Report from Dagstuhl Seminar 25021

Grand Challenges for Research on Privacy Documents

Florian Schaub[*1], Christine Utz[*2], Shomir Wilson[*3], and Lu Xian[†4]

1 **University of Michigan – Ann Arbor, US.** fschaub@umich.edu
2 **Radboud University Nijmegen, NL.** christine.utz@ru.nl
3 **Pennsylvania State University – University Park, US.** shomir@psu.edu
4 **University of Michigan – Ann Arbor, US.** xianl@umich.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 25021 "Grand Challenges for Research on Privacy Documents" held in January 2025. This Dagstuhl Seminar gathered an interdisciplinary group of researchers from privacy, natural language processing, human-computer interaction, public policy, and law to identify and characterize key challenges to research on privacy documents, such as privacy policies, terms of use, cookie policies, and other texts about data practices.

Seminar participants worked together to identify and characterize key challenges in privacy document research with the goal of producing a research roadmap for tackling these challenges. Through a series of perspectives talks and panel discussions, participants exchanged experiences in working with privacy documents in research and learned about associated challenges, as well as interdisciplinary intersections and policy considerations. Through deeper engagement in working groups, participants deeply explored research challenges and research directions across five interconnected topics: (1) formats and standardization; (2) datasets, automation, and analysis methods; (3) usable and useful notice and consent; (4) consumer privacy beyond notice and choice; and (5) cross-stakeholder engagement.

Seminar January 5–10, 2025 – https://www.dagstuhl.de/25021
2012 ACM Subject Classification Applied computing → Law, social and behavioral sciences; Computing methodologies → Natural language processing; Security and privacy → Human and societal aspects of security and privacy; Social and professional topics → Privacy policies
Keywords and phrases Human-Computer Interaction, Machine Learning, Natural Language Processing, Privacy Policy, Public Policy
Digital Object Identifier 10.4230/DagRep.15.1.1
Funding Organizers Schaub and Wilson would like to acknowledge support by the National Science Foundation under Award No. 2105734 and 2105736 ("Collaborative Research: SaTC: CORE: Medium: A Large-Scale, Longitudinal Resource to Advance Technical and Legal Understanding of Textual Privacy Information").

1 Executive Summary

Shomir Wilson (Pennsylvania State University – University Park, US)
Florian Schaub (University of Michigan – Ann Arbor, US)
Christine Utz (Radboud University Nijmegen, NL)

The five-day Dagstuhl Seminar "Grand Challenges for Research on Privacy Documents" gathered an interdisciplinary group of researchers from privacy, natural language processing, human-computer interaction, public policy, and law to identify and characterize key challenges to research on privacy documents, such as privacy policies, terms of use, cookie policies,

[*] Editor / Organizer
[†] Editorial Assistant / Collector

and other texts about data practices. In the status quo, privacy documents primarily serve the compliance needs of companies, while failing to fulfill the needs of other stakeholders in our information society. Although many Internet users have concerns about their privacy, most lack the time, knowledge, and other resources to understand these documents, leaving them underinformed and compromising the goals of the notice and choice paradigm. The needs of other stakeholders, including regulators, researchers, policymakers, and privacy practitioners, are similarly stymied. Although a growing body of research is devoted to analyzing, reconstituting, or otherwise using these documents to satisfy stakeholders' needs, broader interdisciplinary efforts are needed.

The goal of this seminar was to identify and characterize key challenges in privacy document research and to produce a research roadmap of how to tackle them in order to move the field forward. At a high level, the seminar schedule was structured into two stages to produce those outcomes. The first stage consisted of a series of perspectives talks providing background and introductions to relevant disciplines and approaches, as well as thematic panel discussions among participants. In the second stage, participants organized in topical working groups to more deeply explore specific areas and develop elements of the roadmap. Working groups focused on the following themes: (1) document formats and standardization; (2) datasets, automation, and analysis methods; (3) usable and useful notice and consent; (4) consumer privacy beyond notice and choice; and (5) cross-stakeholder engagement.

This report collects abstracts of the three perspective talks and presents research challenges and directions identified in the working groups. Each section describes respective challenges, key research questions, and solution ideas and directions. We present this report to the research community as a resource for discussion and inspiration for future work.

2 Table of Contents

Executive Summary
Shomir Wilson, Florian Schaub, and Christine Utz 1

Overview of Perspective Talks

Legal Perspectives on Privacy Documents
Kirsten Martin 4

NLP/ML Perspectives on Privacy Documents
Sepideh Ghanavati 4

Human-centered Perspectives on Privacy Documents
Simone Fischer-Hübner 4

Working Groups

Formats and Standardization
Christine Utz, Rinku Dewri, Emma Tosch, and Lu Xian 5

Datasets, Automation, and Analysis Methods
Peter Story, Sepideh Ghanavati, Henry Hosseini, Jelena Mitrovic, Tim Samples, Isabel Wagner, and Tianyang Zhao 10

Usable and Useful Notice & Consent
Simone Fischer-Hübner, Kai-Wei Chang, Nico Ebert, Agnieszka Kitkowska, and Shidong Pan 17

Consumer Privacy Beyond Notice and Choice
Noah Apthorpe, Eleanor Birrell, Travis Breaux, Kirsten Martin, Rishab Nithyanand, Sarah Radway, Yan Shvartzshnaider, and Maximiliane Windl 23

Cross-Stakeholder Interaction
Jose M. del Alamo, Soheil Human, Konrad Kollnig, Daniel Smullen, and Kami Vaniea 29

Participants 33

3 Overview of Perspective Talks

3.1 Legal Perspectives on Privacy Documents

Kirsten Martin (Carnegie Mellon University – Pittsburgh, US)

This perspective talk explained why privacy documents have evolved as large, ambiguous documents that do not serve the needs of stakeholders. Privacy laws can be seen as entering a second phase. Where a first phase of privacy regulation focused on the handoff of data to firms as if privacy is relinquished when shared with firms. This placed a heavy burden on privacy notices to ensure individuals choose firms correctly to 'give up' their data. The second phase correctly pivots to recognize that people have privacy interests and rights even when they have shared data with firms.

3.2 NLP/ML Perspectives on Privacy Documents

Sepideh Ghanavati (University of Maine, US)

Privacy policy documents aim to inform users about how their personal information is collected, used, processed, or shared. However, the documents are generally long, contain legal jargon, and include vagueness and ambiguities, which make it hard for the end-users to understand them and make informed decisions. In the last 20 years, researchers leveraged machine learning (ML) and natural language processing (NLP) techniques to create datasets of annotated policies, extract features from the privacy policies, analyze data practices, assess the readability, usability, and utility of privacy policies, and evaluate the consistency and compliance with laws, regulations, and best practices. This talk provided an overview of the state-of-the-art research regarding privacy document analysis with ML/NLP techniques, identifying the key features and contributions, and then provided guidelines and potential research directions for future work.

3.3 Human-centered Perspectives on Privacy Documents

Simone Fischer-Hübner (Karlstad University, Chalmers University of Technology and Gothenburg University, SE)

Privacy by Design can only be achieved if transparency and control functions for users are usable. Therefore, the GDPR is also requiring that data subject rights functionality must be provided in "concise, transparent, intelligible and easily accessible form, using clear and plain language."

This talk first addressed and discussed challenges for human-centered privacy related to privacy documents. These include challenges of notice and consent, the challenge that privacy is only a secondary goal for users, the users' limited rationality and psychological

effects that need to be considered, the challenge that there are no one-size fits all solutions for different types of users and contexts, as well as challenges of explaining privacy-enhancing technologies (PETs) if mentioned in privacy documents.

Secondly, solutions for approaching these challenges and remaining challenges were outlined. It was highlighted that privacy by design of privacy documents needs to be aligned and combined with human-centered and inclusive design. Moreover, various approaches for designing usable privacy notices and usable explanations for PETs were discussed as well as approaches and guidelines for raising the user's attention to essential policy information by engaging them with interactive policy content.

4 Working Groups

4.1 Formats and Standardization

Christine Utz (Radboud University Nijmegen, NL), Rinku Dewri (University of Denver, US), Emma Tosch (Northeastern University – Boston, US), and Lu Xian (University of Michigan – Ann Arbor, US)

Certain firms are legally required to provide privacy documents that notify consumers or end-users of how their data will be collected, used, stored, and transmitted. When firms operate in multiple jurisdictions, they may be required to provide multiple types of privacy documents, each having different information. Furthermore, some firms may wish to provide users with privacy-related information beyond legal requirements. As a result, any discussion of privacy documents involves a many-to-many relationship between heterogeneous stakeholders. This entails studying privacy policies across a variety of sources.

Locating and understanding these different sources of privacy documents in the field is critically important for researchers and legislators performing privacy audits, as well as end-users who have a right to know how their data is being used. Unfortunately, there are no established standards for this information, neither in terms of a canonical location nor format and content: Some websites provide traditional text documents via a clearly identifiable privacy policy link from their landing page, while others include this information in their Terms of Service. Some mobile applications employ buttons and icons to convey privacy information, while still others condense this information into more user-friendly privacy labels. Each of these user-facing formats has strengths while also presenting challenges for satisfying stakeholders' interests.

There are currently no commonly used consensus standards against which firms write their privacy documents. This lack of standardization leads to variability in their location, scope, format, and other features, which in turn creates difficulties for users, legislators, and other stakeholders to find, analyze, and act upon these documents. Of particular relevance is how to both manually and automatically find and evaluate documents for compliance.

4.1.1 Challenges

There are technical, social, and legal challenges to creating consensus around a standard for privacy documents, especially those that are end-user facing. Prior attempts at standardization have largely been seen as failures. We enumerate the challenges currently facing stakeholders and contextualize them in relation to past attempts at standardization.

4.1.1.1 Technical challenges: underspecified or volatile document features

There are a number of technical challenges to standardization. We enumerate the most salient of these below.

In this section we assume that a given product, system, or service requires a privacy document of some kind (i.e., we do not discuss the conditions under which such a product, system, or service would necessitate a privacy document). We abstract over *producers* of a document and *consumers* of a document; when the document feature entails different challenges for different elements of the producer-consumer relation, we ground these actors with specific examples.

