

Breaking Down Walls of Text: How Can NLP Benefit Consumer Privacy?

Abhilasha Ravichander[◇] Alan W Black[◇] Thomas Norton[♣]
Shomir Wilson[♡] Norman Sadeh[◇]

[◇]Carnegie Mellon University, Pittsburgh, PA

[♣]Fordham Law School, New York, NY [♡]Penn State University, University Park, PA

{aravicha, awb, sadeh}@cs.cmu.edu

{shomir}@psu.edu, {tnorton1}@law.fordham.edu

Abstract

Privacy plays a crucial role in preserving democratic ideals and personal autonomy. The dominant legal approach to privacy in many jurisdictions is the “Notice and Choice” paradigm, where privacy policies are the primary instrument used to convey information to users. However, privacy policies are long and complex documents that are difficult for users to read and comprehend. We discuss how language technologies can play an important role in addressing this information gap, reporting on initial progress towards helping three specific categories of stakeholders take advantage of digital privacy policies: consumers, enterprises, and regulators. Our goal is to provide a roadmap for the development and use of language technologies to empower users to reclaim control over their privacy, limit privacy harms, and rally research efforts from the community towards addressing an issue with large social impact. We highlight many remaining opportunities to develop language technologies that are more precise or nuanced in the way in which they use the text of privacy policies.

1 Introduction

Privacy is a fundamental right central to a democratic society, in which individuals can operate as autonomous beings free from undue interference from other individuals or entities (Assembly, 1948). However, certain functions of privacy, such as the power to grant or deny access to one’s personal information, are eroded by modern commercial and business practices that involve vast collection, linking, sharing, and processing of digital personal information through an opaque network, often without data subjects’ knowledge or consent. In many jurisdictions, online privacy is largely governed by “Notice and Choice” (Federal Trade Commission, 1998). Under this framework, data-collecting

and data-processing entities publish privacy policies that disclose their data practices. Theoretically, users are free to make choices about which services and products they use based on the disclosures made in these policies. Thus, the legitimacy of this framework hinges on users reading a large number of privacy policies to understand what data can be collected and how that data can be processed before making informed privacy decisions.

In practice, people seldom read privacy policies, as this would require prohibitive amounts of their time (McDonald and Cranor, 2008; Cate, 2010; Cranor, 2012; Reidenberg et al., 2015; Schaub et al., 2015; Jain et al., 2016). Thus, an opportunity exists for language technologies to bridge this gap by processing privacy policies to meet the needs of Internet and mobile users. NLP has made inroads in digesting large amounts of text in domains such as scientific publications and news (Jain et al., 2020; Cachola et al., 2020; Kang et al., 2018; Rush et al., 2015; See et al., 2017), with several practical tools based on these technologies helping users every day (Cachola et al., 2020; TLDR, 2021; News, 2021). These domains have also received considerable research attention: several benchmark datasets and technologies are based in texts from these domains (Nallapati et al., 2016; See et al., 2017; Narayan et al., 2018; Beltagy et al., 2019). We highlight that the privacy domain can also benefit from increased research attention from the community. Moreover, technologies developed in the privacy domain have potential for significant and large-scale positive social impact—the affected population includes virtually every Internet or mobile user (Sadeh et al., 2013).

Automated processing of privacy policies opens the door to a number of scenarios where language technologies can be developed to support users in the context of different tasks. This includes saving data subjects the trouble of having to read the

entire text of policies when they are typically only concerned about one or a small number of issues (e.g., determining whether they can opt out of some practices or whether some of their data might be shared with third parties). It includes helping companies ensure that they are compliant and that their privacy policies are consistent with what their code actually does. It also includes supporting regulators, as they face the daunting task of enforcing compliance across an ever-growing collection of software products and processes, including sophisticated data collection and use practices. In this work, we conduct an extensive survey of initial progress in applying NLP to address limitations of the Notice and Choice model. We expect our work to serve as a useful starting point for practitioners to familiarize themselves with technological progress in this domain, by providing both an introduction to the basic privacy concerns and frameworks surrounding privacy policies, as well as an account of applications for which language technologies have been developed. Finally, we highlight many remaining opportunities for NLP technologies to extract more precise, more nuanced, and ultimately more useful information from privacy policy text—describing key challenges in this area and laying out a vision for the future.

2 Privacy as a Social Good

In 1890, Warren and Brandeis defined the right to privacy as “the right to be let alone” (Warren and Brandeis, 1890). More recently, Westin defined the right as “the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others” (Westin, 1968). A primary aspiration of privacy is to allow for the separation of individual and society as a means of fostering personal autonomy. To that end, privacy “protects the situated practices of boundary management through which the capacity for self-determination develops,” and further “shelters dynamic, emergent subjectivity from the efforts of commercial and government actors to render individuals and communities fixed, transparent, and predictable” (Cohen, 2012). Privacy, therefore, is “foundational to the practice of informed and reflective citizenship,” and serves as “an indispensable structural feature of liberal democratic political systems” (Cohen, 2012).

When privacy is threatened, we risk losing the chance for critical self-reflection of political pro-

cesses and social norms. Indeed, privacy undergirds the concepts of human dignity and other key values, such as the freedoms of association and speech. For these reasons and others, privacy is regarded as a fundamental human right (Assembly, 1948). In the digital age, privacy is threatened by aggressive, rapid, and largely automated collection, linking, sharing, and processing of digital personal information. Digital privacy is intrinsically linked to the fundamental ethical principles of *transparency*, *fairness* and *agency*.

- *Transparency*: Users have a right to know how information about them is collected and used. Entities collecting user data stay clear of manipulative schemes designed to influence the data subject’s willingness to disclose their data (e.g. overemphasizing benefits while remaining silent about potential risks associated with the disclosure of data in a given context).
- *Fairness*: Users should receive perceived value commensurate to the perceived loss of privacy associated with disclosure and use of their data.
- *Agency*: Users should have a choice about what data is collected about them and how it is used.

The dominant paradigm to address these principles in the United States and most legal jurisdictions around the world, is the ‘Notice and Choice’ regulatory framework (Westin, 1968; Federal Trade Commission, 1998). ‘Notice and Choice’ regimes are based on the presupposition that consumers will adequately manage their privacy, if provided sufficient information about how their data will be collected, used and managed, as well as offered meaningful choices. Today, ‘Notice’ is often practically realized through publishing *privacy policies*, which are long and verbose documents that users are expected to read and understand. ‘Choice’ is often limited to the user clicking ‘I agree’ to the privacy policy, or even interpreting their continued use of the service as some sort of meaningful consent to the terms of the policy.

