

Going Against the (Appropriate) Flow: A Contextual Integrity Approach to Privacy Policy Analysis

Yan Shvartzshnaider
*New York University &
Princeton University*

Noah Apthorpe
Princeton University

Nick Feamster
Princeton University

Helen Nissenbaum
Cornell Tech

Abstract

We present a method for analyzing privacy policies using the framework of contextual integrity (CI). This method allows for the systematized detection of issues with privacy policy statements that hinder readers’ ability to understand and evaluate company data collection practices. These issues include missing contextual details, vague language, and combinatorial possible interpretations of described information transfers. We demonstrate this method in two different settings. First, we compare versions of Facebook’s privacy policy from before and after the Cambridge Analytica scandal. Our analysis indicates that the updated policy still contains fundamental ambiguities that limit readers’ comprehension of Facebook’s data collection practices. Second, we successfully crowdsourced CI annotations of 48 excerpts of privacy policies from 17 companies with 141 crowdworkers. This indicates that regular users are able to reliably identify contextual information in privacy policy statements and that crowdsourcing can help scale our CI analysis method to a larger number of privacy policy statements.

1 Introduction

Federal and state regulations require online services to notify consumers about information collection and sharing practices through privacy policies [9]. In principle, this “notice and choice” framework provides user control and seems to be a fair and transparent process. In practice, privacy policies are confusing [26], notoriously time consuming to read [20], and difficult to comprehend [37, 38]. Furthermore, privacy policies frequently do not conform with users’ expectations of company practices [18].

Researchers have demonstrated that privacy policies are vague and often incomplete, contributing to “misunderstanding among stakeholders, wherein stakeholders have different interpretations regarding the incomplete information” [3, 4]. In this paper, we use the theory of contextual integrity (CI) [22] to synthesize these existing privacy policy evaluation methods and provide a new formal approach for detecting specific types of ambiguities that interfere with read-

ers’ ability to understand the information collection practices described in privacy policies.

Our CI-based analysis method (Section 3) involves identifying and annotating contextual parameters of information flows described in privacy policies, specifically the *senders*, *recipients* and *subjects* of information, information types (*attributes*), and the conditions under which information may be transferred or collected (*transmission principles*). The resulting annotations allow descriptive and normative analyses based on a core principle of contextual integrity: Understanding and assessing the privacy implications of an information flow requires knowing the full context of the flow (i.e., all five contextual parameters). This assertion allows us to evaluate privacy policies for specific issues that hinder understandability, including information flow incompleteness, parameter bloating, and vagueness.

Incomplete information flows, which omit one or more contextual parameters, invite readers to interpret the missing parameters according to their own expectations, which may not match the actual practices of the company [3, 19]. Parameter bloating, or specifying more than one instance of a contextual parameter, increases the cognitive load required for readers to decipher which combinations of five parameters define fully-specified information flows that are actually allowed by the policy [21]. Finally, vague information flows contain language that makes it unclear which actors share the information or under what conditions the data collection practice described by the flow actually takes place.

Analyzing privacy policies on the basis of a consistent set of CI parameters also allows for seamless and rigorous comparison between policy versions and across many policies from different companies. Finally, the use of CI ties our method to an existing body of research using CI for descriptive and normative analyses of privacy implications in other settings [2, 12, 15, 33, 40, 44].

We demonstrate our CI analysis method by applying it in two different settings. First, we show that the technique can help evaluate privacy policies and how they evolve over time. We compare two versions of Facebook’s privacy policy

from before and after the Cambridge Analytica scandal [11] (Section 4). We find that the updated privacy policy fails to provide more clarity to the consumer, despite describing almost as twice as many total information flows as the previous policy. This lack of clarity is due to the updated policy describing more incomplete information flows than the previous policy, and because over 50% of information flows in both policies contain vague language. The updated policy also includes more instances of parameter bloating than the previous version.

Second, we show that crowdworkers are able to perform the annotation component of our CI analysis method by identifying CI parameters in privacy policy statements (Section 5). This indicates that CI analysis could be scaled to a large corpus of privacy policies in future research. We crowdsource the annotation of 48 excerpts of privacy policies from 17 companies with 141 Amazon Mechanical Turk workers. The overall high precision of crowdworker annotations (0.96) shows that regular users are able to reliably identify relevant contextual information in privacy statements.

In summary, this work makes the following contributions:

1. We present a method for annotating privacy policies using the contextual integrity framework (Section 3). The use of a structured framework allows rigorous analysis of difficult information policy statements and is applicable to policies across companies and sectors.
2. We demonstrate a range of analytical methods enabled by our approach through two applications: a comparative analysis of Facebook privacy policy updates (Section 4) and crowdsourced annotations of 48 privacy policy excerpts (Section 5).

To support future research and policymaking efforts, we plan to make the privacy policy annotations performed for this work publicly available to the wider community.

2 Related Work

Prior efforts by the research community have analyzed privacy policies in order to identify statements that are uninformative or potentially confusing to the reader. These works fall into two main categories: 1) Detecting textual ambiguity and vagueness in privacy policies, and 2) Privacy policy annotations.

Ambiguity and Vagueness. In 2016, Bhatia et al. [4] proposed a formal “theory of vagueness for privacy policy statements based on a taxonomy of vague terms” to show that statements with vague language affect readers’ perception of privacy risk from the described data collection practices. More recent work by Bhatia and Breaux [3] used frame semantics [10] to identify incomplete privacy statements that omit relevant contextual information. Textual ambiguity in privacy policies has also been the focus of work performing lexical analysis to extract hypernyms, meronyms, and synonyms in

information type descriptions [5, 8, 14]. These projects have aimed to build a concise ontology of information types described in privacy policies.

Our CI-based analysis benefits from these insights; however, we capture a more complete picture of data collection practices described in privacy policies including and beyond issues of textual ambiguity. We are able to evaluate privacy policy statements with respect to a broader space of issues that make it difficult for readers to assess whether the practices being described respect or violate privacy norms.

Using CI to analyze privacy policies is also supported by recent work showing the importance of contextual factors to users’ privacy expectations. In 2016, Rao et al. compared users’ privacy expectations to existing companies’ practices [24]. A total of 240 participants were asked to state their expectations for the data collection, sharing, and deletion practices of 16 websites. The results showed that users’ privacy expectations depend on the type of website and the type of information being exchanged.

In 2016, Martin and Nissenbaum [19] showed that when confronted with a privacy-related scenario that was missing some contextual information, respondents mentally supplemented the information, essentially generating a different version of the scenario. Martin and Nissenbaum also conducted a survey of 569 respondents presented with 40 scenarios with random combinations of contextual factors. The results showed that the “context of information exchange – how information is used and transmitted, the sender and receiver of the information – all impact the privacy expectations of individuals” [19].

Similar results were reported in 2018 by Bhatia and Breaux [3] in three studies that showed that adding relevant contextual information to the description of a data practice affects user’s perception of privacy risk. Specifically, users’ willingness to share information significantly increased with addition of statements describing the purpose and provision of choice.

Privacy Policy Annotations. In 2016, Wilson et al., [41] (Usable Privacy Project [28]) recruited law students to hand-annotate privacy policies with metadata tags such as “first party collection/use,” “user choice/control,” “data retention,” and “data security.” They then used the hand-labeled policies to train a machine learning algorithm for annotating policies with the same tags. This labelling taxonomy was used in more recent work [13] to train a neural network classifier to automatically annotate segments of privacy policies and to build a Question-Answering system that supports free-form querying of the privacy policy content. Wilson et al., [42, 43] also explored the feasibility of asking crowdworkers to answer questions on data collection practices. The results showed that the answers of the crowdworkers agreed with those of skilled annotators over 80% of the time, indicating that crowdsourcing can be used to identify paragraphs describing specific

practices in privacy policies.

