A Deontological Analysis of Privacy Policy Presentations

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Jack Mingjie Liu

Spring 2023

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Kent A. Wayland, Department of Engineering and Society

Introduction

More than ever before, we are surrounded by devices that are connected to the internet in our daily lives. Many of these devices are equipped with cameras, microphones, and other sensors that companies are willing to use to collect information whether the user is aware or not. For example, recent studies show that some video conferencing applications continuously monitor microphone input even when on "mute" and even transmit the data to remote telemetry servers (Yang et al., 2022). All of this happens behind the scenes with no visible indication to the user. This raises the question of how data collection and data use policies are communicated to users, and how they can be made more transparent.

Privacy policies are the standard way for companies to disclose their practices regarding data collection, usage, and management. While this may cover the necessary legal bases for the service to operate, this does not translate well to the end user. More often than not privacy policies are full of legalese and difficult to parse through. This leads to users ignoring the message and blindly accepting it without fully understanding what the agreement entails. Malicious companies could take advantage of this fact and potentially exploit users for their data. With these concerns, it is important to study the presentation of privacy policies to users and how certain choices can go counter to their intended function.

Background

Many studies have been done on the content and style of privacy policies in order to balance readability and usefulness. Hintze (2017) argues that privacy policies should be drafted to be fully comprehensive, as this ultimately promotes more transparency, even at the cost of being left unread for the end user. This is because other parties like reporters can parse through the document and convey the relevant information to the general public. While it is true that

more detailed statements about all the various sources of data companies collect can aid in transparency, this argument is not satisfying. It only shifts the responsibility of disclosure to a third party and still doesn't guarantee that all users will be able to understand the relevant risks and protections to their data.

On the other hand, Krumay and Klar (2020) investigated the ability to automate quantitative readability measures for privacy policies. They used metrics such as word count, sentence length, and syllables per word to define readability as well as comparing their own ratings to those given by subjects of various ages to see if they align. The researchers found that these methods can help provide more insight in coordination with the qualitative measures, but looking at one criterion alone can be misleading. However, there is still potential for augmenting more qualitative measures with automated tools like the ones mentioned in the paper in order to help companies redesign their policies.

One standard component of modern privacy policies is the inclusion of a specific internet "cookie" policy. Cookies are small pieces of data that are stored in your internet browser that both help websites function but also collect your data. While not necessarily malicious, these cookies can be used to identify a device and track browsing behavior across websites. In 2018, a piece of legislation called the General Data Protection Regulation (GDPR) was passed in the European Union which had a large impact on how internet cookies can be used. One example is the requirement that websites ask for permission for any cookies that a website may use leading to the ubiquitous cookie consent banner we can see today.

With all of the prompts for the user to accept, it can be a challenge to guarantee that users understand what they are agreeing to. But this problem has already been studied in the area of informed consent agreements for scientific experiments involving human subjects which is an

interesting parallel to educating users of privacy policies. These agreements are intended to inform and protect participants, but they can be poorly designed which results in participants accepting without fully grasping their rights. Rossi and Lenzini (2020) developed several information design patterns for researchers so they can more effectively and transparently disclose how information may be used to empower the participants. One example is the use of visuals to help retain the attention of the reader, make abstract concepts more tangible, and aid those with lower literacy. Therefore, it may be useful to take the design patterns proposed by the authors of this paper and see if they can be applied to improve transparency around data collection by companies and technology.

Ethical Framework

In order to study these questions, it may also be helpful to draw upon the theory of deontology, a duty-based branch of ethics. For the philosopher Immanuel Kant, a key tenant of deontology is that of the categorical imperative. One interpretation of this is to never treat others as only a means to an end and to respect their autonomy. In the context of user privacy, the question becomes whether or not companies do enough to allow users to make informed choices for themselves about their data. Moreover, deceptive actions such as tracking users without their knowledge or after they chose to reject cookies may violate this categorical imperative.

Several researchers used this moral framework to create a set of principles and best practices for privacy policies and then evaluated what they saw in the industry at the time (Dean et al., 2016). Some example practices are to respect individual autonomy by allowing consumers to have a choice to provide or withhold information and to respect the value of individuals by encouraging data subjects to check their personal information for accuracy. After sampling a variety of privacy policies found online, they, unfortunately, found that many did not comply

with the principles they derived. While these findings are illuminating, a lot can change in the 9 years since this study due to the face paced nature of technology. Thus, it can be interesting to investigate whether companies have adapted to better fit these policies or if rather they are still sorely lacking.