Document location. There is no standard location where to find disclosures about a company's privacy practices, neither in terms of visual placement nor in terms of file path. For example, on the Web, links to a website's privacy policy are often placed in the website footer, placing a burden on the user to scroll down to find it. In mobile app stores, a link in the app listing should lead to the app's privacy policy, but often only leads to the developer's website, where the user has to conduct further investigation to find the policy. The problem is exacerbated if privacy information is made available in multiple languages. In these cases, the firm or service provider produces the privacy document, while the consumer may be an end-user, a bot or crawler, a lawyer, or even another producer (e.g., a company seeking to refer to a third party's privacy policy).

Dynamically generated web pages and symbolic paths complicate the automated localization of privacy documents. This is especially true for automated agents. That said, there are benefits to dynamic generation: rather than enumerating the full set of data practices, the producer can specialize their content to the particular product, service, jurisdiction, end-user, application, etc. This specialization benefits an end-user consumer, who only sees the relevant information. However, implementing this functionality correctly is quite challenging due to the range of components that the privacy document for an arbitrary producer-consumer pair may require.

Document granularity. The range of necessary information to include in a privacy document for a given producer-consumer pair points to the next challenge in developing standards: document granularity and scope. By "document granularity" we mean what fragment of a company's privacy disclosures are contained in a single document. By "scope" we mean the territorial, personal, and material boundaries the provisions of the privacy document are intended to apply to and the boundaries and extent of information that the policy covers regarding the collection, use, storage, and sharing of personal data. Scope can contribute to issues with granularity: some organizations have a single document named "Privacy Policy"; some incorporate this information into their Terms of Service, while others spread privacy information across multiple documents to make it more easily digestible or to adapt the presented information to different regulatory environments (special audiences, fields, or jurisdictions). Furthermore, organizations may provide separate privacy documents for each of their services. Finally, multiple pieces of software may have individual privacy policies that must be merged into a new service that uses them. There are currently no standard ways to compose these documents, leading to complex and illegible practices.

Document format. Privacy documents also widely differ in qualitative features that an end-user would experience. These include file type (e.g., PDF, HTML, plain text file), media type (text, tables, images, video), and subdivision into subsections and paragraphs. The information in privacy documents can also be visualized, summarized, and synthesized via non-

textual means of presentation, such as icons (e.g., the CCPA icon [1] or "nutrition labels" [2]). Privacy policy text also does not have to be static text but is sometimes dynamically generated. Beyond finding policy text, these differences in format and creation add additional obstacles to extracting privacy policy text for further processing and analysis [3].

Document content, audience, applicable jurisdiction, and scope. Format and presentation are closely related to the content of privacy documents, which in turn is influenced by layered regulatory requirements. Based on the location or jurisdiction alone, there can already be multiple tiers of regulation that apply, including the supranational (EU privacy laws), federal (national data protection laws), and state level (e.g., California Consumer Privacy Act (CCPA), Washington Privacy Act (WPA)). Other requirements hail from special regulations for protected audiences, such as the Children's Online Privacy Protection Act (COPPA), or specific fields or industries, such as the Health Insurance Portability and Accountability Act (HIPAA) for healthcare, the Gramm–Leach–Bliley Act (GLBA) for finance, or Family Educational Rights and Privacy Act (FERPA) for education.

While these regulations specify which content to include in privacy documents (e.g., the CCPA mandates that privacy policies need to state whether personal information is sold or shared for marketing purposes; the GDPR's disclosure requirements in Articles 13 and 14), they usually lack specific guidelines on how such content has to be presented.

4.1.1.2 Challenges of the standardization process

The technical challenges related to document location, granularity and scope, and format all impact any process of standardization. The need for standardization can be felt across different formats of privacy communication to facilitate adoption and address inherent differences in the objective of the communication. They should cover different modalities of communication, as well as the parties on either end of the communication. While the establishment of (privacy communication) standards at multiple points of a product's lifecycle may seem to be the key to richer communication, the fatigue of enforcing such standards could demotivate meaningful implementation. The challenge therefore lies in assessing the effective benefit of standardizing one or more points of communication, and whether the benefits will outweigh the complexity of the induced processes to meet such standards. This is amplified by the potential risk of introducing inconsistencies between different representations of a privacy concept across these multiple points. For example, this is applicable if different document variants are standardized for different stakeholders or when alternative visualizations are used to familiarize consumers about privacy practices. Vocabulary mismatches across standards, as well as in relation to regulations, can become failure causes in standards adoption (e.g., in P3P [7]).

The purpose behind data collection is ever evolving, can be generic in nature, and can be subject to a variety of legal and ethical frameworks. Under such situations, the requirements that a standard should meet are unclear. Standards can and do undergo revisions to incorporate the changing landscape of applicable platforms; however, the requirement to meet frequently revised standards may be seen as an operational burden and hamper adoption.

Furthermore, although it is understood that different stakeholders have different expectations from a privacy document, the precise nature of those expectations is understudied. Hence, what communication a standard should facilitate, and what structure is ideal for such communication, is likely to present itself as a challenge in standards development. Taxonomies help organize content in meaningful, unambiguous, and contained ways. The FPC states 11 Fair Information Practice Principles (FIPPs) for efficient privacy management [4]: consent

and choice; purpose legitimacy and specification; collection limitation; data minimization; use/retention/disclosure limitation; accuracy and quality; openness, transparency and notice; individual participation and access; accountability; information security; and privacy compliance. On the other hand, studies around privacy documents have focused on select categories of information: first party collection/use; third party sharing/collection; user choice/control; user access, edit & deletion; data retention; data security; policy change; Do Not Track; and international & specific audiences. These categories are loosely tied to the FIPPs and have served as a gold standard in data annotation and automated analysis attempts [6]. Nonetheless, the completeness of these categories in capturing all pertinent aspects of privacy communication, especially with respect to different stakeholders, is unknown. The lack of a well-crafted taxonomy creates barriers to standardization, which inherently must find ways to balance expressiveness and verbosity. On a similar note, a standardized vocabulary of privacy-relevant terms is also missing, which may lead to conflicting interpretations across specifications.

4.1.1.3 Sociopolitical challenges

A process of standardization can also be vulnerable to sociopolitical challenges. As learned from the P3P standardization process, this can manifest as waste of time while oscillating between specificity and generality, the introduction of too much transparency as viewed by specific stakeholders, the potential for misuse to generate a false notion of privacy, the introduction of the notion that a specification would substitute legislation, disparate stakeholders hinging on others to take the first step (who goes first – policy writers or policy checkers), and bare minimum implementations of a specification.

4.1.1.4 Technical challenges for firms attempting compliance

As consensus documents produced by formal organizations, privacy standards are typically designed to address legal obligations. These standards must then be translated into technical solutions. Additional challenges arise during this process. For instance, programmers face challenges when interpreting the content of a privacy policy against what a product or service actually does or against what it could do, especially when the product or service predates the policy.

4.1.2 Key research questions

What goals should a standardized privacy communication attempt to reach? Care has to be taken to account for the different expectations from involved groups (regulators, consumers, business owners), and accordingly preserve the indispensable elements. As such, identified objectives will significantly drive the design process of a standardized format, the specificity of the content used to realize a standard, and the feasibility of assessing if the objectives are met. The introduction of stakeholder-specific standards will inevitably introduce subsequent questions on standards mapping and also inform the creation of a comprehensive vocabulary and taxonomy of privacy communication artifacts. Some key research questions to consider include:

- What are the points in the data processing pipeline that could benefit from standards or format templates, and what is the smallest set of privacy data requirements to enable the functionality of a given service beyond that point?

- Where are there opportunities for auto-generating privacy documents, what kinds of formats should an auto-generator produce, and what entities and processes ought to govern the approval of new output formats?
- What inconsistencies can arise when multiple formats for the same concept are created, and how to address interpretive variations in these formats?
- What level of complexity of formal specification is required to capture the minimal expressiveness of different privacy document specifications?
- What quantitative and qualitative methods are needed to measure the conformance of an implementation to a specific standard, and what avenues exist (or need to be created) to integrate such conformance testing into an organization's operational activities? While quantitative methods are aimed at generating measurable insights into an implementation's adherence to concrete requirements of a specification (e.g., a data retention practice must unambiguously indicate a retention period; a purpose definition must be tied to every collected data artifact), qualitative methods are aimed at assessing conformance in terms of coverage, clarity, etc. Integration also includes the notion of feedback to facilitate iterative refinements.
- How to avoid sociopolitical and technical pitfalls in a standards development process? Recommendations such as simplicity and management of expectations are present from the P3P experience, but a deeper discussion is desired to prepare for previously unseen barriers.

4.1.3 Research directions

The specific solutions that resolve or address the challenges we enumerate should be in service of the following concrete desired outcomes:

- Identification of standardization points that serve (preferably) separate sections of an end-to-end privacy communication pipeline
- A systematization of objectives to be met at various standardization points
- A clear statement of requirements that implementation of the standard must meet: necessary, unambiguous, complete, precise, well-structured, consistent, testable [5]
- Formal methods to check for consistency violations in the standards; note: not only across revisions of a specific standard but also across manifestation of the same principle across standardization points
- A baseline implementation with guidelines for extensibility

Next, we discuss specific suggestions for the standardization of certain privacy document features.

Document location. This can be standardized by means of a known semantic endpoint for privacy-related information. Existing similar standards and proposals include:

- `robots.txt` (`https://www.rfc-editor.org/rfc/rfc9309`) – directives for bots which parts of a website (not) to access.
- `security.txt` (`https://www.rfc-editor.org/rfc/rfc9116.html`) – point of contact for security vulnerability notifications, already used by major companies.
- Other proposed standards of metadata files at well-known locations on web servers include `ads.txt`, `human.txt`, and `sellers.json`.
- There is already a proposal for a `privacy.txt` standard (https://privacytxt.dev), but unlike security or ads, the disclosures required in privacy policies are not universal and may vary between jurisdictions, which this proposal does not account for.