These techniques [13, 41–43] aim to make users aware of certain information handling practices, such as third party data collection and use, data retention, and so forth, by labelling relevant paragraphs in privacy policies describing these behaviors. These labels serve as helpful landmarks to navigate users to relevant parts of lengthy policies; however, they still require further interpretation of the text of labeled sections to understand data collection details. In contrast, we use CI to annotate five information flow parameters, rather than a large labelling taxonomy. This allows us to directly evaluate privacy policies for specific properties, such as excessive or missing details that are impossible to detect using previous annotation methodologies.

Additional efforts have relied on NLP and ML techniques to analyze privacy policy text to identify provision of choice statements [30] and opt-out choice statements [29] in order to point the user to relevant parts of privacy policies. While extracting relevant paragraphs saves time for the interested reader, it does not provide a way of identifying issues with the policy itself, such as missing information or combinatorial interpretations due to overloaded contextual descriptions. Our CI approach allows for these analyses.

3 CI Analysis Method

We use the framework provided by CI to identify and annotate information flows and their component parameters described in privacy policy statements.

3.1 CI Overview

In contrast to other existing theories of privacy, Contextual Integrity (CI) defines privacy as the appropriateness of information flows determined by conformance with existing legitimate, informational norms specific to given social contexts [22]. In other words, a person’s privacy is *prima facie* violated when a transfer of information deviates from established norms in a particular context. For example, someone might view sharing Fitbit data with their doctor as appropriate but sharing the same data with an insurance company as a privacy violation. Changing the recipient of the information alters the flow, and as a consequence, could violate a contextual norm. The sources of these contextual norms can vary, ranging from law and regulation to societal beliefs and family values.

To facilitate analysis, CI offers a framework to describe information flows using 5-parameter tuples. These five parameters capture specific actors (*senders*, *recipients*, and *subjects*) involved in an information flow, the type (*attribute*) of information in the flow, and the condition (*transmission principle*) under which the information flow occurs. *Importantly, all five parameters must be specified in order to understand the context of an information flow, and changing even one parameter can affect a flow’s overall appropriateness.* This is a central

premise of CI theory; without stating all parameters characterizing an information flow, its context is underspecified and its implications are ambiguous. In the terms of “informed consent,” past research [3, 19] shows that privacy policy information flow statements which do not clearly state relevant contextual parameters create gaps in users’ understanding of data collection and use.

3.2 Privacy Policy Annotation

We use the following definitions to identify and label information flows and contextual parameters in privacy policy text. These annotations are the raw data for the analyses described in the following section. Annotation can be performed manually or formulated as crowdworking task for scalable application of the CI analysis method.

- **Information Flow.** Any self-contained description of a transfer of information. Information flows are typically single sentences or short paragraphs, but are also presented as bulleted lists in some privacy policy formats.
- **Sender.** Any entity (person, company, website, device, etc.) that transfers or shares information. This may be a pronoun or a specific entity, such as “Company A,” “strategic partners,” or “publisher.”
- **Recipient.** Any entity (person, company, website, device, etc.) that ultimately receives information. This may be a pronoun or a specific entity, such as “third party,” “developer,” “other users,” or “Company B and its affiliates.”
- **Transmission principle.** Any clause describing the “terms and conditions under which [...] transfers ought (or ought not) to occur” [22]. This includes descriptions of how information may be used or collected. Examples include “if the user gives consent,” “when an update occurs,” or “to perform specified functions.”
- **Attribute.** Any description of information type, instance, and/or example, such as “date of birth,” “credit card number,” “photos,” or, more generally, “personal information.”
- **Subject.** Any subjects of information exchanged in a flow. Subjects may be explicitly stated or implicitly described using pronouns and possessives.

For example, the following annotated statement from the Facebook privacy policy describes a single information flow:

We [Facebook]^{recipient} also collect contact information^{attribute} that you^{sender} provide if you upload, sync or import this information (such as an address book) from a device.^{TP}

This flow contains an explicit sender, recipient, attribute, and transmission principle (TP). The subject parameter is not included, but is implicitly the user agreeing to the privacy policy.

3.3 Information Flow & Parameter Analyses

We can use annotated information flows and parameters in privacy policy texts for a variety of analyses, including, but not limited to, the following.

Comparing Privacy Policy Versions. We can compare snapshots of a privacy policy across updates or get an aggregated view across different privacy policies. This offers insights into the general nature of the policy differences, including which parameters were preferentially added, removed, or modified.

Identifying Incomplete Flows. In order to understand the privacy implication of an information flow, it is important to provide a complete description with all five contextual parameters specified. Otherwise, consumers are left uninformed about company behavior [19]. Identifying privacy statements that underspecify information flows can reveal problematic sections of the privacy policies.

Diagnosing Vague Statements. The use of vague and ambiguous terminology in privacy policy statements makes it increasingly difficult for readers to reason about information flow appropriateness and privacy implications. Building on prior work [4, 25], we can use CI annotations to identify specific privacy statements that describe such ambiguous flows. This also makes it easier for regulators and policymakers to monitor the appearance of such statements across privacy policy updates and privacy policies from different companies.

Recognizing CI Parameter Bloating. CI parameter bloating occurs when a single information flow contains two or more semantically different CI parameters of the same type (e.g., two senders or four attributes) without a clear indication of how these parameter instances are related to each other. This creates an information flow with a combinatorial number of possible contexts. It is difficult for readers or regulators to determine which combinations of parameters describe contexts in which information flows actually take place. Previous research indicates that “eliminating connectives that clarify the relationship between ideas makes sentences harder to understand because readers are left to infer the relationship” [21]. CI parameter bloating is a specific example of this phenomenon.

4 Detecting Privacy Policy Ambiguities

Revelations about the misuse of consumer data by Facebook and Cambridge Analytica [11] rekindled the debate around users’ privacy and informed consent on such platforms. In response to public outcry, Facebook worked to rectify the situation by updating its privacy policy (data policy) on April 19, 2018.

We apply our CI analysis technique to the Facebook privacy policy from immediately before and after this update. We used the Brat rapid annotation tool [1], and the annotation guidelines in Section 3 to manually annotate information flows and CI parameters in the previous and updated policy versions.

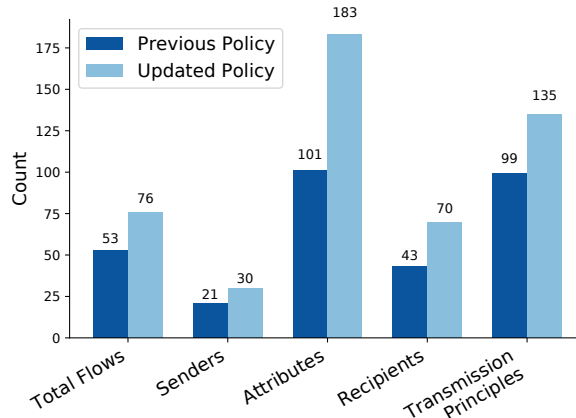


Figure 1: Distribution of unique CI parameters identified in the previous and updated Facebook privacy policies.

Two of the authors separately annotated both versions of the policy and performed statement by statement comparison to produce the final annotation.

From a legal perspective, the new document discloses more about the company’s information sharing practices. However, our CI analysis method reveals fundamental ambiguity issues present in both versions. These issues prevent users from interpreting new details in the updated version to fully understand how their data is being collected and shared.

