

SoK: Technical Implementation and Human Impact of Internet Privacy Regulations

Eleanor Birrell
Pomona College

Jay Rodolitz
Northeastern University

Angel Ding
Wellesley College

Jenna Lee
University of Washington

Emily McReynolds
Future of Privacy Forum

Jevan Hutson
Hintze Law PLLC

Ada Lerner
Northeastern University

Abstract—Growing recognition of the potential for exploitation of personal data and of the shortcomings of prior privacy regimes has led to the passage of a multitude of new privacy regulations. Some of these laws—notably the European Union’s General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA)—have been the focus of large bodies of research by the computer science community, while others have received less attention. In this work, we analyze a set of 24 privacy laws and data protection regulations drawn from around the world—both those that have frequently been studied by computer scientists and those that have not—and develop a taxonomy of rights granted and obligations imposed by these laws. We then leverage this taxonomy to systematize 270 technical research papers published in computer science venues that investigate the impact of these laws and explore how technical solutions can complement legal protections. Finally, we analyze the results in this space through an interdisciplinary lens and make recommendations for future work at the intersection of computer science and legal privacy.

Index Terms—SoK, Privacy Regulations, Data Protection, Usable Privacy, Measurements

1. Introduction

Privacy law is shifting and developing rapidly throughout the world. As of March 2022, 157 nations have enacted privacy laws, with 12 laws enacted in 2021 alone [99]. Over the past decade, a large body of computer science research has emerged that studies these laws and their effects. This work has covered topics including rates of corporate compliance, people’s ability to exercise privacy rights, design patterns that subvert the consent process, and citizens’ expectations, understandings, and actions in the face of laws like the GDPR, which has been called “one of the strictest privacy laws in the world” [237].

What have computer scientists learned about these laws and their effects on privacy? And perhaps more importantly, what directions for future computer science research will be most impactful and effective at informing and driving better privacy laws that create more meaningful and equitable privacy outcomes? This paper sets out to answer these ques-

tions through an interdisciplinary lens by an authorship team consisting of both computer scientists and legal scholars.

We begin with a close reading of 24 privacy and data protection regulations, which provides a broad overview of a landscape of laws far too numerous to discuss individually yet which often share significant commonalities due to the broad influence of the GDPR on legislators worldwide [292]. From this close reading we construct a taxonomy of rights guaranteed and business obligations imposed by current Internet privacy and comprehensive privacy regulations.

We then review and systematize the computer science literature around modern digital privacy laws by organizing studies according to the rights and obligations in our taxonomy that they examine. We find that while the research in this space is extensive, its thoroughness varies dramatically, with the overwhelming majority of papers studying either the EU’s GDPR or California’s CCPA, laws which protect only about 6.3% of the world’s population [48], [49], [268]. Additionally, certain aspects of these laws—e.g., design patterns in consent banners—have been explored deeply, while other aspects—e.g., the right to non-discrimination—have been studied little or not at all.

Building on these analyses, our discussion presents two major arguments. First, we ask how great a limitation our focus on a few specific laws and a few specific rights is to our broad understanding of privacy law writ large. In other words: can we generalize from our deep study of these few contexts? Based on our systematization of the literature and on our interdisciplinary team’s analysis of legal factors that can cause varying privacy outcomes even under similar or identical laws, we argue that we should be wary of generalizing specific results beyond their cultural, temporal, and legal contexts. However, we also find that in combination with scholarship from other fields, this body of work has built a compelling case for certain general ideas about privacy law. Most prominently, we analyze the literature’s body of evidence against privacy self-management as a paradigm for privacy regulation, arguing for the necessity of alternative paradigms in order to produce effective and equitable privacy laws. We thus conclude our discussion by presenting a roadmap for computer science research that includes approaches beyond privacy self-management.

The contributions of this paper are:

- 1) We develop a taxonomy of rights and obligations enacted by modern Internet privacy and comprehensive privacy laws through close readings of global laws.
- 2) We systematize the computer science literature in the privacy law space, characterizing its extent, depth, and skew within our taxonomy.
- 3) We analyze results in this space through an interdisciplinary lens to formulate recommendations for how future computer science research can help guide more effective and equitable privacy regulation.

2. Methodology

This work involved three phases: (1) we identified computer science papers relating to Internet privacy and comprehensive privacy regulations, (2) we developed a taxonomy of rights and obligations imposed by such regulations around the world, and (3) we systematized the computer science literature relating to these regulations within this taxonomy, and (4) we formulated recommendations for future work informed by these results. Our methodology for phases (1) and (2) is described in this section. Our systematization is presented in Section 3. Our recommendations are discussed in Section 4.

2.1. Paper Selection

We identified ten computer science conferences that regularly publish papers about privacy:

- 1) IEEE Symposium on Security and Privacy (“Oakland”)
- 2) ACM Conference on Computer and Communications Security (CCS)
- 3) USENIX Security Symposium
- 4) Privacy Enhancing Technologies Symposium (PETS)
- 5) Symposium on Usable Privacy and Security (SOUPS)
- 6) ACM Conference on Human Factors in Computing Systems (CHI)
- 7) ACM Conference On Computer-Supported Cooperative Work And Social Computing (CSCW)
- 8) Network and Distributed System Security Symposium (NDSS)
- 9) ACM The Web Conference (WWW)
- 10) ACM Conference on Fairness, Accountability, and Transparency (FAccT)

For each conference, one author initially looked at the title and abstract of each paper published in that venue between 2017-2022¹; we also included papers published in the 2023 conferences through August 2023. This generated a preliminary list of 127 computer science papers relating to Internet privacy and comprehensive data protection regulations.

For each paper on our preliminary list, we applied the same analysis to (1) all of the backwards citations (i.e., works cited by that paper) and (2) all of the forward citations (i.e., subsequent papers found on Google Scholar that cited that paper). We also iteratively applied this analysis to any

1. For FAccT, which was founded in 2018, we considered the six years 2018-2023.

additional papers that were published in any of our ten selected venues (e.g., papers that were published prior to 2017 or that were initially excluded). This resulted in a maximal list of 410 candidate papers.

For each candidate paper, one of the authors read through the full paper and made a final determination about whether the paper was in scope for this work. Papers were considered in scope if they were computer science papers that evaluate the impact of an Internet privacy or comprehensive data protection regulation or if they describe or evaluate a tool to complement or enhance privacy under such a regulation. Merely mentioning a relevant law in the introduction was insufficient to be considered in scope. We focused exclusively on laws that directly impact the Internet; other sector-specific privacy laws (e.g., HIPAA) were out of scope. Any papers for which the assigned reader was unsure whether it was in scope were discussed collectively until a consensus was reached. This resulted in a list of 314 in-scope papers.

To validate our scope decisions, we double-coded all 1,357 papers published at USENIX Security between 2017-2023; we found strong inter-rater reliability ($\kappa = .842$).

Due to practical limitations, we removed any papers published outside of major computer science conferences unless they mentioned a relevant regulation or a specific right under such a regulation in the title of the paper. This resulted in a revised full list of 270 papers: 134 papers published in one of our ten selected venues, 19 papers published in other major computer science conferences (e.g., IMC, VLDB), 102 papers published in other computer science venues (e.g., minor conferences, workshops, or journals), 5 whitepapers, and 11 pre-prints posted on arXiv.

2.2. Legal Taxonomy

We organize our systematization around a taxonomy of legal features. To determine which laws to analyze, we considered the top 21 countries by population [294] (75.39% of the global population) and the top 8 countries by GDP [293] (75.90% of global GDP). We also included sub-national and super-national jurisdictions above our inclusion thresholds (e.g., the European Union and California). Our legal scholars then identified which of these jurisdictions have enacted Internet privacy or comprehensive privacy regulations and, for those that have not, whether they have draft regulations. This resulted in a list of 24 current and draft regulations (Table 2). Drawing on a close-reading of these 24 laws, two authors—both practicing privacy attorneys with significant experience in industry, private practice, and academia—developed a taxonomy of rights conferred and obligations imposed by these regulations (Table 1).

2.3. Paper Coding

We deductively coded each of the 270 papers in our revised full list using our taxonomy of legal features. We also inductively coded each paper for: research methodology, system stage, platform, and applicable laws.

| Definitions | Self-Managed Rights | Fundamental Rights | Business Obligations | Applicability | Enforcement |
|--------------------------------------|--|--|---|---------------------------------|--|
| 1. Personal info 2. Anonymization | 1. Right to access 2. Right to portability 3. Right to correct 4. Right to delete 5. Right to opt-out of processing 6. Right to consent to processing | 1. Right to not be subject to automated decisions 2. Prohibitions on certain technologies 3. Prohibitions on certain processing 4. Right to nondiscrim. (based on protected attr.) 5. Right to nondiscrim. (for invoking rights) | 1. Notice and transparency 2. Purpose/processing lim. 3. Data minimization 4. Security requirements 5. Privacy by design 6. Record keeping 7. Cross-border trans. lim. 8. Risk assessment 9. Contracting reqs. 10. Breach notification | 1. Subjects 2. Organizations | 1. Protection auth. 2. Gov. agency 3. Elected official 4. Priv. right of act. 5. Class action 6. Arbitration 7. Civil penalties 8. Criminal penalties |

TABLE 1: Taxonomy of rights and obligations under Internet privacy and data protection regulations around the world.

| Jurisdiction | Law | Status |
|--------------|--|--------|
| Brazil | Lei Geral de Proteção de Dados | Final |
| Bangladesh | Data Protection Act | Draft |
| California | California Consumer Privacy Act | Final |
| Canada | Personal Information Protection and Electronic Documents Act | Final |
| China | Personal Information Protection Law | Draft |
| Congo | Digital Code N°23-010 | Final |
| Egypt | Personal Data Protection Law | Final |
| Ethiopia | Personal Data Protection Proclamation | Draft |
| EU | E-Privacy Directive | Final |
| EU | General Data Protection Regulation | Final |
| India | Digital Personal Data Protection Act | Draft |
| Indonesia | Law regarding Personal Data Protection | Final |
| Japan | Act on the Protection of Personal Information | Final |
| Mexico | Federal Law for the Protection of Personal Data Held by Private Entities | Final |
| Nigeria | Nigeria Data Protection Regulation | Final |
| Pakistan | Personal Data Protection Bill | Draft |
| Philippines | Data Privacy Act | Final |
| Russia | Federal Law No. 152-FZ on Personal Data | Final |
| South Korea | Personal Information Protection Act | Final |
| Turkey | Law on the Protection of Personal Data | Final |
| Thailand | Personal Data Protection Act | Final |
| UK | UK General Data Protection Regulation | Final |
| US | Children’s Online Privacy Protection Act | Final |
| Vietnam | Personal Data Protection Decree | Draft |

TABLE 2: The 24 laws used to develop the legal taxonomy.

3. Systematization

Our legal taxonomy organized features of privacy laws into six general categories based on elements consistent across the texts: (1) definitions, (2) self-managed rights, (3) fundamental rights and prohibitions, (4) obligations, (5) applicability, and (6) enforcement. Papers that did not study any specific legal features—e.g., that observed changes before and after a law was implemented without tying them to a particular requirement—are described in Section 3.6.

Research Methodologies. 160 papers (59.3%) conducted measurement studies to observe the implementation of legal requirements. Of those, 60.0% used automated techniques to conduct large-scale studies and 51.3% conducted manual measurement studies (some papers did both). 20.6% of these paper included a longitudinal study that analyzed changes over time, and 16.3% included a cross-jurisdictional study that analyzed differences between different jurisdictions.

79 papers (29.3%) used HCI methods to investigate the human impact of privacy laws. Of these, 51.9% conducted

an exploratory study, 51.9% conducted a large-scale quantitative study, and 24.1% conducted an experimental study.

45 papers (16.7%) used systems techniques to implement and evaluate a system for implementing or enhancing privacy regulations. 10 papers (3.7%) introduced novel attacks based on legal privacy features. 17 papers (6.3%) used theoretical methods such as cryptography. 15 papers (5.6%) introduced frameworks.

System Stage. Of the 45 papers that were about tools and systems, 6 were deployed with a non-trivial userbase in the real world, 23 had implemented prototypes, 2 were in the design phase, and 3 were still high-level proposals.

Platform. 34.1% of papers investigated privacy features specifically in the context of websites, 18.9% focused on mobile, 9.6% focused on IoT devices, and 13.7% focused on other domains including databases, routers, TVs, social media, blockchains, networks, and cloud services. 33.7% of papers did not apply to any specific platform.

Applicable Laws. 87.0% of papers on our revised final list studied the implementation and impact of GDPR. 17.0% studied CCPA, 9.3% studied COPPA, and just 9.6% studied any other law. Only six papers (2.2%) considered laws outside of the United States and Europe. Two papers looked at Canada’s PIPEDA [300], [183]. One paper included India’s proposed PDPB [235], one included Singapore’s PDPA [200], one paper included the Phillipines’s DPA [221], and one paper included Turkey’s KVKK [133].

3.1. Definitions

Legal interpretation of privacy regulations depends strongly on key definitions. Many works focusing on definitions do so within the scope of a specific legal requirement. However, 8 papers explore definitions of terms that broadly affect the scope of applicable privacy regulations.

Most of the work that focused on definitions focused on GDPR’s definition of anonymization. Work in this area included proposing a statistical framework for GDPR-compliant anonymization [92] and formalizing GDPR’s “singling out” terminology—used to define identifiable data in Recital 26—and establishing the relation between this definition and existing techniques such as differential privacy and k -anonymity [58]. Gruschka et al. [103] analyzed two projects that rely on sensitive large datasets to identify how those projects attempt to protect users’ privacy

as required under GDPR. Other work has been critical. Cohen [57] demonstrated that common anonymization techniques such as k -anonymity might fall short of GDPR's legal standard for de-identifying data. Narayanan and Shmatikov [183] and Kutyłowski et al. [146] both discussed challenges of applying legal definitions of personally-identifiable or personal data versus pseudoanonymized or de-identified data in real world systems.

Less work has focused on other definitions. Gomez et al. [96] explored user perceptions of what constituted sensitive information. The many other definitions critical to the interpretation of privacy regulations—including comparing different definitions and interpretations across different laws—have not been the focus of computer science research.

3.2. Self-Managed Rights

Many Internet privacy and comprehensive data protection regulations confer *self-managed rights*, that is rights that require the individual to take action for the right to be leveraged. These include rights to access, rights to portability, rights to correct, rights to delete, and rights to opt-out of certain types of processing.

3.2.1. Right to Access. 20/24 laws we examined grant people a right to access data about them collected by others, although detailed requirements (e.g., time limits, format specifications, exceptions) vary between laws. This right has a high level of awareness in the EU [143], [197]; people in other jurisdictions are less aware of this right, even jurisdictions with legally-guaranteed rights of access, although many companies honor access requests regardless of user location [197]. The right to access has been extensively evaluated through measurement studies and user studies, and several tools have been designed to enhance this right.

Access Request Mechanisms. While researchers have consistently found that the most common Subject Access Request (SAR) mechanisms are email requests and web form [271], [272], [40], detailed procedures are generally different for every website [40]. Moreover, many websites [40] and child care apps [100] do not accept requests via email, the mechanism people consider most natural [7].

Authentication of Access Requests. Returning (potentially personal) data without authenticating the request can constitute a data breach. However, refusing a genuine request denies people their right to access. This tension makes authentication of SARs critical. However, the 28 EU Data Protection Authorities provide inconsistent guidelines [37].

Several papers have quantified how organizations authenticate SARs [37], [44], [70], [40], [69] with varying results: 10-71% authenticate requests with national ID cards, 15-31% use subject email access, 15-36% use subject account login, 6-22% use secret questions or confidential information, 0-11% use device cookies, and 1-5% call the data subject. Researchers have also examined whether companies consider IP addresses as sufficiently identifying to authenticate a request [2], with negative results.

To evaluate authentication of SARs, several projects have issued experimental, spoofed requests. A 2014 study found 25% of websites returned information in response to requests from email addresses that didn't match the account information [116]. Work conducted after GDPR went into effect produced varying results. One study found that 58/334 popular websites failed to take any steps to authenticate requests [40]. Using publicly information, researchers have had success rates of 10/14 [44], 15/55 [70], and 60/150 [196] at accessing PII in response to spoofed SARs. Social engineering persuaded 2-6% of companies to accept weaker authentication than initially requested [70], [40]. Around half of vulnerable organizations remained vulnerable to the same attack years later [69], and mid-sized organizations and non-profits were most frequently vulnerable [196].

An interview study found that companies do not receive many requests for access (i.e., less than 100 per year at some large companies), that some have not suspected any misuse of SARs, and that some have just 4-5 people with the access required to answer SARs while others handle them across a full customer service department [69].

Compliance with Right to Access. Many papers have measured compliance with SARs by making legitimate requests and analyzing the responses [116], [271], [272], [37], [40], [141], [39], [214].