The ‘Notice and Choice’ framework is fundamentally broken. In practice, users seldom read privacy policies (McDonald and Cranor, 2008; Cate, 2010; US Federal Trade Commission et al., 2012) and it is prohibitively expensive for them to even do so. McDonald and Cranor (2008) estimate that if internet users were to actually read the privacy policies of the websites they visited, they would have to spend roughly 250 hours each year just reading

Challenge	Example												
Ambiguity	We may also use aggregate personal information for regulatory compliance, industry and market analysis, research, demographic profiling, marketing and advertising, and other business purposes.												
Vagueness	[X] collects, or may have a third-party service providers collect, non-personally-identifying information of the sort that mobile applications typically make available , such as the type of device using the Application, the operating system, location information, and aggregated user statistics.												
Modality	If you use our services to make and receive calls or send and receive messages, we may collect call and message log information like your phone number, calling-party number, receiving-party number...												
Negation	No apps have access to contact information, nor do they read or store any contact information												
Lists and Document Structure	We may collect data or ask you to provide certain data when you visit and use our websites, products and services. The sources from which we collect Personal Data include: <ul style="list-style-type: none"> • Data collected directly from you or your device ; • If we link other data relating to you with your Personal Data, we will treat that linked data as Personal Data; and • We may also collect Personal Data from trusted third-party sources.... 												
Tabular Understanding	<table border="1"> <thead> <tr> <th>Reasons we Can Share Your Personal Information</th> <th>Does X share?</th> <th>Can you limit this sharing?</th> </tr> </thead> <tbody> <tr> <td>For our everyday business purposes ...</td> <td>Yes</td> <td>No</td> </tr> <tr> <td>For our everyday marketing purposes ...</td> <td>Yes</td> <td>No</td> </tr> <tr> <td>For joint marketing with other companies</td> <td>No</td> <td>We don't share</td> </tr> </tbody> </table>	Reasons we Can Share Your Personal Information	Does X share?	Can you limit this sharing?	For our everyday business purposes ...	Yes	No	For our everyday marketing purposes ...	Yes	No	For joint marketing with other companies	No	We don't share
Reasons we Can Share Your Personal Information	Does X share?	Can you limit this sharing?											
For our everyday business purposes ...	Yes	No											
For our everyday marketing purposes ...	Yes	No											
For joint marketing with other companies	No	We don't share											

Table 1: Examples of some challenging aspects for language understanding in privacy policies, including reasoning over ambiguity and vagueness, modality, negation (including [scope](#)), lists and document structure, and tables.

privacy policies. A 2014 report from the Presidents Council of Advisors on Science and Technology stated that “only in some fantasy world” were users reading and understanding privacy policies before giving their consent ([of the President’s Council of Advisors on Science and Technology, 2014](#)). Indeed, 91% of people in the U.S have reported feeling like they have lost control over their information ([Madden et al., 2014](#)). Moreover, recent privacy laws such as the EU’s General Data Protection Regulation (GDPR) ([Regulation, 2016](#)) still fail to address the critical limitation of notice and choice: the continued reliance on users to read and understand a large number of privacy policies. Studies have shown that GDPR requirements have actually resulted in longer privacy policies ([Linden et al., 2020](#)), and users still encounter unreadable privacy policies ([Becher and Benoliel, 2019](#)).

The lack of respect for individuals’ rights to privacy also has implications for society. With social platforms in particular having access to an unprecedented scale of information about human behaviour, [Vicario et al. \(2019\)](#) discuss that users’ polarization and confirmation bias can play a role in spreading misinformation on social platforms. [Madden et al. \(2017\)](#) report that particular groups of less-privileged users on the internet are uniquely vulnerable to various forms of surveillance and privacy harms, which could widen existing economic gaps.

[Introna \(1997\)](#) describe privacy as central to human autonomy in social relationships. In this work, we examine the potential of language technologies in enabling people to derive the benefits of their rights to transparency, fairness and agency.

3 Can NLP Help Privacy?

Privacy policies present interesting challenges for NLP practitioners, as they often feature characteristic aspects of language that remain under-examined or difficult to process (Table. 1). For example, while many policies discuss similar issues surrounding how user data is collected, managed and stored, policy silence about certain data practices may carry great weight from a legal, policy, and regulatory perspective.¹ In the privacy policy domain, understanding what has *not* been said in a privacy policy (*policy silence*) is just as important as understanding what *is* said ([Zimmeck et al., 2019a](#); [Marotta-Wurgler, 2019](#)).

Further, though policies tend to feature literal language (compared to more subjective domains like literature or blog posts), processing them ef-

¹For example, in *United States v. Path*, the defendant’s (Path) privacy policy described that its app collected “certain information such as your Internet Protocol (IP) address, your operating system, the browser type.” The Federal Trade Commission found this disclosure to be incomplete and insufficient to provide notice about the collection of users’ contact data ([FTC, 2013](#)).

Task	Goal	Consumer	Regulator	Enterprise
Data Practice Identification (Wilson et al., 2016b)	Annotate segments of privacy policies with described data practices.	✓	✓	✓
Opt-Out Identification (Sathyendra et al., 2017; Bannihatti Kumar et al., 2020)	Extract opt-out choices buried in privacy policy text.	✓		
Compliance Analysis (Zimmeck et al., 2017, 2019a)	Analyze mobile app code and privacy policy to identify potential compliance issues.		✓	✓
Privacy Question-Answering (Ravichander et al., 2019; Ahmad et al., 2020)	Allow consumers to selectively query privacy policies for issues that are important to them.	✓		
Policy Summarization (Zaeem et al., 2018; Keymanesh et al., 2020)	Construct summaries to aid consumers to quickly digest the content of privacy policies.	✓		
Readability Analysis (Massey et al., 2013; Meiselwitz, 2013)	Characterize the ease of understanding or comprehension of privacy policies.		✓	

Table 2: Overview of some applications of NLP to privacy policies, and primary stakeholders they are intended to benefit.

fectively also requires several additional capabilities such as reasoning over vagueness and ambiguity, understanding elements such as lists (including when they are intended to be exhaustive and when they are not (Bhatia et al., 2016)), effectively incorporating ‘co-text’- aspects of web document structure such as document headers that are meaningful semantically to the content of privacy policies (Mysore Gopinath et al., 2018) and incorporating domain knowledge (for example, understanding whether information is *sensitive* requires background knowledge in the form of applicable regulation). Privacy policies also differ from several closely related domains, such as legal texts which are largely meant to be processed by domain experts. In contrast, privacy policies are legal documents with legal effects—generally drafted by experts—that are ostensibly meant to be understood by everyday users. NLP applications in the privacy domain also need to be designed with end user requirements in mind. For example, from a legal standpoint, when generating answers to a user’s question about the content of a privacy policy, it is generally advisable to include disclaimers, but users may prefer to be presented with shorter answers, where disclaimers are kept as short as possible. Challenges are described in more detail in (§4).

We survey current efforts to apply NLP in the privacy domain, discussing both existing task formulations as well as future areas in this domain where language technologies can have impact.²

²Our survey includes relevant papers from major NLP venues, including ACL, EMNLP, NAACL, EACL, COLING, CoNLL, SemEval, TACL, and CL. We supplemented these publications with a review of the literature at venues such as SOUPS, PETS, WWW, ACM, and NDSS. We also included relevant legal venues, such as law reviews and journals.

