Towards Usable Privacy Policies: Semi-automatically Extracting Data Practices From Websites' Privacy Policies

Norman Sadeh[1]*, Alessandro Acquisti[1], Travis D. Breaux[1], Lorrie Faith Cranor[1], Aleecia M. McDonald[2], Joel Reidenberg[3], Noah A. Smith[1], Fei Liu[1], N. Cameron Russell[3], Florian Schaub[1], Shomir Wilson[1], James T. Graves[1], Pedro Giovanni Leon[1], Rohan Ramanath[1], Ashwini Rao[1]

[1]Carnegie Mellon University
Pittsburgh, PA 15213

[2]Stanford University
Stanford, CA 94305

[3]Fordham University
New York, NY 10023

1. MOTIVATION

Natural language privacy policies have become the de facto standard "notice and choice" method on the Web, in order to communicate a website's data practices. Yet, website privacy policies are often complex and difficult to understand. As a result, few users bother to read them [9]. It has been proposed to improve notice and choice mechanisms by making privacy practices machine-readable, e.g., in the Platform for Privacy Preferences (P3P) or Do Not Track (DNT) initiatives. While those initiatives made significant progress in terms of standardizing data collection and usage policies, as well as stipulating dialogue between stakeholders, many website operators are reluctant to adopt such approaches.

In our work, we build on recent advances in natural language processing (NLP), privacy preference modeling, crowdsourcing, and privacy interface design in order to develop a practical framework based on a website's existing natural language privacy policy that empowers users to more meaningfully control their privacy, without requiring additional cooperation from website operators.

2. APPROACH OVERVIEW

Research on user preference modeling suggests that a small number of key features in privacy policies largely determine the user's comfort level and privacy concerns for a visited website [3, 6]. For instance, whether contact or location information is collected, and whether collected information is shared with third parties are important. We leverage crowdsourcing and NLP in order to semi-automatically extract such key features from website's privacy policies and then create a simplified model of the website's stated data practices. Such models facilitate automated analysis of the policy and the creation of more concise privacy notices to be presented to users. Figure 1 provides an overview of our approach. We discuss the main research areas in the following.

2.1 Semi-automated Privacy Policy Feature Extraction

We extract relevant features from privacy policies in a hybrid approach that combines crowdsourcing and NLP. We leverage crowdsourcing to obtain annotations of privacy policies in terms of what information is collected by a website, whether that information is shared with third parties with or without the user's consent, and whether the collected data can be deleted by users.

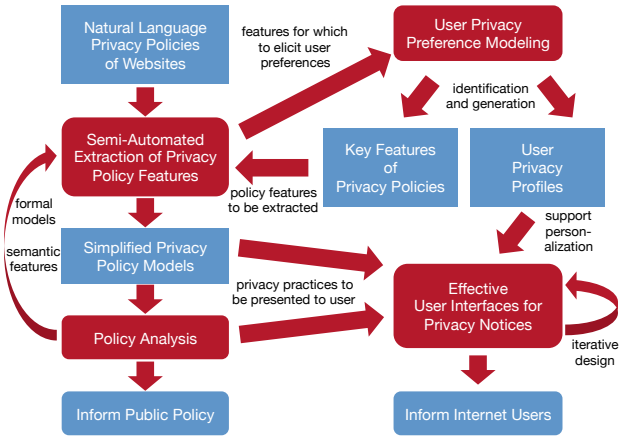

Figure 1: Overview of the proposed approach.

Ensuring that crowdsourcing yields high quality data requires careful task design. Encouraged by our prior results [2], we are experimenting with different task decomposition approaches to enhance annotation quality. Those approaches cover general data practices, such as collection, processing, or sharing with third parties; different information types, such as contact information, current location, or financial information; as well as more fine-grained annotations of recipients of information and purpose statements [4, 5].

The resulting annotation data is used to generate NLP models. For instance, we employ sequence alignment to identify policy segments that likely pertain to the same data practice across different policies [10]. We currently leverage the NLP results to improve annotation interfaces for our crowdsourcing effort and optimizing task scheduling, e.g., by selecting or highlighting parts of the policy, which are potentially relevant for a specific annotation question. However, we plan to extend the expressiveness of our NLP models in order to move towards automated extraction of salient features from privacy policies.

2.2 Privacy Policy Analysis

We use salient information extracted from privacy policies to reason about the website's data practices and conduct extensive privacy policy analysis for multiple purposes. Translating policy features into descriptive logic statements facilitates detection of inconsistencies and contradictions in privacy policies [4].

* Point of contact: Norman Sadeh (sadeh@cs.cmu.edu)

Annotation disagreement among crowd workers further helps identifying potential ambiguities in the policy. Comparing a website's privacy policy with those from similar websites holds the potential to detect likely omissions in the privacy policy. Temporal monitoring of changes in privacy policies facilitates content-based trend analysis. We use policy analysis results to provide more effective and accurate privacy notices to users. Furthermore, we combine reasoning results with legal analysis of privacy policies to study the effectiveness of self-regulation efforts in different sectors and inform public policy. In addition, we plan to make analysis results available to website operators in order to help them improve their privacy policies.

