

# Evaluating Privacy Policy Summarization: An Experimental Study among Japanese Users

Vanessa Bracamonte<sup>1</sup>, Seira Hidano<sup>1</sup>, Welderufael B. Tesfay<sup>2</sup> and Shinsaku Kiyomoto<sup>1</sup>

<sup>1</sup>*KDDI Research, Inc., Saitama, Japan*

<sup>2</sup>*Goethe University Frankfurt, Frankfurt, Germany*

**Keywords:** Privacy Policy, User Perception, Privacy, Risk.

**Abstract:** Summarization and visualization applications can help users understand the content of privacy policies. However, research has focused on English language privacy policies and has not considered users who are not native English speakers nor the potential situation of encountering a privacy policy in a foreign language. In this paper, we contribute to the research on privacy policy summarization by conducting an experimental survey on Japanese users to assess their interest on using such an application, and the influence of this application on their perception. We conducted an experimental survey among Japanese participants, and evaluated their perception on different privacy policy languages (Japanese or English) and risk levels, using PrivacyGuide. We found that PrivacyGuide can increase interest in the contents of the privacy policy for both languages, and can communicate risk level for the English privacy policy. In addition, we found that respondents who indicated interest in using the application mentioned a wide variety of scenarios for its use, while respondents who answered negatively or were hesitant mentioned lack of trust and uncertainty about PrivacyGuide's reputation and accuracy. We discuss these results and offer suggestions for improving adoption of privacy policy summarization tools like PrivacyGuide.

## 1 INTRODUCTION

Privacy policies are often complicated and difficult to read (Proctor et al., 2008; Sunyaev et al., 2015). In an effort to address this situation, regulation like the GDPR (EU, 2016) indicates that information about privacy policies should be written in clear language, but even so, the cost of reading every privacy policy would be too high (McDonald and Cranor, 2008).

Presenting privacy-related information in formats that are easier to understand, such as shorter privacy policies (Gluck et al., 2016) or graphical information regarding privacy risks (Gideon et al., 2006) can help users make better decisions regarding privacy. In order to provide this more understandable information, applications that can automatically analyze existing privacy policies and present the summarized results them have been introduced (Harkous et al., 2018; Tesfay et al., 2018; Wilson et al., 2016).

A limitation of these studies is that they focus on English speaking users and English privacy policies. However, users can access websites all over the

world. In Japan, for example, the top sites accessed as ranked by Alexa<sup>1</sup> include international websites —such as those of well-known companies based in English speaking countries. These websites provide local language versions of their privacy policies, but these are often direct translations of the English version, and therefore have the same shortcomings. And if they do not provide a translated privacy policy, it is possible that foreign users cannot read them at all.

Applications that summarize privacy policies and offer a structured presentation of results may be useful to address these two scenarios, since a standardized format and pre-defined categorizations are suited for translation.

In this paper, we use PrivacyGuide (Tsfay et al., 2018) to evaluate the effects of privacy policy summarization in a non-English language context. We consider the question of whether PrivacyGuide will increase interest in the privacy policy and whether it will effectively communicate risk to users. In addition, we consider the question of

---

<sup>1</sup> <https://www.alexa.com/topsites/countries/JP> (accessed 8.2.18)

whether Japanese users would be interested in using such an application and under which circumstances.

We addressed these questions by conducting an experimental survey among Japanese users, considering two scenarios: Japanese and English language privacy policies. We found that PrivacyGuide results increased interest in the content of a privacy policy in both languages. PrivacyGuide also communicated risk to users for the English language privacy policy. In addition, Japanese users mentioned a variety of possible uses for PrivacyGuide, but also identified barriers to its use, such as uncertainty about its trustworthiness, reputation and accuracy of the results. These results contribute evidence of the potential and challenges for privacy policy summarization for non-English speaking users. To the best of our knowledge, this is the first user perception study on privacy policy analysis tools such as PrivacyGuide. We discuss our findings and possible ways to improve adoption of privacy policy summarizing applications.