- Some companies already seem to acknowledge the problem of privacy documents being hard to find and automatically process, prompting them to provide a version of their privacy policy as a simple plaintext file; one example is Hulu (`https://www.hulu.com/privacy.txt`).

Document content and format. Templates can serve to capture high-level structures (e.g., prescribed headings, paragraphs, placements, etc.) as well as low-level structures (e.g., data items, data collector, purpose, trigger, mechanism, etc.).

Cross-referencing between multiple privacy documents. Privacy documents are hosted in a variety of formats, including privacy policy text, links, nutrition labels, and icons. Existing privacy documents often refer to each other. For example, a section in the general privacy document may refer to another, jurisdiction-specific document (e.g., designated privacy policies for residents of California, who are subject to the CCPA). These cross-references complicate human understanding of the issuer's exact data practices and their associated privacy rights. Designing a standardized path through the documents, links, labels, and icons would create a hierarchical organization of discrete, distributed documents that would aid different stakeholders in finding the applicable privacy information.

References

1 California Department of Justice. *CCPA Privacy Icons*. `https://oag.ca.gov/privacy/ccpa/icons-download`, accessed May 21, 2025.

2 Apple Inc. *App Privacy Details*. `https://developer.apple.com/app-store/app-privacy-details/`, accessed May 21, 2025.

3 H. Hosseini, M. Degeling, C. Utz, and T. Hupperich. *Unifying Privacy Policy Detection*. Proceedings on Privacy Enhancing Technologies (PoPETs), 2021(4), pp. 480–499. `https://doi.org/10.2478/popets-2021-0081`

4 ISO/IEC. *Information technology-Security techniques-Privacy framework*. International Organization for Standardization, Geneva, Switzerland, International Standard ISO/IEC 29100:2011(E), 2011.

5 ETSI. *A Guide to Writing World Class Standards*. `https://www.etsi.org/images/files/Brochures/AGuideToWritingWorldClassStandards.pdf`, accessed May 21, 2025.

6 S. Wilson, F. Schaub, A. A. Dara, F. Liu, S. Cherivirala, P. G. Leon, M. S. Andersen, S. Zimmeck, K. M. Sathyendra, N. C. Russell et al. *The creation and analysis of a website privacy policy corpus*. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 2016, pp. 1330–1340.

7 A. Schwartz. *Why P3P didn't work*. `https://cdt.org/wp-content/uploads/pdfs/P3P_Retro_Final_0.pdf`, accessed May 21, 2025.

4.2 Datasets, Automation, and Analysis Methods

Peter Story (Clark University – Worcester, US), Sepideh Ghanavati (University of Maine, US), Henry Hosseini (Universität Münster, DE & Westfälische Hochschule – Gelsenkirchen, DE), Jelena Mitrovic (Universität Passau, DE & Institute for Artificial Intelligence R&D of Serbia – Novi Sad, RS), Tim Samples (University of Georgia, US), Isabel Wagner (Universität Basel, CH), and Tianyang Zhao (Pennsylvania State University – University Park, US)

4.2.1 Challenges

Concerning datasets, automation, and analysis methods for privacy documents, we have identified four key challenges as well as four meta-challenges. These challenges deliberately focus on the landscape of privacy documents found today, instead of considering possible future improvements or standardization efforts. The rationale behind this focus is that the current state is likely to persist for at least several years, given the typical duration of legislative processes (e.g., the GDPR was first proposed in 2012 and came into effect in 2018) and the current lack of substantial legislative initiatives regarding privacy documents in the EU and elsewhere.

The primary stakeholders in this research area are fellow privacy researchers, as well as regulatory authorities, end users, and developers.

The first challenge is *continuously collecting privacy documents from many companies in many languages*. This is important for providing training data for classifiers, enabling the tracking of changes in policy documents, and ensuring greater inclusivity with respect to languages and jurisdictions.

The second challenge is about *scaling the automated annotation of privacy documents*. For example, regarding the number of annotations, effort, or standard labeling schemes, this approach aims to reduce manual effort and enable more large-scale annotations to train more capable classifiers.

The third challenge is developing *effective computational methods* to analyze privacy documents. This enables the building of tools for stakeholders, including summarization, locating opt-outs, detecting inconsistencies, assigning privacy grades, and identifying topical trends over time.

The fourth challenge is *detecting inconsistencies between software, privacy labels, and privacy and policy documents*. This is important to ensure compliance with laws and avoid misleading users about the privacy practices of systems.

In addition to the four challenges described above, we further identified four meta-challenges, which we describe next, that broadly apply to some or all of the challenges we identified in the area of datasets, automation, and analysis methods. Failing to address these meta-challenges risks limiting the longevity, applicability, impact, and reproducibility of research.

The *FAIR principles* [1] (findability, accessibility, interoperability, and reusability) apply primarily to collected datasets, such as corpora and annotations, but the spirit of the principles can also inform the development of solutions for privacy policy analysis, such as models and tools. Considering the FAIR principles can, for example, influence the choice of file formats. Plain-text formats, such as CSV, should be preferred over proprietary file formats, such as Excel, or formats that require infrastructure, such as an SQL server. The FAIR principles can also inform the instructions given to annotators (e.g., providing codebooks) and the detail and structure of documentation provided with datasets (e.g., including each coder's annotations in addition to the agreed-upon annotations).

Environmental impact and sustainability are rapidly becoming important considerations in science, especially in light of the resource consumption of new machine learning approaches such as large language models (LLMs) and the potential impact on climate change. Large-scale research on privacy documents could lead to substantial resource consumption. Resource consumption should be measured and reported from the start. When the deployment phase of systems is reached, the insights derived from this information can assist in evaluating the balance between task efficiency and the utilization of resources, guiding decisions on potential tradeoffs.

The *selection of downstream tasks* for analyzing privacy documents should be prioritized. Privacy documents contain a wide variety of information, and research has proposed methods to extract and analyze different aspects of this information. These downstream tasks range from broad classification of information categories in privacy policies to identifying narrow information items, such as opt-out statements, to analyzing the consistency of policy statements with data flows. While working on novel downstream tasks may be beneficial for academic indicators of success, it is also crucial to consider which downstream tasks are most beneficial for stakeholders, such as end users.

Although the top privacy publication venues, such as IEEE S&P and USENIX Security, are often novelty-driven, *maintainability and availability of research products* should be given more attention. Past research projects have generated numerous artifacts, such as tools and corpora, that could be useful and applicable to other research projects. However, research projects are often forced to duplicate or repeat previous efforts because tools and datasets are not maintained, unavailable, or insufficiently documented. The challenge is that maintaining research products is not well incentivized in the academic world and is therefore not a natural outcome of academic work compared to industry products. In the interest of open science and reproducibility, research projects should plan for long-term maintainability and availability of their research products from the start, e.g. by selecting repositories and platforms that ensure long-term availability (e.g., preferring Zenodo and Software Heritage over self-hosted websites) or by committing technician and/or student time to maintenance.

4.2.2 Key research challenges and questions

4.2.2.1 Challenge 1: Continuous multilingual large-scale collection, storage, and organization of privacy documents

- RQ1.1: How can we improve the automated collection, storage, and organization of privacy documents?
- RQ1.2: Which privacy document(s) apply to a system? Some documents reference other documents, sometimes from other companies. For example, a mobile app's privacy policy may reference the privacy policy of a third-party analytics library.
- RQ1.3: Which sections of a privacy document apply to a system? For example, a company that offers a wide range of products may write a single privacy policy that covers all of its products. Additionally, certain sections of a privacy document may only apply to children or vulnerable groups, or may only be applicable in specific jurisdictions.
- RQ1.4: How prevalent are non-textual elements in privacy documents? Privacy documents are often reduced to plain text for analysis, while this approach may result in the disordering of textual fragments and the loss of information for certain elements. Examples of such potentially problematic elements include tables, multi-layered policies, lists, etc.
- RQ1.5: Can we generate changelogs to highlight privacy document changes? Some companies visualize document changes using a "diff" (e.g., Google). It might be possible to generate diffs for the policies of other companies. In addition, it may be useful to go beyond a simple diff, perhaps highlighting or summarizing "interesting" changes.

4.2.2.2 Challenge 2: Scaling annotation of privacy documents

- RQ2.1: Can we use large language models (LLMs) to help scale the annotation of privacy policies, with quality comparable to human annotators?
- RQ2.2: How can we improve annotation practices for privacy documents? What are the pros and cons of various annotation tools? How can annotation tools be improved?
- RQ2.3: How can we create annotated corpora in multiple languages, which also meet the regulatory standards in various jurisdictions?

4.2.2.3 Challenge 3: Developing effective computational methods to analyze privacy documents

- RQ3.1: What computational methods offer the greatest performance for downstream tasks? Downstream tasks may involve extracting information from privacy documents, such as opt-out choice links, types of information collected and shared, data retention duration, and data deletion options.
- RQ3.2: Which computational methods balance task performance with other goals? An example of such balance could be reducing resource consumption while improving explainability.
- RQ3.3: What privacy-relevant content is present in documents other than designated privacy policies? For example, terms of service, cookie policies, cookie banners, community guidelines, and FAQs may contain information that is relevant to privacy.

4.2.2.4 Challenge 4: Detecting inconsistencies between software, privacy labels, and privacy documents

- RQ4.1: How can we determine privacy-related behaviors of software systems? These could include, among others, mobile applications, web applications, IoT devices, smart cities, or vehicles.
- RQ4.2: What methodologies can be employed to identify privacy-related behaviors within a given source code?
- RQ4.3: In what ways can we assist developers with limited resources in generating privacy documents, while ensuring alignment and consistency between the written text and the corresponding code?

4.2.3 Research directions

4.2.3.1 Addressing challenge 1

To address RQ1.1, while the toolchain by Hosseini et al. [2] provides a comprehensive solution to collect and preprocess privacy documents automatically on a large scale in 41 languages, the final stage of this toolchain, which uses trained classifiers to distinguish between privacy documents and non-privacy documents, is limited to English and German. Additional classifiers that cover the 39 other languages that the toolchain can collect and preprocess would improve inclusivity.

