the group?
- What is the impact of various forms of media on knowledge and awareness?

There has been some research demonstrating what and how users learn about security from others, such as personal contacts [29] and the media [31]. Yet, few have examined how to measure privacy-related knowledge of communities of people, and how that knowledge can impact individual decisions and outcomes. As mentioned previously, a starting point may be to draw upon research strategies or results in other domains, such as public health, where similar issues of community-oriented awareness and knowledge diffusion have been examined. Similarly, few privacy researchers have investigated interventions to promote knowledge sharing and diffusion within groups as a step towards improving privacy management.

4.5.2.5 Cultural Differences

Privacy values – and consequently privacy regulations – vary across the world. This can perhaps most clearly be seen in the different regulatory approaches taken by the European Union (EU) and the United States. The EU enacted the General Data Protection Regulation (GDPR) nearly a decade ago, specifying a range of privacy protections for citizens and requirements for companies that collect personal data. The U.S. has yet to pass comprehensive federal privacy reforms, relying instead on sector-specific and state laws. Thus, it is important to consider the role that cultural differences play in the design, implementation, and (non)use of technology. However, as Li [13] notes, few privacy studies account for cultural differences in their design. Therefore, we suggest three research questions to guide future work in this space:

- How are the cultural differences in privacy in the offline world related to cultural differences in the online world?
- How are cultural differences related to all of the other research we have outlined in this report?
- What aspects of an organizational privacy culture are related to desired privacy behaviors?

4.5.2.6 Bystander Privacy

An important consideration in privacy is that there are people impacted beyond an individual user. For example, many smart devices are capable of capturing information about "bystanders" – those near the device but who have no interaction with or control of that device [18]. A classic example are always-on doorbell cameras that capture people walking by on a public street, who likely have no awareness of when and by whom they are being recorded. Other example include individuals sharing information about other people, such as a person posting a photo of others, or visiting a friend who has smart home devices [19]. Even when individuals are carefully considering privacy, they may be unaware of the "collateral damage" that their technology usage or decisions can cause. For example, people who have utilized DNA services may inadvertently impact the privacy of current and future relatives. We note that the word "bystander" does not adequately describe many of these situations, and other terms have also been used, such as incidental user or non-user [18]. But what unites all of these is the lack of agency and control, or even awareness, of the way that information is collected, used, and shared about oneself by others. Key research questions related to this area are:

- When and how are people impacted by decisions of others and hence lack agency and/or control?

- How can we empower people in multi-user settings and facilitate privacy negotiations?
- How can we leverage protection from collateral damage caused by other individuals' or groups' privacy decisions?

The problem of bystanders has been widely researched and discussed in relation to particular technologies, such as in the examples mentioned above of photo sharing, smart homes [18, 19], and DNA data sharing [32, 25]. General guidelines have been presented in such research, such as to minimize data collection or increase the transparency of data collection. There have also been a range of specific design interventions proposed, yet many of these are specific to a particular device and use case. For example, there has been extensive research on ways to automatically detect and protect the privacy of bystanders in photos. Yet there has only been limited industry uptake of a few of these, such as visual indicators of recording on smart devices as individuals might not buy too obvious devices [19]. Outside of these examples, exploration and evaluation of solutions is still quite limited, and does not cover the range of privacy harms that can come to bystanders in a variety of contexts. Finally, technical solutions may not always be feasible and instead groups may need to rely on social conventions and conversations to negotiate privacy protections – other research areas discussed in this report.

4.5.3 Research Roadmap

We already have many of the methods and tools to approach the questions discussed above, and we propose that we should start pursuing those right away. These methods, which we already use, largely came out of collaborations with psychologists, who tend to take an individual-centric approach to questions and methods. Given that the questions we map here are societal/group, we recognize that our existing interdisciplinary connections will need to expand. So, to deepen and expand the questions in this document and envision next steps, we will need to strengthen collaboration with scholars from other fields such as sociology and anthropology, which will require new types of incentives and funding to support that. These collaborations will be slow to come to fruition, and thus should also be started immediately.