One focus of this work has been quantifying compliance with GDPR's 30-day deadline for responding to access requests. Estimates of non-compliance rates have varied, with studies reporting 34-41.7% [271], 45% [272], 28.2% [40], and 24.4% [39] of websites failing to respond within 30 days. However, some of those responses stated that no data was found. Kröger et al. performed a similar experiment for mobile apps and found that 19-28% of apps failed to handle the request within 30 days [141]. One study found that sending a reminder email to the company decreased non-compliance to 20% [141]. Another study noted inconsistencies in how companies counted the time limit (from initial request or from the time they received additional information required) [271].

User studies in which website users [39], [197] or smart home users [50] issue requests show people find the process confusing and frustrating, and authentication can be difficult.

Studies have also quantified the distribution of data formats returned in response to SARs. Over half of responses were answered with data in structured formats such as CSV, JSON, or XML [40], [141], [271], but many non-machine-readable formats were received as well, including screenshots, pdfs, raw-text emails, and printed files [40]. Many of these formats are not what users expect to receive (pdfs, word files, or spreadsheets) [7]. Some studies found that data was unintelligible due to obscure labeling or formatting errors [272], [141]. User studies found that people often found the returned data incomprehensible, meaningless, unusable, or not useful [39], [285], [197]. Some users also struggled to open unfamiliar file formats (e.g., JSON files) [280].

Responses to SARs often fall short of user expectations [7], [39], [197]: most users would like responses to

include derived data (82%), data acquired from third-parties (81%), and metadata (73%), but only a minority of responses (39%, 49%, and 4% respectively) included these data, and the returned data was often incomplete [39]. Some companies (22% in 2015 prior to GDPR, 6% in 2019) returned only data types and not actual values [141]. In 2014, only 43% of apps and websites returned data that matched the observed data collected [116]. In 2022, many Android apps did not return data types that could be observed through traffic analysis [214]. However, a series of focus groups still found that people are surprised, shocked, and scared by the level of detail that some data downloads provide [280], [20], and users reported finding the data Twitter and Facebook shared about tracking to be illuminating [285].

Dashboards and Tools. Some companies offer tools or dashboards for directly downloading or interactively exploring data. However, companies rarely provide both an online tool and an opportunity to download formatted data, they often omit data of concern to users, and they do not help users understand what data is collected by companies [269]. A qualitative examination of 10 privacy dashboards determined that none complied with GDPR’s right to access [261]. Nonetheless, a user study found that people were surprised by and changed their attitudes after interacting with the Facebook’s transparency dashboard [20].

Researchers have proposed and developed various new tools to support the right to access. These include a privacy dashboard intended to facilitate user rights under GDPR [203], a tool for efficiently monitoring access request handling using temporal logic [19], a browser extension that enables people without accounts to make CCPA-compliant and GDPR-compliant requests [129], a co-design effort to propose designs, formats, and tools for future data downloads to achieve transparency goals [280], tools for retrofitting access requests into legacy systems [4], [160], a document engineering approach to designing and evaluating a disclosure interface [186], and a static analysis tool that identifies incomplete responses to access requests for WordPress plugins [232].

3.2.2. Right to Portability. 14/24 laws provide a right to portability, the ability to export data from one organization to another. Among GDPR’s self-managed rights, this was the least-known and the most misunderstood [143].

Most research pertaining to portability is comprised of measurement studies specific to GDPR. Three independent projects—one conducted in 2018 [289] and two conducted in 2021 [246], [143]—made GDPR portability requests to hundreds of online services; their results were compatible. 25-30% of companies failed to provide data export within the GDPR-mandated timeframe of 30 days. 40%-50% of file formats received were not compatible with GDPR’s requirement that data be exported in a structured, commonly-used, and machine readable format [289], [290]. Larger companies provide a larger scope of data export, more rigorous procedures for authentication, and are more likely to provide import options [246], but import options

remain rare (about 25% of services) [143]. Portability for IoT devices has only been measured at smaller scales: 4 devices [267] and 34 devices [26]. They found that 0-47% of IoT devices fail to provide data export within 30 days, and several issues were encountered including lack of request authentication, data not structured for machine readability, and lack of documentation and explanation.

Compared to the right to access, few user studies have explored the right to portability. A survey study found that 10% of participants had considered switching between online services, but about two-thirds felt that lack of portability was an obstacle to switching [143]. Usability of provided portability mechanisms has not yet been evaluated.

No tools have been designed specifically for portability, and most general self-management tools—with the exception of Odlaw [160]—do not support portability requests.

3.2.3. Right to Correct. 21/24 privacy regulations provide a right to correct data. The right to correct was one motivation for building VICEROY [129], an authenticated access tool for privacy self-management rights, but no work has focused specifically on the right to correct.

3.2.4. Right to Delete. 20/24 privacy laws include some form of a right to delete, including rights to withdraw consent. This right has been extensively evaluated through measurement studies and user studies, albeit primarily in the context of GDPR and the earlier EU Right to be Forgotten. Challenges have also been extensively explored, particularly in relation to distributed systems, and several tools—both formal tools and user-facing mechanisms—for supporting deletion have been developed.

The Right to Be Forgotten. Researchers have examined trends among delisting requests, including who makes these requests (most commonly law firms and reputation management services [33] about young adult men [297]) and types of information delisted (most commonly personal information or official information relating to illegal activities [33], [297]). While it is possible to use data-driven inferences to find many delisted links, the right to be forgotten does appear to enhance privacy of the delisted material [297].

The Right to Erasure. The right to erasure is the most widely known among GDPR’s self-management rights [143], and it is compatible with users’ expectations and normative beliefs about privacy [77], [175], [221]. Many people assume deleted posts are compromising, although 80% of users have deleted a post, often for innocuous reasons [174].

Measurement studies have found that 52-57% of websites and apps delete accounts upon request [116] and that 27% of websites updated or added privacy policies about deletion after GDPR [110], although many blockchain systems do not address this right in their policy [212].

However, user studies show people are often not aware of or do not use deletion controls [77], [109] because they cannot locate the controls [109], [253] or because they encountered barriers (e.g., authentication requirements, account requirements, paywalls, and dark patterns) [253],

[221]. Some controls are too coarse-grained, e.g., only offering an option to delete a full account [176]. People also report difficulty determining whether removal was successful and find that keeping information off the Internet requires ongoing efforts [253], and one third of attempted account deletion requests are never completed [221].

Challenges to Deletion. Researchers have identified technical challenges to implementing the right to delete, such as deleting all replicas including distributed data centers, backups, and offline copies [220], [229], [226], [64]. High-efficiency lazy deletion algorithms take hours (e.g., on Redis [226]) or months (e.g., 180 days in Google Cloud [229]) to delete data, and systems do not always guarantee that all copies were deleted [229]. A GDPR-compliant version of Redis that deletes data immediately incurred significant overhead [226], as did a SQL-based implementation designed to automate compliance with deletion and retention requirements [223]. However, some researchers argue that state-of-the-art practices can implement these rights [199], [67]. Other research suggests that rights to delete might be fundamentally incompatible with some technologies such as blockchains [212], although Farshid et al. [78] propose a blockchain-like technology that supports data deletion.

Legal exceptions can also pose a barrier to the right to delete: some electronic monitoring apps claim exemption from the right to delete under CCPA citing clauses about retaining data to “comply with a legal obligation” [193].

Tools for Deletion. Several of the tools developed to enhance right to access also support the right to delete [19], [4], [160], [232]. Researchers have also provided formal definitions for deletion [89], [95].

3.2.5. Rights to Opt-out of Processing.

17/24 laws grant a right to opt-out of some processing, with different scope.

The Right to Opt-out of Sale under CCPA. While compliance with the right to opt-out of sale has increased since CCPA went into effect [189], there are still numerous websites that do not provide an opt-out of sale link on their homepage [189], [279]. Opt-out interfaces often exhibit dark patterns and other interface designs that make it harder to opt-out [189], [279]; manipulative designs that frequently occur in the wild—e.g., asymmetric UI, extra clicks, and fillable forms—significantly decrease opt-out of sale rates [189].

The right to opt-out of sale is widely misunderstood: only 30.5-61.1% of Californians can correctly identify which behaviors would be covered by this right [55]. Websites are also inconsistent about how they interpret “sale” [55].

Several projects have focused on enhancing the right to opt-out of sale. Efforts to improve usability include standardized icons and taglines [61], [111]—which were incorporated into the text of the CCPA regulations—and browser extensions that improve visibility of opt-out mechanisms [233] and set browser headers and privacy signals to automatically opt-out of sale [302]. However, while these signals are legally-enforceable, compliance with these signals is low [305], [54].

Other Rights to Opt-out of Processing. Nearly 90% of websites provide opt-out choices for email communications or targeted advertising in their privacy policy, but these choices are often hard to find and comprehend due to poor readability and the lack of standardized wording [110]. Kumar et al. [24] built a corpus and a model to extract and classify opt-out links in privacy policies. They observe variations in kinds and frequencies of opt-outs based on the popularity of sites, and find that many privacy policies have no such links. They also created a browser extension, Opt-Out Easy, which surfaces these links to users. Arfelt et al. [19] express GDPR’s withdrawal of consent (Article 7(3)), right to restrict processing (Article 18), and right to object (Article 21) in temporal logic and demonstrate efficient monitoring. Allegue et al. [9] provide a consent manager for IoT smart homes. Odlaw [160] claims to support the right to object in legacy systems.

3.2.6. Rights to Consent to Processing. 15/24 laws require affirmative consent for certain types of data processing (e.g., sensitive information, children’s information, and/or cookies). For example, GDPR requires consent—defined as a freely-given, affirmative act—for any processing of personal information that is not covered by an alternative legal basis, a model that was also adopted by many subsequently laws.

Consent Interfaces. Many papers have categorized and quantified consent interfaces, primarily cookie banners:

- 1) *Choice Options.* Measurement studies have consistently found that 20-37% of banners present no options or are confirmation-only [121], [65], [171]. Rates are higher in certain countries [65], contexts (e.g., porn websites [277]), or modalities (e.g., mobile [105]). Banners with choices use a variety of mechanisms including binary buttons, sliders, checkboxes, and per-vendor settings. Among choice mechanisms, checkboxes were the most common immediately after GDPR went into effect [65], but options to opt-out of cookies directly in a banner are rare [65], [215], [275], [134].
- 2) *Location.* Most desktop cookie banners (57.9%) are implemented as bars at the bottom of the page [275].
- 3) *Design Elements.* Highlighting, pre-selection, and other forms of nudging and manipulative patterns are common [275], [169], [188], [171], [209].
- 4) *Banner Text.* GDPR requires that cookie purposes be disclosed to the user; this requirement is not always met. Many banners fail to mention a purpose [219], use vague purpose language [219], assign incorrect purposes [35], or use biased text with framing [219].

Computer scientists have identified requirements for cookie consent [217], characterized dark patterns [168], and developed frameworks for evaluating privacy choice interfaces [107], [80]. Overall, estimates suggest that 54-95% of cookie banners fail to meet GDPR’s standard for consent [169], [188], [218], [35], [209], and some people argue that no designs satisfy all of GDPR’s elements [98]. Notifying companies of non-compliant interfaces does not consistently result in improvements to the banner design [115].

Legal requirements for opt-in consent have driven adoption of Consent Management Providers (CMPs). Adoption is highest among mid-market websites, but CMPs are increasingly used by all sorts of sites, with significant jumps in adoption immediately after GDPR and CCPA went into effect [117]. However, CMP-produced banners are not always legally-compliant [263], [242], [209]. The default banner produced by many CMPs uses highlighting to nudge user consent [263], 38% of CMPs do not support any nudge-free designs [242], and some cookie consent libraries support cookie banners with no decline option [65].

Consent interface designs have been quantified in other contexts: (1) mobile app tracking—42.6% are confirmation-only [185]; (2) smart homes—interfaces are manipulative, frustrating, and lack options [50] and withholding or revoking consent is hard [52]; (3) Hybrid Broadcast Broadband TVs—some channels lack consent mechanisms or fail to support revocation [247]; and (4) voice assistants [225].

Impact of Interfaces on Consent Decisions. In 2022, there were 56% more cookie banners in the EU compared to other jurisdictions [202], and many European users report banner fatigue [144]. Nonetheless, interface design—including nudging and dark patterns—can impact consent decisions.

- 1) *Choice Options.* Removing the opt-out button from the first page has the most effect on consent choices, increasing consent rates by 20-30 percentage points [275], [188], [108], [38]. One study found users significantly more likely to consent to cookies when presented with binary-choice compared to finer-grained options [275], but another found the exact number of choices does not have an effect [162].
- 2) *Location.* A large-scale study of European users in the wild found that people are three times as likely to interact with banners in the lower-left than with banners at the top and bottom of a page [275]. However, a later study with predominantly-American MTurk users found that position had no significant effect [32].
- 3) *Design Elements.* Pre-selection or defaults can significantly increase cookie opt-in [275], [162] and result in lower recall and more regret [162]. A color-based nudging bar displaying privacy threat is most effective at nudging users away from default options [32]. However, highlighting does not significantly increase opt-in [275], [32]. In general, defaults and nudging can influence the cognitive decision process, particularly for users with lower privacy concerns [22].
- 4) *Banner Text.* Many textual variations—e.g., the labels on buttons and whether a banner implies that declining cookies will negatively effect experience—have no significant effect on cookie consent [108], [38], but loss versus gain framing for the button labels does [161].

Behavior can also differ significantly between desktop and mobile devices [275], [32], [202]. Manipulative consent interfaces are not constrained to cookie consent; some websites use nudging and dark patterns to steer users towards particular behaviors [167]. Users familiar with computer security are more likely to change defaults [32].

Tools for Consent. Given the many banner-related privacy concerns, several tools have been designed to enhance privacy through automated tools and signals. These tools can automatically answer consent pop-ups based on a user's preferences [187], [202], predict actions required to disable unnecessary cookies [134], and identify legally non-compliant banner language [278]. Hils et al. [118] conducted a longitudinal study of privacy preference signals.

Effect of Consent Requirements on Cookies and Tracking. Researchers have applied two distinct methodologies to measure the effect of consent requirements on cookies and tracking: (1) conducting longitudinal studies before and after a law goes into effect, and (2) quantifying non-consensual processing after a law is in effect.

1. *Longitudinal Studies:* Many projects have conducted longitudinal analyses of the effect of GDPR on tracking and third-parties [65], [240], [121], [126], [273], [120], [128], [63]. Although they observed a reduction in usage of third-party web technology and cookies immediately after GDPR went into effect [128], [121], there was no significant long-term drop [128], [65]. There were some changes towards increased market concentration in third-party services [128], [240], [126], [120], [63], and one study observed a 40% reduction in cookie syncing, although the general shape of the ecosystem did not change [273]. Another study observed an uptick in new third-party cookies placed when GDPR went into effect [120], which might signify adoption of CMPs. However, few developers reported making substantive changes such as adding popup consent dialogues [286].

Rasaii et al. [202] measured the impact of CCPA and LGPD and found that those laws had no effect on cookies.

Research has also looked at tracking by mobile apps. Although there are geodifferences in tracking behavior [145], there was no significant change in tracking by apps after GDPR went into effect [137]. It is common for apps to engage in third-party tracking, and very few of those apps obtain consent [127], [304], [138], [122]. Apps continued to share data with tracking companies prior to user consent after Apple introduced App Tracking Transparency (ATT) [139]. Ret et al. [207] looked at information exposure by IoT devices across jurisdictions; they noted that US devices contact more third-parties compared to UK devices.

Although most users feel negatively about third-parties, many take no action to prevent tracking [60]. Among those who do, most common actions are using a browser extension or manually deleting cookies; fewer people use a privacy-oriented browser or private browsing mode [171].

2. *Processing without consent:* Websites without banners sometimes set cookies [46], [76], [121], [276], and 82.5% of websites that use CMPs set at least one cookie that is not covered by their cookie banner [35]. Websites with banners set cookies prior to obtaining consent [169], [265], [188], [35], [215], [276], [209], a practice that is even more common among EU government websites [213]. However, the practice of setting cookies prior to consent decreased after GDPR [62], [151], [121]. Some websites set cookies even after the user rejects those cookies [121], [35], [169],

[215], some use browser fingerprinting to bypass cookie-consent [195], and 3.8% of top websites use cookie respawning [86]. However, opting out of cookies does reduce cookies [202] and can result in fewer scripts (and corresponding reduced vulnerability to scripting attacks) [135].

Some issues with non-consensual cookies arise from third-party elements imported into a site. Many profiling cookies set prior to user consent are actually set by advertising networks [265] and ghosted cookies—cookies set by hierarchically-imported resources—pose a challenge to consent because the website does not have full control over these first-party cookies [216]. Only 12.8% of third-party cookies are mentioned by cookie policies, and only 5% include a description of the cookie’s purpose in a well-structured table [85].