To address RQ1.2, open web indices, such as the OpenWebSearch.eu Open Web Index (OWI)[1], can be used to identify linking relationships to and from privacy documents [9]. As part of the preprocessing of a privacy document, any referenced privacy documents can be retrieved and inserted into the "main text" of the original privacy document.

[1] https://openwebsearch.eu/open-webindex/

To address RQ1.3, classifiers can be developed to identify privacy-relevant sections in related documents (e.g., terms of use). These sections could be included in the overall privacy analysis of the system.

To address RQ1.4, a measurement study could be conducted to determine the prevalence of problematic elements that hinder machine-readability of privacy documents (e.g., tables, multi-layered policies, enumeration). Data from this study could inform the development of tools for converting problematic elements to easier-to-parse plain text for classifiers.

For RQ1.5, while the GDPR mandates that users be informed about changes in privacy policies that alter the legal basis or purpose of processing, not all companies may comply with this requirement. A simple "diff" can be generated from the plain text versions of privacy policies. However, to make these results useful to users, it is necessary to develop automated methods to filter out insignificant changes (e.g., rephrasing) and to highlight the impactful changes.

4.2.3.2 Addressing challenge 2

To address RQ2.1, approaches that utilize LLMs in an active learning setting have been demonstrated to be effective in numerous NLP tasks [11]. LLM-in-the-loop approaches should be attempted for privacy documents if the format and annotation schema allow. When exploring approaches using LLMs, it is essential to ensure the reproducibility of the results. Sharing source code, prompts, and the checkpoint/version of LLMs is essential, as some LLMs are not open-source and do not always show the latest version (e.g., ChatGPT). It should be standard practice to report the resources used, including running time and computational resources such as GPUs.

Researchers should conduct a comprehensive study that compares the quality of such machine-assisted annotations with that of human-annotated corpora. The results may depend on the type of annotation schema used, so it would be important to study multiple annotation schemes.

To address RQ2.2, researchers can start by creating and maintaining an online resource (e.g., a wiki) describing annotation tools and their pros and cons. It may also be worth improving existing tools, such as using simple NLP methods to highlight negations in document text. An annotation SoK paper could also be published to highlight annotation best practices. For example, it is essential to develop a usable annotation scheme that other labs can apply with low error rates. One approach is to limit the annotation scheme to a single page to limit its complexity. A challenging aspect of annotation is handling disagreements between annotators [3]. Current customs include majority vote and union of labels. With human annotators, we can take the approach of hosting a meeting at the end to resolve disagreements.

Apart from extending the size of the current corpora, the need for regulatory-aware collections of annotated documents arises, especially in Europe, where the GDPR is applicable. For example, in a corpus annotated using a GDPR-based annotation schema [10], what are the relations between specific data processors and data controllers *across the entire corpus*? A possible approach to addressing this question could be the creation of a graph-based collection of documents that contains relations between paragraphs of privacy documents, as well as information on how different labels relate to one another.

RQ2.3 can be explored through large-scale targeted crawling campaigns. These could extend existing web-based corpora (e.g., the English-German corpus by Arora et al. [4]) and enhance already developed crawlers [2], as well as utilize already existing web indices to filter out more task-specific documents and create large, multilingual corpora of privacy documents

that would then still require labeling. Another approach could be to leverage open machine translation software. However, one needs to be aware of the different jurisdictions and create relevant labels for each language. Finally, it is essential to develop taxonomies that are applicable to various research questions and regulatory environments. This would ensure that datasets have lasting value in the face of evolving regulations.

4.2.3.3 Addressing challenge 3

To address RQ3.1, researchers should compare the relative task performance of different computational methods for analyzing privacy documents. For example, comparing the accuracy of text classification using LLMs with classical machine learning algorithms such as logistic regression (LR). It is also worth exploring hybrid techniques, such as symbolic NLP in combination with LLMs. To report task performance, researchers should report at a minimum the accuracy score, the F1 score, and the number of ground truth instances in each category. Other metrics, such as precision and recall, are also beneficial. Researchers can also perform an ablation study to determine which features are the most important.

To address RQ3.2, researchers should report task performance metrics (e.g., accuracy) alongside other quality metrics. Quality metrics may include computational resource consumption, licensing costs, and explainability. Depending on the downstream task, task performance might be more or less important than other quality metrics. For example, if LLMs and LR offer similar task performance, then LR might be preferable due to its lower resource consumption and greater explainability. The predictions of an LR model can be understood by examining the model coefficients. In contrast, if LLMs perform tasks substantially better than LR, LLMs might be chosen despite consuming more computational resources, having higher licensing costs, and offering limited explainability. Of course, the choice between models is context-specific and would depend on the downstream task. Another factor to consider is that the use of closed-weight LLMs (e.g., ChatGPT, Gemini, Claude) may limit the reproducibility of research. Closed-weight LLMs can be modified server-side without notice, or access could be completely revoked. In contrast, open-weight LLMs (e.g., Llama) can be archived, which supports reproducible research. We recommend that researchers perform tasks using open-weight LLMs, perhaps in conjunction with closed-weight LLMs. Access to model weights will also affect the deployment for downstream tasks.

To address RQ3.3, researchers should develop computational methods for documents other than privacy policies. Terms of service, cookie policies, cookie banners, community guidelines, and FAQs may all contain privacy-relevant information. Privacy policies often fail to address users' privacy-related questions. However, answers to users' questions might be found in other documents, such as privacy FAQs. Thus, tools to answer users' questions will be more effective if they can draw from a variety of privacy-related documents.

4.2.3.4 Addressing challenge 4

To address RQ4.1, several directions could be followed. For example, researchers could use crowdsourcing to collect data from users about applications, thereby inferring backend data-handling practices. In addition, reverse engineering techniques could be used to understand the behavior of closed-source applications [5]. Lastly, researchers could advocate for regulatory bodies to provide them with access to source code, thus enhancing the privacy and security of the general public. This is similar to how the EU AI Act envisions access to models, algorithms, and datasets.

To address RQ4.2, researchers should examine the source code of various applications to create a comprehensive taxonomy of privacy behaviors. Past research [6, 7, 8], for example, defined privacy behaviors in terms of four categories of practices (i.e., collection, sharing, processing, and others) and four categories of purposes (i.e., functionality, advertisement, analytics, and others). These categories are limited and do not encompass various cases of privacy behaviors beyond those related to permissions and access. Researchers could begin with the existing taxonomies for policy documents and extend them to source code. Using the expanded taxonomy, researchers should focus on creating ground truth datasets. Creating such datasets can be cumbersome, but with the advancement of LLMs, these models can be used in conjunction with human-in-the-loop (HitL) approaches to create more robust ground-truth datasets. These datasets also require evaluation. Potentially, independent developers could be recruited through crowdsourcing platforms to assess and improve the quality of the dataset.

To address RQ4.3, researchers should leverage, extend, and develop new static or dynamic analysis tools to identify and extract data flows from the source code. It is worth going beyond mobile applications and considering other software, such as backend code. Researchers must determine the most effective way to prepare and represent code for model training and inference. Approaches that focus on pruning and slicing source code to focus on privacy features should be explored. To identify inconsistencies between source code and privacy policies, researchers should focus on mapping and creating traceability between source code and policy documents or labels. Lastly, software engineering research has focused on code summarization and captioning for more than a decade. Researchers in the privacy domain could leverage or adopt some of the techniques from software engineering to automatically generate labels from policy texts and code, and do the translation.

References

1 M. D. Wilkinson et al., *The FAIR Guiding Principles for scientific data management and stewardship*. Sci Data, vol. 3, p. 160018, Mar. 2016, doi: 10.1038/sdata.2016.18.

2 H. Hosseini, C. Utz, M. Degeling, and T. Hupperich. *A Bilingual Longitudinal Analysis of Privacy Policies Measuring the Impacts of the GDPR and the CCPA/CPRA*. Proceedings on Privacy Enhancing Technologies, 2024(2):434—463, February 2024.

3 D. G. Gordon and T. D. Breaux, *The role of legal expertise in interpretation of legal requirements and definitions*. 2014 IEEE 22nd International Requirements Engineering Conference (RE), Karlskrona, Sweden, 2014, pp. 273-282, doi: 10.1109/RE.2014.6912269.

4 S. Arora, H. Hosseini, C. Utz, V. B. Kumar, T. Dhellemmes, A. Ravichander, P. Story, J. Mangat, R. Chen, M. Degeling, T. Norton, T. Hupperich, S. Wilson, and N. Sadeh. *A Tale of Two Regulatory Regimes: Creation and Analysis of a Bilingual Privacy Policy Corpus*. In Proceedings of the 13th Conference on Language Resources and Evaluation, LREC 2022, pages 5460–5472, Paris, France, 2022. ELRA.

5 S. Zimmeck, P. Story, D. Smullen, A. Ravichander, Z. Wang, J. Reidenberg, N. C. Russell, and N. Sadeh. *MAPS: Scaling Privacy Compliance Analysis to a Million Apps*. Proceedings on Privacy Enhancing Technologies, vol. 2019, no. 3, pp. 66–86, Jul. 2019, doi: 10.2478/popets-2019-0037.

6 V. Jain, S. Ghanavati, S. T. Peddinti, and C. McMillan. *Towards Fine-Grained Localization of Privacy Behaviors*. In Proceedings of the 8th IEEE European Symposium on Security and Privacy (Euro S&P'23), Delft, July 3-7, 2023.

7 V. Jain, S. D. Gupta, S. Ghanavati, S. T. Peddinti, and C. McMillan. *PAcT: Detecting and Classifying Privacy Behavior of Android Applications*. In Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '22), ACM, NY, USA, 104–118, 2022.

8 V. Jain, S.D. Gupta, S. Ghanavati, and S.T. Peddinti. *PriGen: Towards Automated Translation of Android Applications' Code to Privacy Captions*. 15th International Conference on Research Challenges in Information Science (RCIS2021), Cypress, 2021.