While researchers may already be using relevant methods, there is a gap when it comes to measures. There are fewer existing measures of group-level privacy used in our research community. For example, an open research question is how to measure privacy outcomes at a societal level. Foundations for such measures may again be inspired from other fields that research organizational or societal phenomena. Another challenge is that creating interventions suffers from many of the same challenges as any multi-user system, with more complex design requirements, implementations, and evaluation protocols.

Funding for this research needs to support overcoming these challenges, and incentivize researchers towards our overarching goals. Funding agencies can encourage interdisciplinary work through specialized programs, post-doctoral opportunities, interdisciplinary symposia, summer schools, and other novel structures to encourage researchers from different communities to share ideas and collaborate on projects.

A goal of all of the above research is to provide practical and actionable guidance for socio-technical designs that empower people to achieve their own and their community's privacy needs. Thus, we also discussed the following two questions for the research community:

- How do we help researchers to influence technology design and policy making? In other words, how can we help researchers to communicate and frame results for interested communities, such as regulators, legislators, and technology creators?
- How can we help people choose research questions and study designs that speak to those different audiences?

Finally, we note that our research can also help communities of users contribute to technology and societal outcomes – empowering people to influence designers and regulators towards meetings privacy needs. Thus, we end with two final research questions:

- What privacy advocacy has an impact on regulations?
- How can we support and promote privacy advocacy?

References

1 Alt, F., Kubitza, T., Bial, D., Zaidan, F., Ortel, M., Zurmaar, B., Lewen, T., Sahami Shirazi, A., & Schmidt, A. (2011). Digifieds: insights into deploying digital public notice areas in the wild. In Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia (MUM '11). Association for Computing Machinery, New York, NY, USA, 165–174. https://doi.org/10.1145/2107596.2107618

2 Altman, I. (1975). The environment and social behavior: privacy, personal space, territory, and crowding.

3 Anthony, D., Campos-Castillo, C., & Horne, C. (2017). Toward a sociology of privacy. Annual review of sociology, 43(1), 249-269. https://doi.org/10.1146/annurev-soc-060116-053643

4 Biczók, G., & Chia, P. H. (2013). Interdependent privacy: Let me share your data. In Financial Cryptography and Data Security: 17th International Conference, FC 2013, Okinawa, Japan, April 1-5, 2013, Revised Selected Papers 17 (pp. 338-353). Springer Berlin Heidelberg. https://link.springer.com/chapter/10.1007/978-3-642-39884-1_29

5 Böhme, R., & Köpsell, S. (2010). Trained to accept? A field experiment on consent dialogs. In Proceedings of the SIGCHI conference on human factors in computing systems (pp. 2403-2406).

6 Boyd, D. (2012). Networked privacy. Surveillance & society, 10(3/4), p.348.

7 Chen, X., Sacré, M., Lenzini, G., Greiff, S., Distler, V. & Sergeeva, A. (2024). The Effects of Group Discussion and Role-playing Training on Self-efficacy, Support-seeking, and Reporting Phishing Emails: Evidence from a Mixed-design Experiment. In Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 829, 1–21. https://doi.org/10.1145/3613904.3641943

8 Das, S., Hyun-Jin Kim, T., Dabbish, L.A., and Hong, J.I. (2014). The effect of social influence on security sensitivity. In 10th Symposium On Usable Privacy and Security (SOUPS 2014), pages 143–157, Menlo Park, CA, July 2014. USENIX Association.

9 Friedman, B. (1996). Value-sensitive design. interactions, 3(6), 16-23.

10 Gerber, N., Zimmermann, V., von Preuschen, A., & Renaud, K. (2025). Unpacking the Social and Emotional Dimensions of Security and Privacy User Engagement. In 21st Symposium on Usable Privacy and Security (SOUPS 2025).