## 2 BACKGROUND

### 2.1 User Perception of Privacy Policy Summarization

Previous research has found that presenting information in more understandable formats can help users better understand the privacy practices of websites. Icons indicating the privacy risk level (Gideon et al., 2006) and detailed information of what data is at risk (Harbach et al., 2014) can help users make more privacy conscious decisions. In particular for privacy policies, standardized information in a format similar to a nutrition label can be more effective in helping users obtain accurate information than a text privacy policy (Kelley et al., 2010) and a shorter text can provide enough information for users to understand risks in privacy policies, compared to longer documents (Gluck et al., 2016).

Although these studies demonstrated possible ways to design privacy policies to effectively communicate risk to users and make them aware of privacy practices, the current situation is that privacy policies are often complex and lengthy text-only documents. However, there have been research efforts to provide applications to automatically analyze existing privacy policies and visually present these results. These applications define a categorization of the privacy policies' contents, although the basis for that categorization may be

different: based on fair information practices and policies (Zimmeck and Bellovin, 2014), on the knowledge of privacy domain experts (Harkous et al., 2018; Wilson et al., 2016) or based on regulation such as the GDPR (Tesfay et al., 2018). However, these studies have not included user validation studies of the effectiveness of these applications for communicating risk.

As far as we could determine, there are no studies on privacy policies in foreign language in the privacy literature. However, we consider that this is a topic that merits research. Anecdotal evidence indicates that Japanese users are concerned when they see English privacy policies. This was the case when the GDPR came into effect on May 2018 and Japanese users received updates to privacy policies in English, which they could not read.

In addition, currently the list of top smartphone apps on Google Play and Apple's App store for Japan regularly includes apps from foreign developers whose websites have only an English language privacy policy. Although in general, very few users check the privacy policies of the websites they visit (Steinfeld, 2016), this is not an ideal situation from the point of view of providing users the information they need to make privacy related decisions.

## 3 METHODOLOGY

### 3.1 PrivacyGuide

In this study, we use of PrivacyGuide (Tesfay et al., 2018), a machine learning-based application for automatically analyzing and summarizing privacy policies written in English. The goal of PrivacyGuide is to support users' understanding of the privacy policy and to elicit interest in the detail of its contents, by providing risk-related information about the privacy policy.

PrivacyGuide classifies the content of the privacy policy into eleven privacy aspects (see Tesfay et al., (2018) for details), which are based on an analysis of criteria from the EU GDPR, and determines a risk level for each privacy aspect. It presents the result of this analysis with the use of icons and colors: one icon for each privacy aspect in green, yellow or red color corresponding to the risk level identified. Privacy guide was considered suitable for the purposes of this experiment due to its structured result presentation format, which could be straightforwardly translated to Japanese language.

In addition, although PrivacyGuide was

developed specifically for the analysis of English language privacy policies following the GDPR, we considered that this approach was compatible for language privacy policies of international websites. Currently, the content and structure of the English privacy policies for these websites have been developed with awareness of the GDPR. And as we could confirm during the development of the experiment's privacy policies, the Japanese versions of international websites' privacy policy are in many cases a direct translation of the English language privacy policy. Therefore, we considered that the use of PrivacyGuide was appropriate for the experiment.

### 3.2 Experiment Design

We designed an experimental survey to address the following research questions:

- Will PrivacyGuide increase users' interest in the content of the privacy policy itself?
- Will the summarized privacy policy result provided by PrivacyGuide correctly communicate risk?
- Will users be interested in using PrivacyGuide and if so, under which circumstances?

In addition, we considered the use case of English language privacy policies for each of these questions.

To address these questions, we conducted an experimental survey where we asked participants to answer an online questionnaire based on their perception of a website page, which included a privacy policy, and a PrivacyGuide result. We defined four experimental conditions resulting from the combination of privacy policy language (English or Japanese) and risk level (low or high) factors.

The online questionnaire instructed participants to imagine a situation where they had found a website and were considering whether or not to register on it, then presented the experiment website page and asked participants to view it as they would in their normal internet use. The participants were not primed to consider privacy in the instructions and we did not ask them to read the privacy policy, which would not have been possible for every participant in the English language conditions.