Beyond website cookies, data sharing and use without opt-in consent occurs in many other contexts. Some websites send marketing emails without opt-in consent, without an option to revoke consent, or after a user revokes consent [142]. In the context of mobile apps, 16.7% of mobile apps share data with third-party trackers prior to user consent, and 1.0% share data after a user explicitly declines consent [185]. 20.9% of apps transmit data without displaying a consent dialogue or privacy notice [136]. 34.4% of apps share personal information with third-party data controllers without opt-in consent, most commonly sending the AAID [184]. Leaks of unresettable user identifiers (UIIs) on Android devices can bypass the permission consent mechanisms [172]. Apps practice deceptive uses of legitimate interest to justify data collection without consent in ways inconsistent with user preferences [147]. In the context of IoT, personal information about other users can be extracted from Amazon Echo devices [88], and Hybrid Broadcast Broadband TV channels track users prior to receiving consent [247].

Misconceptions about consent requirements are common, with many developers believing that having a privacy policy supersedes the need for explicit consent [184]. Many developers rely on third-party app builders or SDKs to make their app compliant and assume that libraries implement compliant behavior. To prevent non-consensual data use, researchers have developed tools that automatically verify and enforce usage consistent with user consent [56], [131].

Effect of Consent Requirements for Sensitive Data. 13/24 laws require opt-in consent for processing sensitive personal information. However, definitions of “sensitive” vary and differ significantly from what users consider sensitive [96].

Work conducted prior to GDPR found that Facebook made significant use of sensitive data for targeting ads and therefore predicted that the regulation should significantly impact data processing by major ad providers [43]. Subsequent work did find significantly lower rates of tracking for sensitive data in Europe compared to other jurisdictions [63], however tracking persisted on sites related to health [123], [63], sexual orientation and preferences [123], [277], religion [123], [63], and politics [63]. 74% of paid apps held the same dangerous permissions as their free versions [113].

Restrictions on the collection of sensitive personal information can also have unintended negative consequences. Interviews with industry experts working on algorithmic fairness for machine learning revealed that GDPR’s restrictions were viewed as prohibitive, and the resulting lack of access to racial data resulted in no longer trying to detect racial bias in their machine learning systems [16].

Effect of Consent Requirements for Children’s Data. 12/24 laws require parental consent for information about children, but the details (e.g., collection versus processing, the age limit for protection, and the requirements for verifying parental consent) vary between regulations.

Several projects have investigated COPPA compliance by mobile apps. This work has consistently found that many apps violate COPPA in a variety of ways, including using SDKs that prohibit use in child-targeted apps because they collect or share PII [208], [6], [79], [100], [204], using libraries without necessary COPPA-compliant parameters [6], not asking for parental consent prior to collection [100], [10], using fingerprinting-alike libraries to bypass parental consent [83], or using parental consent mechanisms such as age gates and knowledge-based questions that do not satisfy the FTC’s requirements for verifiable parental consent [10]. Behaviors that violate COPPA have also been observed in COPPA-approved and child-oriented websites [281] and in voice assistant skills in the “kids” category [150]. Overlapping behavior in free and paid versions of apps might also be indicative of practices that violate COPPA [113]. Xie et al. [296] automatically analyzed use of children’s data by IoT devices and found noncompliance in 8/512 skills.

Failure to comply with restrictions on processing children’s data might be due to barriers to compliance rather than malicious intent. Developers of popular Android children’s apps report issues including a lack of transparency from libraries, a need for data to understand user behavior, and difficulty monetizing apps in age-appropriate ways [72].

Efforts to facilitate COPPA compliance have resulted in tools that leverage dynamic execution and traffic monitoring [287] or machine learning [29], [296] to analyze apps or IoT devices and detect behavior that violates legal regulations. These tools have high accuracy but are not currently in widespread use.

In general, COPPA requirements are consistent with parental expectations and social norms, although younger parents are significantly more accepting of data collection [18]. However, by incentivizing online services to ban users younger than 13, there is also some evidence that COPPA may have reduced privacy for adolescents, who may lie about their age to join platforms, thus aging out of protections for minors before they turn 18 [68].

3.3. Fundamental Rights and Prohibitions

In contrast with self-managed rights, which must be explicitly invoked by individuals, fundamental rights impose general prohibitions on certain types of behaviors. These rights have been rarely studied by computer scientists.

15/24 privacy laws create rights to not be subject to automated decision making, but only four papers have focused on this right. Kaushik et al. [132] found that GDPR's right to not be subject to automated decision making is commonly misunderstood—many people believe users can opt-out in advance—and fails to meet expectations for transparency. Krishna et al. [140] formulate the problem of implementing the right to explanation in the context of automated decisions as an optimization problem robust against model updates to accommodate deletion requests. Other work has considered how this and other fundamental rights might be implemented for databases [228], [229].

GDPR and LGPD grant freedom from discriminatory processing on the basis of sensitive personal information. Four laws—CPRA, GDPR, LGPD, and China's PIPA—grant a right to non-discrimination if a user invokes their rights to privacy self-management. These rights have not been explored in the computer science literature.

3.4. Obligations

Obligations are procedural requirements that must be satisfied when collecting or processing data. Examples include transparency requirements, legal basis requirements, data minimization requirements, and privacy by design. Some of these legal aspects—specifically transparency and consent as a legal basis—have been the focus of significant bodies of work by computer scientists. Others have not.

3.4.1. Notice and Transparency. Transparency was a key motivation behind many privacy laws, and 22/24 privacy laws have specific transparency requirements. For example, GDPR calls for data to be processed “in a transparent manner in relation to the data subject” [5(1)(a)], and CCPA features the “right to know” as the first right granted to consumers. Despite some challenges [173], a lot of work has focused on privacy policies as the most common form of notice; some work has also looked at other transparency mechanisms in the context of legal requirements.

Impact of Regulations on Privacy Policies. Several longitudinal studies looked at privacy policies before and after GDPR went into effect to understand its impact [65], [110], [286], [154], [299], [13], [3]. 4.9% of apps added privacy policies and 50% updated pre-existing policies [65]—the most widespread changes in the last decade [13], [3]—with many adding new content (e.g., options regarding deletion [110], information about privacy self-management rights [3], and information about protections for children's data [299]) to meet particular GDPR requirements. 38% of Android developers reported adding or updating the privacy policy for their app [286]. Overall, privacy policy sensitivity increased between 1997-2018, with spikes corresponding to the enactment of privacy regulations [158]. However, not all companies had privacy policies even after GDPR [65], [145], [163]. Cross-jurisdictional analyses have also identified jurisdictional differences between policies in the EU and the US [145], [21], with some evidence suggesting that the GDPR reduced data collection [21].

The impact of regulations on the transparency of privacy policies was mixed. Using NLP techniques, researchers found an increase in terms related to GDPR rights [65], [154], [3] and in granularity of disclosures [21]; most apps with different policies in different countries were due to additional clauses relating to GDPR or CCPA [145]. Some longitudinal studies have found that policies became more specific with improved presentation [154] and simpler and more regularized [158]. However, many privacy policies showed no improvement in any of 10 privacy measures [299]. Other studies have found that privacy policies also became longer [154], [13], [3] and harder to read [13]. Some policies covered more data use practices at the cost of reduced specificity [154], and information about tracking and information sharing with third parties were still frequently missing [13].

Compliance with GDPR Transparency Requirements. According to analyses, privacy policies can violate GDPR in five ways: (1) omitting required information [59], [153] (Fan et al. [75] identify six required notice categories; Vanezi et al. [298] identify a list of 89 terms across 7 groups that should be included in privacy policies), (2) describing prohibited data practices [59], [153], (3) using unclear language [59], [153], [191], (4) not providing the privacy policy in an accessible location, and (5) inaccurately or incompletely disclosing data practices. A large body of work has also been devoted to developing tools that use machine learning to automatically analyze whether privacy policies comply with GDPR's transparency requirements [59], [179], [271], [272], [14], [15], [75], [85], [205], [298], [191], [156], [73], [102], [201], [200], [155], [192]. Many of these tools have been applied to corpuses of post-GDPR privacy policies to determine rates of compliance with GDPR.

- 1) *Omitting Required Information.* Between 8.3% [298] and 23.7% [75] of website privacy policies are missing at least one required category. Observed compliance levels differ significantly for different disclosure requirements imposed by GDPR [201]: more than 90% of policies disclose categories of data collected and purposes of data processing (albeit sometimes bundled [176]), but only 15.3% disclose how personal information is used for automated decision making or profiling [267]. 43% of child care apps do not mention processing sensitive data about children [100]. 56% of privacy policies for browser extensions omit one or more pieces of information required by GDPR [155]. Users struggle to identify policy excerpts relevant to GDPR's articles [179].
- 2) *Illegal Data Use Practices.* None of the papers we systematized measured prevalence of illegal practices.
- 3) *Unclear Language.* An estimated 1.4% of websites have readable privacy policies [191]. Many use vague language [176], [122]. Moreover, 7.6-18.1% of privacy policies for Android apps contain contradictions that may be indicative of misleading statements [14], [15], [42]; many are contradictions due to inconsistencies with GDPR's definition of personal information. GDPR

requires that privacy notices be understandable to children if a company processes children's data; however, studies with children reveal safety concerns and a lack of awareness of the value of data [66], children do not understand the technical terms that appear in privacy policies [178] and misconceptions about data practices remain pervasive among children [245], suggesting that current policies fall short of meeting this legal standard.

- 4) *Inaccessible Policy.* Most websites provide a privacy policy in an accessible location [191], but there still are some without privacy policies [276]. Compliance can be lower in other contexts: 74% of IoT producers' websites [267], [163], 62.2% of Google Assistant actions [150], 50.5-55% of Android App Store pages [304], [113], 32% of browser extensions [155], 27.7% of Amazon Alexa skills [150], 16% of porn sites [277], and only 9% of Smart Home devices [163]. Moreover, there are no clear standards for what constitutes an accessible location for apps and smart home devices [304], [163].
- 5) *Inaccurate Policy.* App behavior is not always consistent with notices provided by websites and apps [206], [15], [127], [306], [75], [42], [10]: 42.4-77.9% of apps exhibit at least one behavior inconsistent with their privacy policy [15], [75], [295], 17-18% share information with third parties without disclosing it [306], [79], and app behavior is frequently inconsistent with app privacy labels [136], [295]. 35% of European websites collect data not disclosed by their privacy policy [192]. 11/165 popular trackers with opt-out choices exhibited data practices inconsistent with their privacy policy [41]. IoT skills for Amazon and Google devices have also been found to be inconsistent with their privacy policies [106], [296], and 31% of IoT companion apps share data without disclosing it. Independent work developing a testbed for IoT devices found that half of 11 sample devices collected data not disclosed by their privacy policy [244], and almost half of browser extensions have inconsistencies between privacy policies and actual data practices [155].

Overall, early estimates suggested up to a third of the privacy policies for large companies were not compliant with GDPR [59], but later work found that number could be as high as 97% [201], with many privacy policies having multiple compliance issues [156]. Compliance is higher for top tracking companies, most of which meet the minimum legal requirements set out by GDPR [272].

Compliance with Other Transparency Requirements. Three projects focused on the transparency requirements imposed by COPPA on apps that target children under 13. The first found that only 10.8% of these apps targeted at young children provided a privacy policy in their Google Play Store page in 2013 even though half collected personal information [152]. Five years later, only half of parental control apps clearly informed users of their data practices, and only 24% provided a complete list of third-parties with which they share information [79]. Some companies violate

COPPA by failing to address children's data in their privacy policy [204] or by not disclosing data practices in the privacy policy [304]. In 2022, many child care apps still failed to disclose use of trackers or processing of sensitive data about children [100].

Two papers evaluated CCPA's transparency requirements. Chen et al. found that while almost all U.S. websites describe their data sharing practices, only 24.4% of those list every category of personal information that is shared with each category of recipient as required by CCPA [55]. Musa et al. [180] proposed a technique for automatically inferring data sharing relationships and validated it against CCPA's data broker registry, suggesting a means to verify accuracy of CCPA disclosures.

The only paper that considered transparency requirements outside Europe and the U.S. was Qamar et al. [200], who developed a tool that measures similarity between the text of a law and a privacy policy; their tool supports compliance evaluation for Singapore's PDPA, but they have not published compliance rates.

Tools to Improve Transparency. One approach to facilitating compliance with transparency requirements is to automatically generate policies that comply with the various laws. Such tools have been developed for COPPA [152], [303], GDPR [90], [12], [303], [27], [231], CCPA [303], and CalOPPA [303]. There are also questionnaire-based generators that create privacy policies for mobile apps [17], [125], [256], [257]; three of these claim to generate policies that are compliant with COPPA, GDPR, and CCPA. However, an analysis conducted in January 2021 found that policies generated by available tools were only compliant with GDPR [303]; all generated policies violated multiple transparency-related requirements imposed by COPPA and by CCPA.

An alternative approach is to enhance transparency with summarization or annotation. Several independent projects applied NLP to automatically identify and highlight parts of privacy policies relevant to GDPR's requirements [258], [259], [179], [53], [21]. Mustapha et al. [181] provide an improved tool for automatically annotating privacy policies.

Other tools include formal languages for expressing GDPR's transparency requirements [19], [101], expressing and verifying privacy policies [260], and modeling inter-process communication to audit policies [27]. Wang et al. [284] used cryptography and trusted execution environments to automate compliance with privacy regulations.

3.4.2. Purpose or Processing Limitations. 17/24 privacy laws impose purpose or processing limitations as a business obligation. These include both legal bases for processing and explicit purpose limitations.

Legal Bases for Processing. Many laws accept user consent as one possible legal basis for processing, effectively turning a business obligation into a self-managed right. Consent as a legal basis has been extensively studied by computer scientists (Section 3.2.6), but only three papers have studied other legal bases for processing. Arfelt et al. [19] express

GDPR's legal basis requirement (Article 6(1)) in temporal logic and find it can be efficiently monitored. Kutylowski et al. [146] discuss challenges that can arise from GDPR's legal basis requirement in cases where a processor ceases to exist but personal data are still stored by a third-party storage provider. Han et al. [114] argue that the "Legitimate Interests" clause in the GDPR covers use of data collected prior to GDPR; they build a system to provide sequential recommendations by training a global model using pre-GDPR data then fine-tuning that model locally with more recent data.

Purpose Limitation. Few projects have looked at purpose limitations. One found that after GDPR went into effect, Android apps declared fewer dangerous permissions and that many apps reduced the number of times they accessed permissions [177], suggesting that GDPR's purpose limitation requirement might have significantly enhanced privacy. Another study found that additional metadata required to enforce purpose limitation and other GDPR requirements imposed a 2-5x performance slowdown on three widely-used database systems [228].

To facilitate compliance with purpose limitation requirements, Wolf et al. [288] developed HivePBAC, an adaptation of purpose-based access control designed to ensure purpose limitation for message-oriented architectures, and Karami et al. [131] developed a programming language that generates runtime errors if data are used for purposes other than those for which they were collected or if they are not deleted after their purpose is complete.

3.4.3. Data Minimization. 12/24 laws impose a data minimization obligation, however this obligation has been rarely studied. One study found that personalization can be relatively robust to global minimization but that quality loss is significant for some users [34]. Another found that few apps enable data-minimization SDK settings [139] and that 16/31 IoT devices had at least one unnecessary data flow [164]. Senarath et al. [224] found that developers struggle to implement data minimization due to uncertainty about how data could potentially be used, and that developers are inconsistent in how they apply data minimization.

Efforts to support data minimization requirements are also few. They include expressing data minimization requirements in temporal logic to enable efficient monitoring [19], developing a framework for implementing data minimization in machine learning systems by iteratively estimating the system performance curve and use of personal data when a performance-based stopping criteria is achieved [227], and developing a tool for automatically comparing privacy policies between counterparts and identifying overly-broad data practices (a subsequent analysis flagged 48.3% of privacy policies as overly broad) [301].

3.4.4. Security Requirements. 16/24 laws impose security requirements such as encryption for personal information, but only a few projects have evaluated the impact of this requirement, and all focused on GDPR. Longitudinal studies

observed a 9% decrease in plaintext transmission of personal data after GDPR went into effect, but 39% of top Android apps still transmitted plaintext data [127], [75]. Many mHealth apps that attempted to encrypt data contained at least one error [75]. Although some apps contain geodifferences in security settings, those differences do not correspond to privacy regulation jurisdictions [145].

In some cases, compliance with legally-mandated encryption requirements can impose a significant performance overhead [226]. Marjanov et al. [165] analyzed fines imposed for violations of GDPR's security requirements and identified common failings and danger points.