Of course, Facebook’s privacy policy was unlikely written with contextual integrity in mind. We therefore intend the following analysis not as a criticism of Facebook per se, but as an opportunity to demonstrate our method and to point out issues common across privacy policies from many companies.

4.1 Information Flow Updates

We used our CI annotations to compare numbers (Figure 1) and specifics of each information flow parameter described in the previous and updated Facebook privacy policies. The updated Facebook privacy policy has about 50% more information flows than the previous policy (Figure 1). However, more information flows does not necessarily equal less confusion. Our analysis shows that many of the newly introduced information flows are incomplete (Section 4.2), are overloaded with CI parameters (Section 4.3) and/or use vague language (Section 4.4).

Sender. The updated policy offers a more detailed account of the sources of information transfer. It elaborates on categories from the previous privacy policy and also includes several new senders, such as “WhatsApp,” “connected TV,” and “a business,” which were not specified in the previous policy. Not surprisingly, the most frequent senders in both policies are Facebook and the user (Table 1).

Recipient. Similarly to the sender parameter, the updated version introduces new recipients, such as “people and busi-

CI Param	Version	Instances (frequency)
Recipients	Previous	we [Facebook] (25), Third party service, vendors, partners (25)
	Updated	we/us [Facebook] (37), Third party service, vendors, partners (41)
Senders	Previous	we [Facebook] (13), you (11)
	Updated	we [Facebook] (18), you (17)
Attributes	Previous	information (9), information about you (2), information we have (5), non-personally identifiable information only (2), data (2)
	Updated	information (18), content (5), information about you (4), information that we have (6), public information (4), communications (2), shipping and contact details (2).

Table 1: The most frequent recipients, senders, and attributes mentioned in the previous and updated Facebook privacy policies.

nesses outside the audience that you shared with,” “content creators,” “page admin,” “Instagram business profiles,” and “companies that aggregate.” As expected, the most common recipients in both versions are “Facebook,” and “third party service, vendors, partners” (Table 1).

Attribute. When describing the types of information being transferred or collected, the updated policy contains more attributes (183) than the previous policy (101). However, we note that some attributes from the previous policy were omitted in the update. The updated policy does not mention “user id” (opting for “username” instead), or “age range” (instead providing the example “. . . ad was seen by a woman between the ages of 25 and 34”). Generally, the updated policy describes new types of information and/or elaborates on information that was previously generic or abstract (Table 2). For example, the updated policy provides significantly more details about the type of content that is being collected about the user, including “racial or ethnic origins,” “health,” “events attended,” “interests,” “religious views,” “general demographics,” “political views,” “trade union membership,” and “philosophical beliefs.” Furthermore, the updated policy describes attributes not discussed in the previous policy, such as “connected TVs,” and “information about nearby Wi-Fi access points, beacons, and cell towers.”

Transmission Principle. When specifying conditions under which information transfer may be performed, the updated policy includes all conditions and information flow constraints in the previous policy. In addition, the updated policy also contains new transmission principles, such as “whether or not you have a Facebook account or are logged in to Facebook,” “to recognise you in photos, videos and camera experiences,”

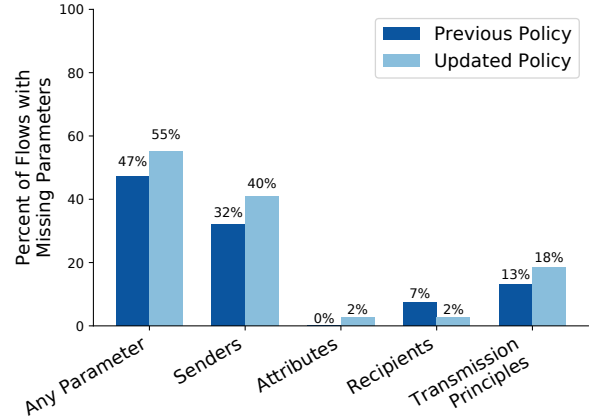


Figure 2: Percentage of incomplete information flows in Facebook’s previous and updated privacy policies with missing CI parameters.

“reshared or downloaded through APIs,” “to have lawful rights to collect, use and share your data before providing any data to us,” and many others (Table 2).

Subject. The subject of nearly all flows in both Facebook policy versions is the user. However, privacy policies from other companies may refer to non-implicit subjects in their information flow statements, especially for technologies targeted to specific populations, such as minors or other dependants of the user. Subject parameters are particularly important in such policies, as they may need to be explicitly disclosed to comply with privacy regulation, such as the U.S. Children’s Online Privacy Protection Act (COPPA).

4.2 Incomplete Information Flows

Our analysis of the Facebook privacy policy versions finds many described information flows with missing (non-specified) parameters (Figure 2). In the previous privacy policy, 47% (25/53) of flows are missing one or more parameters. In the updated policy, this number increases to 55% (43/77), including 16 incomplete flows from the previous policy and 27 new incomplete flows.

Missing Recipient. Table 3 lists the flows from both policies with missing recipient parameters. The previous policy only has three flows without an explicit recipient while the updated policy has two. Not stating information recipients forces users to infer what entities will have access to their information from other sources, often leading to incorrect notions of company behavior [19, 38]. Identifying the recipient can sometimes be difficult, as in the flow “*We are able to suggest that your friend tags you in a picture by comparing your friend’s pictures to information we’ve put together from your profile pictures and the other photos in which you’ve been tagged.*”

CI Parameter	Previous Policy	Updated Policy
Sender	people you share and communicate with	specific friends or accounts, friends and followers, other people using Facebook and Instagram, people
	devices, phones, computers, devices where you install or access our Services	connected TVs, web-connected devices you use that integrate with our Products
Recipient	family of companies that are part of Facebook	Facebook companies, Facebook company products
	people you share and communicate	audience they choose, specific friends or accounts, those you connect and share with around the world, people in your networks, friends and followers, people and businesses outside the audience that you shared with, anyone who can see the other person's content, anyone on or off our products
	partners conducting academic research, partners conducting surveys	research partners, research partners who we collaborate with, academics
	third-party companies who help us provide and improve our services or who use advertising or related products	websites that integrate with our products, other services that integrate with our products, companies that aggregate
	N/A	systems, devices and operating systems providing native versions of Facebook and Instagram (i.e. where we have not developed our own first-party apps), anyone on or off our product, content creator, seller, page admins, regulators, network
Attribute	information about how you use our services, how you use and interact with our services	information about any of your Instagram followers, the ads you see and how you use their services, other web-connected devices you use that integrate with our products, when you last used our products, whether a window is foregrounded or backgrounded, when you're using and have last used our products, identifiers from apps or accounts that you use, actions that you have taken on our products
	content about you	the features you use, life events, racial or ethnic origin, activities, where you live, what games you play, information about your interests actions and connections, who you are "interested in", your health, events you attend, interests, preferences, your religious views, general demographic, the places you like to go and the businesses and people you're near, whether you are currently active on Instagram messenger or Facebook, check-ins, websites you visit, other information about your Facebook friends from you, political views, trade union membership, philosophical beliefs
	information about the reach and effectiveness of their advertising	reports about the kinds of people seeing their ads, which Facebook ads led you to make a purchase or take an action with an advertiser, ads you see, family device ids
	Device information	information about operations and behaviours performed on the device, other identifiers unique to Facebook company products associated with the same device or account, available storage space
	N/A	information about nearby wi-fi access points, beacons, and cell towers
Transmission Principle	N/A	to detect when someone needs help, to recognise you in photos videos and camera experiences, help you stream a video from your phone to your tv, combat harmful conduct, can help distinguish humans from bots, to aid relief efforts, whether or not you have a Facebook account or are logged in to Facebook, reshared or downloaded through APIs, to have lawful rights to collect, use and share your data before providing any data to us and many others.