Another example of deontology used in the study of privacy policies is by Irena Pollach (2005). In the study, Pollach utilized four different ethical frameworks (virtue ethics, deontology, teleology, and justice) and applied them to the central idea of informed consent concerning accepting privacy policies. For the methods, Pollach used linguistics to study the privacy policies of several large retailers and travel agencies. Some specific aspects she looked at are the vocabulary used, verb tenses, and passive vs. active voice. For example, passive voice is commonly used by companies when discussing the data collection process which Pollach argues is done to form a degree of separation and obscure the company's responsibility. Overall, this work corroborated other research that privacy policy language is commonly opaque and vague to the point of non-transparency to the reader.

Methods

The methods for this study consisted of two stages. The first revolved around drawing from ethical frameworks such as deontology and virtue ethics in order to lay a theoretical foundation for the study. These frameworks are useful as educating users of privacy policies can be thought of as the duty of a company, and it shifts the focus away from consequences which can be difficult and unclear in the context of privacy policies. Inspiration was drawn from the work of Dean, Payne, and Landry which used previous frameworks and legislative documents to define a set of principles that privacy policies should follow. A special focus will be on the most

recent developments in privacy legislation including the GDPR law which has significantly influenced the privacy landscape since its enactment in 2018.

The second phase of the methods consisted of evaluating the principles defined previously on privacy policies found today. For this, it was beneficial to focus on a select few policies rather than a shallow pass through many. In order to get a representative sample, this study used privacy policies from some of the larger technology companies in different domains namely Facebook, Google, and Apple. Not only are these companies of note because of their size, but they have come under scrutiny for their privacy policy in the past, and Apple stands out as a company that is also invested in hardware in addition to software. Finally, a smaller company called Chess24 was also included as the company has close ties with the EU and also to see how much standardization there is between companies of different scales.

In order to judge these documents, they were compared with the principles formulated in phase one to determine the degree of compliance. To support this analysis, quantitative measures can be used to gather a degree of readability of the sources. The Flesch Reading Ease and Flesch-Kincaid grade level readability are two measures that use the average length of sentences and syllables per word to determine a readability score for any given text. Qualitative measures in the design and presentation of privacy policies were also analyzed in this study.

Results

Dean, Payne, and Landry (2016) defined a list of proposed practices for privacy policies using the principles of deontology. This is summarized below in Table 1, and notably, they organize the practices under the three principles of universalizability, respect for individual value, and individual autonomy. For this study, some practices will not be as studied such as whether the policies comply with existing laws and whether they support industry accountability.

It is assumed that large corporations are legitimate and are not actively breaking laws as this would presumably have been brought to attention already. Furthermore, the focus of this study is on the impact of privacy policies on the individual user.

However, there is a note that some other principles require additional scrutiny as we've seen in previous sources whether a privacy policy respects individual autonomy depends on more than merely a choice being offered. It should be clear what each choice entails and that the result is what can be reasonably assumed.

Table 1

Proposed practices for online privacy policies based on deontological principles

Philosophical rule	Proposed practices for online privacy policies
Universally consistent actions	- Comply with the letter of all applicable law, including providing notice, security, and data integrity, particularly with sensitive information - Support legislation that allows for a reasonable system of information collection and sharing - Support industry efforts to self-police, using the spirit of the law
Respect individuals as inherently valuable	- Encourage data subjects to access their personal information to check for accuracy, relevance, comprehensiveness, and timeliness - Support industry efforts to self-police, using the spirit of the law - Encourage and engage in discussions about why people share personal information, what constitutes personal information - Encourage firm and employee pride in self and community through clearly alerting data collectors and users as to their policies and adherence to same
Respect autonomy of all rational beings	- Allow consumers to have a choice to provide or withhold information - Use only truthful, candid privacy policies and notices regarding collection, storage, security and sharing of personal information

The privacy policies themselves were found on the respective websites of the appropriate companies. There was a range of formats ranging from interactive websites with embedded images and videos to plain text documents. Furthermore, Facebook and Chess24 had separate

documents for their cookie policy while Google and Apple included this component in the main privacy policy itself. There was also an explicit cookie banner on the Chess24 that all visitors to the website must interact with to enable certain cookies on the site which will also be analyzed in addition to the policy documents.

As mentioned before, the Flesch Reading Ease and Flesch-Kincaid grade level readability scores were calculated for each of the bodies of text. Only the primary privacy policy was analyzed for each website, and the results are summarized in the table below.