3.4.5. Privacy by Design. Among the laws we reviewed, only GDPR (and the UK GDPR) includes a privacy by design obligation, although such obligations are also present in laws beyond the scope of our review (e.g., Australia's Privacy Act, Kenya's Data Protection Act). A few projects have considered this requirement. Research has analyzed the mobile ecosystem through the lens of privacy by design [47], conducted case studies [94], and studied developer perspectives on privacy by design to identify barriers [8]. Recommendations include more guidance for developers [47], changes at individual and organization levels [8], and development of tools targeted at the engineering mindset [166].

Researchers have developed tools for facilitating privacy by design using domain-specific languages [90], [91], semantic models and automated verification [56], and formal modeling and interactive theorem proving [130]. Deshpande [67] proposed an architecture for private-by-design database systems. Tamò-Larrieux et al. [254] analyzed privacy by design as a stepping stone toward a right to customize data processing.

3.4.6. Record Keeping. 19/24 laws impose a record-keeping requirement, but this obligation has been studied only in limited contexts such as measuring overhead incurred by adding synchronous logging to Redis [226] or by a multi-level logging scheme [264]. To facilitate compliance with this obligation, Ryan et al. [211] proposed DPCat, a standardized representation for the collection and transfer of Register of Processing Activities (ROPA) information.

3.4.7. Cross-Border Transfer Limitations. 17/24 laws impose restrictions on cross-border transfers. Analyses of the mobile tracking ecosystem prior to GDPR predicted a significant impact [204]. However, subsequent work found that tracking flows did not change significantly [123] and that compliance rates were low, with 93% of websites embedding third parties located in regions outside the Privacy Shield [270] and 66% of apps including cross-border transfers that were not accurately disclosed in their privacy policy [104]. IoT companion apps also transmit data across regions in ways that could violate GDPR [182].

3.4.8. Risk Assessments. 13/24 laws require risk assessments such as Data Protection Impact Assessments, Privacy Impact Assessments, or Algorithmic Impact Assessments.

Three projects have briefly considered this requirement in the context of big data systems [103], database systems [229], and smart homes [94], but it has not been an area of significant focus.

3.4.9. Contracting Requirements. 12/24 laws have contracting requirements, e.g., for service providers or third-parties who process personal information. Amaral et al. [11] developed an NLP approach to automating compliance checking for GDPR data processing agreements. No other work has looked at this obligation.

3.4.10. Breach Notification Requirements. 18/24 laws have breach notification requirements. Shastri et al.'s work on GDPR-compliant databases briefly mentions this obligation [228], [229], but no work explicitly focuses on it.

3.5. Applicability and Enforcement

The 24 laws we analyzed include a range of different scopes of applicability, both in terms of which people are granted these protections (criteria include residence, citizenship, physical location, age, and employee status) and which organizations are subject to the regulation (criteria include number of users, company revenue, revenue from selling personal information, organization's country of registration, jurisdictions in which a company does business, and non-profit or governmental status). They also include a range of different enforcement mechanisms. While some papers conducted cross-jurisdictional measurements or studies, none looked explicitly at the impact of legal applicability or enforcement.

3.6. General Papers

While we were able to position most papers within our legal taxonomy, some work approached privacy regulations from a more general perspective.

Several projects evaluate the effect of GDPR on specific things without tying their results to any particular legal requirement. These include the impact of GDPR on cross-library data harvesting [283], on WHOIS records [159], on feasibility of large-scale vulnerability notifications [241], and on universities' cloud migration [84]. Wong et al. [291] evaluated the impact of both GDPR and CCPA on business risks identified in investor documents.

Many projects explored generally what types of architectural and design changes would be required to bring systems into compliance with GDPR [93], [194], [148], [119], [229], [222], [28], [45], and one did so for the proposed Indian Personal Data Protection Bill [235]. Other general projects explored how various technologies could facilitate GDPR compliance generally. These included developed or proposed formal languages [25], [19], [262], [36] or tools [74], [81], [198], [82], [1], [36], [11], [255], [124] for facilitating and automating GDPR compliance. The ISO 27001 standards [157] and blockchains [266] were also explored.

Some work explored users' [230], [300] and developers' [236], [87], [51], [248], [10], [251], [50], [274], [133] attitudes towards, awareness of, and understanding of various privacy regulations. Since legal compliance is often framed as the developer's choice and responsibility [252], research has also explored advice or guidance available to (and used by) developers [252], [251], [47], [250], [149], [10] and identified barriers to compliance from the developer's perspective [251], [248], [10], [112], [30], [243], [276], [133], [249], [5]. However, Utz et al. [276] found that notifications about privacy issues were less well-received than notifications about security issues.

4. Discussion

Looking at the 270 papers systematized in this work through an interdisciplinary lens, we see a deep, active, and highly impactful body of scholarship. However, systematizing this volume of work also illuminates patterns in how computer scientists currently approach the problem of analyzing the implementation and impact of privacy regulations. Based on our analysis of these patterns and our systematization of existing research in this space, we formulate recommendations about directions for future computer science research at the intersection of privacy and law.

Recommendation 1: *Computer science researchers should expand our efforts to evaluate and amplify non-self-management aspects of privacy and data protection regulations*

A majority of the work we systematized (51.5%) focused on measuring, evaluating, or enhancing privacy self-management features such as access, deletion, opting-in, or opting-out. Another 27.8% focused on the right to transparency, most of which was framed within the context of notice and consent, implicitly interpreting this requirement through the lens of privacy self-management. Most of this work is comprised of criticisms of privacy self-management, such as quantifying designs that deter users from invoking their rights or measuring the (poor) usability of existing implementations of self-management rights.

These critiques of self-management contribute to a growing interdisciplinary consensus that privacy self-management is inherently unworkable. Enhanced transparency is intended to empower users to make informed decisions about whether to consent to data practices, but the resulting disclosures are still unreadable [13], [191] and omit critical information [75], [298]. Implementations of self-management rights frequently deter people from invoking their rights by leveraging cognitive biases in so-called "dark patterns" [71], [282], [97], [275], [189]. Moreover, self-management simply doesn't scale to the number of companies with which users regularly interact and the difficulty of identifying the many third parties with access to personal data [170], [238], [210], [239].

Non-self-management aspects, including fundamental rights and obligations, have the potential to overcome the limitations of privacy self-management. However, these aspects have been under-studied. We recommend that com-

puter science researchers shift our efforts towards evaluating non-self-management features of existing regulations—e.g., evaluating the impact of non-discrimination requirements through measurements and user studies—and that we work to develop tools that complement these regulatory efforts—e.g., tools for auditing data minimization.

Recommendation 2: *Computer science researchers should extend our efforts to evaluate the implementation and impact of currently under-studied aspects of privacy regulations.*

Much of the work we systematized focused on a small number of legal aspects such as right to access (34 papers), consent interfaces (41 papers), and right to transparency (75 papers). By contrast, other aspects of modern privacy laws have gone largely unstudied by our community. Some of these understudied aspects challenge and enhance common approaches to self-management or take users out of the self-management loop. Examples of under-studied legal aspects that could benefit from additional research include right to correct (1 current paper), data minimization (9 papers), anti-discrimination requirements (0 papers), prohibitions on facial recognition (1 paper), and rights not to be subject to automated decisions (4 papers).

Recommendation 3: *Computer science researchers should prioritize longitudinal and cross-cultural work at the intersection of technology and privacy law.*

The measurement work systematized in this paper focused on the time immediately after or immediately around the date when a law went into effect. However, the legal community considers post-enactment guidelines and case law to provide essential interpretations of laws that define the rights and obligations imposed by those laws. For example, what constitutes a valid opt-out of sale mechanism under California law has evolved since January 2020 via new guidelines, new interpretations and legal actions, and subsequent legal amendments. To provide meaningful evaluations of privacy regulations, future work will need to conduct longitudinal measurement studies over longer periods after a law goes into effect, specifically including the evaluation of post-enactment events that can change legal interpretation.

Similar sounding legal rights can also result in substantively different legal realities in different jurisdictions due to differing interpretation and case law. For example, many national privacy laws provide special protection for “sensitive personal data”, but Singapore’s PDPA provides no statutory definition or protection for sensitive data over other types of data [234]. However, guidance from Singapore’s Personal Data Protection Commission indicates that sensitivity of data is a factor for consideration, thus introducing the potential for sensitive data protections under a law which otherwise lacks such protections [190]. These differences can result in results that don’t generalize between different jurisdictions even if they have similar legal aspects and regulatory language. For example, consider an illustrative example from the research on dark patterns and nudging: two studies investigated the effects of the on-page location of consent banners on interactions and user

consent decisions [275], [32]. A study with Europeans found that participants were three times more likely to interact with banners in the lower-left than in other positions; a study with a predominantly-American population found that banner position had no effect on interaction rates [32]. While there are many possible reasons for these inconsistent results (including differences in timing, recruitment, and methodology), one possibility is that a genuine or psychological difference exists between these two populations, which may have been induced by cultural factors, differing experiences with privacy and consent, or the impact that the GDPR (and the changes to corporate behavior it triggered, e.g., the frequency of cookie banners) might have had on Europeans’ behavior. To provide meaningful results about the impact of privacy regulations globally, researchers should conduct cross-jurisdictional studies and should replicate results from single-jurisdictional work to validate whether the results generalize to other legal jurisdictions.

Recommendation 4: *Computer science researchers should explore how technical expertise and methodology might be applied to evaluate proposed future regulations and regulatory approaches in addition to laws currently in effect.*

Existing work provides critiques of extant laws. This can (and has) impacted subsequent interpretation and enforcement. However, the impact of such work could be amplified if computer scientists are able to develop validated techniques for empirically evaluating proposed regulations or regulatory regimes. Certain upcoming and proposed laws contain absolute clauses beyond the scope of existing privacy guarantees, for example the District of Columbia’s Stop Discrimination by Algorithms Act of 2021 and Section 207 (“Civil Rights and Algorithms”) of the previously proposed US omnibus privacy legislation, the American Data Privacy and Protection Act (ADPPA). Such laws are beginning to address a key facet of privacy regulation, which is the ways in which it can be made to enforce equity in privacy across axes of marginalization and class. Studies of these aspects of these laws would be valuable for understanding how non-discrimination approaches impact privacy outcomes generally as well as on privacy equity.

Finally, we encourage researchers to push in the direction of studying privacy regulation approaches which have only been theorized in academic law literature. Such approaches would often be deeply transformative or rely on major shifts in not only laws and regulations but also the ways that societies construct and value privacy. Paradigms that might be rich for this type of study include data fiduciary approaches [23]—under which companies collecting personal data would have fiduciary duties toward data subjects, including confidentiality, care, and loyalty, and companies would be obligated to act in the best interests of data subjects—and civil rights approaches [31]—under which companies collecting personal data would be subject to an array of data protection laws that incorporate explicit legal safeguards against direct discrimination as well as indirect harms (i.e. disparate impact) to marginalized communities.

Acknowledgements

This work was supported by in part by NSF grant 2317115 and by internal funds from Pomona College.

References

- [1] Abdulaziz Aborujilah, Abdulaleem Z Al-Othmani, Zalizah Awang Long, Nur Syahela Hussien, and Dahlan Abdul Ghani. Conceptual model for automating GDPR compliance verification using natural language approach. In *2022 International Conference on Intelligent Technology, System and Service for Internet of Everything (ITSS-IoE)*, pages 1–6, 2022.
- [2] Supriya Adhatara, Cédric Lauradoux, and Cristiana Santos. Why IP-based subject access requests are denied? *arXiv preprint arXiv:2103.01019*, 2021.
- [3] Andrick Adhikari, Sanchari Das, and Rinku Dewri. Evolution of composition, readability, and structure of privacy policies over two decades. *Proceedings on Privacy Enhancing Technologies*, 3:138–153, 2023.
- [4] Archita Agarwal, Marilyn George, Aaron Jeyaraj, and Malte Schwarzkopf. Retrofitting GDPR compliance onto legacy databases. *Proceedings of the VLDB Endowment*, 15(4):958–970, 2021.
- [5] Abdulrahman Alhazmi and Nalin Asanka Gamagedara Arachchilage. Why are developers struggling to put GDPR into practice when developing privacy-preserving software systems? *arXiv preprint arXiv:2008.02987*, 2020.
- [6] Suzan Ali, Mounir Elgharabawy, Quentin Duchaussoy, Mohammad Mannan, and Amr Youssef. Betrayed by the guardian: Security and privacy risks of parental control solutions. In *Annual Computer Security Applications Conference*, pages 69–83, 2020.
- [7] Fatemeh Alizadeh, Timo Jakobi, Alexander Boden, Gunnar Stevens, and Jens Boldt. GDPR reality check-claiming and investigating personally identifiable data from companies. In *5th European Workshop on Usable Security*, 2020.
- [8] Sami Alkhatib, Jenny Waycott, George Buchanan, Marthie Grobler, and Shuo Wang. Privacy by design in aged care monitoring devices? Well, not quite yet! In *32nd Australian Conference on Human-Computer Interaction*, pages 492–505, 2020.
- [9] Sahar Allegue, Mouna Rhahla, and Takoua Abdellatif. Toward GDPR compliance in IoT systems. In *Service-Oriented Computing—ICSO 2019 Workshops: WESOACS, ASOCA, ISYCC, TBCE, and STRAPS, Toulouse, France, October 28–31, 2019, Revised Selected Papers 17*, pages 130–141, 2020.
- [10] Noura Alomar and Serge Egelman. Developers say the darnedest things: Privacy compliance processes followed by developers of child-directed apps. *Proceedings on Privacy Enhancing Technologies*, 4:250–273, 2022.
- [11] Orlando Amaral, Muhammad Ilyas Azeem, Sallam Abualhaija, and Lionel C Briand. NLP-based automated compliance checking of data processing agreements against GDPR. *IEEE Transactions on Software Engineering*, 2023.
- [12] David Restrepo Amariles, Aurore Clément Troussel, and Rajaa El Hamdani. Compliance generation for privacy documents under GDPR: A roadmap for implementing automation and machine learning. *arXiv preprint arXiv:2012.12718*, 2020.
- [13] Ryan Amos, Gunes Acar, Elena Lucherini, Mihir Kshirsagar, Arvind Narayanan, and Jonathan Mayer. Privacy policies over time: Curation and analysis of a million-document dataset. In *The Web Conference*, pages 2165–2176, 2021.
- [14] Benjamin Andow, Samin Yaseer Mahmud, Wenyu Wang, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Tao Xie. PolicyLint: Investigating internal privacy policy contradictions on Google Play. In *28th USENIX Security Symposium*, pages 585–602, 2019.
- [15] Benjamin Andow, Samin Yaseer Mahmud, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Serge Egelman. Actions speak louder than words: Entity-sensitive privacy policy and data flow analysis with PoliCheck. In *29th USENIX Security Symposium*, pages 985–1002, 2020.
- [16] McKane Andrus, Elena Spitzer, Jeffrey Brown, and Alice Xiang. What we can’t measure, we can’t understand: Challenges to demographic data procurement in the pursuit of fairness. In *ACM Conference on Fairness, Accountability, and Transparency*, pages 249–260, 2021.
- [17] App privacy policy generator. <https://app-privacy-policy-generator.nisrulz.com/>.
- [18] Noah Apthorpe, Sarah Varghese, and Nick Feamster. Evaluating the contextual integrity of privacy regulation: Parents’ IoT toy privacy norms versus COPPA. In *28th USENIX Security Symposium*, pages 123–140, 2019.
- [19] Emma Arfelt, David Basin, and Søren Debois. Monitoring the GDPR. In *24th European Symposium on Research in Computer Security*, pages 681–699, 2019.
- [20] Patricia Arias-Cabarcos, Saina Khalili, and Thorsten Strufe. ‘Surprised, shocked, worried’: User reactions to Facebook data collection from third parties. *Proceedings on Privacy Enhancing Technologies*, 1:384–399, 2023.
- [21] Siddhant Arora, Henry Hosseini, Christine Utz, Vinayshekhar K. Bannihatti, Tristan Dhellemmes, Abhilasha Ravichander, Peter Story, Jasmine Mangat, Rex Chen, Martin Degeling, et al. A tale of two regulatory regimes: Creation and analysis of a bilingual privacy policy corpus. In *LREC Proceedings*, 2022.
- [22] Paritosh Bahirat, Martijn Willemsen, Yangyang He, Qizhang Sun, and Bart Knijnenburg. Overlooking context: How do defaults and framing reduce deliberation in smart home privacy decision-making? In *CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2021.
- [23] Jack M Balkin. The fiduciary model of privacy. *Harv. L. Rev. F.*, 134:11–33, 2020.
- [24] Vinayshekhar Bannihatti Kumar, Roger Iyengar, Namita Nisal, Yuanyuan Feng, Hana Habib, Peter Story, Sushain Cherivirala, Margaret Hagan, Lorrie Cranor, Shomir Wilson, et al. Finding a choice in a haystack: Automatic extraction of opt-out statements from privacy policy text. In *The Web Conference*, pages 1943–1954, 2020.
- [25] Adam Barth, Anupam Datta, John C Mitchell, and Helen Nissenbaum. Privacy and contextual integrity: Framework and applications. In *IEEE Symposium on Security and Privacy*, 2006.
- [26] Marlene Barth. A case study on data portability. *Datenschutz und Datensicherheit-DuD*, 45(3):190–197, 2021.
- [27] David Basin, Søren Debois, and Thomas Hildebrandt. On purpose and by necessity: Compliance under the GDPR. In *Financial Cryptography and Data Security: 22nd International Conference, FC 2018, Nieuwpoort, Curaçao, February 26–March 2, 2018, Revised Selected Papers 22*, pages 20–37, 2018.
- [28] Daniel Bastos, Fabio Giubilo, Mark Shackleton, and Fadi El-Moussa. GDPR privacy implications for the internet of things. In *4th Annual IoT Security Foundation Conference*, pages 1–8, 2018.
- [29] Kanad Basu, Suha Sabi Hussain, Ujjwal Gupta, and Ramesh Karri. COPPTCHA: COPPA tracking by checking hardware-level activity. *IEEE Transactions on Information Forensics and Security*, 15:3213–3226, 2020.