11 Nina Gerber and Karola Marky. The nerd factor: The potential of S&P adepts to serve as a social resource in the user's quest for more secure and Privacy-Preserving behavior. In Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022), pages 57–76, Boston, MA, August 2022. USENIX Association. https://www.usenix.org/conference/soups2022/presentation/gerber

12 Horst, H., & Sinanan, J. (2021). Digital housekeeping: Living with data. New Media & Society, 23(4), 834-852.

13 Li, Y. (2022). Cross-Cultural Privacy Differences. In: Knijnenburg, B.P., Page, X., Wisniewski, P., Lipford, H.R., Proferes, N., Romano, J. (eds) Modern Socio-Technical Perspectives on Privacy. Springer, Cham. https://doi.org/10.1007/978-3-030-82786-1_12

14 Tamara Lopez, Thein Tun, Arosha Bandara, Levine Mark, Bashar Nuseibeh, and Helen Sharp. An anatomy of security conversations in stack overflow. In 2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS), pages 31–40. IEEE, 2019. https://doi.org/10.1109/ICSE-SEIS.2019.00012

15 Mansour, A., & Francke, H. (2021). Collective privacy management practices: A study of privacy strategies and risks in a private Facebook group. Proceedings of the ACM on Human-Computer Interaction, 5(CSCW2), 1-27. https://doi.org/10.1145/3479504

16 Mantelero, A. (2017). From Group Privacy to Collective Privacy: Towards a New Dimension of Privacy and Data Protection in the Big Data Era. In: Taylor, L., Floridi, L., van der Sloot, B. (eds) Group Privacy. Philosophical Studies Series, vol 126. Springer, Cham. https://doi.org/10.1007/978-3-319-46608-8_8

17 Masur, P. K., Epstein, D., Quinn, K., Wilhelm, C., Baruh, L., & Lutz, C. (2025). Comparative privacy research: Literature review, framework, and research agenda. The Information Society, 1-22. https://doi.org/10.1080/01972243.2025.2451863

18 Marky, K., Voit, A., Stöver, A., Kunze, K., Schröder, S., & Mühlhäuser, M. (2020, October). "I don't know how to protect myself": Understanding Privacy Perceptions Resulting from the Presence of Bystanders in Smart Environments. In Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society (pp. 1-11).

19 Marky, K., Gerber, N., Pelzer, M. G., Khamis, M., & Mühlhäuser, M. (2022). "You offer privacy like you offer tea": Investigating mechanisms for improving guest privacy in IoT-equipped households. Proceedings on Privacy Enhancing Technologies.

20 McDonald, N., & Forte, A. (2020, April). The politics of privacy theories: Moving from norms to vulnerabilities. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (pp. 1-14).

21 Memarovic, N., Elhart, I., & Langheinrich, M. (2011, December). FunSquare: First experiences with autopoiesic content. In Proceedings of the 10th International Conference on Mobile and Ubiquitous Multimedia (pp. 175-184).https://doi.org/10.1145/2107596.2107619

22 Murtezaj, D., Paneva, V., Distler, V., & Alt, F. (2024, September). Public Security User Interfaces: Supporting Spontaneous Engagement with IT Security. In Proceedings of the New Security Paradigms Workshop (pp. 56-70).

23 Müller, J., Walter, R., Bailly, G., Nischt, M., & Alt, F. (2012, May). Looking glass: a field study on noticing interactivity of a shop window. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 297-306).

24 Nissenbaum, H. (2009). Privacy in Context: Technology, Policy, and the Integrity of Social Life. Stanford Law Books.

25 Niu, Y., Meng-Schneider, N., Qiu, W., & Kokciyan, N. (2025, April). "I am not the primary focus"-Understanding the Perspectives of Bystanders in Photos Shared Online. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (pp. 1-23).

26 Palen, L., & Dourish, P. (2003, April). Unpacking "privacy" for a networked world. In Proceedings of the SIGCHI conference on Human factors in computing systems (pp. 129-136).

27 Petronio, S. (2002). Boundaries of privacy: Dialectics of disclosure. SUNY Press.

28 Pfeffer, K., Mai, A., Weippl, E., Rader, E., & Krombholz, K. (2022). Replication: Stories as informal lessons about security. In Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022) (pp. 1-18).