After answering questions on their perception of the website, the survey introduced PrivacyGuide as a privacy policy analysis and summary application and instructed participants to take some time to check the PrivacyGuide results. The participants then answered the remaining questions.

The survey was conducted online, using a third-party online survey company. The survey company distributed the call for participation among their subscribers. The sample is therefore a convenience sample, although the recruitment process targeted a pool with an equal distribution of gender and an age distribution similar to that of the Japanese population. Each participant was randomly assigned to one of the four conditions, and viewed only one version of the website page—therefore, only one privacy policy—and only one PrivacyGuide result screen. All participants answered the same questions.

### 3.3 Website Page

We developed a non-interactive website sign-in page for the experiment. We used a fictional company to control for reputation effects.

The page consisted of a simple online registration form with first and last name, email address and password input fields, a scrollable text area with the experiment privacy policy, and a check box to indicate agreement to the privacy policy. Besides the privacy policy content, the design of the website page was the same for all conditions.

We developed the page in Japanese and therefore did not conduct a translation process, but we reviewed the language and interface design naturalness.

### 3.4 Privacy Policies

We constructed four privacy policies, corresponding to the combination of languages and risk levels determined for the experiment. The PrivacyGuide result screen consists of eleven privacy aspects, whose respective risk levels are indicated by icons in three different colors: red, yellow and green, depending on the content of the privacy policy. However, we determined that it was not feasible to test all combinations of privacy aspects and risk levels, nor was it the goal of this experiment to measure the effect of specific privacy aspects.

In order to facilitate the distinction, we established two risk levels—low and high—for the experiment. These levels were not intended to represent an absolute scale, but rather to approximate a risk level that users might realistically encounter in normal circumstances.

We used the following procedure to construct the low and high privacy policies for the experiment. First, we obtained a list of the top 50 websites accessed from Japan from the Alexa website (Alexa

Internet, Inc., 2018) on August 8, 2018. From this list, we selected those websites which provided an English language privacy policy.

We identified ten privacy policies using these criteria, which were then analyzed using PrivacyGuide. The PrivacyGuide result for each privacy policy consisted of a combination of privacy aspects and corresponding risk levels. We assigned a value to each level (green=1; yellow=2; red=3) and calculated a total risk value for each combination. Higher values were considered to indicate a higher risk privacy policy. We removed the privacy policies with the highest and lowest risk values, and took the next values as initial candidates for low and high risk.

We then counted the frequency of each risk level for each privacy aspect and identified a base pattern: a group of privacy aspects with the same risk level for most privacy policies in the list. To this base pattern, and corresponding risk value, we added the remaining privacy aspect levels at a risk level that would help reach a value corresponding to a low and a high-risk level privacy policy.

To create the text of the experiment privacy policies we used fragments from privacy policies of existing websites. We obtained these privacy policies from websites which had been used to develop PrivacyGuide and which had both a Japanese and English language privacy policy, for which the Japanese version was a translation of the English version. A person fluent in Japanese and English verified that the fragments included the same content in both languages. We removed references to the original website and put together the fragments to create the English and Japanese low and high-risk privacy policies for the experiment.

To minimize any possible influence due to differences in privacy policy length, we verified that the privacy policies created had a similar number of lines, and word count (for the English version) or character count (for the Japanese version).

### 3.5 PrivacyGuide Result

In order to construct the PrivacyGuide result screen for the experiment, we first translated the PrivacyGuide interface and privacy aspects to Japanese, as follows. A person fluent in English and Japanese translated the text elements of the interface, such as page title, instructions and button labels. A native Japanese speaker fluent in English, with expert knowledge of Japanese and European privacy regulation, translated the privacy aspects name and description. A second native Japanese

speaker reviewed the translated PrivacyGuide interface elements and privacy aspects, focusing on understandability and naturalness. We addressed the translation issues identified in this stage.