[30] Kathrin Bednar, Sarah Spiekermann, and Marc Langheinrich. Engineering privacy by design: Are engineers ready to live up to the challenge? *The Information Society*, 35(3):122–142, 2019.

[31] Alvaro M Bedoya. Privacy as civil right. *NML Rev.*, 50:301, 2020.

[32] Carlos Bermejo Fernandez, Dimitris Chatzopoulos, Dimitrios Papadopoulos, and Pan Hui. This website uses nudging: Mturk workers’ behaviour on cookie consent notices. *Human-Computer Interaction*, 5(CSCW2):1–22, 2021.

[33] Theo Bertram, Elie Bursztein, Stephanie Caro, Hubert Chao, Rutledge Chin Feman, Peter Fleischer, Albin Gustafsson, Jess Hemerly, Chris Hibbert, Luca Invernizzi, et al. Five Years of the Right to be Forgotten. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 959–972, 2019.

[34] Asia J Biega, Peter Potash, Hal Daumé, Fernando Diaz, and Michèle Finck. Operationalizing the legal principle of data minimization for personalization. In *International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 399–408, 2020.

[35] Dino Bollinger, Karel Kubicek, Carlos Cotrini, and David Basin. Automating cookie consent and GDPR violation detection. In *31st USENIX Security Symposium*, 2022.

[36] Piero A Bonatti, Sabrina Kirrane, Iliana M Petrova, and Luigi Sauro. Machine understandable policies and GDPR compliance checking. *KI-Künstliche Intelligenz*, 34:303–315, 2020.

[37] Coline Boniface, Imane Fouad, Natalia Bielova, Cédric Lauradoux, and Cristiana Santos. Security analysis of subject access request procedures. In *Annual Privacy Forum*, pages 182–209, 2019.

[38] Elijah Robert Bouma-Sims, Megan Li, Yanzi Lin, Adia Sakur-Lemessy, Alexandra Nisenoff, Ellie Young, Eleanor Birrell, Lorrie Faith Cranor, and Hana Habib. A US-UK usability evaluation of consent management platform cookie consent interface design on desktop and mobile. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–36, 2023.

[39] Alex Bowyer, Jack Holt, Josephine Go Jefferies, Rob Wilson, David Kirk, and Jan David Smeddinck. Human-GDPR interaction: Practical experiences of accessing personal data. In *CHI Conference on Human Factors in Computing Systems*, pages 1–19, 2022.

[40] Luca Bufalieri, Massimo La Morgia, Alessandro Mei, and Julinda Stefa. GDPR: When the right to access personal data becomes a threat. In *IEEE International Conference on Web Services*, pages 75–83, 2020.

[41] Duc Bui, Brian Tang, and Kang G. Shin. Do opt-outs really opt me out? In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 425–439, 2022.

[42] Duc Bui, Yuan Yao, Kang G. Shin, Jong-Min Choi, and Junbum Shin. Consistency analysis of data-usage purposes in mobile apps. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 2824–2843, 2021.

[43] José González Cabañas, Ángel Cuevas, and Rubén Cuevas. Unveiling and quantifying Facebook exploitation of sensitive personal data for advertising purposes. In *27th USENIX Security Symposium*, pages 479–495, 2018.

[44] Matteo Cagnazzo, Thorsten Holz, and Norbert Pohlmann. Gdpirated-stealing personal information on-and offline. In *European Symposium on Research in Computer Security*, pages 367–386, 2019.

[45] Daniel Prett Campagna, Altigran Soares da Silva, and Vanessa Braganholo. Achieving GDPR compliance through provenance: An extended model. In *Anais do XXXV Simpósio Brasileiro de Bancos de Dados*, pages 13–24, 2020.

[46] Claudio Carpineto, Davide Lo Re, and Giovanni Romano. Automatic assessment of website compliance to the European cookie law with CooLCheck. In *Workshop on Privacy in the Electronic Society*, pages 135–138, 2016.

[47] C Castelluccia, S Guerses, M Hansen, JH Hoepman, J van Hoboken, B Vieira, et al. Privacy and data protection in mobile applications: A study on the app development ecosystem and the technical implementation of GDPR. ENISA, the European Union Agency for Network and Information Security, 2017.

[48] Central Intelligence Agency. European Union. <https://www.cia.gov/the-world-factbook/countries/european-union/#people-and-society>, August 2022.

[49] Central Intelligence Agency. World. <https://www.cia.gov/the-world-factbook/countries/european-union/#people-and-society>, August 2022.

[50] George Chalhoub and Ivan Flechais. Data protection at a discount: Investigating the ux of data protection from user, designer, and business leader perspectives. *Proceedings of the ACM on Human-computer Interaction*, 6(CSCW2):1–36, 2022.

[51] George Chalhoub, Ivan Flechais, Norbert Nthala, and Ruba Abu-Salma. Innovation inaction or in action? The role of user experience in the security and privacy design of smart home cameras. In *16th Symposium on Usable Privacy and Security*, pages 185–204, 2020.

[52] George Chalhoub, Martin J Kraemer, Norbert Nthala, and Ivan Flechais. “It did not give me an option to decline”: A longitudinal analysis of the user experience of security and privacy in smart home products. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2021.

[53] Cheng Chang, Huaxin Li, Yichi Zhang, Suguo Du, Hui Cao, and Haojin Zhu. Automated and personalized privacy policy extraction under GDPR consideration. In *International Conference on Wireless Algorithms, Systems, and Applications*, pages 43–54, 2019.

[54] Jan Charatan and Eleanor Birrell. Two steps forward and one step back: The right to opt-out of sale under CPRA. *Proceedings on Privacy Enhancing Technologies*, 2:91–105, 2024.

[55] Rex Chen, Fei Fang, Thomas Norton, Aleecia M. McDonald, and Norman Sadeh. Fighting the fog: Evaluating the clarity of privacy disclosures in the age of CCPA. In *Workshop on Privacy in the Electronic Society*, pages 73–102, 2021.

[56] Tek Raj Chhetri, Anelia Kurteva, Rance J DeLong, Rainer Hilscher, Kai Korte, and Anna Fensel. Data protection by design tool for automated GDPR compliance verification based on semantically modeled informed consent. *Sensors*, 22(7):2763, 2022.

[57] Aloni Cohen. Attacks on deidentification’s defenses. In *31st USENIX Security Symposium*, 2022.

[58] Aloni Cohen and Kobbi Nissim. Towards formalizing the GDPR’s notion of singling out. *Proceedings of the National Academy of Sciences*, 117(15):8344–8352, 2020.

[59] Giuseppe Contissa, Koen Docter, Francesca Lagioia, Marco Lippi, Hans-W Micklitz, Przemysław Pałka, Giovanni Sartor, and Paolo Torroni. Claudette meets GDPR: Automating the evaluation of privacy policies using artificial intelligence. Available at SSRN 3208596, 2018.

[60] Kovila P.L. Coopamootoo, Maryam Mehrnezhad, and Ehsan Toreini. “I feel invaded, annoyed, anxious and i may protect myself”: Individuals’ feelings about online tracking and their protective behaviour across gender and country. In *31st USENIX Security Symposium*, 2022.

[61] Lorrie Faith Cranor, Hana Habib, Yixin Zou, Alessandro Acquisti, Joel Reidenberg, Norman Sadeh, and Florian Schaub. Design and evaluation of a usable icon and tagline to signal an opt-out of the sale of personal information as required by CCPA, 2020. Retrieved September 13, 2022.

[62] Adrian Dabrowski, Georg Merzdovnik, Johanna Ullrich, Gerald Sendera, and Edgar Weippl. Measuring cookies and web privacy in a post-GDPR world. In *International Conference on Passive and Active Network Measurement*, pages 258–270, 2019.

[63] Savino Dambra, Iskander Sanchez-Rola, Leyla Bilge, Davide Balzarotti, Yixin Zou, Alison McDonald, Julia Narakornpichit, Nicola Dell, Thomas Ristenpart, Kevin Roundy, et al. When Sally met trackers: Web tracking from the users' perspective. In *31st USENIX Security Symposium*, 2022.

[64] Javam de Castro Machado and Paulo Roberto Pessoa Amora. How can DB systems be ready for privacy regulations. In *Anais do XXXV Simpósio Brasileiro de Bancos de Dados*, pages 235–240, 2020.

[65] Martin Degeling, Christine Utz, Christopher Lentzsch, Henry Hosseini, Florian Schaub, and Thorsten Holz. We value your privacy... now take some cookies: Measuring the GDPR's impact on web privacy. In *26th Network and Distributed System Security Symposium*, 2019.

[66] John Dempsey, Gavin Sim, and Brendan Cassidy. Designing for GDPR-investigating children's understanding of privacy: A survey approach. In *Proceedings of the 32nd International BCS Human Computer Interaction Conference 32*, pages 1–13, 2018.

[67] Amol Deshpande. Sypse: Privacy-first data management through pseudonymization and partitioning. In *Conference on Innovative Data Systems Research*, 2021.

[68] Ratan Dey, Yuan Ding, and Keith W Ross. Profiling high-school students with Facebook: How online privacy laws can actually increase minors' risk. In *Internet Measurement Conference*, pages 405–416, 2013.

[69] Mariano Di Martino, Isaac Meers, Peter Quax, Ken Andries, and Wim Lamotte. Revisiting identification issues in GDPR 'right of access' policies: A technical and longitudinal analysis. *Proceedings on Privacy Enhancing Technologies*, 2022(2):95–113, 2022.

[70] Mariano Di Martino, Pieter Robyns, Winnie Weyts, Peter Quax, Wim Lamotte, and Ken Andries. Personal information leakage by abusing the GDPR 'right of access'. In *15th Symposium on Usable Privacy and Security*, pages 371–385, 2019.

[71] Nora A. Draper and Joseph Turow. The corporate cultivation of digital resignation. *New Media & Society*, 21(8):1824–1839, 2019.

[72] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. "Money makes the world go around": Identifying barriers to better privacy in children's apps from developers' perspectives. In *CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2021.

[73] Rajaa El Hamdani, Majd Mustapha, David Restrepo Amariles, Aurore Troussel, Sébastien Meeùs, and Katsiaryna Krasnashchok. A combined rule-based and machine learning approach for automated GDPR compliance checking. In *International Conference on Artificial Intelligence and Law*, pages 40–49, 2021.

[74] Lavanya Elluri and Karuna Pande Joshi. A knowledge representation of cloud data controls for EU GDPR compliance. In *IEEE World Congress on Services*, pages 45–46, 2018.

[75] Ming Fan, Le Yu, Sen Chen, Hao Zhou, Xiapu Luo, Shuyue Li, Yang Liu, Jun Liu, and Ting Liu. An empirical evaluation of GDPR compliance violations in Android mHealth apps. In *IEEE 31st International Symposium on Software Reliability Engineering*, pages 253–264, 2020.

[76] Sabrina Fang. Investigating GDPR compliance across consumer-related websites: Are businesses telling consumers the truth about data collection? <https://www.ideals.illinois.edu/items/109189>, 2018.

[77] Florian M. Farke, David G. Balash, Maximilian Golla, Markus Dürmuth, and Adam J. Aviv. Are privacy dashboards good for end users? Evaluating user perceptions and reactions to Google's My Activity. In *30th USENIX Security Symposium*, pages 483–500, 2021.

[78] Simon Farshid, Andreas Reitz, and Peter Roßbach. Design of a forgetting blockchain: A possible way to accomplish GDPR compatibility. <https://scholarspace.manoa.hawaii.edu/items/69463b3e-4203-452b-b2b5-1b02173bc828>, 2019.

[79] Álvaro Feal, Paolo Calciati, Narseo Vallina-Rodríguez, Carmela Troncoso, Alessandra Gorla, et al. Angel or devil? A privacy study of mobile parental control apps. *Proceedings on Privacy Enhancing Technologies*, 2:314–335, 2020.

[80] Yuanyuan Feng, Yaxing Yao, and Norman Sadeh. A design space for privacy choices: Towards meaningful privacy control in the internet of things. In *CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2021.

[81] Pietro Ferrara and Fausto Spoto. Static analysis for GDPR compliance. In *ITASEC*, 2018.

[82] Mafalda Ferreira, Tiago Brito, José Fragoso Santos, and Nuno Santos. RuleKeeper: GDPR-aware personal data compliance for web frameworks. In *IEEE Symposium on Security and Privacy*, pages 2817–2834, 2023.

[83] Christof Ferreira Torres and Hugo Jonker. Investigating fingerprinters and fingerprinting-alike behaviour of Android applications. In *European Symposium on Research in Computer Security*, pages 60–80, 2018.

[84] Tobias Fiebig, Seda Gürses, Carlos H Gañán, Erna Kotkamp, Fernando Kuipers, Martina Lindorfer, Menghua Prisse, and Taritha Sari. Heads in the clouds? Measuring universities' migration to public clouds: Implications for privacy & academic freedom. *Proceedings on Privacy Enhancing Technologies*, 2:117–150, 2023.

[85] Imane Fouad, Cristiana Santos, Feras Al Kassar, Natalia Bielova, and Stefano Calzavara. On compliance of cookie purposes with the purpose specification principle. In *IEEE European Symposium on Security and Privacy Workshops*, pages 326–333, 2020.

[86] Imane Fouad, Cristiana Santos, Arnaud Legout, and Natalia Bielova. My cookie is a phoenix: Detection, measurement, and lawfulness of cookie respawning with browser fingerprinting. *Proceedings on Privacy Enhancing Technologies*, 3:79–98, 2022.

[87] M da C Freitas and Miguel Mira da Silva. GDPR compliance in SMEs: There is much to be done. *Journal of Information Systems Engineering & Management*, 3(4):30, 2018.

[88] Eoghan Furey and Juanita Blue. Can I trust her? Intelligent personal assistants and GDPR. In *2019 International Symposium on Networks, Computers and Communications (ISNCC)*, pages 1–6, 2019.

[89] Sanjam Garg, Shafi Goldwasser, and Prashant Nalini Vasudevan. Formalizing data deletion in the context of the right to be forgotten. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 373–402, 2020.

[90] Armin Gerl, Nadia Bennani, Harald Kosch, and Lionel Brunie. Lpl, towards a GDPR-compliant privacy language: formal definition and usage. In *Transactions on Large-Scale Data-and Knowledge-Centered Systems XXXVII*, pages 41–80. 2018.

[91] Armin Gerl and Bianca Meier. The layered privacy language Art. 12–14 GDPR extension–privacy enhancing user interfaces. *Datenschutz und Datensicherheit-DuD*, 43(12):747–752, 2019.

[92] Matteo Giomi, Franziska Boenisch, Christoph Wehmeyer, and Borbála Tasnádi. A unified framework for quantifying privacy risk in synthetic data. *Proceedings on Privacy Enhancing Technologies*, 2:312–328, 2023.

[93] Harald Gjermundrød, Ioanna Dionysiou, and Kyriakos Costa. privacytracker: a privacy-by-design GDPR-compliant framework with verifiable data traceability controls. In *International Conference on Web Engineering*, pages 3–15, 2016.

[94] Olga Gkotsopoulou, Elisavet Charalambous, Konstantinos Limniotis, Paul Quinn, Dimitris Kavallieratos, Gohar Sargsyan, Stavros Shiaeles, and Nicholas Kolokotronis. Data protection by design for cybersecurity systems in a smart home environment. In *IEEE Conference on Network Softwarization*, pages 101–109, 2019.

[95] Jonathan Godin and Philippe Lamontagne. Deletion-compliance in the absence of privacy. In *2021 18th International Conference on Privacy, Security and Trust*, pages 1–10, 2021.

[96] Alejandra Gómez Ortega, Jacky Bourgeois, and Gerd Kortuem. What is sensitive about (sensitive) data? Characterizing sensitivity and intimacy with Google assistant users. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2023.