3.1 Data Practice Identification

Initial efforts in applying NLP in the privacy domain have largely focused on discovering or identifying data practice categories in privacy policies (Costante et al., 2012a; Ammar et al., 2012; Costante et al., 2012b; Liu et al., 2014b; Ramanath et al., 2014a; Wilson et al., 2016b). Automating the identification of such data practices could potentially support users in navigating privacy policies more effectively³, as well as automate analysis for regulators who currently do not have techniques to assess a large number of privacy policies. Wilson et al. (2016b) create a corpus of 115 website privacy policies annotated with detailed information of the privacy policies described. The corpus and associated taxonomy have been of utility in the development of several subsequent privacy-enhancing language technologies (Mysore Sathyendra et al., 2017a; Zimmeck et al., 2017; Ravichander et al., 2019; Ahmad et al., 2020).

3.2 Choice Identification

Studies have shown that consumers desire control over the use of their information for marketing communication, and object to the use of their information for web tracking or marketing purposes including targeted advertising (Cranor et al., 2000; Turow et al., 2009; Ur et al., 2012; Bleier and Eisenbeiss, 2015). However, McDonald and Cranor (2010) find that many people are unaware of the opt-out choices available to them. These choices are often buried in policy text, and thus there has been interest in applying NLP to extract choice language. Mysore Sathyendra et al. (2017b) automatically identify choice instances within a privacy

³For example, through the data exploration tool developed by the Usable Privacy Policy Project: <https://explore.usableprivacy.org/?view=machine>




Figure 1: The results from Opt-Out Easy, a browser extension to extract opt-out choices from privacy policies, for Overleaf.com (Bannihatti Kumar et al., 2020).

policy, labeling different types of opt-out choices, with a particular emphasis on extracting actionable choices in the policy, i.e. those associated with hyperlinks. Bannihatti Kumar et al. (2020) develop a web-browser extension to present extracted choice instances to users (Figure. 1), finding that the tool can considerably increase awareness of choices available to users and reduce the time taken to identify actions the users can take.

3.3 Compliance Analysis

In 2012, six major mobile app stores entered into an agreement with the California Attorney General, where they agreed to adopt privacy principles that require mobile apps to have privacy policies (Justice, 2012). Regulations such as the the EU General Data Protection Directive (GDPR) and the California Consumer Protection Act (CCPA) impose further requirements on what entities collecting and using personal data need to disclose in their privacy policies and what rights they need to offer to their users (e.g. privacy controls, option to request deletion of one’s data). However, regulators lack the necessary resources to systematically check that these requirements are satisfied. In fact, even app stores lack the resources to systematically check that disclosures made in privacy policies are consistent with the code of apps and comply with relevant regulatory requirements. Thus, there has been interest in developing technologies to automatically identify potential compliance issues (Enck et al., 2014; Zimmeck et al., 2017; Wang et al., 2018; Libert, 2018a; Zimmeck et al., 2019b).

A first application of language technologies to

aid compliance analysis is detailed by Zimmeck et al. (2017), including results of a systematic analysis of 17,991 apps using both natural language processing and code analysis techniques. Classifiers are trained to identify data practices based on the OPP-115 ontology (Wilson et al., 2016b), and static code analysis techniques are employed to extract app’s privacy behaviors. The results from the two procedures are compared to identify potential compliance issues. The system was piloted with personnel at the California Office of the Attorney General. Users reported that the system could significantly increase productivity, and decrease the effort and time required to analyze practices in apps and audit compliance. Zimmeck et al. (2019b) review 1,035,853 apps from the Google Play Store for compliance issues. Their system identifies disclosed privacy practices in policies using classifiers trained on the APP-350 corpus (Story et al., 2019), and static code analysis techniques to identify apps’ privacy behaviors. Results of the analysis of this large corpus of privacy policies revealed a particularly large number of potential compliance problems, with a subset of results shared with the Federal Trade Commission. The system was also reported to have been used by a large electronics manufacturer to verify compliance of legacy mobile apps prior to the introduction of GDPR.

3.4 Policy Summarization

Due to the lengthy and verbose nature of privacy policies, it is appealing to attempt to develop automated text summarization techniques to generate short and concise summaries of a privacy policy’s contents (Liu et al., 2015). Tomuro et al. (2016) develop an extractive summarization system that identifies important sentences in a privacy policy along five categories: purpose, third parties, limited collection, limited use and data retention. Zaeem et al. (2018, 2020) identify ten questions about privacy policies, and automatically categorize ‘risk levels’ associated with each of the questions, as shown in Table. 3. Keymanesh et al. (2020) focus on extractive summarization approaches to identify ‘risky sections’ of the privacy policy, which are sentences that are likely to describe a privacy risk posed to the end-user. However, while automated summarization seems like a promising application of language technologies, identifying which parts of a policy should be shown to users is exceedingly difficult, and studies by privacy experts have shown

#	Question	Green Risk Level	Yellow Risk Level	Red Risk Level
(1)	How well does this website protect your email address?	Not asked for	Used for intended service	Shared w/ third parties
(2)	How well does this website protect your credit card information and address?	Not asked for	Used for intended service	Shared w/ third parties
(3)	How well does this website handle your social security number?	Not asked for	Used for intended service	Shared w/ third parties
(4)	Does this website use or share your PII for marketing purposes?	PII not used for marketing	PII used for marketing	PII shared for marketing
(5)	Does this website track or share your location?	Not tracked	Used for intended service	Shared w/ third parties
(6)	Does this website collect PII from children under 13?	Not collected	Not mentioned	Collected
(7)	Does this website share your information with law enforcement?	PII not recorded	Legal docs required	Legal docs not required
(8)	Does this website notify or allow you to opt-out after changing their privacy policy?	Posted w/ opt-out option	Posted w/o opt-out option	Not posted
(9)	Does this website allow you to edit or delete your information from its records?	Edit/delete	Edit only	No edit/delete
(10)	Does this website collect or share aggregated data related to your identity or behavior?	Not aggregated	Aggregated w/o PII	Aggregated w/ PII

Table 3: Ten privacy questions used for summarization, and associated ‘risk levels’ from (Zaem et al., 2018).

that such ‘one-size-fits-all’ approaches are unlikely to be effective (Gluck et al., 2016; Rao et al., 2016).

3.5 Privacy Question-Answering

A desire to move away from ‘one-size-fits-all’ approaches has led to increased interest in supporting automated privacy question-answering (QA) capabilities. If realized, such functionality will help users selectively and iteratively explore issues that matter most to them. Table 4 lists current efforts to develop resources for privacy question-answering. Amongst the initial explorations in this area, Harkous et al. (2018) examine privacy questions asked by Twitter users to companies, with answers annotated by the paper’s authors. Ravichander et al. (2019) collect questions asked by crowdworkers about a mobile app without seeing the app’s privacy policy, and hire legal experts to identify sentences in the privacy policy relevant for each question. (Ahmad et al., 2020) provide ‘skilled annotators’ with privacy policy segments drawn from the OPP-115 corpus (Wilson et al., 2016b), and ask them to construct questions based on the provided span of text. Ravichander et al. (2019) and Ahmad et al. (2020) both find that current QA baselines based on pretrained language models (Devlin et al., 2019) are inadequate for answering privacy questions. Ahmad et al. (2020) indicate that identifying longer evidence spans are challenging and describe transfer learning as a potential direction to improve performance. Ravichander et al. (2019) examine unanswerability as a challenge to privacy QA systems, highlighting the many facets of unanswerable questions that can be asked. It is worth noting that all three resources formulate ground truth based in the text of the privacy policy, but policy language is difficult for non-experts to understand (Reidenberg et al., 2015). Future QA dataset architects

could consider abstractive answers as ground truths, which are validated by legal experts for correctness and evaluated by users for helpfulness. It may also be desirable for benchmarks to aim for ecological validity (de Vries et al., 2020), with users asking questions, and legal experts constructing answers.

