Natural Language Processing: Enhancing Transparency in Privacy Policies of Connected Devices

CS4991 Capstone Report, 2025

John DeFranco
Computer Science
The University of Virginia
School of Engineering and Applied Science
Charlottesville, Virginia USA
qbk3xy@virginia.edu

ABSTRACT
Privacy policies for connected devices are often unclear and confusing, making it difficult for non-technical users to understand how their personal data is being collected, used or shared. To address this problem, I propose using a system that summarizes privacy policies in simple, understandable language which highlights data collection processes, usage and sharing. The proposed system involves creating a natural language processing (NLP) model and integrating the model into an application or web-extension which provides users with an easily understandable summary showing how their data is being used. The use of the model will lead to a reduction in user confusion regarding privacy policies and data collection practices in connected devices. Future work includes testing the system's effectiveness and improving its accuracy.

1. INTRODUCTION

The rapid growth of connected devices, from smartphones to home appliances, has caused concerns regarding user's privacy and security. Each of these devices has its own individual privacy policy, which explains how a user's data is collected and potentially analyzed or shared. However, these privacy policies can be complex and long, which can make it difficult for users that come from non-technical backgrounds to fully understand and comprehend what they are agreeing to. Some users may skim or ignore these policies as a result of their intricacy, which may potentially expose themselves to the risks that come with sharing their personal data.

In order to address the issue of non-technical users not being fully aware of what happens when they agree to a privacy policy for connected devices, I have explored automated techniques that simplify and summarize privacy policies. One of these techniques is natural language processing (NLP), which can be used to create concise summaries of complex documents, revealing key data collection practices. By reducing the cognitive burden and prior knowledge needed to read and comprehend these lengthy privacy policies, NLP can give users the power to make more informed decisions about how their personal data is being collected and used. This paper proposes an NLP model that provides clear summaries and takeaways of the privacy policies associated with connected devices, thus shortening the divide between the dense language in policies and user comprehension.

2. RELATED WORKS

The simplification of privacy policies has been a topic of contention for many years, and there have been studies that have emphasized the need to make the complex language commonly included in these policies more accessible to general users that come from non-technical backgrounds. For example, a study from 2008 showed that if an average American read the privacy policy for each unique website they accessed, they would spend between 181 and 304 hours per year reading these policies (McDonald & Cranor, 2008). Although this study is old, both the number of websites on the internet and the amount of connected devices the average American uses every year has significantly increased since 2008, so one could assume that the time it would take for a user to read every privacy policy for every digital service they use to be even higher in 2025.

There have been other attempts to simplify privacy policies for non-technical users. In 2018, researchers created an automatic framework called *Polisis* in order to analyze and simplify privacy policies. The framework utilizes neural networks to detect and label specific data practices within privacy policies. The researchers also created an application called *PriBot*, which was the first free form question answering system for privacy policies (Harkous et al., 2018). Similarly, another group of researchers have developed a web-based annotation tool called *Corpus*. The researchers created an annotation scheme with ten different data practice categories, and the tool allows skilled annotators to apply the scheme to select privacy policies. After manually annotating a select group of privacy policies, the researchers were able to partly automate the annotation process (Wilson et al., 2016).

3. PROPOSAL DESIGN

The proposed design consists of two sections: design and evaluation. The design section outlines the technical components of the system, and the evaluation section describes how the performance of the system is assessed.

3.1 Design

The proposed design contains two main components: summarization of the data via natural language processing and the implementation of a user interface.

First, the model utilizes named entity recognition (NER) to extract important entities (such as organizations, locations, and dates) from a user-inputted privacy policy. NER analyzes a text, and can be used to classify words and sections into predefined categories (“A Comprehensive Guide to Named Entity Recognition”, n.d). Next, the model utilizes an abstractive transformer model (ATM) to summarize the key points from the privacy policies. Using an ATM to summarize each privacy policy is key because it can produce simple human-sounding summaries, thus allowing non-technical users to understand its output (Nnadi & Bertini, 2024).

Finally, the model is integrated into a user interface. A web-based interface allows users to operate the model from any internet connected device. The interface has a text input area where users can insert a privacy policy. After inputting the policy, the model summarizes the text, displays important entities, and shows the user metrics to help them understand the effectiveness of the model.