[97] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. The dark (patterns) side of ux design. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pages 1–14, 2018.

[98] Colin M Gray, Cristiana Santos, Natalia Bielova, Michael Toth, and Damian Clifford. Dark patterns and the legal requirements of consent banners: An interaction criticism perspective. In *CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2021.

[99] Graham Greenleaf. Now 157 countries: Twelve data privacy laws in 2021/22. *176 Privacy Laws & Business International Report 1*, 3-8, March 2022.

[100] Moritz Gruber, Christian Höfig, Maximilian Golla, Tobias Urban, and Matteo Große-Kampmann. “We may share the number of diaper changes”: A privacy and security analysis of mobile child care applications. *Proceedings on Privacy Enhancing Technologies*, 3:394–414, 2022.

[101] Elias Grünewald and Frank Pallas. Tilt: A GDPR-aligned transparency information language and toolkit for practical privacy engineering. In *ACM Conference on Fairness, Accountability, and Transparency*, pages 636–646, 2021.

[102] Elias Grünewald, Paul Wille, Frank Pallas, Maria C Borges, and Max-R Ulbricht. TIRA: An OpenAPI extension and toolbox for GDPR transparency in RESTful architectures. In *IEEE European Symposium on Security and Privacy Workshops*, pages 312–319, 2021.

[103] Nils Gruschka, Vasileios Mavroeidis, Kamer Vishi, and Meiko Jensen. Privacy issues and data protection in big data: A case study analysis under GDPR. In *IEEE International Conference on Big Data*, pages 5027–5033, 2018.

[104] Danny S. Guamán, Jose M. Del Alamo, and Julio C. Caiza. GDPR compliance assessment for cross-border personal data transfers in Android apps. *IEEE Access*, 9:15961–15982, 2021.

[105] Johanna Gunawan, Amogh Pradeep, David Choffnes, Woodrow Hartzog, and Christo Wilson. A comparative study of dark patterns across mobile and web modalities. In *Conference on Computer-Supported Cooperative Work and Social Computing*, volume 5, 2021.

[106] Zhiyu Guo, Zijin Lin, Pan Li, and Kai Chen. SkillExplorer: Understanding the behavior of skills in large scale. In *29th USENIX Security Symposium*, pages 2649–2666, 2020.

[107] Hana Habib and Lorrie Cranor. Evaluating the usability of privacy choice mechanisms. In *Symposium on Usable Privacy and Security*, 2022.

[108] Hana Habib, Megan Li, Ellie Young, and Lorrie Cranor. “Okay, whatever”: An evaluation of cookie consent interfaces. In *CHI Conference on Human Factors in Computing Systems*, pages 1–27, 2022.

[109] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. “It’s a scavenger hunt”: Usability of websites’ opt-out and data deletion choices. In *CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2020.

[110] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. An empirical analysis of data deletion and opt-out choices on 150 websites. In *15th Symposium on Usable Privacy and Security*, pages 387–406, 2019.

[111] Hana Habib, Yixin Zou, Yaxing Yao, Alessandro Acquisti, Lorrie Cranor, Joel Reidenberg, Norman Sadeh, and Florian Schaub. Toggles, dollar signs, and triangles: How to (in)effectively convey privacy choices with icons and link texts. In *CHI Conference on Human Factors in Computing Systems*, pages 1–25, 2021.

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

[113] Catherine Han, Irwin Reyes, Álvaro Feal, Joel Reardon, Primal Wijesekera, Narseo Vallina-Rodriguez, Amit Elazar, Kenneth A Bamberger, and Serge Egelman. The price is (not) right: Comparing privacy in free and paid apps. *Proceedings on Privacy Enhancing Technologies*, 2020(3), 2020.

[114] Jialiang Han, Yun Ma, Qiaozhu Mei, and Xuanzhe Liu. DeepRec: On-device deep learning for privacy-preserving sequential recommendation in mobile commerce. In *The Web Conference*, pages 900–911, 2021.

[115] Anne Hennig, Heike Dietmann, Franz Lehr, Miriam Mutter, Melanie Volkamer, and Peter Mayer. Your cookie disclaimer is not in line with the ideas of the GDPR. Why? In *International Symposium on Human Aspects of Information Security and Assurance*, pages 218–227, 2022.

[116] Dominik Herrmann and Jens Lindemann. Obtaining personal data and asking for erasure: Do app vendors and website owners honour your privacy rights? *arXiv preprint arXiv:1602.01804*, 2016.

[117] Maximilian Hils, Daniel W Woods, and Rainer Böhme. Measuring the emergence of consent management on the web. In *Internet Measurement Conference*, pages 317–332, 2020.

[118] Maximilian Hils, Daniel W Woods, and Rainer Böhme. Privacy preference signals: Past, present and future. *Proceedings on Privacy Enhancing Technologies*, 2021(4):249–269, 2021.

[119] Kalle Hjerppe, Jukka Ruohonen, and Ville Leppänen. The General Data Protection Regulation: Requirements, architectures, and constraints. In *IEEE 27th International Requirements Engineering Conference*, pages 265–275, 2019.

[120] Xuehui Hu, Guillermo Suarez de Tangil, and Nishanth Sastry. Multi-country study of third party trackers from real browser histories. In *IEEE European Symposium on Security and Privacy*, pages 70–86, 2020.

[121] Xuehui Hu and Nishanth Sastry. Characterising third party cookie usage in the EU after GDPR. In *ACM Conference on Web Science*, pages 137–141, 2019.

[122] Irene Ioannidou and Nicolas Sklavos. On General Data Protection Regulation vulnerabilities and privacy issues, for wearable devices and fitness tracking applications. *Cryptography*, 5(4):29, 2021.

[123] Costas Iordanou, Georgios Smaragdakis, Ingmar Poese, and Niko Laoutaris. Tracing cross border web tracking. In *Internet Measurement Conference 2018*, pages 329–342, 2018.

[124] Zsolt István, Soujanya Ponnappalli, and Vijay Chidambaram. Software-defined data protection: Low overhead policy compliance at the storage layer is within reach! *Proceedings of the VLDB Endowment*, 14(7):1167–1174, 2021.

[125] iubenda. <https://www.iubenda.com/en/>.

[126] Sorensen Jannick Kirk, Hilde Van den Bulck, and Sokol Kosta. Privacy policies caught between the legal and the ethical: European media and third party trackers before and after GDPR. In *TPRC47: The 47th Research Conference on Communication, Information and Internet Policy*, 2019.

[127] Qiwei Jia, Lu Zhou, Huixin Li, Ruoxu Yang, Suguo Du, and Hacjin Zhu. Who leaks my privacy: Towards automatic and association detection with GDPR compliance. In *International Conference on Wireless Algorithms, Systems, and Applications*, pages 137–148, 2019.

[128] Garrett Johnson, Scott Shriver, and Samuel Goldberg. Privacy & market concentration: Intended & unintended consequences of the GDPR. *Available at SSRN 3477686*, 2021.

[129] Scott Jordan, Yoshimichi Nakatsuka, Ercan Ozturk, Andrew Paverd, and Gene Tsudik. VICEROY: GDPR-/CCPA-compliant enforcement of verifiable accountless consumer requests. In *Network and Distributed System Security*, 2023.

[130] Florian Kammueler. Formal modeling and analysis of data protection for GDPR compliance of IoT healthcare systems. In *IEEE International Conference on Systems, Man, and Cybernetics*, pages 3319–3324, 2018.

[131] Farzane Karami, David Basin, and Einar Broch Johnsen. DPL: A language for GDPR enforcement. In *IEEE Computer Security Foundations Symposium*, pages 112–129, 2022.

[132] Smirity Kaushik, Yaxing Yao, Pierre Dewitte, and Yang Wang. “How I know for sure”: People’s perspectives on solely automated decision-making (SADM). In *17th Symposium on Usable Privacy and Security*, pages 159–180, 2021.

[133] Dilara Keküllüoglu and Yasemin Acar. “We are a startup to the core”: A qualitative interview study on the security and privacy development practices in Turkish software startups. In *IEEE Symposium on Security and Privacy*, pages 2015–2031, 2023.

[134] Rishabh Khandelwal, Asmit Nayak, Hamza Harkous, and Kassem Fawaz. Automated cookie notice analysis and enforcement. In *32nd USENIX Security Symposium*, pages 1109–1126, 2023.

[135] David Klein, Marius Musch, Thomas Barber, Moritz Kopmann, and Martin Johns. Accept all exploits: Exploring the security impact of cookie banners. In *Proceedings of the 38th Annual Computer Security Applications Conference*, pages 911–922, 2022.

[136] Simon Koch, Malte Wessels, Benjamin Altpeter, Madita Olvermann, and Martin Johns. Keeping privacy labels honest. *Proceedings on Privacy Enhancing Technologies*, 4:486–506, 2022.

[137] Konrad Kollnig, Reuben Binns, Max Van Kleek, Ulrik Lyngs, Jun Zhao, Claudine Tinsman, and Nigel Shadbolt. Before and after GDPR: Tracking in mobile apps. *Internet Policy Review*, 10(4), 2021.

[138] Konrad Kollnig, Pierre Dewitte, Max Van Kleek, Ge Wang, Daniel Omeiza, Helena Webb, and Nigel Shadbolt. A fait accompli? An empirical study into the absence of consent to third-party tracking in Android apps. In *17th Symposium on Usable Privacy and Security*, pages 181–196, 2021.

[139] Konrad Kollnig, Anastasia Shuba, Max Van Kleek, Reuben Binns, and Nigel Shadbolt. Goodbye tracking? Impact of iOS app tracking transparency and privacy labels. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 508–520, 2022.

[140] Satyapriya Krishna, Jiaqi Ma, and Himabindu Lakkaraju. Towards bridging the gaps between the right to explanation and the right to be forgotten. In *International Conference on Machine Learning*, pages 17808–17826, 2023.

[141] Jacob Leon Kröger, Jens Lindemann, and Dominik Herrmann. How do app vendors respond to subject access requests? A longitudinal privacy study on iOS and Android apps. In *International Conference on Availability, Reliability and Security*, pages 1–10, 2020.

[142] Karel Kubicek, Jakob Merane, Carlos Cotrini, Alexander Stremitzer, Stefan Bechtold, and David Basin. Checking websites’ GDPR consent compliance for marketing emails. *Proceedings on Privacy Enhancing Technologies*, 2022(2):282–303, 2022.

[143] Sophie Kuebler-Wachendorff, Robert Luzsa, Johann Kranz, Stefan Mager, Emmanuel Syrmoudis, Susanne Mayr, and Jens Grossklags. The right to data portability: Conception, status quo, and future directions. *Informatik Spektrum*, 44(4):264–272, 2021.

[144] Oksana Kulyk, Nina Gerber, Annika Hilt, and Melanie Volkamer. Has the GDPR hype affected users’ reaction to cookie disclaimers? *Journal of Cybersecurity*, 6(1):1–14, 2020.

[145] Renuka Kumar, Apurva Virkud, Ram Sundara Raman, Atul Prakash, and Roya Ensafi. A large-scale investigation into geodifferences in mobile apps. In *31st USENIX Security Symposium*, 2022.

[146] Mirosław Kutyłowski, Anna Lauks-Dutka, and Moti Yung. GDPR–Challenges for reconciling legal rules with technical reality. In *25th European Symposium on Research in Computer Security*, pages 736–755, 2020.

[147] Lin Kyi, Sushil Ammanaghata Shivakumar, Cristiana Teixeira Santos, Franziska Roesner, Frederike Zufall, and Asia J. Biega. Investigating deceptive design in GDPR’s legitimate interest. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2023.

[148] Clément Labadie and Christine Legner. Understanding data protection regulations from a data management perspective: A capability-based approach to EU-GDPR. In *14th International Conference on Wirtschaftsinformatik*, 2019.

[149] Tianshi Li, Elizabeth Louie, Laura Dabbish, and Jason Hong. How developers talk about personal data and what it means for user privacy: A case study of a developer forum on Reddit. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW):1–28, 2021.

[150] Song Liao, Christin Wilson, Long Cheng, Hongxin Hu, and Huixing Deng. Measuring the effectiveness of privacy policies for voice assistant applications. In *Annual Computer Security Applications Conference*, pages 856–869, 2020.

[151] Timothy Libert, Lucas Graves, and Rasmus Kleis Nielsen. Changes in third-party content on European news websites after GDPR. <https://ora.ox.ac.uk/objects/uuid:5a5d4eea-6e74-49b4-8c77-71ec6760f127>, 2018.

[152] Ilaria Liccardi, Monica Bulger, Hal Abelson, Daniel J. Weitzner, and Wendy Mackay. Can apps play by the COPPA rules? In *12th Annual International Conference on Privacy, Security and Trust*, 2014.

[153] Ruta Liepina, Giuseppe Contissa, Kasper Drazewski, Francesca Lagioia, Marco Lippi, H-W Micklitz, Przemyslaw Palka, Giovanni Sartor, and Paolo Torroni. GDPR privacy policies in Claudette: Challenges of omission, context and multilingualism. In *3rd Workshop on Automated Semantic Analysis of Information in Legal Texts*, volume 2385, 2019.

[154] Thomas Linden, Rishabh Khandelwal, Hamza Harkous, and Kassem Fawaz. The privacy policy landscape after the GDPR. *Proceedings on Privacy Enhancing Technologies*, 2020(1):47–64, 2020.

[155] Yuxi Ling, Kailong Wang, Guangdong Bai, Haoyu Wang, and Jin Song Dong. Are they toeing the line? Diagnosing privacy compliance violations among browser extensions. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, pages 1–12, 2022.

[156] Shuang Liu, Baiyang Zhao, Renjie Guo, Guozhu Meng, Fan Zhang, and Meishan Zhang. Have you been properly notified? Automatic compliance analysis of privacy policy text with GDPR Article 13. In *The Web Conference*, pages 2154–2164, 2021.

[157] Isabel Maria Lopes, Teresa Guarda, and Pedro Oliveira. How ISO 27001 can help achieve GDPR compliance. In *2019 14th Iberian Conference on Information Systems and Technologies*, pages 1–6, 2019.

[158] Juniper Lovato, Philip Mueller, Parisa Suchdev, and Peter Dodds. More data types more problems: A temporal analysis of complexity, stability, and sensitivity in privacy policies. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 1088–1100, 2023.

[159] Chaoyi Lu, Baojun Liu, Yiming Zhang, Zhou Li, Fenglu Zhang, Haixin Duan, Ying Liu, Joann Qiongna Chen, Jinjin Liang, Zai Feng Zhang, et al. From WHOIS to WHOVAS: A Large-Scale Measurement Study of Domain Registration Privacy under the GDPR. In *28th Network and Distributed System Security Symposium*, 2021.

[160] Connor Luckett, Andrew Crotty, Alex Galakatos, and Ugur Cetintemel. Odlaw: A tool for retroactive GDPR compliance. In *IEEE International Conference on Data Engineering*, pages 2709–2712, 2021.

[161] Eryn Ma and Eleanor Birrell. Prospective consent: The effect of framing on cookie consent decisions. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–6, 2022.

[162] Dominique Machuletz and Rainer Böhme. Multiple purposes, multiple problems: A user study of consent dialogs after GDPR. *Proceedings on Privacy Enhancing Technologies*, 2:481–498, 2020.

[163] Sunil Manandhar, Kaushal Kafle, Benjamin Andow, Kapil Singh, and Adwait Nadkarni. Smart home privacy policies demystified: A study of availability, content, and coverage. In *31st USENIX Security Symposium*, pages 3521–3538, 2022.

[164] Anna Maria Mandalari, Daniel J. Dubois, Roman Kolcun, Muhammad Talha Paracha, Hamed Haddadi, and David Choffnes. Blocking without breaking: Identification and mitigation of non-essential IoT traffic. *Proceedings on Privacy Enhancing Technologies*, 4:369–388, 2021.

[165] Tina Marjanov, Maria Konstantinou, Magdalena Józwiak, and Dayana Spagnuelo. Data security on the ground: Investigating technical and legal requirements under the GDPR. *Proceedings on Privacy Enhancing Technologies*, 3:405–417, 2023.

[166] Yod-Samuel Martin and Antonio Kung. Methods and tools for GDPR compliance through privacy and data protection engineering. In *IEEE European Symposium on Security and Privacy Workshops*, pages 108–111, 2018.

[167] Arunesh Mathur, Gunes Acar, Michael J. Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. Dark patterns at scale: Findings from a crawl of 11k shopping websites. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–32, 2019.

[168] Arunesh Mathur, Mihir Kshirsagar, and Jonathan Mayer. What makes a dark pattern... dark? Design attributes, normative considerations, and measurement methods. In *CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2021.

[169] Célestin Matte, Natalia Bielova, and Cristiana Santos. Do cookie banners respect my choice? Measuring legal compliance of banners from IAB Europe’s transparency and consent framework. In *IEEE Symposium on Security and Privacy*, pages 791–809, 2020.