Table 2

Quantitative readability scores for gathered online privacy policies

Website	Word Count	Flesch Reading Ease	Flesch-Kincaid Grade Level
Apple	4,054	36.9	13.1
Google	7,532	43.0	12.5
Facebook	14,841	53.0	10.1
Chess24	2,078	40.2	13.6

Discussion

Before analyzing the contents of the privacy policies themselves, it's important to take a moment to identify the accessibility of the documents themselves. Across the board, there is standardization in this aspect. For each of the considered websites, the privacy policy is linked directly at the footer of the page without needing to search for it. That being said, the links are in a much smaller font compared to the rest of the page making it difficult to notice if you are not actively searching for them. This can be seen as an expression of the universalizability principle of deontology where companies act similarly in the inclusion and styling of their privacy policies. While it may not be the most visible, it is a rule to include a link to the privacy policy

and to not actively obscure it by making it less conspicuous to other links. This has the benefit that any privacy-conscious user is able to obtain relevant policies without too much effort.

Looking at the readability metrics calculated in Table 2 above, we can see that the privacy documents analyzed have a range of lengths with respect to word count, but are fairly consistent in their Flesch Reading Ease and Flesch-Kincaid Grade Level scores. In general, documents meant for the general public should aim to have a Flesch Reading Ease score of around 60.0 to 70.0, or about an 8th-grade reading level, so that it can be read by about 80% of the population (Readable, 2021). However, the privacy documents analyzed largely were at a college reading level far above this benchmark. Even the best scoring document from Facebook was still at a 10th-grade reading level and considering that teens actively use all of these services there is some obvious disconnect. Beyond simply being hard to read, using language that is not accessible to readers can impact their autonomy to make informed decisions. Even if a privacy agreement does not actively include false information, an individual cannot make a rational decision if they are unable to comprehend it. In order to address this issue, revisions must be made to improve readability or supplemental sources must be included that are more accessible to the general public.

On the qualitative side, there is a distinction between the privacy documents of larger and smaller companies. Google, Facebook, and Apple all have interactive privacy documents that allow readers to see an overview of key points and expand each section as necessary. In addition, Google and Facebook feature embedded videos and pictures which reflect the design patterns of Rossi and Lenzini (2020) aimed to help retain attention and make abstract concepts more comprehensible. This is in direct contrast with the privacy documents of Chess24 which follows the more traditional format of a plaintext document. Even though it is the shortest of all the

privacy policies considered, it also feels the densest because of its format. This potentially shows the difference between larger organizations that can dedicate more resources towards designing a friendlier privacy document and smaller organizations that view it more as a contractual obligation.

Concerning the actual contents of the privacy documents, all feature a similar set of section headings that include what data is collected, why it's being done, and how much user control there is over their practices. Following the principles set forth by Dean et al. in Table 1, these features help satisfy the deontological rule of respect for persons. This is especially for the interactive policies of Facebook and Google where there are embedded links to download your personal data for review as well as to account management pages to change what is collected. By being upfront about company policies, it helps to mitigate the problem of treating users as simply a means for obtaining data. While it is true that these companies take advantage of their users for data, they also allow the user to decide whether or not they want to use the service. However, this is contingent on the fact that they are not coercive or misleading in delivering this information. As touched on previously, companies across the board tend to shift the spotlight away from privacy information making such controls often difficult to find in the first place. This is especially a concern when many services operate on an opt-out regime where companies will only limit their data collection if explicitly asked by the user.

One notable aspect of Chess24's privacy initiatives is a cookie consent banner that all visitors to the site must interact with before continuing to the rest of the page. This screen is shown below in Figure 1. As mentioned before, these kinds of banners have been made more common as a result of the GDPR law in the EU. But, by looking at the design choices made in the banner, we can uncover the hidden values of the company that go beyond the word of the

law. The first note is that by default only the cookies necessary for the site to function are toggled on which is a good baseline for general users. However, by looking at the confirmation buttons at the bottom of the prompt we can see evidence of some deceptive design choices.