9 M. Granitzer, S. Voigt, N.A. Fathima, M. Golasowski, C. Guetl, T. Hecking, G. Hendriksen, D. Hiemstra, J. Martinovič, J. Mitrović, I. Mlakar, S. Moiras, A. Nussbaumer, P. öster, M. Potthast, M. Srdič Senčar, M. Sharikadze K. Slaninová, B. Stein, A. de Vries, V. Vondrák, A. Wagner, and S. Zerhoudi. *Impact and development of an Open Web Index for open web search*. Journal of the Association for Information Science and Technology, vol. 75, no.5, pp. 512-520, 2024, doi:10.1002/asi.24818.

10 H. Darji, J. Mitrović, and M. Granitzer. *German BERT Model for Legal Named Entity Recognition*. Proceedings of the 15th International Conference on Agents and Artificial Intelligence (ICAART) vol. 3, pp. 723-728, 2023. doi:10.5220/0011749400003393.

11 N. Kholodna, S. Julka, M. Khodadadi, M.N. Gumus, and M. Granitzer. *LLMs in the Loop: Leveraging Large Language Model Annotations for Active Learning in Low-Resource Languages*. In Machine Learning and Knowledge Discovery in Databases. Applied Data Science Track. ECML PKDD 2024. Lecture Notes in Computer Science, vol. 14950. pp. 397-412. doi:10.1007/978-3-031-70381-2_25.

4.3 Usable and Useful Notice & Consent

Simone Fischer-Hübner (Karlstad University, Chalmers University of Technology and Gothenburg University, SE), Kai-Wei Chang (UCLA, US), Nico Ebert (ZHAW – Winterthur, CH), Agnieszka Kitkowska (Jönköping University, SE), and Shidong Pan (Australian National University – Canberra, AU)

4.3.1 Challenges

4.3.1.1 Challenge 1: Limitations of human-centered and inclusive approaches to the design of notice and consent

Better stakeholder engagement. In the design of notice and consent, not only end-users but various other stakeholder groups that engage with privacy practices and consent mechanisms need to be actively included and involved in the development and design process. These can be lawyers, software engineers, and other stakeholder groups, such as journalists or any marginalized or vulnerable groups.

Better integration of diverse communication channels, formats, and media. A holistic and inclusive approach is needed that also considers a variety of communications and interaction channels and formats. Due to emerging technology (e.g., AI-based conversational agents, voice assistants, autonomous vehicles, and VRs that intrusively collect large amounts of more sensitive personal data), new ways for effective, usable and useful notice and consent are required (e.g., audio-assisted consent). With an increasing emphasis on inclusivity also additional design aspects are important, for instance, complementary easy-accessible audio formats of privacy information should be available for people with visual impairments. Similarly, the intersectional lens could be considered, where, in the given example, also easy to comprehend language is used in the audio-notice to ensure the inclusion of people with lower cognitive functions or migration backgrounds. For example, a privacy notice in audio

format might use straightforward terms such as "we use your email to send updates about our products" instead of complex legal jargon like "we process your contact information to disseminate product-related notifications".

Better personalization and contextualization of notice and consent (including language). There is a growing need for better personalization in notice and consent to reflect the individual demands and preferences of stakeholders. People from diverse backgrounds may rely on different terminologies or expressions. For instance, while a lawyer might understand a term like "third-party", it would likely be incomprehensible to a regular user. Adapting language and presentation to suit the broader audience is crucial for improving comprehension. Moreover, individuals typically may have different preferences regarding policy content that are of interest or relevance to them. To ensure improved personalization, there is a need for a stronger focus on user context to make privacy notice and consent more relevant and engaging for the different groups of users. While theories (e.g., contextual integrity) and new approaches that respect contexts have been developed (e.g., contextual privacy notices [16], just-in-time notices [19]), further research is needed, in particular considering the contexts within the emerging technologies (e.g., VR, Metaverse, or smart technologies).

In summary, there is a need for a human-centered and inclusive design approach that considers a variety of stakeholders and their preferences, channels/formats and personalization.

4.3.1.2 Challenge 2: A lack of a risk-oriented approach to the design of notice and consent

What is presented in a notice and consent process typically refers, in general terms, to all possible data processing practices, data controllers, data processors, and other third parties ("procedural transparency"). Users might be confronted with too much, too irrelevant, or deceiving information. At the same time, potential risks of data processing might not be transparent and clear to them. As dealing with privacy information is often only a secondary task for users, a focus on privacy risks could help to make communication more effective ("risk transparency"). A risk-based approach to communication has been proven effective in other areas, such as safety, and indications for their effectiveness in privacy exist, too [1]. It can be seen as complementary to traditional approaches that want to give a complete picture of data processing practices and will remain relevant for specific stakeholder groups (e.g., privacy-interested users or professionals).

A risk-oriented approach would, however, require identifying what privacy risks exist and how they can be categorized. It would also require new tools (e.g., based on AI) that help to quickly predict risks and allow dynamic risk communications (e.g., based on changes in the privacy policy that introduce new or unexpected risks and in the context of dynamic consent).

Implementing a risk-oriented approach in privacy notices could foster end-user trust in the respective organization, if it is also clearly communicated if and how risks are adequately mitigated. A challenge is communicating residual risks in relation to the benefits of disclosing data, e.g., when using a service. This could not only provide increased transparency for users but also for controllers, who could see benefits in a transparent risk-based approach that can also create trust and, therefore, would support it.

4.3.1.3 Challenge 3: A lack of standardized methods to evaluate the usability and usefulness of notice and consent

Although researchers typically evaluate new design artifacts, it is difficult to compare the results of different studies with regard to general criteria. Standardized design criteria and evaluation methods could help to assess the benefits of artifacts. A concrete example is a validated user-survey instrument comparable to the system usability scale [11] that would allow the community to compare different artifacts (e.g., nutrition labels vs. short privacy notices with complementing privacy icons). There are also no standardized methods for the evaluation across contexts or channels/formats (e.g., privacy information presented in a car vs. on a mobile phone). Standardized methods could range from general criteria for "good" design of notice and consent to validated instruments to survey users' perceptions of the usability and usefulness of notice and consent.

Another issue exacerbating the problem of the lack of standardized evaluation methods in the space of design of notice and consent is lack of reproducibility (when the measurement can be obtained with stated precision by a different team, a different measuring system, or in a different location on multiple trials), replicability (when the measurement can be obtained with stated precision by a different team using the same measurement procedure, the same measuring system, under the same operating conditions, in the same or a different location on multiple trials) and repeatability (when the measurement can be obtained with stated precision by the same team using the same measurement procedure, the same measuring system, under the same operating conditions, in the same location on multiple trials) of research. Not only are such studies scarce within the field, but also difficulties related to how such research is defined should be addressed, considering contextual factors surrounding privacy (e.g., temporality). For example, while some researchers may claim the effectiveness of newly proposed notice and consent artifacts, other researchers may have no incentives, capability, or means to reproduce the results.

Also, structural issues (e.g., lack of outlets accepting replication studies) may prevent replication. Furthermore, the introduction of AI components adds another layer of complexity, as the inherent uncertainty of AI-driven systems can make reproducibility even more challenging. Replication, reproduction, and repetition, however, are necessary to create an actual body of knowledge in the discipline. Adding to this challenge are the often little interest and encouragement from publication venues towards publishing replication studies, as well as requiring researchers to make the datasets, artifacts, and other research-relevant materials publicly available.

4.3.2 Key research questions

The overarching goal to achieve in the next decade is to define guidelines enabling the holistic human-centered approach to the design of notice and consent. Specifically, we have to address the following:

- RQ1: How should a holistic human-centered approach to the design of consent and notice look like – and how can it address key design challenges?
- RQ2: How can privacy risks be identified and communicated in a way that benefits users and other stakeholders and fosters reliable trust?
- RQ3: How can we standardize methods for the evaluation of usability and usefulness of privacy policies across research?

Table 1 draws connections between the identified challenges, research questions, and suggested approaches.

4.3.3 Research Directions

In order to develop solutions we could adapt approaches from related areas to our problem domain.

4.3.3.1 Approach 1: Human-centered and inclusive privacy by design and default

Usability is an important prerequisite for privacy by design and by default, and at the same time privacy by design and by default provides means for enhancing usability. Moreover, solutions for human-centric privacy by design and by default have to be inclusive and adapt to the needs and values of different types of users and other relevant stakeholders. This approach helps to balance different design requirements in the design space for privacy notices [19]. Hence, new methodologies and approaches for privacy by design (for privacy notices and consent solutions) based on and combined with human-centric and inclusive design approaches should be researched and developed [8].

4.3.3.2 Approach 2: Risk perception and risk communication

Risk research has a long tradition of studying risk perception and risk communication in the non-digital world that could also be applied in traditionally digital domains [22, 21]. Also in digital domains, for example, "saliency" of privacy risks plays a key role in improving risk perception [6] and promoting protective user behavior [7]. Other ideas from risk research have not yet found their way into privacy research but might be beneficial in developing risk-oriented approaches in the domain of notice and consent. For example, risk researchers hypothesized that risk mitigation measures may lead to an increased risk acceptance level of individuals ("risk compensation" [9]). For instance, explanations of privacy-enhancing technologies (PETs) could introduce an unexpected level of personal data sharing on the side of the end-user when all risks are assumed to be mitigated. The opposite effect could occur if the core protection functionality of PETs is misunderstood. Further, different theoretical lenses could be applied to identification of privacy risks, particularly among marginalized populations. One approach could be critical feminist frameworks, such as intersectionality, that have been used in the field of human-computer interaction (HCI) [20].

4.3.3.3 Approach 3: AI support and tools

Artificial intelligence (e.g., LLMs) may provide means that can be beneficial in all phases of the design process, from creation to evaluation of artifacts, but also as a functional part of proposed design artifacts. In the design phase, an LLM can be used to simulate different personas with different privacy preferences or groups of people to quickly evaluate privacy design in an early stage. In the consent interpretation stage, recent years have seen the development of machine learning-based Personalized Privacy Assistants (PPAs) that can, based on an analysis of the users' previous privacy decisions (e.g., related to setting or rejecting privacy permissions for Android or IoT systems [12, 2, 23, 3]), predict the users' preferred choices and subsequently assist users in privacy decision-making with suitable recommendations. Nonetheless, PPAs can only help to semi-automate privacy decisions like consent, which according to the GDPR requires an affirmative action by the user and thus cannot be fully automated [15]. Still, in the future, PPAs could be developed more broadly for other technical areas (e.g., for IoT Trigger Action Platforms, cloud environments) and also assist users, e.g., with extracting or highlighting core information to meet their personal interests in addition to the personalized recommended decisions. In the consent

Table 1 Research Questions, Challenges, and Approaches.