[170] Aleecia M. McDonald and Lorrie Faith Cranor. The cost of reading privacy policies. *IJS: A Journal of Law and Policy for the Information Society*, 4:540–565, 2008.

[171] Maryam Mehrnezhad, Kovila Coopamootoo, and Ehsan Toreini. How can and would people protect from online tracking? *Proceedings on Privacy Enhancing Technologies*, 1:105–125, 2022.

[172] Mark Huasong Meng, Qing Zhang, Guangshuai Xia, Yuwei Zheng, Yanjun Zhang, Guangdong Bai, Zhi Liu, Sin G Teo, and Jin Song Dong. Post-GDPR threat hunting on android phones: Dissecting OS-level safeguards of user-unresettable identifiers. In *Network and Distributed System Security Symposium*, 2023.

[173] Abraham Mhaidli, Selin Fidan, An Doan, Gina Herakovic, Mukund Srinath, Lee Matheson, Shomir Wilson, and Florian Schaub. Researchers’ experiences in analyzing privacy policies: Challenges and opportunities. *Proceedings on Privacy Enhancing Technologies*, 4:287–305, 2023.

[174] Mohsen Minaei, Mainack Mondal, and Aniket Kate. Empirical understanding of deletion privacy: Experiences, expectations, and measures. In *31st USENIX Security Symposium*, 2022.

[175] Reham Ebada Mohamed and Sonia Chiasson. Online privacy and aging of digital artifacts. In *14th Symposium on Usable Privacy and Security*, pages 177–195, 2018.

[176] Jayashree Mohan, Melissa Wasserman, and Vijay Chidambaram. Analyzing GDPR compliance through the lens of privacy policy. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB 2019 Workshops, Poly and DMAH, Los Angeles, CA, USA, August 30, 2019, Revised Selected Papers 5*, pages 82–95, 2019.

[177] Nurul Momen, Majid Hatamian, and Lothar Fritsch. Did app privacy improve after the GDPR? *IEEE Security & Privacy*, 17(6):10–20, 2019.

[178] Charlotte Moremen, Jordan Hoogsteden, and Eleanor Birrell. Generational differences in understandings of privacy terminology. *Proceedings on Privacy Enhancing Technologies*, 2024.

[179] Najmeh Mousavi Nejad, Simon Scerri, and Jens Lehmann. Knight: Mapping privacy policies to GDPR. In *European Knowledge Acquisition Workshop*, pages 258–272, 2018.

[180] Maaz Bin Musa and Rishab Nithyanand. ATOM: Ad-network to-mography. *Proceedings on Privacy Enhancing Technologies*, 4:295–313, 2022.

[181] Majd Mustapha, Katsiaryna Krasnashchok, Anas Al Bassit, and Sabri Skhiri. Privacy policy classification with xlnet (short paper). In *International Workshop on Data Privacy Management*, pages 250–257, 2020.

[182] Yuhong Nan, Xueqiang Wang, Luyi Xing, Xiaojing Liao, Ruoyu Wu, Jianliang Wu, Yifan Zhang, and XiaoFeng Wang. Are you spying on me? Large-Scale analysis on IoT data exposure through companion apps. In *32nd USENIX Security Symposium*, pages 6665–6682, 2023.

[183] Arvind Narayanan and Vitaly Shmatikov. Myths and fallacies of “personally identifiable information”. *Communications of the ACM*, 53(6):24–26, 2010.

[184] Trung Tin Nguyen, Michael Backes, Ninja Marnau, and Ben Stock. Share first, ask later (or never?) Studying violations of GDPR’s explicit consent in Android apps. In *30th USENIX Security Symposium*, 2021.

[185] Trung Tin Nguyen, Michael Backes, and Ben Stock. Freely given consent? studying consent notice of third-party tracking and its violations of GDPR in android apps. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, pages 2369–2383, 2022.

[186] Chris Norval, Kristin Cornelius, Jennifer Cobbe, and Jatinder Singh. Disclosure by design: Designing information disclosures to support meaningful transparency and accountability. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 679–690, 2022.

[187] Midas Nouwens, Rolf Bagge, Janus Bager Kristensen, and Clemens Nylandsted Klokmos. Consent-o-matic: Automatically answering consent pop-ups using adversarial interoperability. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–7, 2022.

[188] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In *CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2020.

[189] Sean O’Connor, Ryan Nurwono, Aden Siebel, and Eleanor Birrell. (Un)clear and (In)conspicuous: The right to opt-out of sale under CCPA. In *Workshop on Privacy in the Electronic Society*, pages 59–72, 2021.

[190] Personal Data Protection Commission of Singapore. Advisory guidelines on key concepts in the personal data protection act. <https://www.pdpc.gov.sg/guidelines-and-consultation/2020/03/advisory-guidelines-on-key-concepts-in-the-personal-data-protection-act>.

[191] Junhyoung Oh, Jinhyoung Hong, Changsoo Lee, Jemin Justin Lee, Simon S Woo, and Kyungho Lee. Will EU’s GDPR act as an effective enforcer to gain consent? *IEEE Access*, 9:79477–79490, 2021.

[192] Haoran Ou, Yong Fang, Yongyan Guo, Wenbo Guo, and Cheng Huang. Viopolicy-detector: An automated approach to detecting GDPR suspected compliance violations in websites. In *Proceedings of the 25th International Symposium on Research in Attacks, Intrusions and Defenses*, pages 409–430, 2022.

[193] Kentrell Owens, Anita Alem, Franziska Roesner, and Tadayoshi Kohno. Electronic monitoring smartphone apps: An analysis of risks from technical, human-centered, and legal perspectives. In *31st USENIX Security Symposium*, 2022.

[194] Harshvardhan J. Pandit, Declan O’Sullivan, and Dave Lewis. Queryable provenance metadata for GDPR compliance. *Procedia Computer Science*, 137:262–268, 2018.

[195] Emmanouil Papadogiannakis, Panagiotis Papadopoulos, Nicolas Kourtellis, and Evangelos P Markatos. User tracking in the post-cookie era: How websites bypass GDPR consent to track users. In *The Web Conference*, pages 2130–2141, 2021.

[196] James Pavur, Casey Knerr, and LTD Dionach. GDPArrrrr: Using privacy laws to steal identities. In *Blackhat*, 2019.

[197] Justin Petelka, Elisa Oreglia, Megan Finn, and Janaki Srinivasan. Generating practices: Investigations into the double embedding of GDPR and data access policies. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–26, 2022.

[198] Luca Piras, Mohammed Ghazi Al-Obeidallah, Andrea Pratano, Aggeliki Tsouhou, Haralampos Mouratidis, Beatriz Gallego-Nicasio Crespo, Jean Baptiste Bernard, Marco Fiorani, Emmanouil Magkos, Andres Castillo Sanz, et al. DEFeND architecture: A privacy by design platform for GDPR compliance. In *International Conference on Trust and Privacy in Digital Business*, pages 78–93, 2019.

[199] Eugenia Politou, Eftimios Alepis, and Constantinos Patsakis. Forgetting personal data and revoking consent under the GDPR: Challenges and proposed solutions. *Journal of Cybersecurity*, 4(1):tyy001, 2018.

[200] Ayesha Qamar, Tehreem Javed, and Mirza Omer Beg. Detecting compliance of privacy policies with data protection laws. *arXiv preprint arXiv:2102.12362*, 2021.

[201] Tamjid Al Rahat, Tu Le, and Yuan Tian. Automated detection of GDPR disclosure requirements in privacy policies using deep active learning. *arXiv preprint arXiv:2111.04224*, 2021.

[202] Ali Rasaii, Shivani Singh, Devashish Gosain, and Oliver Gasser. Exploring the cookieverse: A multi-perspective analysis of web cookies. In *International Conference on Passive and Active Network Measurement*, pages 623–651, 2023.

[203] Philip Raschke, Axel Küpper, Olha Drozd, and Sabrina Kirrane. Designing a GDPR-compliant and usable privacy dashboard. In *IFIP International Summer School on Privacy and Identity Management*, pages 221–236, 2017.

[204] Abbas Razaghpanah, Rishab Nithyanand, Narseo Vallina-Rodriguez, Srikanth Sundaresan, Mark Allman, Christian Kreibich, Phillipa Gill, et al. Apps, trackers, privacy, and regulators: A global study of the mobile tracking ecosystem. In *The 25th Annual Network and Distributed System Security Symposium*, 2018.

[205] Ronak Razavisousan and Karuna P Joshi. Analyzing GDPR compliance in cloud services’ privacy policies using textual fuzzy interpretive structural modeling (tfism). In *IEEE International Conference on Services Computing*, pages 89–98, 2021.

[206] Joel Reardon, Álvaro Feal, Primal Wijesekera, Amit Elazari Bar On, Narseo Vallina-Rodriguez, and Serge Egelman. 50 ways to leak your data: An exploration of apps’ circumvention of the android permissions system. In *28th USENIX Security Symposium*, 2019.

[207] Jingjing Ren, Daniel J Dubois, David Choffnes, Anna Maria Mandalaris, Roman Kolcun, and Hamed Haddadi. Information exposure from consumer IoT devices: A multidimensional, network-informed measurement approach. In *Proceedings of the Internet Measurement Conference*, pages 267–279, 2019.

[208] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, Serge Egelman, et al. “Won’t somebody think of the children?” examining COPPA compliance at scale. In *The 18th Privacy Enhancing Technologies Symposium*, 2018.

[209] Gerard Reynolds and Seamus Dowling. An analysis of Ireland’s homecare companies’ cookie practices in terms of GDPR compliance. In *2022 Cyber Research Conference-Ireland*, pages 1–7, 2022.

[210] Neil Richards and Woodrow Hartzog. The pathologies of digital consent. *Wash. UL Rev.*, 96:1461, 2018.

[211] Paul Ryan, Rob Brennan, and Harshvardhan J Pandit. DPCat: Specification for an interoperable and machine-readable data processing catalogue based on GDPR. *Information*, 13(5):244, 2022.

[212] Rahime Belen Sağlam, Çağrı Burak Aslan, Shujun Li, Lisa Dickson, and Ganna Pogrebna. A data-driven analysis of blockchain systems’ public online communications on GDPR. In *IEEE International Conference on Decentralized Applications and Infrastructures*, pages 22–31, 2020.

[213] Nayanamana Samarasinghe, Aashish Adhikari, Mohammad Mannan, and Amr Youssef. Et tu, Brute? Privacy analysis of government websites and mobile apps. In *Proceedings of the ACM Web Conference*, pages 564–575, 2022.

[214] Nikita Samarin, Shayna Kothari, Zaina Siyed, Oscar Bjorkman, Reena Yuan, Primal Wijesekera, Noura Alomar, Jordan Fischer, Chris Hoofnagle, and Serge Egelman. Lessons in VCR repair: Compliance of Android app developers with the California Consumer Privacy Act (CCPA). *Proceedings on Privacy Enhancing Technologies*, 3:103–121, 2023.

[215] Iskander Sanchez-Rola, Matteo Dell’Amico, Platon Kotzias, Davide Balzarotti, Leyla Bilge, Pierre-Antoine Vervier, and Igor Santos. Can I opt out yet? GDPR and the global illusion of cookie control. In *Asia Conference on Computer and Communications Security*, pages 340–351, 2019.

[216] Iskander Sanchez-Rola, Matteo Dell’Amico, Davide Balzarotti, Pierre-Antoine Vervier, and Leyla Bilge. Journey to the center of the cookie ecosystem: Unraveling actors’ roles and relationships. In *IEEE Symposium on Security and Privacy*, 2021.

[217] Cristiana Santos, Natalia Bielova, and Célestin Matte. Are cookie banners indeed compliant with the law? Deciphering EU legal requirements on consent and technical means to verify compliance of cookie banners. *Technology and Regulation*, pages 91–135, 2020.

[218] Cristiana Santos, Midas Nouwens, Michael Toth, Natalia Bielova, and Vincent Roca. Consent management platforms under the GDPR: Processors and/or controllers? In *Annual Privacy Forum*, pages 47–69, 2021.

[219] Cristiana Santos, Arianna Rossi, Lorena Sanchez Chamorro, Kerstin Bongard-Blanchy, and Ruba Abu-Salma. Cookie banners, what’s the purpose? Analyzing cookie banner text through a legal lens. In *Workshop on Privacy in the Electronic Society*, pages 187–194, 2021.

[220] Subhadeep Sarkar, Jean-Pierre Banatre, Louis Rilling, and Christine Morin. Towards enforcement of the EU GDPR: Enabling data erasure. In *IEEE International Conference on Internet of Things and IEEE Green Computing and Communications and IEEE Cyber, Physical and Social Computing and IEEE Smart Data*, pages 222–229, 2018.

[221] Brennan Schaffner, Neha A. Lingareddy, and Marshini Chetty. Understanding account deletion and relevant dark patterns on social media. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–43, 2022.

[222] Malte Schwarzkopf, Eddie Kohler, M. Frans Kaashoek, and Robert Morris. Position: GDPR compliance by construction. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB 2019 Workshops, Poly and DMAH, Los Angeles, CA, USA, August 30, 2019, Revised Selected Papers 5*, pages 39–53, 2019.

[223] Nick Scope, Alexander Rasin, Ben Lenard, Karen Heart, and James Wagner. Harmonizing privacy regarding data retention and purging. In *Proceedings of the 34th International Conference on Scientific and Statistical Database Management*, pages 1–12, 2022.

[224] Awanthika Senarath and Nalin Asanka Gamagedara Arachchilage. Understanding software developers’ approach towards implementing data minimization. *arXiv preprint arXiv:1808.01479*, 2018.

[225] William Seymour, Mark Coté, and Jose Such. Legal obligation and ethical best practice: Towards meaningful verbal consent for voice assistants. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–16, 2023.

[226] Aashaka Shah, Vinay Banakar, Supreeth Shastri, Melissa Wasserman, and Vijay Chidambaram. Analyzing the impact of GDPR on storage systems. In *11th USENIX Workshop on Hot Topics in Storage and File Systems*, 2019.

[227] Divya Shanmugam, Fernando Diaz, Samira Shabanian, Michele Finck, and Asia Biega. Learning to limit data collection via scaling laws: A computational interpretation for the legal principle of data minimization. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 839–849, 2022.

[228] Supreeth Shastri, Vinay Banakar, Melissa Wasserman, Arun Kumar, and Vijay Chidambaram. Understanding and benchmarking the impact of GDPR on database systems. *Proceedings of the VLDB Endowment*, 13(7).

[229] Supreeth Shastri, Melissa Wasserman, and Vijay Chidambaram. The seven sins of personal-data processing systems under GDPR. In *11th USENIX Workshop on Hot Topics in Cloud Computing*, 2019.

[230] Swapneel Sheth, Gail Kaiser, and Walid Maalej. Us and them: a study of privacy requirements across North America, Asia, and Europe. In *36th International Conference on Software Engineering*, 2014.

[231] Faysal Hossain Shezan, Yingjie Lao, Minlong Peng, Xin Wang, Mingming Sun, and Ping Li. NL2GDPR: Automatically develop GDPR compliant Android application features from natural language. In *IEEE Conference on Communications and Network Security*, pages 1–9, 2022.

[232] Faysal Hossain Shezan, Zihao Su, Mingqing Kang, Nicholas Phair, Patrick William Thomas, Michelangelo van Dam, Yinzhi Cao, and Yuan Tian. CHKPLUG: Checking GDPR compliance of WordPress plugins via cross-language code property graph. In *Network and Distributed System Security Symposium*, 2023.

[233] Aden Siebel and Eleanor Birrell. The impact of visibility on the right to opt-out of sale under CCPA. *arXiv preprint arXiv:2206.10545*, 2022.

[234] The Statutes Of The Republic Of Singapore. Personal Data Protection Act 2012, 2020 revised edition. <https://sso.agc.gov.sg/Act/PDPA2012>.

[235] Ram Govind Singh and Sushmita Ruj. A technical look at the Indian personal data protection bill. *arXiv preprint arXiv:2005.13812*, 2020.

[236] Sean Sirur, Jason RC Nurse, and Helena Webb. Are we there yet? Understanding the challenges faced in complying with the General Data Protection Regulation (GDPR). In *Proceedings of the 2nd International Workshop on Multimedia Privacy and Security*, pages 88–95, 2018.

[237] Daniel Solove. Beyond GDPR: The challenge of global privacy compliance—an interview with Lothar Determann. <https://teachprivacy.com/challenge-of-global-privacy-compliance/>, [<https://perma.cc/4956-Q6TK>], 11 2017.

[238] Daniel J Solove. Introduction: Privacy self-management and the consent dilemma. *Harv. L. Rev.*, 126:1880, 2012.

[239] Daniel J Solove. The myth of the privacy paradox. *Geo. Wash. L. Rev.*, 89:1, 2021.