Table 2: List of notable CI parameters introduced or refined between the previous and updated Facebook privacy policies.

Information Flow	Version
<ul style="list-style-type: none"> For example, people may share a photo of you mention or tag you at a location in a post or share information about you that you shared with them. Bear in mind that information that others have shared about you is not part of your account and will not be deleted when you delete your account. You can manage the content and information you share when you use Facebook through the Activity Log tool. 	Previous
<ul style="list-style-type: none"> You can choose to provide information in your Facebook profile fields or life events about your religious views, political views, who you are “interested in” or your health. This and other information (such as racial or ethnic origin, philosophical beliefs or trade union membership) could be subject to special protections under the laws of your country For example, people can share a photo of you in a story or mention, tag you at a location in a post or share information about you in their posts or messages 	Updated

Table 3: Information flows in the previous and updated Facebook privacy policies with missing recipient parameters. In all the above case, the reader must infer who will end up receiving the information.

Missing Sender. The sender parameter is not specified in 17 (32%) flows in the previous policy nor in 31 (40%) flows in the updated policy. Many of the statements with missing senders describe “use-of-data,” i.e., they inform the consumer how the collected information will be used but not from where it is collected. Missing senders can easily lead to misinterpretations and false privacy expectations. For example, the source of the information in the following statement is unclear: “We collect information about the people, Pages, accounts, hashtags and groups you are connected to and how you interact with them.” Without knowing which of Facebook’s various services collect and send this information, users are unable to take specific action to avoid this data collection or adjust their behaviour on the platform.

Missing Transmission Principle. We identified 7 information flows in the previous policy where the transmission principle is missing. For example, the statement “We share information we have about you within the family of companies that are part of Facebook” does not specify under what conditions/constraints the information is being shared. Likewise, the statement “We also collect information about how you use our Services, such as the types of content you view or engage with or the frequency and duration of your activities” does not contain any transmission principles. Previous research [19]

Advertisers, app developers and publishers^{senders} can send us^{recipient} information through Facebook Business Tools that they use, including our social plug-ins (such as the Like button), Facebook Login, our APIs and SDKs or the Facebook pixel^{TP}. These partners provide information about your^{subject} activities off Facebook including information about your device, websites you visit, purchases you make, the ads you see and how you use their services^{attributes} whether or not you have a Facebook account or are logged in to Facebook.^{TP}

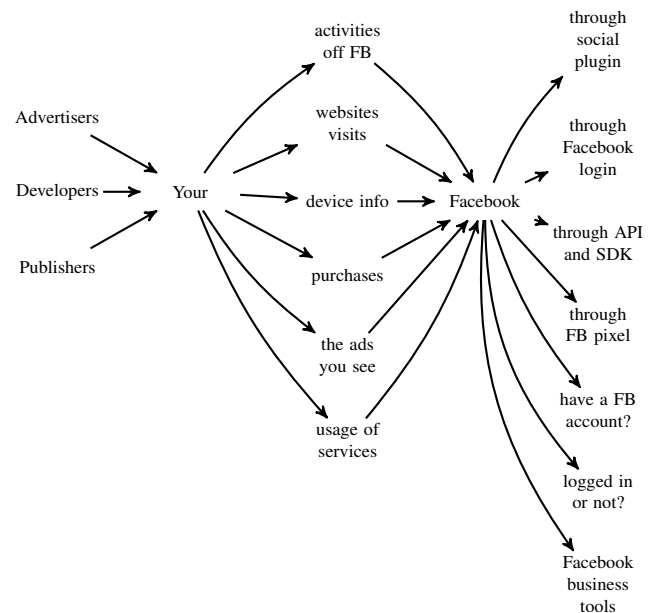


Figure 3: Example of CI parameter bloating in privacy policy text (top) and mapped into possible interpretations (bottom).

shows that in these instances consumers will end up guessing to guess when and for what reason information is collected.

The updated policy contains even more (14) flows with missing transmission principles. Without a transmission principle, flows like “We also receive information about your online and offline actions and purchases from third-party data providers who have the rights to provide us with your information” become ambiguous because it is not clear when or why this information is being collected.

4.3 CI Parameter Bloating

We used our annotations of the Facebook policy versions to identify information flows that suffer from CI parameter bloating, including the flow in Figure 3. At first glance, this statement seems transparent and informative. It explicitly specifies the type of information that is being exchanged, among what actors (sender, recipient, subject) and under what conditions. However, this is an example of CI parameter bloating. Taking




Figure 4: Extent of CI parameter bloating in privacy policy statements with multiple instances of at least two different CI parameters. *Not shown*: one outlier flow with 180 possible permutations in the previous policy and one outlier flow 492 possible permutations in the updated policy.

into account all the possible permutations results in total of $3 \text{ (senders)} \times 1 \text{ (subject)} \times 6 \text{ (attributes)} \times 1 \text{ (recipient)} \times 7 \text{ (TPs)} = 126$ possible flows.

How should the consumer reason about this privacy policy statement? Do all listed senders transfer all of these information types to Facebook or does each particular sender transmit a specific information type? Do flows with each sender/attribute pair occur under each listed TP or only specific ones? Even technically-savvy users will have difficulty reasoning about the many possible information flows with all combinations of each parameter type.

Our CI annotation analysis identifies several statements in both previous and updated policies that suffer from parameter bloating. The previous policy has 15 statements (28% of all flows) with multiple instances of two or more CI parameters. These statements have up to 4 senders, 20 attributes, 10 recipients, and 7 transmission principles and describe 4 to 180 total information flow permutations each (Figure 4). The updated policy has 30 statements (39% of all flows) with multiple instances of two or more CI parameters. These statements have up to 7 senders, 41 attributes, 8 recipients, and 8 transmission principles and describe 4 to 492 total information flow permutations each (Figure 4).

Given that an average consumer today spends little to no time reading privacy policies [31], it is unreasonable to assume that the even the most privacy-concerned citizen will dissect all possible combinations of this many multi-parameter flows. Instead, we believe that privacy policies should list all prescribed information flows explicitly, with each including all five parameters. This will increase the length of the policy and might be initially be construed to decrease readability [36]. However, adopting a regular 5-tuple structure for all policy statements will increase machine interpretability and allow user interfaces that can provide “different notices for different audiences” [32] by automatically parsing, filtering, and categorizing privacy policy statements.

4.4 Vague Information Flows

We used the annotations to identify flows which are prescribed using one or more combinations of vague terms (See

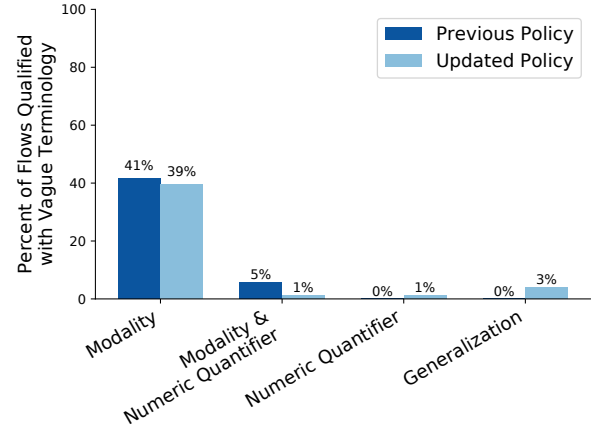


Figure 5: Percentage of information flows in Facebook’s previous and updated privacy policies qualified with various categories of vague terminology (as defined in Appendix A).