Research Question	RQ/Challenge	Approach
What immediate research questions need to be answered?		
How can relevant stakeholders be identified and engaged in the phases of requirement elicitation and the generation, design, and evaluation of privacy notices and consent?	RQ1/C1	A1
What are more effective user interactions with policy content for raising the user's attention and awareness of the relevant policy information, esp. in the context of consent (e.g., via interactive voice communication)?	RQ1/C1	A1
What are relevant risks from a user perspective, and how can risk communication be best personalized with the help of AI tools?	RQ2/C2	A2
What tools do developers need to be better supported in the development of privacy notices and consent?	RQ3/C3	A4
What best practices can be proposed for replication studies on usable privacy related to privacy notice and consent?	RQ3/C3	A4
What are the challenges and questions to be answered in the next three years?		
How can we apply AI to facilitate the design and produce personalized, useful privacy notices for users?	RQ1/C1	A3
What are usable and inclusive forms of policy communication utilizing multiple channels?	RQ1/C1	A1
How can transparency about risks and perceived consequences be achieved without unnecessarily scaring users?	RQ2/C2	A2
How can risks, risk mitigation and the residual risks be made more transparent?	RQ2/C2	A2
How can personal or societal benefits achieved from disclosing personal data compared to the residual risk be communicated?	RQ2/C2	A2
How can risks be dynamically detected and used to inform the users and obtain dynamic consent?	RQ2/C2	A3
How effective is a risk-based approach to privacy communication compared to traditional approaches?	RQ2/C2	A3
What are challenges and questions to be answered within the next decade?		
What are the criteria for the evaluation of the usability and usefulness of privacy policy notice and consent?	RQ3/C3	A4
What "evaluation checklist" and tool support for privacy notice and consent should be provided to developers?	RQ3/C3	A4

informing stage, AI-based tools or PPAs could be used to create personalized and dynamic privacy information that adapts to users' context and information needs [12]. For instance, if a weather app starts to use location data not only for the purpose of showing the local weather but for location-based advertising or sharing data with location data brokers, LLMs can easily generate a new consent prompt to describe the data practice, and users could be notified and requested to consent dynamically. At the same time, such mechanisms need to be designed to be privacy- and values-preserving, legally compliant, and ethical [15, 18, 17]. This is particularly important given the widespread criticism of AI techniques for their frequent hallucinations (in the case of LLMs) and lack of transparency.

4.3.3.4 Approach 4: Systematic reviews, replication studies, and mega-studies

To enable the creation of standardized evaluation criteria that could be applied in the research on privacy notice and consent design, there is a need for the production of more empirical evidence within the field. However, to identify what empirical investigations are more urgent in the fast-changing technology landscape, more systematic reviews and meta-analysis studies are needed. Only based on such studies can researchers pursue the right research problems.

Moreover, considering the evaluation criteria and assessment of the effectiveness of design, the field could pursue methods used in the social sciences, such as mega-studies [4]. Such studies were shown to be successful in making behavioral science findings more successful in their applicability, particularly in the context of behavior change through design in the context of nudging [5, 10, 13, 14]. Mega-studies themselves, although challenging to conduct, could also serve as quantitative evaluation criteria.

In order to develop solutions, we would also require additional skills and collaboration with other fields. Research on a holistic human-centered approach to the design of privacy notice and consent will require interdisciplinary expertise and the cooperation of a broad range of stakeholders. Most importantly, the research needs to include experts from the fields of computer science, information systems, social science, economics, psychology, risk research, safety, and law and should also be based on requirements collected from stakeholders elicited from end users with different demographic backgrounds, representatives from organizations including management and DPOs, as well as regulators.

References

1 Almuhimedi, Hazim, et al. "Your location has been shared 5,398 times! A field study on mobile app privacy nudging." Proceedings of the 33rd annual ACM conference on human factors in computing systems. 2015.
2 Bahirat, Paritosh, et al. "A Data-Driven Approach to Developing IoT Privacy-Setting Interfaces", 23rd International Conference on Intelligent User Interfaces, pp. 165-176, March 2018, [online] Available: https://dl.acm.org/doi/10.1145/3172944.3172982.
3 Das, Anupam, et al. "Personalized privacy assistants for the internet of things: Providing users with notice and choice." IEEE Pervasive Computing 17.3 (2018): 35-46.
4 Duckworth, Angela L., and Katherine L. Milkman. "A guide to megastudies." PNAS nexus 1.5 (2022): pgac214.
5 Duckworth, Angela L., et al. "A national megastudy shows that email nudges to elementary school teachers boost student math achievement, particularly when personalized." Proceedings of the National Academy of Sciences 122.13 (2025): e2418616122.
6 Ebert, Nico, et al. "Bolder is better: Raising user awareness through salient and concise privacy notices." Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. 2021: 1–12.

7 Ebert, Nico, et al. "When information security depends on font size: how the saliency of warnings affects protection behavior." Journal of Risk Research 26.3 (2022): 233–255.

8 Fischer-Hübner, Simone, and Karegar, Farzaneh. "Addressing Challenges: A Way Forward." The Curious Case of Usable Privacy: Challenges, Solutions, and Prospects. Cham: Springer International Publishing, 2024. 133-160.

9 Hedlund, James. "Risky business: safety regulations, risk compensation, and individual behavior." Injury prevention 6.2 (2000): 82-89.

10 Kuan, Robert, et al. "Behavioral nudges prevent loan delinquencies at scale: A 13-million-person field experiment." Proceedings of the National Academy of Sciences 122.4 (2025): e2416708122.

11 Lewis, James R. "The system usability scale: past, present, and future." International Journal of Human–Computer Interaction 34.7 (2018): 577-590.

12 Liu, Bin, et al. "Follow my recommendations: A personalized privacy assistant for mobile app permissions." Twelfth Symposium on Usable Privacy and Security (SOUPS 2016). 2016.

13 Milkman, Katherine L., et al. "Megastudies improve the impact of applied behavioural science." Nature 600.7889 (2021): 478-483.

14 Milkman, Katherine L., et al. "A megastudy of text-based nudges encouraging patients to get vaccinated at an upcoming doctor's appointment." Proceedings of the National Academy of Sciences 118.20 (2021): e2101165118.

15 Morel, Victor, and Fischer-Hübner, Simone. "Automating privacy decisions-where to draw the line?" 2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW). IEEE, 2023.

16 Pan, Shidong, et al. "A NEW HOPE: Contextual Privacy Policies for Mobile Applications and An Approach Toward Automated Generation" 33rd USENIX Security Symposium (USENIX Security 24). 2024.

17 Pan, Shidong, et al. "Is It a Trap? A Large-scale Empirical Study And Comprehensive Assessment of Online Automated Privacy Policy Generators for Mobile Apps." 33rd USENIX Security Symposium (USENIX Security 24). 2024.

18 Morel, Victor, et al. "AI-driven Personalized Privacy Assistants: a Systematic Literature Review." (2025), `https://arxiv.org/abs/2502.07693`

19 Schaub, Florian, et al. "A design space for effective privacy notices." Eleventh Symposium on Usable Privacy and Security (SOUPS 2015).

20 Schlesinger, Ari, et al. "Intersectional HCI: Engaging identity through gender, race, and class." In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (pp. 5412-5427).

21 Siegrist, Michael, & Árvai, Joseph (2020). Risk perception: Reflections on 40 years of research. Risk Analysis 40(S1) (2020), pp. 2191-2206.

22 Slovic, Paul, et al. Why study risk perception?. Risk Analysis, 2(2) (1982), pp. 83-93.

23 Smullen, Daniel, et al. "The Best of Both Worlds: Mitigating Tradeoffs Between Accuracy and User Burden in Capturing Mobile App Privacy", Proceedings on Privacy Enhancing Technologies, 2020(1), pp. 195-215.

4.4 Consumer Privacy Beyond Notice and Choice

Noah Apthorpe (Colgate University – Hamilton, US), Eleanor Birrell (Pomona College – Claremont, US), Travis Breaux (Carnegie Mellon University – Pittsburgh, US), Kirsten Martin (Carnegie Mellon University – Pittsburgh, US), Rishab Nithyanand (University of Iowa – Iowa City, US), Sarah Radway (Harvard University – Allston, US), Yan Shvartzshnaider (York University – Toronto, CA), and Maximiliane Windl (LMU München, DE)

4.4.1 Challenges

The notice and choice regime that appears in US and EU privacy law has dominated how online privacy is regulated. Websites, apps, IoT devices, and other technologies post privacy policies describing (to some degree) data practices, and people are expected to choose services that meet their privacy needs.

However, decades of research have consistently shown that these documents are long, vague, and rarely read, thereby undermining the premise that consumer choice is informed. Data flows – what data is collected, how that data is used and shared, what inferences are created from that data – are complicated, and vulnerabilities emanating from these data flows are difficult to identify. In addition, consumers are often not provided an authentic choice due to the market power of an organization (e.g., web search), opt-out as the default choice (e.g., behavioral advertising), dark patterns aimed at minimizing opt-out and otherwise influencing decision-making, or consumer choices simply being ignored.

Problems with notice and choice are not limited to implementation, but are inherent in the notice and choice regime. Privacy notices serve multiple purposes – providing legal disclosures, serving as a legally-enforceable data-use contract, and providing transparency to users – so there is no notice length or level of precision in the notice that meets the needs of each of these purposes. Moreover, the time required to read these documents does not scale to the number of services with which users interact. In addition, the choice to opt-in by one data subject can impact other subjects, e.g., when the choice to disclose covers personal information from more than one subject. Finally, notice and choice inherently places all responsibility on the user to understand data practices and make an informed decision. Yet the ever-changing data flows and nuanced implications of data practices render consumers poorly-positioned to ensure organizations' data practices meet the needs of individuals or society.