3.6 Other Applications

In this section, we survey further tasks where NLP has been applied to consumer privacy, including analyzing privacy policy *readability*, with the goal of aiding writers of privacy policies (Fabian et al., 2017; Massey et al., 2013; Meiselwitz, 2013; Ermakova et al., 2015), and understanding data practice categories are described in a policy, known as measuring policy *coverage* (Linden et al., 2020; Shvartzshnaider et al., 2020). A significant amount of recent work has also focused on information extraction from privacy policies (Costante et al., 2012a). Shvartzshnaider et al. (2018); Shvartzshnaider et al. (2019, 2020) identify contextual integrity parameters (Nissenbaum, 2004) in policy text. Studies have also tried to extract other, more specific kinds of information from policies, such as third party entities (Libert, 2018b; Bokaie Hosseini et al., 2020) and information about regulated information types (Bhatia et al., 2016; Evans et al., 2017) as well as their similarity (Hosseini et al., 2016). There have also been efforts to analyze vague statements in privacy policies (Liu et al., 2016b; Lebanoff and Liu, 2018), and explore how benchmarks in this domain can be constructed through crowdsourcing (Ramanath et al., 2014b; Wilson et al., 2016c; Audich et al., 2018). Lastly, there has been research focused on identifying header information in privacy policies (Mysore Gopinath et al., 2018) and generating them (Gopinath et al., 2020). Techniques to

Dataset	#Questions	Question Scenario	Legal Expert Annotator	Asker Cannot See Evidence	Unanswerable Questions	Non-Contiguous Answer
Polisis (Harkous et al., 2018)	120	Twitter users ask questions to a company.	✗	✓	✗	✗
PrivacyQA (Ravichander et al., 2019)	1750	Crowdworkers ask questions about a mobile app.	✓	✓	✓	✓
PolicyQA (Ahmad et al., 2020)	714	Skilled annotators are shown a text span and data practice, and asked to construct a question.	✗	✗	✗	✗

Table 4: Comparison of Polisis (Harkous et al., 2018), PrivacyQA (Ravichander et al., 2019) and PolicyQA (Ahmad et al., 2020) QA datasets. *Question Scenario* describes conditions under which the questions were generated. ‘Asker Cannot See Evidence’ indicates the asker of the question was not shown evidence from the document when formulating questions. Unanswerable questions indicates if the corpus includes unanswerable questions. ‘Non Contiguous Answer’ indicates the answers are allowed to be from non-adjacent segments of the privacy policy.

process privacy policies have largely followed successful approaches elsewhere in NLP, starting from feature-based approaches (Sathyendra et al., 2017; Zimmeck et al., 2019a), training domain-specific word embeddings (Kumar et al., 2019) and fine-tuning pretrained language models on privacy policies (Nejad et al., 2020; Mustapha et al., 2020).

3.7 Towards New Tasks and Formulations

We discuss a vision of future applications of NLP in aiding consumer privacy. We believe these applications present interesting opportunities for the community to develop technologies, both because of the technical challenges they offer and the impact they are likely to have.

Detecting surprising statements: Since users do not read privacy policies, their expectations for the data practices of services might not align with services’ actual practices. These mismatches may result in unexpected privacy risks which lead to loss of user trust (Rao et al., 2016). Identifying such ‘surprising’ statements will require understanding social context and domain knowledge of privacy information types. For example, it is natural for a banking website to collect payment information, but not health information. Moreover, understanding what statements will be surprising for each individual user requires understanding their personal, social and cultural backgrounds (Rao et al., 2016). We speculate that NLP can potentially be leveraged to increase transparency by identifying discordant statements within privacy policies.

Detecting missing information: In contrast to detecting surprising statements, privacy policies may be *underspecified*. Story et al. (2018) find that many policies contain language appearing in unrelated privacy policies, indicating that policy writers

may use privacy policy generators not suited to their application, potentially resulting in missing information. Techniques from compliance analysis could help in flagging some of these issues (Zimmeck et al., 2017, 2019a).

Generating privacy nutrition labels: One proposal to overcome the gap in communicating privacy information to users has been the privacy ‘nutrition label’ approach (Kelley et al., 2009, 2013), as shown in Fig. 2. The proposal draws from industries such as nutrition, warning and energy labeling where information has to be communicated to consumers in a standardized way. Recently, Apple announced that developers will be required to provide information for these labels (Campbell, 2020), which disclose to the user the information a company and third parties collect.⁴ This approach could potentially be helpful to users to understand privacy information at a glance, but presents challenges to both developers and app platforms. Developers need to ensure their nutrition label is accurate and platforms need to enforce compliance to these requirements. Potentially, early successes of language technologies in compliance systems can be extended to analyzing a specified nutrition label, policy and application code. NLP may also be used to generate nutrition labels which developers inspect, as opposed to the more costly process of developers specifying nutrition labels from scratch which may hinder adoption (Fowler, 2021).

Personalized privacy summaries: One approach to mitigating inadequacies of policy summarization—where generic summaries may not be sufficiently complete—is personalized summarization (Díaz and Gervás, 2007; Hu et al.,

⁴An example of such a nutrition label can be found in Appendix. A

2012). In this formulation, policies are summarized for each user based on issues that matter most to them. This formulation may alleviate some downsides of QA approaches, which require the user know how to manage their privacy by asking the right questions. Personalized summarization systems would benefit from modeling users' level of knowledge, as well as their beliefs, desires and goals. In NLP, there has been effort towards addressing similar challenges for personalized learning in intelligent tutoring (McLaren et al., 2006; Malpani et al., 2011).

Assistive Policy Writing: We speculate advances in natural language generation and compliance analysis techniques may jointly be leveraged to help app developers create more accurate privacy policies, rather than relying on policy generators (Story et al., 2018). Privacy policies generally cover a known set of data practices (Wilson et al., 2016a), providing potential statistical commonalities to aid natural language generation. Code analysis can be leveraged to constrain generation to accurately describe data practices of a service.

4 Progress and Challenges

Although privacy policies have legal effects for most Internet users, these types of texts constitute an underserved domain in NLP. NLP has the potential to play a role in easing user burden in understanding salient aspects of privacy policies, help regulators enforce compliance and help developers enhance the quality of privacy policies by reducing the effort required to construct them. Yet, the privacy domain presents several challenges that require specialized resources to deal with effectively. We describe some of these distinctive challenges, as well as the capabilities that will need to be developed to process policies satisfactorily.