[240] Jannick Sørensen and Sokol Kosta. Before and after GDPR: The changes in third party presence at public and private European websites. In *The World Wide Web Conference*, pages 1590–1600, 2019.

[241] Wissem Soussi, Maciej Korczynski, Sourena Maroofi, and Andrzej Duda. Feasibility of large-scale vulnerability notifications after gdpr. In *IEEE European Symposium on Security and Privacy Workshops*, pages 532–537, 2020.

[242] Alina Stöver, Nina Gerber, Christin Cornel, Mona Henz, Karola Marky, Verena Zimmermann, and Joachim Vogt. Website operators are not the enemy either-analyzing options for creating cookie consent notices without dark patterns. *Mensch und Computer 2022-Workshopband*, 2022.

[243] Alina Stöver, Nina Gerber, Henning Pridöhl, Max Maass, Sebastian Bretthauer, I Spiecker, M Hollick, and D Herrmann. How website owners face privacy issues: Thematic analysis of responses from a covert notification study reveals diverse circumstances and challenges. *Proceedings on Privacy Enhancing Technologies*, 2023.

[244] Alanoud Subahi and George Theodorakopoulos. Ensuring compliance of iot devices with their privacy policy agreement. In *IEEE 6th International Conference on Future Internet of Things and Cloud*, pages 100–107, 2018.

[245] Kaiwen Sun, Carlo Sugatan, Tanisha Afnan, Hayley Simon, Susan A Gelman, Jenny Radesky, and Florian Schaub. “They see you’re a girl if you pick a pink robot with a skirt”: A qualitative study of how children conceptualize data processing and digital privacy risks. In *CHI Conference on Human Factors in Computing Systems*, pages 1–34, 2021.

[246] Emmanuel Sirmoudis, Stefan Mager, Sophie Kuebler-Wachendorff, Paul Pizzinini, Jens Grossklags, and Johann Kranz. Data portability between online services: an empirical analysis on the effectiveness of GDPR Art. 20. *Proceedings on Privacy Enhancing Technologies*, 2021(3):351–372, 2021.

[247] Carlotta Tagliaro, Florian Hahn, Riccardo Sepe, Alessio Aceti, and Martina Lindorfer. I still know what you watched last sunday: Privacy of the HbbTV protocol in the European smart tv landscape. In *Network and Distributed System Security Symposium*, 2023.

[248] Mohammad Tahaei, Alisa Frik, and Kami Vaniea. Deciding on personalized ads: Nudging developers about user privacy. In *17th Symposium on Usable Privacy and Security*, pages 573–596, 2021.

[249] Mohammad Tahaei, Alisa Frik, and Kami Vaniea. Privacy champions in software teams: Understanding their motivations, strategies, and challenges. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2021.

[250] Mohammad Tahaei, Tianshi Li, and Kami Vaniea. Understanding privacy-related advice on Stack Overflow. *Proceedings on Privacy Enhancing Technologies*, 1:18, 2022.

[251] Mohammad Tahaei, Kopo M Ramokapane, Tianshi Li, Jason I Hong, and Awais Rashid. Charting app developers’ journey through privacy regulation features in ad networks. *Proceedings on Privacy Enhancing Technologies*, 1:24, 2022.

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

[253] Keiji Take, Kevin Gallagher, Andrea Forte, Damon McCoy, and Rachel Greenstadt. “It feels like whack-a-mole”: User experiences of data removal from people search websites. *Proceedings on Privacy Enhancing Technologies*, 1:20, 2022.

[254] Aurelia Tamò-Larrieux, Zaira Zihlmann, Kimberly Garcia, and Simon Mayer. The right to customization: Conceptualizing the right to repair for informational privacy. In *Annual Privacy Forum*, pages 3–22, 2021.

[255] Alain Tchana, Raphael Colin, Adrien Le Berre, Vincent Berger, Benoît Combemale, and Ludovic Pailler. rgpdOS: GDPR enforcement by the operating system. In *53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks—Supplemental Volume (DSN-S)*, pages 100–104, 2023.

[256] Termly. <https://termly.io/products/privacy-policy-generator/>.

[257] TermsFeed. <https://www.termsfeed.com/>.

[258] Welderufael B. Tesfay, Peter Hofmann, Toru Nakamura, Shinsaku Kiyomoto, and Jetzabel Serna. I read but don't agree: Privacy policy benchmarking using machine learning and the EU GDPR. In *The Web Conference*, pages 163–166, 2018.

[259] Welderufael B. Tesfay, Peter Hofmann, Toru Nakamura, Shinsaku Kiyomoto, and Jetzabel Serna. PrivacyGuide: Towards an implementation of the EU GDPR on Internet privacy policy evaluation. In *ACM International Workshop on Security and Privacy Analytics*, pages 15–21, 2018.

[260] Shukun Tokas, Olaf Owe, and Toktam Ramezanifarkhani. Static checking of GDPR-related privacy compliance for object-oriented distributed systems. *Journal of Logical and Algebraic Methods in Programming*, 125:100733, 2022.

[261] Jan Tolsdorf, Michael Fischer, and Luigi Lo Iacono. A case study on the implementation of the right of access in privacy dashboards. In *Annual Privacy Forum*, pages 23–46, 2021.

[262] Damiano Torre, Ghanem Soltana, Mehrdad Sabetzadeh, Lionel C Briand, Yuri Auffinger, and Peter Goes. Using models to enable compliance checking against the GDPR: An experience report. In *ACM/IEEE 22nd International Conference on Model Driven Engineering Languages and Systems*, pages 1–11, 2019.

[263] Michael Toth, Natalia Bielova, and Vincent Roca. On dark patterns and manipulation of website publishers by CMPs. In *22nd Privacy Enhancing Technologies Symposium*, 2022.

[264] Casey Tran, Reza Tourani, Satyajayant Misra, Travis Machacek, and Gaurav Panwar. Analyzing GDPR compliance of named data networking. In *Conference on Information-Centric Networking*, pages 107–117, 2021.

[265] Martino Trevisan, Stefano Traverso, Eleonora Bassi, and Marco Mellia. 4 years of EU cookie law: Results and lessons learned. *Proceedings on Privacy Enhancing Technologies*, 2019(2):126–145, 2019.

[266] Nguyen Binh Truong, Kai Sun, Gyu Myoung Lee, and Yike Guo. GDPR-compliant personal data management: A blockchain-based solution. *IEEE Transactions on Information Forensics and Security*, 15:1746–1761, 2019.

[267] Sarah Turner, July Galindo Quintero, Simon Turner, Jessica Lis, and Leonie Maria Tanczer. The exercisability of the right to data portability in the emerging Internet of Things (IoT) environment. *new media & society*, 23(10):2861–2881, 2021.

[268] United States Census Bureau. QuickFacts California. <https://www.census.gov/quickfacts/fact/table/CA/PST045221>, July 2021.

[269] Tobias Urban, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. “Your hashed IP address: Ubuntu.” Perspectives on transparency tools for online advertising. In *Annual Computer Security Applications Conference*, pages 702–717, 2019.

[270] Tobias Urban, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. Beyond the front page: Measuring third party dynamics in the field. In *The Web Conference*, pages 1275–1286, 2020.

[271] Tobias Urban, Dennis Tatang, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. The unwanted sharing economy: An analysis of cookie syncing and user transparency under GDPR. *arXiv preprint arXiv:1811.08660*, 2018.

[272] Tobias Urban, Dennis Tatang, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. A study on subject data access in online advertising after the GDPR. In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, pages 61–79. 2019.

[273] Tobias Urban, Dennis Tatang, Martin Degeling, Thorsten Holz, and Norbert Pohlmann. Measuring the impact of the GDPR on data sharing in ad networks. In *Asia Conference on Computer and Communications Security*, pages 222–235, 2020.

[274] Christine Utz, Sabrina Amft, Martin Degeling, Thorsten Holz, Sascha Fahl, and Florian Schaub. Privacy rarely considered: Exploring considerations in the adoption of third-party services by websites. *Proceedings on Privacy Enhancing Technologies*, 1:5–28, 2023.

[275] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 973–990, 2019.

[276] Christine Utz, Matthias Michels, Martin Degeling, Ninja Marnau, and Ben Stock. Comparing large-scale privacy and security notifications. *Proceedings on Privacy Enhancing Technologies*, 3:173–193, 2023.

[277] Pelayo Vallina, Álvaro Feal, Julien Gamba, Narseo Vallina-Rodriguez, and Antonio Fernández Anta. Tales from the porn: A comprehensive privacy analysis of the web porn ecosystem. In *Internet Measurement Conference*, pages 245–258, 2019.

[278] Marieke Van Hofslot, Almila Akdag Salah, Albert Gatt, and Crisitana Santos. Automatic classification of legal violations in cookie banner texts. In *Proceedings of the Natural Legal Language Processing Workshop 2022*, pages 287–295, 2022.

[279] Maggie Van Nortwick and Christo Wilson. Setting the bar low: Are websites complying with the minimum requirements of the CCPA? *Proceedings on Privacy Enhancing Technologies*, 2022(1):608–628, 2022.

[280] Sophie Veys, Daniel Serrano, Madison Stamos, Margot Herman, Nathan Reitinger, Michelle L Mazurek, and Blase Ur. Pursuing usable and useful data downloads under GDPR/CCPA access rights via Co-Design. In *17th Symposium on Usable Privacy and Security*, pages 217–242, 2021.

[281] Natalija Vlajic, Marmara El Masri, Gianluigi M. Riva, Marguerite Barry, and Derek Doran. Online tracking of kids and teens by means of invisible images: COPPA vs. GDPR. In *International Workshop on Multimedia Privacy and Security*, pages 96–103, 2018.

[282] Ari Ezra Waldman. Cognitive biases, dark patterns, and the ‘privacy paradox’. *Current Opinion in Psychology*, 31:105–109, 2020.

[283] Jice Wang, Yue Xiao, Xueqiang Wang, Yuhong Nan, Luyi Xing, Xiaoqing Liao, JinWei Dong, Nicolas Serrano, Haoran Lu, XiaoFeng Wang, et al. Understanding malicious cross-library data harvesting on android. In *30th USENIX Security Symposium*, pages 4133–4150, 2021.

[284] Lun Wang, Usmann Khan, Joseph Near, Qi Pang, Jithendaraa Subramanian, Neel Somani, Peng Gao, Andrew Low, and Dawn Song. PRIVGUARD: Privacy regulation compliance made easier. In *31st USENIX Security Symposium*, 2022.

[285] Miranda Wei, Madison Stamos, Sophie Veys, Nathan Reitinger, Justin Goodman, Margot Herman, Dorota Filipczuk, Ben Weinsel, Michelle L Mazurek, and Blase Ur. What twitter knows: Characterizing ad targeting practices, user perceptions, and ad explanations through users’ own twitter data. In *29th USENIX Security Symposium*, pages 145–162, 2020.

[286] Charles Weir, Ben Hermann, and Sascha Fahl. From needs to actions to secure apps? The effect of requirements and developer practices on app security. In *29th USENIX Security Symposium*, pages 289–305, 2020.

[287] Primal Wijesekera, Abbas Razaghpanah, Joel Reardon, Irwin Reyes, Narseo Vallina-Rodriguez, Serge Egelman, and Christian Kreibich. “Is our children’s apps learning?” Automatically detecting COPPA violations. In *Workshop on Technology and Consumer Protection*, 2017.

[288] Karl Wolf, Frank Pallas, and Stefan Tai. Messaging with purpose limitation–privacy-compliant publish-subscribe systems. In *IEEE 25th International Enterprise Distributed Object Computing Conference*, pages 162–172, 2021.

[289] Janis Wong and Tristan Henderson. How portable is portable? Exercising the GDPR’s right to data portability. In *International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*, pages 911–920, 2018.

[290] Janis Wong and Tristan Henderson. The right to data portability in practice: Exploring the implications of the technologically neutral GDPR. *International Data Privacy Law*, 9(3):173–191, 2019.

[291] Richmond Y Wong, Andrew Chong, and R Cooper Aspren. Privacy legislation as business risks: How GDPR and CCPA are represented in technology companies’ investment risk disclosures. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1):1–26, 2023.

[292] Mike Woodward. 16 countries with gdpr-like data privacy laws. <https://securityscorecard.com/blog/countries-with-gdpr-like-data-privacy-laws>, July 2021.

[293] World Bank. GDP (current US\$). <https://data.worldbank.org/indicator/NY.GDP.MKTP.CD>. Accessed on November 17, 2022.

[294] World Bank. Population, total. <https://data.worldbank.org/indicator/SP.POP.TOTL>. Accessed on November 17, 2022.

[295] Yue Xiao, Zhengyi Li, Yue Qin, Xiaolong Bai, Jiale Guan, Xiaojing Liao, and Luyi Xing. Lalaine: Measuring and characterizing non-compliance of Apple privacy labels. In *32nd USENIX Security Symposium*, pages 1091–1108, 2023.

[296] Fuman Xie, Yanjun Zhang, Chuan Yan, Suwan Li, Lei Bu, Kai Chen, Zi Huang, and Guangdong Bai. Scrutinizing privacy policy compliance of virtual personal assistant apps. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, pages 1–13, 2022.

[297] Minhui Xue, Gabriel Magno, E Landulfo Teixeira P Cunha, Virgilio Almeida, and Keith W Ross. The right to be forgotten in the media: A data-driven study. *Proceedings on Privacy Enhancing Technologies*, 2016(4):389–402, 2016.

[298] Alexandros Yeratziotis and George A Papadopoulos. CompLicy: Evaluating the GDPR alignment of privacy policies-a study on web platforms. In *Research Challenges in Information Science*, volume 415, pages 152–168, 2021.

[299] Razieh Nokhbeh Zaeem and K Suzanne Barber. The effect of the GDPR on privacy policies: Recent progress and future promise. *ACM Transactions on Management Information Systems*, 12(1):1–20, 2020.

[300] Leah Zhang-Kennedy and Sonia Chiasson. “Whether it’s moral is a whole other story”: Consumer perspectives on privacy regulations and corporate data practices. In *17th Symposium on Usable Privacy and Security*, pages 197–216, 2021.

[301] Lu Zhou, Chengyongxiao Wei, Tong Zhu, Guoxing Chen, Xiaokuan Zhang, Suguo Du, Hui Cao, and Haojin Zhu. POLICYCOMP: Counterpart comparison of privacy policies uncovers overbroad personal data collection practices. In *32nd USENIX Security Symposium*, pages 1073–1090, 2023.

[302] Sebastian Zimmeck and Kuba Alicki. Standardizing and implementing do not sell. In *Workshop on Privacy in the Electronic Society*, pages 15–20, 2020.

[303] Sebastian Zimmeck, Rafael Goldstein, and David Baraka. PrivacyFlash Pro: automating privacy policy generation for mobile apps. In *28th Network and Distributed System Security Symposium*, 2021.

[304] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel R. Reidenberg, N. Cameron Russell, and Norman Sadeh. MAPS: Scaling privacy compliance analysis to a million apps. *Proceedings on Privacy Enhancing Technologies*, 2019:66–86, 2019.

[305] Sebastian Zimmeck, Oliver Wang, Kuba Alicki, Jocelyn Wang, and Sophie Eng. Usability and enforceability of Global Privacy Control. *Proceedings on Privacy Enhancing Technologies*, 2:1–17, 2023.

[306] Sebastian Zimmeck, Ziqi Wang, Lieyong Zou, Roger Iyengar, Bin Liu, Florian Schaub, Shomir Wilson, Norman Sadeh, Steven M. Bellovin, and Joel Reidenberg. Automated analysis of privacy requirements for mobile apps. In *24th Network & Distributed System Security Symposium*, 2017.

Appendix A.

Meta-Review

The following meta-review was prepared by the program committee for the 2024 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

A.1. Summary of Paper

This SoK paper presents an analysis of 24 privacy and data protection regulations to derive a taxonomy of rights and obligations defined in those laws, which is then used to systematize 270 research papers that investigate the impact of those laws.

A.2. Scientific Contributions

- Provides a Valuable Step Forward in an Established Field

A.3. Reasons for Acceptance

- 1) The paper contributes a well-written comprehensive overview of research on effects of privacy laws and regulations. It does a great job of systematizing prior work and pointing out future research directions.
- 2) The paper presents a detailed analysis of the existing literature, with comprehensive coverage of research efforts studying effects of privacy laws and regulations since 2017. The discussion of research conducted with respect to the different topics in the taxonomy is well done.
- 3) The paper presents interesting outcomes, particularly highlighting the importance of focusing on under-studied aspects of privacy regulations. Another key finding is the lack of focus on cross-cultural practices in most of the research.
- 4) The paper is written well.

Noteworthy Concerns

A mapping between the identified rights/obligations and the laws in which they are described would have been useful to enable readers to more clearly trace the sources of those rights.