Investigating Privacy Policies using PolicyLint Tool

Tricia Camaya
Norfolk State University

Investigating Privacy Policies using PolicyLint tool

Tricia Camaya

Norfolk State University

COVA CCI Undergrad Cyber Research

December 9, 2022

Abstract

Organizations essentially inform clients about data collection and sharing practices through privacy policies. Recent research has proposed tools to help users better comprehend these lengthy and intricate legal documents that summarize collection and sharing. However, these instruments have a significant flaw. They overlook the possibility of contradictions within a particular policy. This paper introduces PolicyLint, a tool for analyzing privacy policies that simultaneously considers negation and varying semantic levels of data objects and entities. PolicyLint accomplishes this by using sentence-level natural language processing to automatically create ontologies from a large corpus of privacy policies and capturing both positive and negative statements regarding data collection and sharing. Using PolicyLint, I examined the policies of 300 apps and found that some contained contradictions that could indicate false statements. I manually check 100 contradictions, spotting troubling patterns like the use of misleading presentation, attempts to redefine terms that are commonly understood, and tracking information that is made possible by sharing or collecting data that can be used to derive sensitive information. As a result, automated privacy policy analysis is significantly improved by PolicyLint.

Keywords: semantic, negation, entities, contradictions, ontologies

Investigating Privacy Policies using PolicyLint tool

Some of the most private information about users such as private communications, precise location data, and even health measurements are collected, managed, and transmitted by mobile apps. These applications frequently send this data to first or third parties. If it is outlined in the app's privacy policy, such data collection and sharing is frequently deemed legal. Privacy policies are complex legal documents that are typically lengthy, ambiguous, and challenging to understand for novices, experts, and algorithms. As a result, it's hard to tell if app developers follow privacy policies, which can help app markets and different examiners recognize protection infringement or assist with finishing clients pick more-security amicable applications. Recent research has begun to investigate whether app behavior matches privacy policy statements. However, the previous research did not consider privacy policies' contradictions; these inconsistencies might prompt incorrect understanding of sharing and assortment rehearses. There are two main obstacles that must be overcome before contradictions can be found. To start with, security approaches allude to data at various semantic granularities. For instance, a policy may initially discuss its practices using broad terms like "personal information" but later use more specific terms like "email address" to describe them. This issue has previously been addressed by crowdsourcing data object ontologies; however, such crowdsourced information is neither complete nor accurate, nor is it simple to collect. Second, previous methods haven't been able to accurately detect negative statements because they have relied on bi-grams (like "not share") or only verb modifiers (like "will share X except Y"), leaving out more complicated statements like "will share X except Y." To determine the correct meaning of a policy statement (such as "not sharing" versus "sharing" information), negative statements must be modelled. To fully describe contradictions, one must address the two challenges that came before it. PolicyLint, a tool for automatically detecting potential contradictions between software privacy policies' sharing and collection practices. Policies become unclear as a result of contradictions, causing confusion for both humans and any automated system that relies on interpreting the policies. PolicyLint defines two contradiction groupings based on these use cases. Logical contradictions are policy statements that contradict one another and are more likely to harm people if analysts and users don't know about them. A policy that initially states that it does not collect personal information later reveals in the fine print that it does collect a user's name and email address for advertisers is one example. Automated techniques that reason over policy statements may be harmed by narrow definitions. To make decisions that are wrong or inconsistent, which could lead to policies that are unclear. The first tool with the sophistication required to reason about both negative sentiments and statements with varying levels of specificity, which is necessary for locating contradictions, is PolicyLint. Not every lint finding is necessarily a real bug, just like with any static approach. An external control or another context that the tool cannot verify could, for instance, mitigate potential bug conditions. The results of a lint finding can frequently only be verified by an individual. In the case of PolicyLint, privacy policies are intricate legal documents that have the potential to be intentionally evasive, ambiguous, or deceptive even for the human eye. Regardless of these difficulties, PolicyLint gathers long, convoluted strategies into a little arrangement of competitor issues important to human or algorithmic examination.

Producing ontologies from privacy policies automatically

PolicyLint extracts ontologies for both data objects and entities from a large corpus of privacy policies by employing an expanded set of Hearst patterns, such as “W such as X, Y, and Z.” PolicyLint is more thorough and scalable than crowdsourcing efforts.