- *Disagreeable privacy policies:* Privacy policies are complex, but are the most important source of information about how user data is collected, managed and used. Reidenberg et al. (2015) find that sometimes discrepancies can arise in the interpretation of policy language, even between experts. This additional complexity should be taken into consideration by those developing language technologies in this domain.
- *Difficulty or validity of collecting annotations:* Privacy policies are legal documents that have legal effects on how user data is collected and

used. While crowdworkers have been found to provide non-trivial annotations for some tasks in this domain (Wilson et al., 2016c), individual practitioners constructing applications must carefully consider the consequences of sourcing non-expert annotations in the context of their task and the impacted stakeholders, and not rely on crowdsourced annotation simply because it is cheaper or easier to scale.

- *Difficult for users to articulate their needs and questions:* Developing effective privacy QA functionality will require understanding the kinds of questions users ask and quantifying to what extent privacy literacy affects users' ability to ask the right questions. Ravichander et al. (2019) find many questions collected from crowdworkers were either incomprehensible, irrelevant or atypical. Understanding these factors could lead to the development of more proactive QA functionality- for example, rather than wait for users to form questions, the QA system could prompt users to reflect on certain privacy issues.
- *Challenges to QA:* Additionally, privacy question-answering systems themselves will require several capabilities in order to have larger impact. These systems must be capable of doing question-answering iteratively, working with the user towards resolving information-seeking needs. They will also need to consider unanswerability (Rajpurkar et al., 2018; Ravichander et al., 2019; Asai and Choi, 2020) as a *graded* problem, recognizing to what extent the privacy policy contains an answer and communicating both what is known and what is not known to the user. QA systems must also consider what kinds of answers are useful, identifying appropriate response format and tailoring answers to the user's level of knowledge and individual preferences.
- *Domain Knowledge:* It remains an open question how to best incorporate expert knowledge into the processing of privacy policies. Although privacy policies are intended to be read by everyday users, experts and users often disagree on their interpretations (Reidenberg et al., 2015).
- *Combining Disparate Sources of Information:* While privacy policies are the single most important source of information about collection and sharing practices surrounding user data, technologies to address users' personalized concerns could leverage additional sources of information- such as analyzing the code of a given technology

such as a mobile app, news articles, or background knowledge of a legal, technical or statistical nature. For example, when the policy is silent on an issue- a QA system could report the practices of other similar services to the user, or if a user asks about the likelihood of a data breach, the QA system could refer to news sources for information about the service.

- *User Modeling*: Personalized privacy approaches will also need to model individual user's personal, social and cultural contexts to deliver impact. This could include information about the issues likely to matter most to users, their background knowledge, privacy preferences and expectations (Liu et al., 2014a; Lin et al., 2014; Liu et al., 2016a).
- *Accessibility*: Efforts to help users understand privacy policies by breaking through walls of text to identify salient aspects, are expected to help users with a range of visual impairments navigate their privacy. Future work would conduct user studies to determine the extent to which developed technologies ease visually-impaired users' accessibility to learn about the content of policies, related to their interests or concerns.

5 Ethical Considerations

While NLP has the potential to benefit consumer privacy, we emphasize there are also ethical considerations to be taken in account. These include:

Bias of agent providing technology: A factor that must be considered in the practical deployment of NLP systems in this domain is the incentives of the entity creating or providing the technology. For example, the incentives of a company that develops a QA system to answer questions about its own privacy policy may not align with those of a trusted third-party privacy assistant that reviews the privacy policies of many different companies. This information also needs to be communicated in an accurate and unbiased fashion to users.

User Trust: While NLP systems have the potential to digest policy text and present information to users, NLP systems are seldom completely accurate, and therefore it is important that users be appropriately informed of these limitations. For example, if a QA system communicates a data practice incorrectly in response to a users' question and the user encounters privacy harms contrary to their expectations as a result, they may lose trust in the

system. It is important to also identify appropriate disclaimers to accompany NLP systems to manage user expectations.

Discriminatory Outcomes: It is possible that different populations will benefit to different extents from the developed technologies, and we are yet unable to anticipate precisely where the benefits will accrue. For example, users with higher degrees of privacy literacy may be able to take better advantage of a developed QA system.

Technological Solutionism: It is important to consider that while language technologies have the potential to considerably alleviate user burden in reading privacy policies, they are unlikely to completely resolve the issue that users are unable to read and review a multitude of privacy policies everyday. Advances toward addressing the limitations of notice and choice will also require progress in regulation and enforcement by regulatory bodies to ensure that enterprises are more accurate in their disclosures and use clearer language, in tandem with creative technological solutions.

6 Conclusion

Privacy is about the right of people to control the collection and use of their data. Today privacy relies on the 'Notice and Choice' framework, which assumes that people actually read the text of privacy policies. This is a fantasy as users do not have the time to do so. In this article, we summarize how language technologies can help overcome this challenge and support the development of solutions that assist customers, technology providers and regulators. We reviewed early successes and presented a vision of how NLP could further help in the future. We hope this article will motivate NLP researchers to contribute to this vision and empower people to regain control over their privacy.

Acknowledgements

This research was supported in part by grants from the National Science Foundation Secure and Trustworthy Computing program (CNS-1330596, CNS-1330214, CNS-15-13957, CNS-1801316, CNS-1914486, CNS-1914444) and DARPA(FA8750-15-2-0277). Part of the work summarized in this paper was conducted by the Usable Privacy Policy Project(<https://usableprivacy.org>). The authors would like to thank Siddhant Arora, Rex Chen and Aakanksha Naik for valuable discussion.