29 Rader, E., Wash, R., & Brooks, B. (2012, July). Stories as informal lessons about security. In Proceedings of the Eighth Symposium on Usable Privacy and Security (pp. 1-17).

30 Rader, E., & Wash, R. (2015). Identifying patterns in informal sources of security information. Journal of Cybersecurity, 1(1), 121-144.

31 Raphael, M. M., Kanta, A., Seebonn, R., Dürmuth, M., & Cobb, C. (2024). Batman Hacked My Password: A Subtitle-Based Analysis of Password Depiction in Movies. In Twentieth Symposium on Usable Privacy and Security (SOUPS 2024) (pp. 199-218).

32 Saqib, E., He, S., Choy, J., Abu-Salma, R., Such, J., Bernd, J., & Javed, M. (2025). Bystander Privacy in Smart Homes: A Systematic Review of Concerns and Solutions. ACM Transactions on Computer-Human Interaction.

33 Solove, D. J. (2004). The digital person: Technology and privacy in the information age (Vol. 1). NYU Press.

34 Solove, D. J. (2025). Artificial intelligence and privacy. Fla. L. Rev., 77, 1.

35 Tolmie, P., Crabtree, A., Rodden, T., Greenhalgh, C., & Benford, S. (2007, September). Making the home network at home: Digital housekeeping. In ECSCW 2007: Proceedings of the 10th European Conference on Computer-Supported Cooperative Work, Limerick, Ireland, 24-28 September 2007 (pp. 331-350). London: Springer London.

36 Westin, A. F. (1968). Privacy and freedom. Washington and Lee Law Review, 25(1), 166.

37 Wood, W., & Neal, D. T. (2016). Healthy through habit: Interventions for initiating & maintaining health behavior change. Behavioral Science & Policy, 2(1), 71-83.

Participants

- Florian Alt
LMU München, DE
- Zinaida Benenson
Universität Erlangen-Nürnberg, DE
- Bettina Berendt
TU Berlin, DE
- Benjamin Berens
KIT – Karlsruher Institut für Technologie, DE
- Nataliia Bielova
INRIA – Sophia Antipolis, FR
- Camille Cobb
University of Illinois – Urbana-Champaign, US
- Ha Dao
MPI für Informatik – Saarbrücken, DE
- Cori Faklaris
University of North Carolina – Charlotte, US
- Simone Fischer-Hübner
Karlstad University, SE
- Nina Gerber
TU Darmstadt, DE
- Sophie Grimme
OFFIS – Oldenburg, DE
- Andreas Gutmann
Ofcom – London, GB
- Dominik Herrmann
Universität Bamberg, DE
- Adam Jenkins
King's College London, GB
- Bailey Kacsmar
University of Alberta – Edmonton, CA
- Apu Kapadia
Indiana University – Bloomington, US
- Farzaneh Karegar
Karlstad University, SE
- Agnieszka Kitkowska
Jönköping University, SE
- Marc Langheinrich
USI – Lugano, CH
- Karola Marky
Ruhr-Universität Bochum, DE
- Mainack Mondal
Indian Institute of Technology – Kharagpur, IN
- Simran Munot
MPI für Informatik – Saarbrücken, DE
- Alena Naiakshina
Universität Köln, DE
- Sameer Patil
University of Utah – Salt Lake City, US
- Maija Poikela
Charité – Berlin, DE
- Sören Preibusch
BfR – Berlin, DE
- Elissa Redmiles
Georgetown University – Washington, DC, US
- Heather Richter Lipford
University of North Carolina – Charlotte, US
- Arianna Rossi
Sant'Anna School of Advanced Studies – Pisa, IT
- Cristiana Santos
Utrecht University, NL
- Anastasia Sergeeva
University of Luxembourg, LU
- William Seymour
King's College London, GB
- Jose Such
Technical University of Valencia, ES & King's College London, GB
- Jessica Vitak
University of Maryland – College Park, US
- Mark Warner
University College London, GB
- Daricia Wilkinson
Arizona State University – Mesa, US
- Yixin Zou
MPI-SP – Bochum, DE