Appendix A for the vague terms taxonomy defined by Bhatia et al. [4]). As discussed in Section 2, vague information flows affect readers’ ability to accurately interpret whether the described data collection practice violates or respects their privacy. Figure 5 shows the percentage of vague information flows in Facebook’s previous and updated policies. In both policies, “modality” vagueness dominates, occurring in close to 40% of all flows. The updated policy does not represent a reduction in vague terminology from the previous version. Rather, the percentage of flows with vague terminology remains the same. This supports our initial claim that the updated policy does not contribute to clarity. The widespread occurrence of flows qualified by vague terminology further supports the problem that privacy policies are too often “obtuse and noncommittal [and] make it difficult for people to know what information a site collects and how it will be used” [37].

5 Crowdsourcing CI Privacy Policy Analysis

We also test our method to see whether crowdworkers are able to identify CI parameters in privacy policy statements. Specifically, we created an Amazon Mechanical Turk (AMT) Human Intelligence Task (HIT) to annotate 48 privacy policy excerpts. These included 16 excerpts from the Google privacy policy circa October 2017 and 16 pairs of excerpts from the privacy policies of 16 well-known companies¹ before and after May 2018 updates for the European General Data Privacy Regulation (GDPR). These excerpt pairs describe information flows with differences in parameters between the pre-GDPR and post-GDPR versions. The excerpts are also self-contained and do not require additional information from the policy to

¹Amazon, Fitbit, Indiegogo, LinkedIn, The New York Times, Microsoft, Shapeways, Slack, Spotify, Steam, Stripe, Tinder, Twitter, Uber, WhatsApp, Yelp

correctly annotate. The excerpts range from 21 to 113 words² and from 1 to 4 sentences for a total of 2621 words over 103 sentences.

We compared aggregated crowdworker annotations to ground-truth annotations from the authors (Section 5.4). The crowdworker annotations had an average precision of 0.96 across CI information flow parameters, indicating that the crowdworkers understood the relatively complex notion of information flow parameters and were able to correctly identify them in real privacy policy text. These results show that crowdworking can be an effectual tool for scaling CI annotation. We will release the crowdworker annotations as a public dataset for further analysis upon publication.

5.1 Annotation Task Design

We developed the annotation task as a Qualtrics [23] survey deployed on AMT. The task was designed to optimize annotation accuracy while minimizing cost.

Consent and Instructions. The first page of the annotation task was a consent form. Participants who did not consent were prevented from proceeding. The annotation task collected no personal information about crowdworkers and was approved by our university’s Institutional Review Board.

The task next presented annotation instructions (Appendix B), including a description of each information flow parameter that should be annotated (sender, attribute, recipient, and transmission principle) and an example annotated flow. The information flow parameter descriptions matched those described in Section 3.2.

Screening Questions. Each crowdworker was asked to annotate (highlight and label) all words and phrases corresponding to CI information flow parameters in three privacy policy excerpts (Figure 6). These excerpts served as screening questions to identify workers who are able to perform high-accuracy annotations. Workers whose annotations had an F_1 score of at least 0.7 compared to ground-truth expert annotations on the first screening question (for which the correct answer is given) and either of the next two screening questions were allowed to proceed with the task. Workers whose annotations did not meet this accuracy threshold did not proceed. This helped limit the effect and cost of workers who did not understand the task or who attempted to “cheat” by performing minimal annotations (e.g., highlighting just the first word in each excerpt).

Annotations. Each worker who passed the screening questions was asked to annotate 5 excerpts selected randomly from the 48 excerpts of interest. The format of the annotation questions was equivalent to the screening questions (Figure 6). The instructions were also repeated at the top of the page for workers to refer to if they wished.

Annotations of all excerpts from multiple workers were

²Mean: 55 words/excerpt, SD: 23 words/excerpt

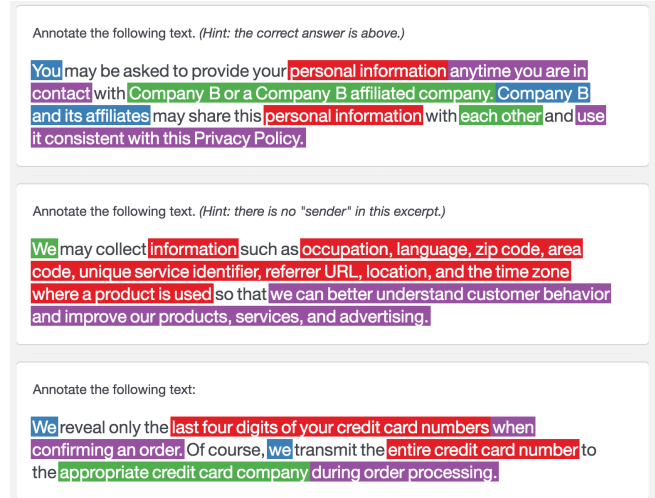


Figure 6: Screening questions to identify AMT workers who were able to perform high accuracy annotations. The ground truth annotations are shown with sender in blue, recipient in green, attribute in red, and transmission principle in purple.

collected, analyzed, and processed into the final crowdworker annotation for each privacy policy (Section 5.3).

5.2 Task Deployment

We first tested the annotation task on UserBob [39], a usability-testing service where users narrate their experience while performing tasks. We collected seven UserBob responses. All UserBob workers completed the task in less than 15 minutes. We used the UserBob responses to adjust task instructions to ameliorate worker confusion. Performing such “cognitive interviews” is common practice in survey design and development [35].

We deployed the annotation task as a HIT on AMT using TurkPrime [17], an online tool for researchers to easily manage AMT tasks. We limited the HIT to AMT workers in the United States with a HIT approval rating of 90–100% and at least 100 HITs approved. We did not collect or place any other criteria on the demographics or technical background of the AMT workers. 141 total workers accepted the HIT. Of these workers, 99 passed the screener questions. All 48 excerpts were annotated by between 7 and 12 workers (mean 10.2). AMT workers who did not pass the screening questions were automatically reimbursed \$0.25. AMT workers who passed the screening test and completed the entire annotation task were reimbursed \$1.50. Collecting all responses took approximately 4 hours from HIT launch until completion and cost a total of \$198 (including AMT fees).

5.3 Majority Vote Annotations

We are ultimately interested in acquiring the single highest-accuracy annotation for each privacy policy excerpt independent of individual workers. We therefore combine multiple

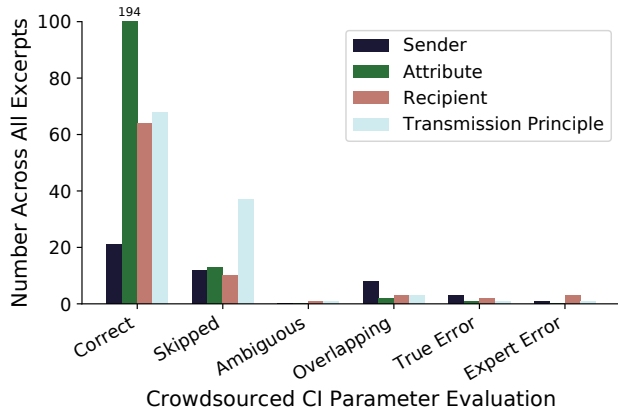


Figure 7: Comparison of crowdworker majority vote annotations to expert ground truth. Correct parameters are labeled in both annotations. Skipped parameters are only labeled by the expert. All other categories are described in Section 5.5.

annotations of each privacy policy excerpt into a “majority vote” annotation, which assigns each word in an excerpt to the CI parameter annotated by at least 50% of the workers presented with that excerpt. If fewer than 50% of workers labeled a word with the same parameter, then the word is given no label in the majority vote annotation.