User choice is not the only possible mechanism for determining whether organizations' data practices are responsible or harm individuals or society. Other industries rely on regulations to provide minimum standards as well as reporting and auditing requirements that serve to hold organizations accountable for their business practices, provide standards of appropriate behavior, and enforce societal requirements and needs on businesses. Businesses and markets retain legitimacy not only through ensuring consumers are informed and make authentic choices but also through the work of auditors, regulators, industry groups, etc. to hold organizations accountable to societal standards.

Here, we propose a new framework that describes how responsible data practices could be enforced beyond the notice and choice regime. We first identify the features necessary in a world beyond the notice and choice regime. We then propose how current and future market and regulatory incentives and stakeholder roles can help hold organizations accountable for their data practices. This includes roles internal to the organization, such as the chief

executive officer, privacy risk and compliance officers, and software engineers, among others, as well as new roles external to the organization, such as auditors and insurers. These roles work together in a "web" to help organizations be accountable for their data practices to regulators. We then outline a possible approach to define privacy standards to guide organizational data practices, potential incentives for organizations to adopt the proposed framework, as well as potential obstacles to framework implementation and how those obstacles may be avoided or overcome. Finally, we present a roadmap for how the framework implementation – beyond notice and choice – could be achieved and what documents would be necessary to support this framework.

4.4.2 Features of an effective privacy protection paradigm

In order to move beyond notice and choice, it is important to articulate the key features of an effective privacy protection paradigm – features that are supposedly being achieved through the notice and choice regime. In this section, we identify four such features – (1) transparency, (2) responsible data practices, (3) minimization of regulatory burden, and (4) legitimacy of the data market – and we outline how the current framework of the notice and choice regime falls short of these requirements. Table 2 provides an overview.

Table 2 Comparison of Privacy Protection Principles: Traditional vs. New Paradigm.

Principles of Privacy Protection	Notice and Choice	New Paradigm
Transparency	NO	YES
Responsible data practices	NO	YES
Minimization of regulatory burden	YES and NO	YES and NO
Legitimacy of the data market	NO	YES

4.4.2.1 Transparency

The current notice and choice regime aims to provide transparency over organizational practices through privacy notices. To this end, organizations share information about their data practices in a way that is supposed to be accessible and understandable to all stakeholders. To meet this standard, various stakeholders require different information presented in different ways. For the data subjects, or individuals about whom data is collected or used, we observe that transparency efforts should focus on conveying information to users in a manner that is clearly written, not overly technical, and easy to read in a reasonably short amount of time. For example, the General Data Protection Regulation (GDPR) imposes similar requirements on notices as described in Article 12. However, there are few metrics to measure whether notices meet these requirements. Visualization tools or interactive interfaces may be helpful. For regulators, technical implementation details about data practices are necessary to evaluate compliance with legal requirements.

The current notice and choice framework provides the same privacy notices to both users and regulators, leaving both parties dissatisfied. Organizations are incentivized to create long, opaque or vague, and overly broad privacy documents to limit liability; as a result, these documents are hard for users to understand and the effort necessary to read the policies does not scale to the number of services with which a user interacts. An effective

privacy protection framework would need to provide transparency of an organization's data practices by ensuring documents are (1) comprehensive, (2) specifically tailored to the needs of different stakeholders, and (3) accurate representations of actual practices.

4.4.2.2 Responsible data practices

Organizational data practices should be consistent with laws and standards, which often encode societal values. Notice and choice regimes place responsibility on the consumer to recognize risks to their personal privacy and to be the sole decision maker about whether that risk is acceptable. Consumer choice acts as the sole force to ensure organizations' data practices are responsible and meet the needs and values of individuals and society. However, the shortcomings of notice and choice – including unreadable policies, incomprehensible implications of data practices, unscalable user burden, and limited choices – ensure that, in practice, users cannot always reasonably choose to use only services with responsible data practices. This often results in adoption of services that are inconsistent with user values, such as those that repurpose user data. To be effective, a privacy protection framework must employ incentives that go beyond the status quo. This includes new actor roles to hold organizations accountable and ensure responsible data practices, while incorporating robust enforcement mechanisms to penalize practices that violate legal standards or social values without relying only on consumer choice.

4.4.2.3 Minimization of regulatory burden

For a regulatory approach to be effective, enforcement mechanisms must be practical and enforceable. Regulators have finite resources that impose practical bounds on their actions; if misbehavior is less likely to be detected, organizations have less incentive to comply. The notice and choice regime minimizes regulatory burden by placing the primary responsibility for ensuring responsible data practices on the data subjects making the "correct" choice, which is criticized as increasingly being an "illusion of choice." On the other hand, the lack of transparency in current notices places a significant burden on regulators, who must investigate current practices without access to detailed internal information or technical details. Rather than rely on regulators to conduct random investigations, a new framework that requires organizations to disclose risks to privacy and non-compliance can guide regulators toward which organizations are likely to be non-compliant. The consequence of such disclosures can further motivate organizations to be more responsible. An effective privacy protection framework should provide incentives, structure, and access that enable effective external regulation without imposing undue regulatory burden.

4.4.2.4 Legitimacy of the data market

A successful privacy framework should engender trust in the marketplace, ensuring that any data transaction is made without fraud, manipulation, or deception. In the notice and choice regime, fully-informed choice presumes accurate information being received by data subjects and ensures individual autonomy, thereby legitimizing transactions over personal data. The broader market's legitimacy is thus determined through the summation of the legitimacy of these transactions. However, this regime assumes data subjects have the time, technical knowledge, and understanding of the broader technology ecosystem to make sense of privacy notices. Research has consistently shown that these requirements are not met, rendering this regime ineffective. An effective privacy protection paradigm should ensure legitimacy of the data market through a multi-faceted approach.

4.4.3 Research directions

In the light of these challenges and key features of an effective privacy protection paradigm, we identify directions for future research, shown in Table 3.

Table 3 Research directions.

Topic	Immediate	Medium-term (3 yrs)	Long-term (10 yrs)
Laying out the privacy protection paradigm *(Econ, Policy, Law, CS, Business)*	Identifying features and expectations from an effective privacy protection paradigm. Enumerating the different market forces, stakeholder roles and responsibilities that may be operationalized to motivate responsible data markets and compliance within them.	Providing specifications and examples of documents needed to support the paradigm. *(HCI, Business, Law)* Drafting regulation in support of the new paradigm.	
Defining responsible data practices *(Econ, Policy, Law, CS, Business)*	Developing models and frameworks for defining responsible data practices. How do we evaluate and assess which data practices are consistent with societal values in different contexts?	Develop mechanisms for assessing fines and incentives within the proposed frameworks. Developing mechanisms to evaluate the effectiveness of different frameworks for responsible data practices.	Developing mechanisms for transitioning frameworks to practice and assessing their performance. How do we assess harms and gains made from violating proposed frameworks?
Transparency needs and responsibilities *(Econ, Policy, Law, CS, Business)*	What are the transparency needs of different stakeholders? Does existing HCI research already suggest the best ways to address these needs? What documents are needed to provide effective transparency, and what should these documents look like?	Additional HCI research to address dark patterns in transparency documents and other transparency gaps/challenges.	
Professional obligations *(Econ, Policy, Law, CS, Business)*	Identify the professional obligations of a licensed CS engineer. Design document formats that professionalized engineers can use to communicate technical details and risks to internal privacy risk officers.	Design the conditions under which a licensed computer scientist is required.	Developing licensing bodies and mechanisms.
Whistleblower Employees *(Policy, Law, Business)*	Research when whistleblowers are needed based on regulatory and reporting requirements.	Identify the conditions and protections for whistleblowers in tech.	

Topic	Immediate	Medium-term (3 yrs)	Long-term (10 yrs)
Audits *(Econ, Policy, Law, CS, Business)*	Design auditing requirements for government contracting and special industries. Design documents that outline specific auditing requirements.	Expand the conditions requiring data audits including for privacy risk reporting for SEC.	Expand conditions requiring data audits to include organizations with individualized data.
SEC/Shareholder reporting *(Econ, Policy, Law, Business)*	Increase risk of high fines for privacy violations through state and federal regulations/policies. Design document formats for communicating organization behavior to regulators to enable meaningful compliance auditing.	Design reporting obligations for privacy risk assessments for publicly traded companies to SEC. Design the role of a privacy risk officer with obligations to report privacy risk based on potential fines and organizations' data practices.	
Insurance *(Econ, Policy, Law, Business)*	Increase risk of high fines for privacy violations (see above).	Design insurance market for organizations wishing to mitigate privacy risk.	

4.5 Cross-Stakeholder Interaction

Jose M. del Alamo (Polytechnic University of Madrid, ES), Soheil Human (Wirtschaftsuniversität Wien, AT), Konrad Kollnig (Maastricht University, NL), Daniel Smullen (CableLabs – Louisville, US), and Kami Vaniea (University of Waterloo, CA)

4.5.1 Challenges

Without noticing it much, we are all subject to a wealth of different privacy documents every day. These include privacy policies, FAQs, and the text shown in consent pop-ups of the various online services that we all use. The breadth of document types is impressive, encompassing terms of service, cookie policies, privacy disclosures, engineering specifications, and more. These documents also include legal statutes, implementing and delegated acts, and regulatory guidelines that underpin data processing. Often, laws from multiple countries apply to a single data processing operation of the same company, given the global scale of the Internet.

Instead of end-users (i.e., the public writ large), legal professionals – foremost lawyers and judges – are those who are currently most empowered by the broad share of currently used privacy document types. Perhaps the most common type of privacy document, privacy policies, in theory are to inform end-users about how their data is used. However, these documents are drafted mainly for lawyers and judges, thereby leaving end-users struggling to understand legal language in order to understand what happens with their data. Similarly, governments are those who enact privacy laws but legal documents are written in a language specific to law, which can be challenging for engineers to understand in terms of computer code and/or system implementation needs. Similarly, these documents are extremely difficult for end-users to interpret, even though they are the data subjects that the documents are intended to enshrine rights to.