References

- Wasi Ahmad, Jianfeng Chi, Yuan Tian, and Kai-Wei Chang. 2020. [PolicyQA: A reading comprehension dataset for privacy policies](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 743–749, Online. Association for Computational Linguistics.
- Waleed Ammar, Shomir Wilson, Norman Sadeh, and Noah A Smith. 2012. Automatic categorization of privacy policies: A pilot study. *Technical Report CMU-LTI-12-019, Carnegie Mellon University*.
- Akari Asai and Eunsol Choi. 2020. Challenges in information seeking qa: Unanswerable questions and paragraph retrieval. *arXiv preprint arXiv:2010.11915*.
- UN General Assembly. 1948. Universal declaration of human rights. *UN General Assembly*, 302(2).
- Dhiren A Audich, Rozita Dara, and Blair Nonnecke. 2018. Privacy policy annotation for semi-automated analysis: A cost-effective approach. In *IFIP International Conference on Trust Management*, pages 29–44. Springer.
- Vinayshekhar Bannihatti Kumar, Roger Iyengar, Namita Nisal, Yuanyuan Feng, Hana Habib, Peter Story, Sushain Cherivirala, Margaret Hagan, Lorrie Cranor, Shomir Wilson, et al. 2020. Finding a choice in a haystack: Automatic extraction of opt-out statements from privacy policy text. In *Proceedings of The Web Conference 2020*, pages 1943–1954.
- Shmuel I Becher and Uri Benoliel. 2019. Law in books and law in action: the readability of privacy policies and the gdpr. In *Consumer Law and Economics*, pages 179–204. Springer.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Jaspreet Bhatia, Morgan C Evans, Sudarshan Wadkar, and Travis D Breaux. 2016. Automated extraction of regulated information types using hyponymy relations. In *2016 IEEE 24th International Requirements Engineering Conference Workshops (REW)*, pages 19–25. IEEE.
- Alexander Bleier and Maik Eisenbeiss. 2015. The importance of trust for personalized online advertising. *Journal of Retailing*, 91(3):390–409.
- Mitra Bokaie Hosseini, Pragyant K C, Irwin Reyes, and Serge Egelman. 2020. [Identifying and classifying third-party entities in natural language privacy policies](#). In *Proceedings of the Second Workshop on Privacy in NLP*, pages 18–27, Online. Association for Computational Linguistics.
- Isabel Cachola, Kyle Lo, Arman Cohan, and Daniel Weld. 2020. [TLDR: Extreme summarization of scientific documents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4766–4777, Online. Association for Computational Linguistics.
- Ian Carlos Campbell. 2020. [Apple will require apps to add privacy ‘nutrition labels’ starting december 8th](#).
- Fred H Cate. 2010. The limits of notice and choice. *IEEE Security & Privacy*, 8(2).
- Julie E Cohen. 2012. What privacy is for. *Harv. L. Rev.*, 126:1904.
- Elisa Costante, Jerry den Hartog, and Milan Petković. 2012a. What websites know about you. In *Data Privacy Management and Autonomous Spontaneous Security*, pages 146–159. Springer.
- Elisa Costante, Yuanhao Sun, Milan Petković, and Jerry den Hartog. 2012b. [A machine learning solution to assess privacy policy completeness: \(short paper\)](#). In *Proceedings of the 2012 ACM Workshop on Privacy in the Electronic Society, WPES ’12*, page 91–96, New York, NY, USA. Association for Computing Machinery.
- Lorrie Faith Cranor. 2012. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *J. on Telecomm. & High Tech. L.*, 10:273.
- Lorrie Faith Cranor, Joseph Reagle, and Mark S Ackerman. 2000. Beyond concern: Understanding net users’ attitudes about online privacy. *The Internet upheaval: raising questions, seeking answers in communications policy*, pages 47–70.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alberto Díaz and Pablo Gervás. 2007. User-model based personalized summarization. *Information Processing and Management: an International Journal*, 43(6):1715–1734.
- William Enck, Peter Gilbert, Seungyeop Han, Vasant Tendulkar, Byung-Gon Chun, Landon P Cox, Jaeyeon Jung, Patrick McDaniel, and Anmol N Sheth. 2014. Taintdroid: an information-flow tracking system for realtime privacy monitoring on smartphones. *ACM Transactions on Computer Systems (TOCS)*, 32(2):1–29.

- Tatiana Ermakova, Benjamin Fabian, and E. Babina. 2015. Readability of privacy policies of health-care websites. In *12. Internationale Tagung Wirtschaftsinformatik*.
- Morgan C Evans, Jaspreet Bhatia, Sudarshan Wadkar, and Travis D Breaux. 2017. An evaluation of constituency-based hyponymy extraction from privacy policies. In *2017 IEEE 25th International Requirements Engineering Conference (RE)*, pages 312–321. IEEE.
- Benjamin Fabian, Tatiana Ermakova, and Tino Lentz. 2017. [Large-scale readability analysis of privacy policies](#). In *Proceedings of the International Conference on Web Intelligence, WI '17*, page 18–25, New York, NY, USA. Association for Computing Machinery.
- Federal Trade Commission. 1998. Privacy online: A report to congress. *Washington, DC, June*, pages 10–11.
- Geoffrey Fowler. 2021. [I checked apple’s new privacy ‘nutrition labels.’ many were false](#).
- FTC. 2013. Path social networking app settles ftc charges it deceived consumers and improperly collected personal information from users’ mobile address books. <https://www.ftc.gov/news-events/press-releases/2013/02/path-social-networking-app-settles-ftc-charges-it-deceived>.
- Joshua Gluck, Florian Schaub, Amy Friedman, Hana Habib, Norman Sadeh, Lorrie Faith Cranor, and Yuvraj Agarwal. 2016. How short is too short? implications of length and framing on the effectiveness of privacy notices. In *12th Symposium on Usable Privacy and Security (SOUPS)*, pages 321–340.
- Abhijith Athreya Mysore Gopinath, Vinayshekhar Bannihatti Kumar, Shomir Wilson, and Norman Sadeh. 2020. Automatic section title generation to improve the readability of privacy policies.
- Hamza Harkous, Kasseem Fawaz, Rémi Lebret, Florian Schaub, Kang G Shin, and Karl Aberer. 2018. Polisis: Automated analysis and presentation of privacy policies using deep learning. *arXiv preprint arXiv:1802.02561*.
- Mitra Bokaei Hosseini, Sudarshan Wadkar, Travis D Breaux, and Jianwei Niu. 2016. Lexical similarity of information type hypernyms, meronyms and synonyms in privacy policies.
- Po Hu, Donghong Ji, Chong Teng, and Yujing Guo. 2012. [Context-enhanced personalized social summarization](#). In *Proceedings of COLING 2012*, pages 1223–1238, Mumbai, India. The COLING 2012 Organizing Committee.
- Lucas D Introna. 1997. Privacy and the computer: why we need privacy in the information society. *Metaphilosophy*, 28(3):259–275.
- Priyank Jain, Manasi Gyanchandani, and Nilay Khare. 2016. Big data privacy: a technological perspective and review. *Journal of Big Data*, 3(1):25.
- Sarthak Jain, Madeleine van Zuylen, Hannaneh Hajishirzi, and Iz Beltagy. 2020. Scirex: A challenge dataset for document-level information extraction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7506–7516.
- California Department of Justice. 2012. [Attorney general kamala d. harris secures global agreement to strengthen privacy protections for users of mobile applications](#).
- Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy, and Roy Schwartz. 2018. [A dataset of peer reviews \(PeerRead\): Collection, insights and NLP applications](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1647–1661, New Orleans, Louisiana. Association for Computational Linguistics.
- Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. 2009. A nutrition label for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, page 4. ACM.
- Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2013. [Privacy as Part of the App Decision-Making Process](#), page 3393–3402. Association for Computing Machinery, New York, NY, USA.
- Moniba Keymanesh, Micha Elsner, and Srinivasan Parthasarathy. 2020. Toward domain-guided controllable summarization of privacy policies. In *Natural Legal Language Processing Workshop*. KDD.
- Vinayshekhar Bannihatti Kumar, Abhilasha Ravichander, Peter Story, and Norman Sadeh. 2019. Quantifying the effect of in-domain distributed word representations: A study of privacy policies. In *AAAI Spring Symposium on Privacy Enhancing AI and Language Technologies: PAL 2019*.
- Logan Lebanoff and Fei Liu. 2018. [Automatic detection of vague words and sentences in privacy policies](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3508–3517, Brussels, Belgium. Association for Computational Linguistics.
- Timothy Libert. 2018a. [An automated approach to auditing disclosure of third-party data collection in website privacy policies](#). In *Proceedings of the 2018 World Wide Web Conference, WWW '18*, page 207–216, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.