2.3 Privacy Preference Modeling

The major goal of our approach is to make privacy policies more usable and accessible for website users. Thus, an important aspect of our work is the identification of those key features in privacy policies that are relevant to users in order to guide the semi-automated extraction of privacy policy features and the development of improved privacy notices. For this purpose, we have been conducting numerous user studies on privacy concerns, perceptions, and preferences, for example, in relation to online behavioral advertising [7]. Furthermore, we strive to gain a deeper understanding of cognitive biases that may negatively affect individuals' privacy decisions, in order to inform how users can be made aware of privacy risks in an effective manner [1].

In addition, we employ crowdsourcing and machine learning to collect users' privacy preferences on a large scale in order to generate preference profiles for different user groups [8], which can then be leveraged to personalize privacy notices and interfaces to individual user requirements and preferences.

2.4 Effective Privacy User Interfaces

Features extracted from privacy policies as well as results from privacy policy analysis and privacy preference modeling inform our design of user interfaces for privacy notices. The goal is to make those policy features that users care about more accessible, for instance, with nutrition label inspired privacy notices [6] or privacy icons symbolizing data practices. We are also investigating the potential of just-in-time notices that highlight data practices when they become relevant for the individual user. For instance, data practices concerning the collection and sharing of contact or financial information may only be relevant when the user creates an account or makes a purchase. Privacy notices can further be designed as privacy nudges, which aim to mitigate cognitive and behavioral biases in order to help users make better privacy decisions [1]. We are in the process of designing browser extensions that leverage policy extraction results and offer notices to users independently of website operators. We follow a user-centric iterative design process to enhance and evaluate the effectiveness of developed privacy interfaces in user studies.

3. CONCLUSIONS

Our work aims to improve the usability of privacy policies by extracting relevant data practices and making them more accessible to users. We follow an interdisciplinary approach, in which results inform and shape research in different areas. For instance, privacy preference modeling informs semi-automated policy analysis; respective results inform privacy notice design; user evaluation and studies on cognitive biases inform interface design; and overall results inform public policy.

In contrast to prior work, our outlined approach does not require any effort or cooperation by website operators. By making the content of privacy policies more salient and accessible, we hope to also nudge companies towards improving their privacy policies by reducing ambiguities. Policies that already provide relevant information in their policies concisely facilitate semi-automatic extraction and privacy policy analysis, which results in more concrete information provided to users through our process.

Our approach currently focuses predominantly on data practices revealed in natural language privacy policies, because they constitute the basis for user consent to data practices. While this approach has limitations in terms of detecting covert data practices and violations of a website's own privacy policy, it can be further complemented with analysis of a company's actual data practices or data flow analysis.

Acknowledgments

This work is supported by the National Science Foundation under Grant No. CNS 13-30596.

4. REFERENCES

[1] A. Acquisti. Nudging privacy: The behavioral economics of personal information. *IEEE Security & Privacy*, 7(6):82–85, 2009.

[2] W. Ammar, S. Wilson, N. Sadeh, and N. A. Smith. Automatic categorization of privacy policies: A pilot study. Tech report CMU-ISR-12-114, 2012.

[3] M. Benisch, P. G. Kelley, N. Sadeh, and L. F. Cranor. Capturing location-privacy preferences: quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing*, 15(7):679–694, Oct. 2011.

[4] T. D. Breaux, H. Hibshi, and A. Rao. Eddy, a formal language for specifying and analyzing data flow specifications for conflicting privacy requirements. *Requirements Engineering*, pages 1–27, 2013.

[5] T. D. Breaux and F. Schaub. Scaling Requirements Extraction to the Crowd: Experiments with Privacy Policies. In Proc. RE ’14. IEEE, 2014.

[6] P. G. Kelley, J. Bresee, L. F. Cranor, and R. W. Reeder. A ”Nutrition label” for privacy. In *Proc. SOUPS ’09*. ACM, 2009.

[7] P. G. Leon, B. Ur, Y. Wang, M. Sleeper, R. Balebako, R. Shay, L. Bauer, M. Christodorescu, and L. F. Cranor. What matters to users?: Factors that affect users’ willingness to share information with online advertisers. In *Proc. SOUPS ’13*. ACM, 2013.

[8] B. Liu, J. Lin, and N. Sadeh. Reconciling mobile app privacy and usability on smartphones: Could user privacy profiles help? In *Proc. WWW ’14*. ACM, 2014.

[9] A. M. McDonald and L. F. Cranor. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*, 4(3):543–568, 2008.

[10] R. Ramanath, F. Liu, N. Sadeh, and N. A. Smith. Unsupervised alignment of privacy policies using hidden Markov models. In *Proc. ACL ’14*, 2014.