While the list of stakeholders has yet to be fully enumerated, and identifying the broad set of documents that relate to them is an open research problem, we can already conclude some basic facts about all privacy documents. Among all stakeholders, from legal to end-user to engineer and beyond, the main purpose of privacy documents is the communication of privacy concepts from one stakeholder group to another. This underscores that the study of privacy documents is the study of communication between pairs of stakeholder groups, whose individual and collective needs must be considered when trying to make any such communication work. In practice, privacy documents often serve multiple stakeholder groups, with each of their unique needs. Too often, privacy documents are designed around a narrow group of stakeholders but used in a way that requires stakeholders outside of this scope to consume them.

There are many other stakeholders beyond data subjects and legal professionals, such as engineers in data-processing companies or competitors, members of civil society organizations, and parents of children who are subject to data processing. These, too, face the general challenge that they are not empowered by the privacy documents that other stakeholders use to communicate with them. The consequences are numerous; stakeholders are inadequately informed by the privacy documents they consume. The documents are at the wrong level of abstraction for their needs. The documents are intended to be understood in a way that may be familiar to the producer but is unfamiliar to the document consumer. In general, the misalignment of documents with their stakeholders exhibits many problems related to their application and suitability for use in numerous real-world scenarios that deserve further study. As we see in our previously cited examples, such as lawyers communicating with end-users (via privacy policies), there is an obvious but poorly understood mismatch in the communication these documents facilitate between the various pairs of stakeholders (i.e., document producers and consumers) and their intent. The result are a lack of transparency, which results in information asymmetry, and all the downstream problems that arise from under- or over-specification and misunderstanding. We view the overarching challenge in this space as one that arises from a lack of shared expertise between stakeholder groups. Without an interdisciplinary effort to improve communication between these groups, each stakeholder runs the risk of ineffective communication.

Figure 1 provides an abstract vignette, notionally illustrating how privacy documents serve as a means of communication between a few pairs of imaginary stakeholders. This illustration also captures the common scenario we see in the real world – that the same document is often used to facilitate the communication between different stakeholders in a way which is well-intentioned but falls short of achieving its intent. When companies use the same documents (i.e., privacy policies) to communicate with both regulators and users, the mismatch between documents intended for companies and regulators to communicate results in a communication failure.

Better communication arises from interdisciplinary work, such as where experts in law and technology work together to more holistically study the problems in this space and develop better solutions to address them. These solutions often revolve around the ideal state of having prescribed formats, structures, or schemata for communicating privacy information and concepts. Prominent examples are the Platform for Privacy Preferences Project (P3P) and Privacy Nutrition Labels, in theory enabling rapid or even automated communication. However, these solutions have rarely been adopted into practice or sustained so that they would meaningfully empower the involved stakeholders – end-users in particular. This situation reveals a key challenge: Too often, one stakeholder group (e.g., computer science academia, advertising industry, privacy regulators, browser developers) comes up

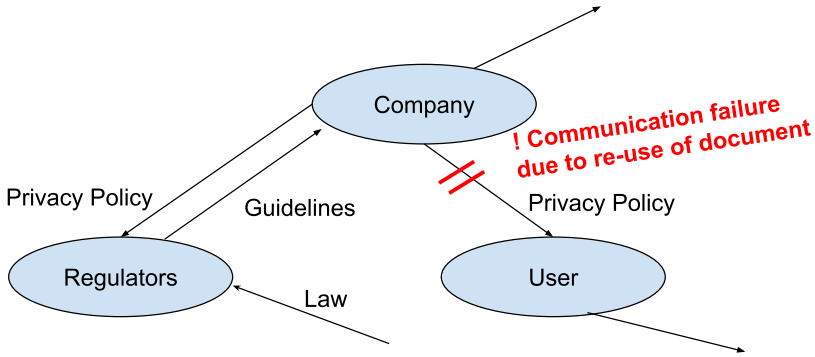

Figure 1 Notional example of how privacy documents serve as a means for communication between a few different imaginary stakeholders. Ineffective communication is common.

with solutions that generalize well within their own stakeholder group or a narrow slice of other groups but do not adequately address the needs of all stakeholders involved. Thus, work towards the adoption of those solutions by others falls short of meeting that goal. They can fail to take into account more pluralistic approaches and miss opportunities for better tailored approaches intended for specific stakeholder pairs.

Addressing a wide range of stakeholders requires interdisciplinary research, which everyone is in favor of in principle. But in practice incentives are missing: access to funding, venues for publication, means for long-term development of tooling, recognition in academic promotion, and so on. As a result, research ends up focusing on narrow parts of the inter-stakeholder relationships, where it aligns with the existing incentives.

In sum, we identify three main challenges:

- Stakeholder imbalance in needs around privacy documents in relation to motivation, information needs, empowerment, specificity, and clarity.
- Researcher needs to conduct interdisciplinary research that aims to connect several disciplines to understand problems and find solutions.
- Robust implementation strategies for how solutions will be developed, adopted, and enforced, as well as an understanding of how the strategies might be resisted.

4.5.2 Key research questions

- What mechanisms can be developed to ensure that theoretical research on privacy documents is translated into practical, industry-wide applications?
- How can research outputs in the design and analysis of privacy documents be made more accessible and relevant to policymakers and industry practitioners?
- What are the barriers to adopting academic innovations in privacy policy tools and practices within organizational settings?

4.5.3 Research directions

Addressing these challenges requires a multifaceted strategy. It is essential to enhance the dialogue between researchers and practitioners to ensure that research outputs are designed with practical constraints and needs in mind. Developing flexible, scalable solutions that can be easily integrated into existing systems and processes is crucial. Furthermore, fostering a regulatory environment that supports and promotes the enforcement of these novel solutions can provide the necessary impetus for their broader adoption.

4.5.3.1 Mapping the ecosystem of privacy management

A comprehensive understanding of the complex network of stakeholders, their interrelationships, and the channels through which communication and enforcement occur is vital. This involves mapping out the ecosystem of privacy management, identifying how different stakeholders interact, and recognizing the influence they exert on each other. Such an understanding can help in tailoring solutions that are not only technically sound and legally compliant but also socially and organizationally feasible.

4.5.3.2 Motivating privacy for stakeholders

Educating all stakeholders about the potential benefits and long-term gains of implementing advanced privacy solutions will be key to overcoming resistance and achieving widespread acceptance. This holistic approach is paramount to successfully translating privacy research into effective, enduring practices.

4.5.3.3 Communicating challenges in domain-specific language

Existing research already aims to understand the needs of different stakeholders and stakeholder pairs. Extensive research in the social sciences, for example, looks at the needs of vulnerable groups. Similarly, mapping the different actors in privacy enforcement is likely studied by many researchers. The challenge is that this work is currently presented in ways that are appropriate for the domains it was conducted in. More effort is needed to enable cross-domain presentation of research in ways that are accessible and help researchers better understand the challenges. In particular, survey or summarization type research would be quite valuable to provide alternative presentations of that are meant to be consumed by other disciplines.

4.5.3.4 Researcher incentivization and support

Addressing the lack of incentives and support for researchers in interdisciplinary fields requires structural and cultural changes within academic and research institutions, as well as funding bodies. Some solutions and ideas to move towards that end are:

- *Promote funding.* Both the National Science Foundation (NSF) and Horizon Europe (HEU) have initiatives to fund interdisciplinary research and promote collaboration, yet more targeted (or open-ended) programs can encourage collaboration across disciplines. In Europe, COST Actions and HEU Marie Curie networks can be helpful instruments to that end. The former can be used to create professional societies/chapters or online platforms to connect interdisciplinary researchers. The latter supports developing programs that train students and early-career researchers in methods and theories from multiple disciplines.
- *Revise academic incentive structures.* The metrics for academic advancement should include recognition of interdisciplinary work, e.g., by considering DORA-like criteria in assessing scientific quality and promotion. Identifying and sharing (or setting up if not available) reputed venues for interdisciplinary publications and/or cross-disciplinary outreach will enable this recognition by peers, further supporting networking and collaboration.

Participants

- Noah Apthorpe
Colgate University – Hamilton, US
- Eleanor Birrell
Pomona College – Claremont, US
- Travis Breaux
Carnegie Mellon University – Pittsburgh, US
- Kai-Wei Chang
UCLA, US
- Jose M. del Alamo
Polytechnic University of Madrid, ES
- Rinku Dewri
University of Denver, US
- Nico Ebert
ZHAW – Winterthur, CH
- Simone Fischer-Hübner
Karlstad University, SE
- Sepideh Ghanavati
University of Maine, US
- Henry Hosseini
Universität Münster, DE & Westfälische Hochschule – Gelsenkirchen, DE
- Soheil Human
Wirtschaftsuniversität Wien, AT
- Agnieszka Kitkowska
Jönköping University, SE
- Konrad Kollnig
Maastricht University, NL
- Kirsten Martin
University of Notre Dame, US
- Jelena Mitrovic
Universität Passau, DE & Institute for Artificial Intelligence R&D of Serbia – Novi Sad, RS
- Rishab Nithyanand
University of Iowa – Iowa City, US
- Shidong Pan
Australian National University – Acton, AU
- Sarah Radway
Harvard University – Allston, US
- Tim Samples
University of Georgia, US
- Florian Schaub
University of Michigan – Ann Arbor, US
- Yan Shvartzshnaider
York University – Toronto, CA
- Daniel Smullen
CableLabs – Louisville, US
- Peter Story
Clark University – Worcester, US
- Emma Tosch
Northeastern University – Boston, US
- Christine Utz
Radboud University Nijmegen, NL
- Kami Vaniea
University of Waterloo, CA
- Isabel Wagner
Universität Basel, CH
- Shomir Wilson
Pennsylvania State University – University Park, US
- Maximiliane Windl
LMU München, DE
- Lu Xian
University of Michigan – Ann Arbor, US
- Tianyang Zhao
Pennsylvania State University – University Park, US