- Timothy Libert. 2018b. An automated approach to auditing disclosure of third-party data collection in website privacy policies. In *Proceedings of the 2018 World Wide Web Conference*, pages 207–216.
- Jiali Lin, Bin Liu, Norman Sadeh, and Jason I Hong. 2014. Modeling users’ mobile app privacy preferences: Restoring usability in a sea of permission settings. In *10th Symposium On Usable Privacy and Security ({SOUPS} 2014)*, pages 199–212.
- Thomas Linden, Rishabh Khandelwal, Hamza Harkous, and Kassem Fawaz. 2020. The privacy policy landscape after the gdpr. *Proceedings on Privacy Enhancing Technologies*, 2020(1):47–64.
- Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhiemedi, SA Zhang, Norman Sadeh, Alessandro Acquisti, and Yuvraj Agarwal. 2016a. Follow my recommendations: A personalized privacy assistant for mobile app permissions. In *Symposium on Usable Privacy and Security*.
- Bin Liu, Jiali Lin, and Norman Sadeh. 2014a. Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help? In *Proceedings of the 23rd international conference on World wide web*, pages 201–212.
- Fei Liu, Nicole Lee Fella, and Kexin Liao. 2016b. Modeling language vagueness in privacy policies using deep neural networks. In *2016 AAAI Fall Symposium Series*.
- Fei Liu, Jeffrey Flanigan, Sam Thomson, Norman Sadeh, and Noah A Smith. 2015. Toward abstractive summarization using semantic representations. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1077–1086.
- Fei Liu, Rohan Ramanath, Norman Sadeh, and Noah A. Smith. 2014b. [A step towards usable privacy policy: Automatic alignment of privacy statements](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 884–894, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Mary Madden, Michele Gilman, Karen Levy, and Alice Marwick. 2017. Privacy, poverty, and big data: A matrix of vulnerabilities for poor americans. *Wash. UL Rev.*, 95:53.
- Mary Madden, Lee Rainie, Kathryn Zickuhr, Maeve Duggan, and Aaron Smith. 2014. Public perceptions of privacy and security in the post-snowden era. *Pew Research Center*, 12.
- Ankit Malpani, Balaraman Ravindran, and Hema A Murthy. 2011. Personalized intelligent tutoring system using reinforcement learning.
- Florencia Marotta-Wurgler. 2019. [Does “notice and choice” disclosure regulation work? an empirical study of privacy policies.](#)
- Aaron K Massey, Jacob Eisenstein, Annie I Antón, and Peter P Swire. 2013. Automated text mining for requirements analysis of policy documents. In *2013 21st IEEE International Requirements Engineering Conference (RE)*, pages 4–13. IEEE.
- Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *ISJLP*, 4:543.
- Aleecia M McDonald and Lorrie Faith Cranor. 2010. Americans’ attitudes about internet behavioral advertising practices. In *Proceedings of the 9th annual ACM workshop on Privacy in the electronic society*, pages 63–72.
- Bruce M McLaren, Sung-Joo Lim, France Gagnon, David Yaron, and Kenneth R Koedinger. 2006. Studying the effects of personalized language and worked examples in the context of a web-based intelligent tutor. In *International Conference on Intelligent Tutoring Systems*, pages 318–328. Springer.
- Gabriele Meiselwitz. 2013. Readability assessment of policies and procedures of social networking sites. In *International Conference on Online Communities and Social Computing*, pages 67–75. Springer.
- Majd Mustapha, Katsiaryna Krasnashchok, Anas Al Bassit, and Sabri Skhiri. 2020. Privacy policy classification with xlnet (short paper). In *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, pages 250–257. Springer.
- Abhijith Athreya Mysore Gopinath, Shomir Wilson, and Norman Sadeh. 2018. [Supervised and unsupervised methods for robust separation of section titles and prose text in web documents](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 850–855, Brussels, Belgium. Association for Computational Linguistics.
- Kanthashree Mysore Sathyendra, Shomir Wilson, Florian Schaub, Sebastian Zimmeck, and Norman Sadeh. 2017a. Identifying the provision of choices in privacy policy text. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2774–2779, Copenhagen, Denmark. Association for Computational Linguistics.
- Kanthashree Mysore Sathyendra, Shomir Wilson, Florian Schaub, Sebastian Zimmeck, and Norman Sadeh. 2017b. [Identifying the provision of choices in privacy policy text](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2774–2779, Copenhagen, Denmark. Association for Computational Linguistics.

- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. **Abstractive text summarization using sequence-to-sequence RNNs and beyond**. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. **Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Brussels, Belgium. Association for Computational Linguistics.
- Najmeh Mousavi Nejad, Pablo Jabat, Rostislav Nedelchev, Simon Scerri, and Damien Graux. 2020. Establishing a strong baseline for privacy policy classification. IFIP International Conference on ICT Systems Security and Privacy Protection.
- Sansa News. 2021. Sansa news. <https://sansa.news/>.
- Helen Nissenbaum. 2004. Privacy as contextual integrity. *Wash. L. Rev.*, 79:119.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. **Know what you don’t know: Unanswerable questions for SQuAD**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Rohan Ramanath, Fei Liu, Norman Sadeh, and Noah A. Smith. 2014a. Unsupervised alignment of privacy policies using hidden markov models. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 605–610, Baltimore, Maryland. Association for Computational Linguistics.
- Rohan Ramanath, Florian Schaub, Shomir Wilson, Fei Liu, Norman Sadeh, and Noah Smith. 2014b. Identifying relevant text fragments to help crowdsource privacy policy annotations. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 2.
- Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. 2016. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 77–96.
- Abhilasha Ravichander, Alan W Black, Shomir Wilson, Thomas Norton, and Norman Sadeh. 2019. **Question answering for privacy policies: Combining computational and legal perspectives**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4947–4958, Hong Kong, China. Association for Computational Linguistics.
- General Data Protection Regulation. 2016. Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46. *Official Journal of the European Union (OJ)*, 59(1-88):294.
- Joel R Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T Graves, Fei Liu, Aleecia McDonald, Thomas B Norton, and Rohan Ramanath. 2015. Disagreeable privacy policies: Mismatches between meaning and users’ understanding. *Berkeley Tech. LJ*, 30:39.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. **A neural attention model for abstractive sentence summarization**. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 379–389, Lisbon, Portugal. Association for Computational Linguistics.
- Norman Sadeh, Ro Acquisti, Travis D Breaux, Lorrie Faith Cranor, Aleecia M Mcdonalda, Joel R Reidenberg, Noah A Smith, Fei Liu, N Cameron Russellb, Florian Schaub, et al. 2013. The usable privacy policy project: Combining crowdsourcing, machine learning and natural language processing to semi-automatically answer those privacy questions users care about. *Technical Report CMU-ISR-13-119, Carnegie Mellon University*.
- Kanthashree Mysore Sathyendra, Abhilasha Ravichander, Peter Garth Story, Alan W Black, and Norman Sadeh. 2017. Helping Users Understand Privacy Notices with Automated Query Answering Functionality: An Exploratory Study. Technical report.
- Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. 2015. A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 1–17.
- Executive Office of the President’s Council of Advisors on Science and Technology. 2014. Big data and privacy: A technological perspective.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. **Get to the point: Summarization with pointer-generator networks**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Yan Shvartzshnider, Ananth Balashankar, Thomas Wies, and Lakshminarayanan Subramanian. 2018. **RECIPE: Applying open domain question answering to privacy policies**. In *Proceedings of the Workshop on Machine Reading for Question Answering*, pages 71–77, Melbourne, Australia. Association for Computational Linguistics.