The majority vote method reduces the influence of unreliable or adversarial crowdworkers who pass the screening questions. Assuming that such crowdworkers are a minority of those assigned to an excerpt, their annotations (or lack thereof) will not affect the final annotation.

5.4 Crowdworker Annotation Accuracy

Two of the authors annotated all excerpts prior to seeing the crowdworker results. The authors compared their independent annotations and manually resolved minor differences to create a single set of ground truth expert annotations.

We then found all discrepancies between the crowdworker and expert annotations and divided them into six categories: correct parameters, skipped parameters, ambiguous parameters, overlapping parameters, true errors, and expert errors (Figure 7, Section 5.5). Two of the authors performed this comparison manually to ensure accuracy and avoid the need for string matching heuristics. Categorizing the discrepancies allowed us to count the number of true positives (correct parameters), false negatives (skipped parameters), and false positives (true errors) and compute precision, recall, and F₁ scores³ for the crowdworker annotations (Table 4).

Overall, the high precision of the majority vote crowdworker annotations indicates that the majority of crowdworkers understood the CI annotation task and were able to correctly identify and highlight CI parameters in short privacy policy excerpts. A closer look at the flows where the majority

	Precision	Recall
Attribute	0.99	0.94
Sender	0.88	0.64
Recipient	0.97	0.86
TP	0.99	0.65

Table 4: Precision and recall scores of crowdworker majority vote annotations for each CI parameter across all excerpts.

of crowdworkers missed some parameters (Section 5.5) provides interesting insight into the reasons for the moderately lower recall numbers.

5.5 Evaluating Annotation Discrepancies

Analyzing the crowdworker annotations raises the question “What causes particular excerpts or CI parameters to be more difficult for crowdworkers to annotate than others?” We evaluated the discrepancies between crowdworker and expert annotations to better understand their underlying causes.

Ambiguity. The annotated excerpts include the various types of ambiguities found in the Facebook evaluation (Section 3). 32 excerpts describe incomplete information flows, 20 excerpts describe bloated information flows, and 27 excerpts include vague language. We used the Mann-Whitney *U* test to compare excerpts with and without incomplete information flows, parameter bloating, and vague language. We found no significant difference in F₁ scores based on these conditions ($p > 0.05$). This supports using crowdworking to scale CI analysis of privacy policies, because it indicates that crowdworkers can identify individual CI parameters even in privacy policy excerpts with semantic ambiguities that hinder interpretation of complete information flows, allowing post-annotation analysis to detect and evaluate these ambiguities.

Readability. We calculated Spearman correlations of the crowdworker majority vote annotation F₁ scores for each excerpt versus word count, Flesch-Kincaid Reading Ease [16], FOG Index [16], and number of CI parameters. However, all of the resulting correlation coefficients had absolute values < 0.5 and $p \gg 0.05$, indicating no significant correlations with F₁ score. This suggests that crowdworker difficulties with annotating certain excerpts were due to more nuanced factors than length or readability, which we explore by looking at each category of discrepancy in more detail.

Skipped Parameters. The most common type of discrepancy occurred when the crowdworkers simply neglected to annotate some or all instances of a given parameter. These discrepancies were the primary contributor to lowering recall scores without affecting precision.

The skipped parameters offer a glimpse into how a majority of the crowdworkers interpret the privacy policy excerpts. For example, we noted that the majority did not annotate a

³Precision = $\frac{TP}{TP+FP}$, Recall = $\frac{TP}{TP+FN}$, $F_1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$

sender in the information flow beginning with “*We may display your Profile name...*” presumably because they don’t see an “act of displaying” as sharing information. Additionally, in the information flow “*We collect information when you sync non-content like your email address book, mobile device contacts, or calendar with your account,*” both the expert and the crowdworkers labeled “email address book,” “mobile device contacts,” and “calendar” as attributes. However, the expert also labeled “information” as an attribute, while the majority of crowdworkers did not. From the CI analysis perspective, it is important to label “information” as an attribute because it acts as a superset, while the provided examples are merely selected instances. This is another type of privacy policy ambiguity that we would like to investigate in future work.

Alternatively, the crowdworkers may have found a few instances of each parameter and then moved on to the next excerpt without double-checking to ensure that none were missed. The crowdworkers may also have intentionally skipped parameters. This could be due to cognitive fatigue or the fact that crowdworkers are incentivized to finish the annotations as quickly as possible to optimize their hourly compensation rate.

Ambiguous Parameters. Ambiguous parameter discrepancies occurred when a CI parameter was mislabeled compared to the expert annotation, but the correct labeling is ultimately open to interpretation. Consider the sentence “*If you want to take full advantage of the sharing features we offer, we might also ask you to create a publicly visible Google Profile, which may include your name and photo.*” In this sentence, “publicly” could be interpreted as a recipient, i.e. the public would receive the data in the Google Profile. However, “publicly” could also be interpreted as a transmission principle i.e. the flow is from “you” to your “Google Profile” and the condition on the flow is that it is public. The expert labeled “publicly” as a recipient, while the crowdworker majority did not. We only identified 2 such ambiguous parameter discrepancies, indicating that CI information flow descriptions map naturally to privacy policy texts.

Overlapping Parameters. Overlapping parameter discrepancies occurred when a CI parameter was mislabeled compared to the expert annotation, but the text in question is part of two or more CI parameters simultaneously. We identified 16 overlapping parameters. Consider the excerpt “*When you use our services or view content provided by Google, we automatically collect and store certain information in server logs.*” The first clause (before the comma) could be interpreted as a single transmission principle, but the “you” could also be a sender. Variations on this issue were the primary cause of discrepancies for the “sender” parameter, i.e. the expert annotated an entire clause as a transmission principle but the majority vote annotation instead labeled a single word in the clause as a sender.

The presence of overlapping parameter discrepancies is due to a tradeoff in our implementation of the CI annotation task. We chose to allow only one CI parameter annotation per word in each excerpt to simplify the task for workers. Future work could instead ask each crowdworker to annotate only a single CI parameter type, simplifying the task from multi-class classification to binary classification. However, this would require more crowdworkers per policy and could lead to higher rates of false positives if crowdworkers are not forced to discriminate between different parameters.

True Errors. True errors occurred when the crowdworkers unambiguously misannotated a CI parameter. Fortunately, we only observed 7 true errors across all annotations. This implies that when a label makes it into the majority vote annotation (with sufficient workers contributing to the vote), it is most likely correct. The low frequency of true errors indicates that, with improvements to reduce the number of skipped parameters, crowdworking can be a high-accuracy method of obtaining CI annotations of privacy policies.

Expert Errors. Finally, we identified 5 cases where the crowdworker majority vote annotation was correct while the “ground-truth” expert annotation was incorrect. Most of these cases were due to the expert annotation missing a one-word sender or recipient, e.g. “we.” We did not adjust recall or precision scores to reflect the incorrect expert annotations, as these judgments were made after, and could have been influenced by, viewing the crowdworker annotations.

6 Discussion & Future Work

Privacy policies generally follow regulations devised by U.S. Federal Trade Commission (FTC) and drawn from the Code of Fair Information Practice Principles (FIPPs). This has resulted in an approach described as “notice and choice” [9], in which companies use privacy policies to notify consumers about their information collection and sharing practices and obtain consent, usually implicitly, when users continue to engage with a service in question. However, companies have found that legalistic language and vague terminology can produce privacy policies that adhere to the letter, but not the spirit, of these regulations. This affords companies leeway to collect and distribute large amounts of data while users remain ignorant of these practices, either because they understandably choose not to read complicated policies or because the policies do not provide enough specifics for even experts to understand exact company behavior [27].