Figure 1

Chess24.com cookie consent banner

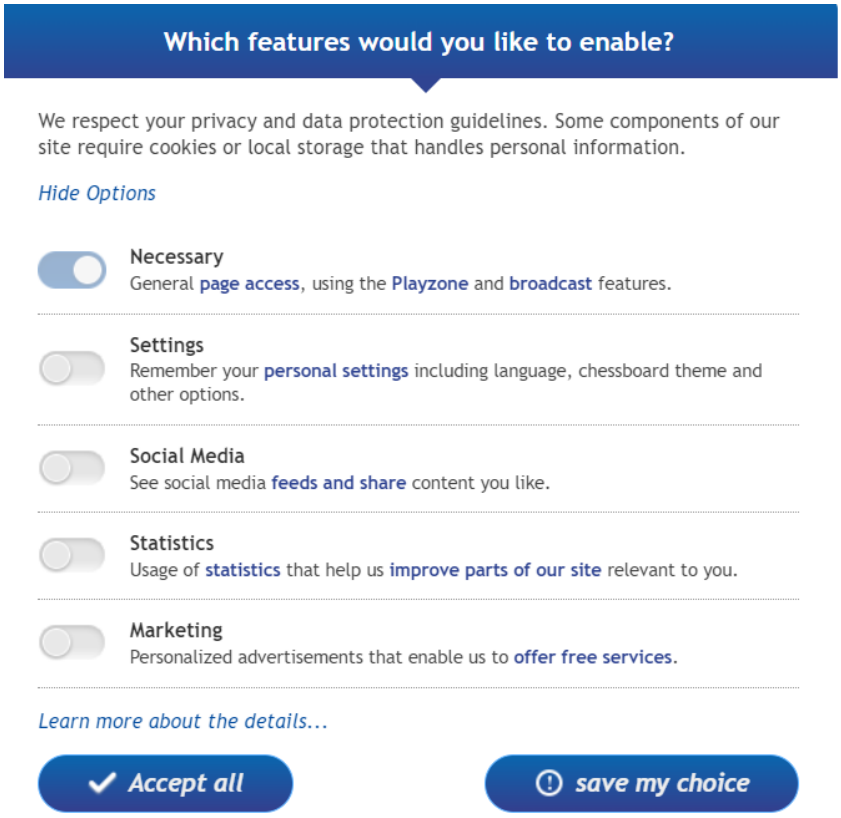

Traditionally, prompts that ask for approval or rejection have the positive choice on the left and the negative on the right. However, for this banner, the button on the left will accept all cookies regardless of what has been selected by the toggle switches above, while only the button on the right respects the user's choices. This is further made an issue by the choice of iconography with a checkmark on the “accept all” button and a typical warning icon on the right. This can be seen as a deceptive act by the company by giving the users the illusion of choice while tricking them into inadvertently accepting all cookies instead. Together, these choices may abide by the

regulations placed on data privacy but still violates individual autonomy by not respecting the values of the user. It goes to show that we must consider intent as well as action when passing regulations.

Conclusion

Privacy policies are more important than ever in our digital world as a way to inform and empower individuals on how their data is being collected by the services they use. Fortunately, some companies have made improvements towards making these policies less dense and full of legalese by adding interactivity and videos to reach a larger audience. However, we've seen that they are still inaccessible for a large portion of the user base and this can undermine individual autonomy when it comes to deciding what they are willing to share with companies. This problem is only exacerbated by deceptive design practices by some companies.

To address these issues, there are still a few areas of research that should be looked into. Because this study focused on analyzing privacy policies from the deontological point of view, there was less emphasis on the consequences of the policies. Specifically, it is important to understand to what degree companies follow the standards that they set, as well as to gather information from users to gauge the effectiveness of privacy policies. Ultimately, this study is only one step towards promoting awareness of the privacy implications of the services that we all interact with every day.

References

Dean, M. D., Payne, D. M., & Landry, B. J. L. (2016). Data mining: An ethical baseline for online privacy policies. *Journal of Enterprise Information Management*, *29*(4), 482–504. https://doi.org/10.1108/JEIM-04-2014-0040

Hintze, M. (2017). In Defense of the Long Privacy Statement. *Maryland Law Review*, *76*(4), 1044–1084.

Krumay, B., & Klar, J. (2020). Readability of Privacy Policies. In A. Singhal & J. Vaidya (Eds.), *Data and Applications Security and Privacy XXXIV* (pp. 388–399). Springer International Publishing. https://doi.org/10.1007/978-3-030-49669-2_22

Pollach, I. (2005). A Typology of Communicative Strategies in Online Privacy Policies: Ethics, Power and Informed Consent. *Journal of Business Ethics*, *62*(3), 221–235. https://doi.org/10.1007/s10551-005-7898-3

Readable. (2021, July 9). *Flesch Reading Ease and the Flesch Kincaid Grade Level*. Readable. https://readable.com/readability/flesch-reading-ease-flesch-kincaid-grade-level/

Rossi, A., & Lenzini, G. (2020). Transparency by design in data-informed research: A collection of information design patterns. *Computer Law & Security Review*, *37*, 105402. https://doi.org/10.1016/j.clsr.2020.105402

Yang, Y., West, J., Thiruvathukal, G. K., Klingensmith, N., & Fawaz, K. (2022). Are You Really Muted?: A Privacy Analysis of Mute Buttons in Video Conferencing Apps. *Proceedings on Privacy Enhancing Technologies*, *2022*(3), 373–393. https://doi.org/10.56553/popets-2022-0077