Automatic extraction of privacy practices at the sentence level

PolicyLint captures data collection and sharing as a four-tuple by employing sentence-level NLP and type-dependency information: entity, actor, action, and data object. "We [actor] share [action] personal information [data object] with advertisers [entity]," for instance, the correct identification of negative statements necessitates sentence-level NLP.

Automated investigation of privacy policies' contradictions

An algorithmic approach to identifying contradictory policy statements is provided by our formal modeling of nine different kinds of contradictions caused by semantic relationships between terms. The foundation for ensuring the soundness of automated policy tools and identifying potentially misleading policy statements is provided by our groupings of narrowing definitions and logical contradictions.

Analyzing contradictions manually to identify trends

Surprisingly, there are a lot of policy contradictions. These inconsistencies include making broad claims that personal information will be protected early on in a policy but later making exceptions for data that the authors try to define as not personal, that could be used to derive sensitive information (like IP addresses and location), or that some regulators consider sensitive but not others.

Usage for PolicyLint

There are four main uses for PolicyLint. First, PolicyLint can be utilized by policy writers to lessen the likelihood of publishing misleading policies. Second, PolicyLint's definition of logical contradictions can be used by regulators to identify deceptive policies. Thirdly, PolicyLint can be used in a similar way on app stores like Google Play to make sure that apps' privacy policies don't contain false claims. They can also use PolicyLint's extraction of policy statements to automatically generate privacy labels that users can see on the market to encourage them to use apps that don't invade their privacy. Finally, PolicyLint's fine-grained extraction of policy statements, formalization of logical contradictions, and narrowing of definitions can help ensure the soundness of automated methods for analyzing privacy policies.

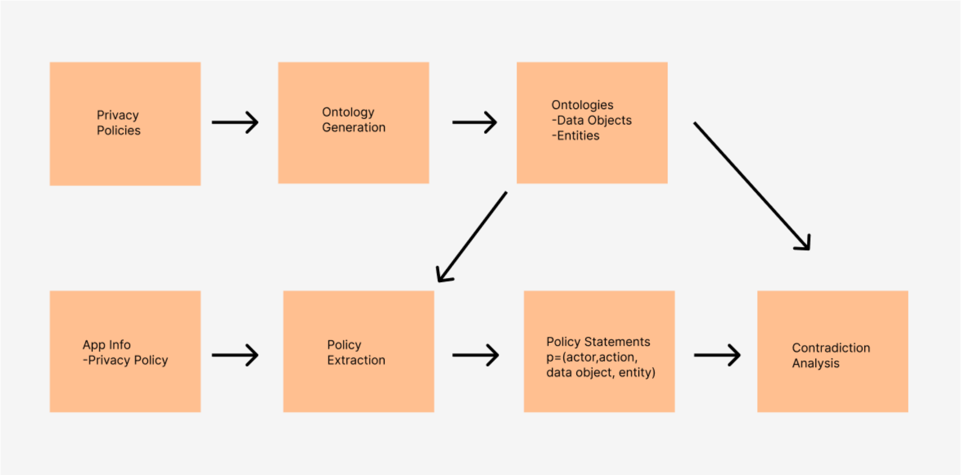

A Quick Look at PolicyLint

PolicyLint

PolicyLint aims to find contradictions in software-specific privacy policies. Like software lint tools, it issues privacy warnings based on contradictory sharing and collection statements in policies. However, manual verification of these warnings is required. "Candidate contradictions" within policies are found by PolicyLint. When considered in the most conservative (i.e., context-insensitive) way, a candidate contradiction is a pair of policy statements that contradict one another. Analyst-validated candidate contradictions are referred to as "validated contradictions.” "Due to the fundamental issues of ambiguity in interpreting the meaning of sentences in natural language (multiple interpretations of the same sentence), manual verification is required.” Take, for instance, a popular recipe app's privacy policy (com.omniluxtrade.allrecipes). "We do not collect personally identifiable information from our users" is stated in one section of the policy. This sentence makes it abundantly clear that the app does not collect any personal data. However, further down the policy, it is stated that "We may collect your email address in order to send information, respond to inquiries, and other requests or questions. "Since email addresses are considered private information, this sentence clearly contradicts the previous one. This is the root of the underlying contradiction. Even though PolicyLint is not the first NLP tool to analyze privacy policies, locating contradictions necessitates addressing two major obstacles.