In this paper, we argue that the notion of information flow appropriateness in the CI framework lends itself well to data collection practices described in privacy policies. Requiring that privacy policies have distinct five-parameter information flow descriptions for all data collection practices would complement ongoing efforts to improve readers’ understanding of privacy implications, move towards an efficient auditing of devices and services, and understand how privacy policies

relate to societal privacy norms.

6.1 Auditing Privacy Policies

The FTC and other regulatory bodies recommend that privacy policies include specific components, including the type of information collected, the entities that receive or store the information, uses of the information, and the conditions governing data acquisition and handling [7]. Our CI analysis method would enable a scalable auditing technique to check whether such requirements on the information flow descriptions in privacy policies were followed. The CI analysis method would also simplify continued auditing of privacy policies across updates by only requiring annotation of the differences between versions rather than each version in its entirety. This would indicate the CI parameter and information flow changes between versions, providing enough information for detecting ambiguous flows while requiring minimal annotation overhead.

Our CI analysis method can also enable auditing online service and device behavior in addition to privacy policies themselves. Recent studies have found several examples of technologies violating their own privacy policies [6], but such audits must often be performed manually because of the effort required to interpret individual privacy policies and compare their stipulations to observed behavior. Extracting information flow descriptions from privacy policies using our annotation technique could be the first step in an automated auditing pipeline. Information transfers from devices or online services could be observed using techniques such as network traffic analysis or taint tracing and automatically compared against CI parameters in their privacy policies.

6.2 Comparing Privacy Policies to Norms

Our analysis method adopts the notions of contextual integrity. On one hand, privacy policy statements made by a company should be compliant with existing regulation and legal statutes. On the other hand, they need to be informed by the context in which they operate. In other words, it becomes not just about being compliant with the law but also respecting users' privacy expectations and societal privacy norms. This challenge is particularly relevant in modern technosocial systems and platforms that operate in myriad of social contexts. Fortunately, the research community has already made steps towards addressing this challenge that can be furthered by our CI privacy policy analysis method.

Previous efforts [2, 34] have used the CI framework as a practical tool to discover privacy norms. These works in conjunction with CI annotations of policies will allow to determine whether the practices described in privacy policies align with users' privacy expectations and societal norms in general. This combination of privacy policy CI annotation and survey data could be used by companies to inform their behavior, as data collection practices more in line with user

norms are less likely to cause consumer backlash. It will also enable longitudinal ethnographic insight into how user norms are changing vis-a-vis privacy policies over time.

7 Conclusion

We present a privacy policy analysis method, based on the theory of contextual integrity, for detecting specific ways that privacy policies make it impossible for readers to assess whether the practices being described respect or violate privacy norms (Section 3).

We demonstrated the utility of the method in two settings: First, we analyzed versions of Facebook's privacy policy from before and after the Cambridge Analytica incident in April 2018 (Section 4).

Our analysis shows that the updated policy describes more information flows than the previous version, but that the updates do not improve the percentage of flows that contain vague language, omit parameters, or allow for many possible interpretations by including several parameters of the same type. Second, we show that non-expert users can help scale the CI analysis method by successfully crowdsourcing annotations of 48 privacy policy excerpts from 17 companies (Section 5).

In summary, our method complements existing privacy policy research and offers a new, scalable, approach to help study and protect user privacy.

References

- [1] Brat Rapid Annotation Tool. <http://brat.nlplab.org>.
- [2] Noah Apthorpe, Yan Shvartzshnaider, Arunesh Mathur, Dillon Reisman, and Nick Feamster. Discovering smart home internet of things privacy norms using contextual integrity. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT/UbiComp)*, 2018.
- [3] Jaspreet Bhatia and Travis Breaux. Semantic incompleteness in privacy policy goals. In *2018 IEEE 26th International Requirements Engineering Conference (RE)*, pages 159–169. IEEE, 2018.
- [4] Jaspreet Bhatia, Travis D Breaux, Joel R Reidenberg, and Thomas B Norton. A theory of vagueness and privacy risk perception. In *Requirements Engineering Conference (RE), 2016 IEEE 24th International*, pages 26–35. IEEE, 2016.
- [5] Jaspreet Bhatia, Morgan C Evans, Sudarshan Wadkar, and Travis D Breaux. Automated extraction of regulated information types using hyponymy relations. In *2016 IEEE 24th International Requirements Engineering Conference Workshops (REW)*, pages 19–25. IEEE, 2016.

- [6] Gordon Chu, Noah Apthorpe, and Nick Feamster. Security and privacy analyses of internet of things children’s toys. *IEEE Internet of Things Journal*, 2018.
- [7] Federal Trade Commission et al. Protecting consumer privacy in an era of rapid change: Recommendations for businesses and policymakers. *Washington, DC: Federal Trade Commission*, 2012.
- [8] Morgan C Evans, Jaspreet Bhatia, Sudarshan Wadkar, and Travis D Breaux. An evaluation of constituency-based hyponymy extraction from privacy policies. In *Requirements Engineering Conference (RE), 2017 IEEE 25th International*, pages 312–321. IEEE, 2017.
- [9] Federal Trade Commission. Privacy online: A report to congress. *Washington, DC, June*, pages 10–11, 1998.
- [10] Charles J Fillmore. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences*, 280(1):20–32, 1976.
- [11] Sarah Frier. Facebook Updates Policies After Privacy Outcry, Limits Data Use. <https://www.bloomberg.com/news/articles/2018-04-04/facebook-updates-policies-after-privacy-outcry-limits-data-use>, 2018.
- [12] Audrey Guinchard. Contextual integrity and eu data protection law: Towards a more informed and transparent analysis. *SSRN*, 2017.
- [13] H Harkous, K Fawaz, R Lebret, F Schaub, KG Shin, and K Aberer. Polisis: Automated analysis and presentation of privacy policies using deep learning. In *27th USENIX Security Symposium (USENIX Security 18)*, 2018.
- [14] Mitra Bokaei Hosseini, Sudarshan Wadkar, Travis D Breaux, and Jianwei Niu. Lexical similarity of information type hypernyms, meronyms and synonyms in privacy policies. In *2016 AAAI Fall Symposium Series*, 2016.
- [15] Gordon Hull, Heather Richter Lipford, and Celine Latulipe. Contextual gaps: privacy issues on Facebook. *Ethics and information technology*, 13(4):289–302, 2011.
- [16] Peter Kincaid, Robert Fishburne Jr, Richard Rogers, and Brad Chissom. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. Technical report, Naval Technical Training Command Millington TN Research Branch, 1975.
- [17] Leib Litman, Jonathan Robinson, and Tzvi Abberbock. Turkprime. com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior research methods*, 49(2):433–442, 2017.
- [18] Kirsten Martin. Privacy notices as tabula rasa: An empirical investigation into how complying with a privacy notice is related to meeting privacy expectations online. *Journal of Public Policy & Marketing*, 34(2):210–227, 2015.
- [19] Kirsten Martin and Helen Nissenbaum. Measuring privacy: an empirical test using context to expose confounding variables. *Colum. Sci. & Tech. L. Rev.*, 18:176, 2016.
- [20] Aleecia McDonald and Lorrie Faith Cranor. The cost of reading privacy policies. *ISJLP*, 4:543, 2008.
- [21] Anca Micheti, Jacquelyn Burkell, and Valerie Steeves. Fixing broken doors: Strategies for drafting privacy policies young people can understand. *Bulletin of Science, Technology & Society*, 30(2):130–143, 2010.
- [22] Helen Nissenbaum. *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press, 2010.
- [23] Qualtrics. www.qualtrics.com, 2018.
- [24] Ashwini Rao, Florian Schaub, Norman Sadeh, Alessandro Acquisti, and Ruogu Kang. Expecting the unexpected: Understanding mismatched privacy expectations online. In *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*, pages 77–96, Denver, CO, 2016. USENIX Association.
- [25] Joel Reidenberg, Jaspreet Bhatia, Travis Breaux, and Thomas Norton. Ambiguity in privacy policies and the impact of regulation. *The Journal of Legal Studies*, 45(S2):S163–S190, 2016.
- [26] Joel Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James Graves, Fei Liu, Aleecia McDonald, Thomas Norton, and Rohan Ramanath. Disagreeable privacy policies: Mismatches between meaning and users’ understanding. *Berkeley Tech. LJ*, 30:39, 2015.
- [27] Joel Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James Graves, Fei Liu, Aleecia McDonald, Thomas Norton, and Rohan Ramanath. Disagreeable privacy policies: Mismatches between meaning and users’ understanding. *Berkeley Tech. LJ*, 30:39, 2015.
- [28] Norman Sadeh, Alessandro Acquisti, Travis D Breaux, Lorrie Faith Cranor, Aleecia McDonald, Joel Reidenberg, Noah Smith, Fei Liu, Cameron Russell, Florian Schaub, et al. The usable privacy policy project. Technical report, CMU-ISR-13-119, Carnegie Mellon University, 2013.