- Yan Shvartzshnaider, Noah Apthorpe, Nick Feamster, and Helen Nissenbaum. 2019. Going against the (appropriate) flow: a contextual integrity approach to privacy policy analysis. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, pages 162–170.
- Yan Shvartzshnaider, Ananth Balashankar, Vikas Patidar, Thomas Wies, and Lakshminarayanan Subramanian. 2020. Beyond the text: Analysis of privacy statements through syntactic and semantic role labeling. *arXiv preprint arXiv:2010.00678*.
- Peter Story, Sebastian Zimmeck, Abhilasha Ravichander, Daniel Smullen, Ziqi Wang, Joel Reidenberg, N Cameron Russell, and Norman Sadeh. 2019. Natural language processing for mobile app privacy compliance. In *AAAI Spring Symposium on Privacy Enhancing AI and Language Technologies: PAL 2019*.
- Peter Story, Sebastian Zimmeck, and Norman Sadeh. 2018. Which apps have privacy policies? In *Annual Privacy Forum*, pages 3–23. Springer.
- Auto TLDR. 2021. Auto tldr. <http://autotldr.io/>.
- Noriko Tomuro, Steven Lytinen, and Kurt Hornsborg. 2016. Automatic summarization of privacy policies using ensemble learning. In *Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy, CODASPY '16*, page 133–135, New York, NY, USA. Association for Computing Machinery.
- Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. Americans reject tailored advertising and three activities that enable it. Available at SSRN 1478214.
- Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. 2012. Smart, useful, scary, creepy: perceptions of online behavioral advertising. In *proceedings of the eighth symposium on usable privacy and security*, pages 1–15.
- FTC US Federal Trade Commission et al. 2012. Protecting consumer privacy in an era of rapid change: Recommendations for businesses and policymakers. *FTC Report*.
- Michela Del Vicario, Walter Quattrociocchi, Antonio Scala, and Fabiana Zollo. 2019. Polarization and fake news: Early warning of potential misinformation targets. *ACM Trans. Web*, 13(2).
- Harm de Vries, Dzmitry Bahdanau, and Christopher Manning. 2020. Towards ecologically valid research on language user interfaces. *arXiv preprint arXiv:2007.14435*.
- Xiaoyin Wang, Xue Qin, Mitra Bokaei Hosseini, Rocky Slavin, Travis D. Breaux, and Jianwei Niu. 2018. Guileak: Tracing privacy policy claims on user input data for android applications. In *Proceedings of the 40th International Conference on Software Engineering, ICSE '18*, page 37–47, New York, NY, USA. Association for Computing Machinery.
- Samuel D Warren and Louis D Brandeis. 1890. The right to privacy. *Harvard law review*, pages 193–220.
- Alan F Westin. 1968. Privacy and freedom. *Washington and Lee Law Review*, 25(1):166.
- S Wilson, F Schaub, A Dara, F Liu, S Cherivirala, P G Leon, M S Andersen, S Zimmeck, K Sathyendra, N C Russell, T B Norton, E Hovy, J R Reidenberg, and N Sadeh. 2016a. The creation and analysis of a website privacy policy corpus. In *Annual Meeting of the Association for Computational Linguistics, Aug 2016*. ACL.
- Shomir Wilson, Florian Schaub, Aswath Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Schaarup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N Cameron Russell, et al. 2016b. The creation and analysis of a website privacy policy corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1330–1340.
- Shomir Wilson, Florian Schaub, Rohan Ramanath, Norman Sadeh, Fei Liu, Noah A Smith, and Frederick Liu. 2016c. Crowdsourcing annotations for websites' privacy policies: Can it really work? In *Proceedings of the 25th International Conference on World Wide Web*, pages 133–143.
- Razieh Nokhbeh Zaeem, Safa Anya, Alex Issa, Jake Nimergood, Isabelle Rogers, Vinay Shah, Ayush Srivastava, and K Suzanne Barber. 2020. Privacycheck v2: A tool that recaps privacy policies for you. In *29th ACM International Conference on Information and Knowledge Management (CIKM)*. ACM. To appear.
- Razieh Nokhbeh Zaeem, Rachel L German, and K Suzanne Barber. 2018. Privacycheck: Automatic summarization of privacy policies using data mining. *ACM Transactions on Internet Technology (TOIT)*, 18(4):1–18.
- Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel Reidenberg, N Cameron Russell, and Norman Sadeh. 2019a. Maps: Scaling privacy compliance analysis to a million apps. *Proceedings on Privacy Enhancing Technologies*, 2019(3):66–86.
- Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel R. Reidenberg, N. Russell, and N. Sadeh. 2019b. Maps: Scaling privacy compliance analysis to a million apps. *Proceedings on Privacy Enhancing Technologies*, 2019:66 – 86.

Sebastian Zimmeck, Ziqi Wang, Lieyong Zou, Roger Iyengar, Bin Liu, Florian Schaub, Shomir Wilson, Norman M Sadeh, Steven M Bellovin, and Joel R Reidenberg. 2017. Automated analysis of privacy requirements for mobile apps. In *NDSS*.

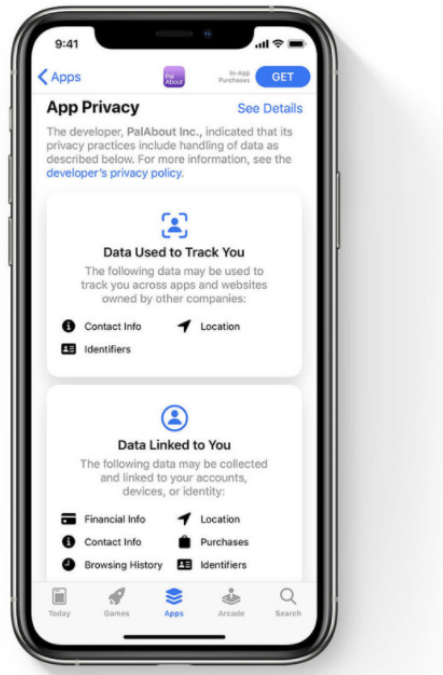


Figure 2: Example of privacy nutrition labels, disclosing information collected by companies and third parties through an application. Source: Apple.

A Privacy Nutrition Labels

Figure.2 includes an example of a privacy nutrition label, intended to disclose to a user the information a company and any third parties collect through an app. Apple requires developers to self-report the information for these nutrition labels.