3.2 Evaluation

In order to evaluate the model, Recall-Oriented Understudy for Gisting Evaluation (ROUGE-L) score will be used. ROUGE-L score measures how similar a

machine-translated text is to a reference text, and will be used to compare the summarized privacy policy to the original policy. A higher ROUGE-L score indicates a higher quality summarization (Chiusano 2023). Additionally, Flesch-Kincaid Reading Ease score (FKRES) will be used to evaluate how easy the summary is to understand. FKRES analyzes sentence length, word difficulty, and text cohesion. The higher the FKRES a summary receives, the easier it is to read and comprehend (Readable, 2023). Also, the difference in length between the two policies, expressed as a percentage (summarized policy length / original policy length), will be used to evaluate the effectiveness of the model as well.

4. RESULTS

The main goal of the model is to make privacy policies easier to understand for non-technical users. Privacy policies are typically thousands of words long, and the model is able to summarize them concisely while maintaining their essential details. Across a dataset of 115 real-world privacy policies from the *OPP-115 Corpus* dataset (Wilson et al., 2016), the model generated summaries that reduced the original word count by an average of 90.20%, making it significantly easier for users to parse the lengthy original text.

In addition to shortening each privacy policy, the model was able to simplify the often complex language used. The FKRES of the summaries was, on average, 3.52 points higher than that of the original policies. The simplification of the language ensures that non-technical users can understand each policy, without the need for preexisting knowledge of legal expertise.

Content preservation was also measured using ROUGE-L score, which had an average of 0.1353 across the dataset. Although the average ROUGE-L score was not as high as anticipated, it suggests that the summaries retained the key concepts and structure from the original policies.

The model allows users to receive a clear, structured summary of what personal data is collected, whether said data is shared with third parties, and whether users are able to opt-out of data collection. This transparency gives users the ability to make informed decisions about what data they share, and will reduce instances of users unknowingly agreeing to intrusive and hazardous privacy policies.

5. CONCLUSION

The sharp increase of connected devices in everyday life has made it extremely important for users to understand how their personal data is being collected and used. However, the complexity and length of the privacy policies for these devices presents a barrier to users understanding what they are agreeing to. The proposed system addresses that problem by using natural language processing to summarize privacy policies in clear concise language.

The system leverages named entity recognition and an abstractive transformer model to summarize and shorten text. Recall-Oriented Understudy for Gisting Evaluation and Flesch-Kincaid Reading Ease scores are then used to evaluate the effectiveness of the model. The system is of great value to users, since it allows them to easily understand the important parts of lengthy and complex privacy policies without needing any preexisting technical or legal knowledge. This allows them to make informed decisions regarding their personal privacy when using connected devices.

6. FUTURE WORK

The next step for this model would be to conduct tests with beta users to further evaluate the system's accuracy and effectiveness. This will include both technical and non-technical users following a beta testing script and then providing feedback on the system.

Looking further into the future, additional work would likely include refining the accuracy and quality of the model to ensure higher FKRES and ROUGE-L scores, along with adding support for multiple languages. Restructuring the system as a browser extension or mobile application could increase accessibility to new users and make it more convenient to use.

REFERENCES

Chiusano, F. (2023). *Two minutes NLP-learn the Rouge metric by examples*. Medium. https://medium.com/nlplanet/two-minutes-nlp-learn-the-rouge-metric-by-examples-f179cc285499

Harkous, H., Fawaz, K., Lebret, R., Schaub, F., & Aberer, K. (2018). Polisis: Automated analysis and presentation of privacy policies using deep learning. *In 27th USENIX Security Symposium* (pp. 531–548). Baltimore, MD: USENIX Association.

McDonald, A., & Cranor, L. (2008). The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*

Nnadi, G. O., & Bertini, F. (2024). *Survey on abstractive text summarization: Dataset, models, and metrics*. ArXiv.org. https://arxiv.org/abs/2412.17165

Readable. (2023). *Flesch Reading Ease and the Flesch Kincaid Grade Level*. Readable. https://readability/flesch-reading-ease-flesch-kincaid-grade-level/

Turing.com. (n.d.). *A comprehensive guide to named entity recognition (NER)*. www.turing.com. https://www.turing.com/kb/a-comprehensive-guide-to-named-entity-recognition

Wilson, S., Schaub, F., Dara, A., Liu, F., Cherivirala, S., Giovanni Leon, P., Schaarup Andersen, M., Zimmeck, S., Sathyendra, K., Russell, N., Norton, T., Hovy, E., Reidenberg, J., & Sadeh, N. (2016). The creation and evaluation of a website privacy policy corpus. *In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 1330–1340). New York, NY: Association for Computing Machinery.