- [29] Kanthashree Mysore Sathyendra, Florian Schaub, Shomir Wilson, and Norman Sadeh. Automatic extraction of opt-out choices from privacy policies. In *AAAI Fall Symposium on Privacy and Language Technologies*, 2016.
- [30] Kanthashree Mysore Sathyendra, Shomir Wilson, Florian Schaub, Sebastian Zimmeck, and Norman Sadeh. Identifying the provision of choices in privacy policy text. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2764–2769, 2017.
- [31] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. Designing effective privacy notices and controls. *IEEE Internet Computing*, 2017.
- [32] Florian Schaub, Rebecca Balebako, Adam L Durity, and Lorrie Faith Cranor. A design space for effective privacy notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, pages 1–17, 2015.
- [33] Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. Learning privacy expectations by crowdsourcing contextual informational norms. In *Fourth AAI Conference on Human Computation and Crowdsourcing*, 2016.
- [34] Yan Shvartzshnaider, Schrasing Tong, Thomas Wies, Paula Kift, Helen Nissenbaum, Lakshminarayanan Subramanian, and Prateek Mittal. Learning privacy expectations by crowdsourcing contextual informational norms. In *Fourth AAI Conference on Human Computation and Crowdsourcing*, 2016.
- [35] Seymour Sudman, Norman M Bradburn, Norbert Schwarz, and Terri Gullickson. Thinking about answers: The application of cognitive processes to survey methodology. *Psychocritiques*, 42(7):652, 1997.
- [36] The Center for Information Policy Leadership. Ten Steps to Develop a Multilayered Privacy Notice. https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/ten_steps_to_develop_a_multilayered_privacy_notice_white_paper_march_2007_.pdf, 2007.
- [37] Joseph Turow, Michael Hennessy, and Amy Bleakley. Consumers’ understanding of privacy rules in the marketplace. *Journal of consumer affairs*, 42(3):411–424, 2008.
- [38] Joseph Turow, Michael Hennessy, and Nora Draper. Persistent Misperceptions: Americans? Misplaced Confidence in Privacy Policies, 2003–2015. *Journal of Broadcasting & Electronic Media*, 62(3):461–478, 2018.
- [39] UserBob. <https://userbob.com/>, 2018.
- [40] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. Android permissions remystified: A field study on contextual integrity. In *USENIX Security Symposium*, pages 499–514, 2015.
- [41] Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Schaarup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N Cameron Russell, et al. The creation and analysis of a website privacy policy corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1330–1340, 2016.
- [42] Shomir Wilson, Florian Schaub, Frederick Liu, Kanthashree Mysore Sathyendra, Daniel Smullen, Sebastian Zimmeck, Rohan Ramanath, Peter Story, Fei Liu, Norman Sadeh, et al. Analyzing privacy policies at scale: From crowdsourcing to automated annotations. *ACM Transactions on the Web (TWEB)*, 13(1):1, 2018.
- [43] Shomir Wilson, Florian Schaub, Rohan Ramanath, Norman Sadeh, Fei Liu, Noah A. Smith, and Frederick Liu. Crowdsourcing annotations for websites’ privacy policies: Can it really work? In *Proceedings of the 25th International Conference on World Wide Web, WWW ’16*, pages 133–143, Republic and Canton of Geneva, Switzerland, 2016. International World Wide Web Conferences Steering Committee.
- [44] Michael Zimmer. Privacy on planet google: Using the theory of contextual integrity to clarify the privacy threats of google’s quest for the perfect search engine. *J. Bus. & Tech. L.*, 3:109, 2008.

Appendices

A Summary of Vagueness Categories as Defined by Bhatia et al. [4]

Category	Definition	Example Terms
Conditionality	it is not clear what is the condition associated with information transfer	“as needed”, “as necessary”, “as appropriate”, “depending”, “sometimes”, “as applicable”, “otherwise reasonably determined”, “from time to time”
Generalization	action or information types are too abstract or vague	“typically”, “normally”, “often” , “general”, “usually”, “generally”, “commonly ”, “among other things”, “widely”, “primarily”, “largely”, “mostly”
Modality	Hard to estimate the possibility of occurrence	“likely”, “may”, “can”, “could” “would”, “might”, “could”, “possibly”
Numeric Quantifier	Vague numeric quantifier	“certain”, “most”, “majority”, “many”, “some” “few”

B CI Annotation Task Instructions

Annotation Task Directions

The following questions contain short excerpts from website privacy policies. These sentences or paragraphs describe how the websites collect and transfer personal information about their users.

Your task is to identify and highlight the following elements in these excerpts:

Type of information: The type of information that is being collected or transferred. Examples include "date of birth," "credit card number," "photos," or, more generally, "personal information."

Sender of information: The entity (person, company, website, device, etc.) that transfers or shares the information. This may be a pronoun (e.g. "we") or a specific entity, such as "Company A," "strategic partners," or "publisher."

Recipient of information: The entity (person, company, website, device, etc.) that ultimately receives the information. This may be a pronoun (e.g. "we") or a specific entity, such as "third party," "developer," "other users," or "Company B and its affiliates."

When or why the information is collected or how it is used: Any condition that determines if or when the information is collected or transferred, or how the information may be used. These often, but not always, start with "if," "when," "for," or "in order to." Examples include "if the user gives consent," "when an update occurs," or "to perform specified functions."

Additional Instructions

Some privacy policy excerpts will not contain all 4 elements. For example, an excerpt may describe the sender of certain information but not the recipient. Excerpts may also have multiple of the same element. For example, an excerpt may describe several types of information. Do your best to identify as many elements as accurately you can.

Example

Use the first question below to practice highlighting and labeling. Highlight text then click the appropriate element type to add a label. You can re-highlight an existing label and click "remove" to remove it. The correct labeling and a video demonstration are below:

You may be asked to provide your **personal information** **anytime you are in contact** with **Company B or a Company B affiliated company.** **Company B and its affiliates** may share this **personal information** with **each other** and **use it consistent with this Privacy Policy.**