Informational references can be expressed at a variety of semantic levels

Ontologies are used in previous research to capture subsumptive (also known as "is-a") relationships between terms; However, these ontologies are crowdsourced, and the authors manually define subsumptive relationships, raising concerns regarding comprehensiveness, and scalability. For instance, prior work uses only 50 and 30 policies, respectively, to construct their ontology. While large-scale crowdsourcing is impractical due to limited resources, crowdsourced ontologies could be comprehensive with unlimited time and manpower. In addition, the specific relationships that are required to reason about the data types and entities referred to in privacy policies are not all captured by the general-purpose ontologies that are currently in use.

Statements against collection and negative sharing are included in privacy policies

Most of the previous research operates at the paragraph level and is unable to capture negative sharing statements. Complex statements, such as "will share personal information except your email address" are missed by prior research that does capture negative statements. Such earlier work extricates coarse-grained synopses of strategy articulations, document level and can never unequivocally demonstrate negative statements or the elements in question. Semantics is influenced by sentence structure: Statements of sharing and collection typically adhere to a set of templates that can be learned. These templates are used by PolicyLint to extract a four-tuple from statements like these: entity, actor, action, and data object. "We [actor] share [action] personal information [data object] with advertisers [entity]," for instance, additionally, the sentence structure provides a deeper understanding of more intricate negative sharing. For instance, "We share personal information with advertisers, with the exception of your email address." PolicyLint builds on top of existing parts-of-speech and dependency parsers to extract such semantics from policy statements. Ontologies are encoded in privacy policies: The legal nature of privacy policies necessitates those general terms be defined in terms of examples or their components. PolicyLint automatically generates an ontology for policies (one for data objects and one for entities) by processing many privacy policies.

Conclusion

Natural language processing is used in the privacy policy analysis tool PolicyLint to find contradictory sharing and collection practices. By automatically generating domain ontologies, PolicyLint explains contradictory policy statements that occur at various semantic granularity levels. A plethora of concerning issues with privacy policies upon further examination, including misleading presentations and redefining common terms. The foundation for ensuring the soundness of automated policy analysis and identifying potentially deceptive policies is provided by PolicyLint's fine-grained extraction techniques, formalization of narrowing definitions, and identification of logical contradictions.

References

Benandow. (n.d.). Benandow/Privacypolicyanalysis: This repository holds the code for Policylint and PoliCheck, which identifies internal contradictions within privacy policies and analyzes data flow to ensure privacy policy consistency. GitHub. Retrieved November 27, 2022, from https://github.com/benandow/PrivacyPolicyAnalysis

Federal Trade Commission Act Section 5: Unfair or Deceptive Acts or Practices. (2016). https://www.federalreserve.gov/boarddocs/supmanual/cch/ftca.pdf

Google. (n.d.). All recipes full - apps on Google Play. Google. Retrieved November 27, 2022, from https://play.google.com/store/apps/details?id=com.omniluxtrade.allrecipes&gl=US

Manandhar, S., Kafle, K., Andow, B., Singh, K., & Nadkarni, A. (n.d.). Smart Home Privacy Policies Demystified: A Study of Availability, Content, and Coverage. https://www.adwaitnadkarni.com/downloads/manandhar-sec22.pdf

Slavin, R., Wang, X., Hosseini, M., Hester, J., Krishnan, R., Bhatia, J., Breaux, T., & Niu, J. (n.d.). Toward a Framework for Detecting Privacy Policy Violations in Android Application Code. Retrieved December 8, 2022, from https://www.cs.cmu.edu/~./breaux/publications/slavin-icse16.pdf

Stamey, J. W., & Rossi, R. A. (2009). Automatically identifying relations in privacy policies. Proceedings of the 27th ACM International Conference on Design of Communication - SIGDOC '09. https://doi.org/10.1145/1621995.1622041

Wang, X., Qin, X., Hosseini, M., Slavin, R., Breaux, T., & Niu, J. (n.d.). GUILeak: Identifying Privacy Practices on GUI-Based Data. Retrieved December 8, 2022, from https://cdn.stmarytx.edu/wp-content/uploads/2020/10/GUILeak-Identifying-Privacy-Practices-on-GUI-Based-Data.pdf

Zaeem, R. N., German, R. L., & Barber, K. S. (2018). PrivacyCheck: Automatic Summarization of Privacy Policies Using Data Mining. Undefined. https://www.semanticscholar.org/paper/PrivacyCheck%3A-Automatic-Summarization-of-Privacy-Zaeem-German/351cc7c7699c9fce80f77c86e8572bc7edd5253a