We then constructed two versions of the PrivacyGuide result screen, corresponding to the low and high risk levels described in the previous subsection. During this development, we simplified the presentation of PrivacyGuide results as a consequence of feedback received from the Japanese reviewer. PrivacyGuide shows the privacy aspect description and a fragment of the analyzed privacy policy that corresponds to the risk assessment when the user hovers over an icon. For the experiment, the non-interactive version of the result screen that we constructed included icons as well as the privacy aspect name and description of the identified risk level for the privacy aspect. We did not include the original privacy policy fragment since we did not ask participants to read the privacy policy and because it was not possible to ask them to compare the accuracy of those results to the English language privacy policy.

### 3.6 Questionnaire

We addressed perception of the website through three constructs: behavioral intention, risk perception and privacy concern using items adapted from previous research on user perception of websites (Kim et al., 2008). We used a six-point Likert scale for the items' score; the scales ranged from Completely Disagree to Completely Agree, with the exception of risk items, which ranged from Very Safe to Very Risky. In addition, we included an open-ended question asking participants if they were interested in using PrivacyGuide to analyze a website's privacy policy, and if so, which one.

We also included Likert-style questions about the participants' normal privacy policy-related behavior, whether they would use websites in English and whether they would read privacy policies in English. We included these two last questions as proxy for measuring self-perceived English ability related to these tasks, in order to avoid self-effacing responses.

We also measured the number of times a participant viewed the experiment website and PrivacyGuide result, the time they spent viewing those screens and the total time spend on the survey.

We initially developed the survey questionnaire in English and then translated it to Japanese with the following procedure. First, a native Japanese speaker fluent in English forward translated the questionnaire. Then, a second native Japanese

speaker—a person familiar with privacy research—reviewed the translation with a focus on accuracy, identifying and modifying any inaccuracies. A third native Japanese speaker conducted an additional review of the translation, which focused on understandability and naturalness. The translated questionnaire was then compared with the original English one, by a native Japanese speaker and a person fluent in English and Japanese.

At every stage, identified issues were discussed and addressed by the translators and reviewers until there was agreement about the questionnaire text.

### 3.7 Limitations

The methodology we used had the following limitations. We used a convenience sample, obtained from a limited pool of users that had subscribed to participate in online surveys conducted by the third-party survey company. This may have introduced bias in the analysis; however, as the results in the next section show, the sample age distribution followed the Japanese population age distribution and we obtained an equal number of male and female respondents. In addition, the website page developed for the experiment as well as the PrivacyGuide result screen were non-interactive, which limited the authenticity of the scenario proposed to the participants. Finally, when we assigned the privacy policies' total risk values we assumed that all privacy aspects had equal importance, but users may have different priorities and may consider a particular privacy aspect more important than another.

## 4 ANALYSIS AND RESULTS

### 4.1 Data Validation

The responses were collected from August 30 until August 31, 2018. We obtained a total of 1040 participant responses, with 260 participants in each group, as predefined by the survey process.

We first analyzed these data to identify suspicious response patterns. We used the criterion of no variability of extreme responses—where the answer to every question was either 1 or 6—to select the initial candidates for elimination, and assessed the total survey response time for these cases. All cases identified by the no variability criterion showed a low total survey response time. The total sample after eliminating these cases was 984.

Construct validity was evaluated using confirmatory factor analysis (CFA) using maximum likelihood estimation with robust standard errors and a Satorra-Bentler scaled test statistic to correct for nonnormality (Curran et al., 1996; Rosseel, 2012). Items loaded on their respective constructs with a standardized loading higher than 0.7 and the model showed good fit: RMSEA=0.06, CFI=0.97, TLI=0.97, SRMR=0.03 (Kline, 2005). All constructs showed good internal consistency, with a minimum Cronbach's alpha value of 0.87. We then constructed composite variables by summing the validated items for each construct. We used non-parametric statistical tests for the analysis due to non-normality and we used the Benjamini–Hochberg procedure to control for false positives (Benjamini and Hochberg, 1995) due to the multiple statistical tests. The adjusted p-values are reported for each test.

Regarding the sample characteristics, there was a similar distribution of gender: 490 male (49.8%) and 494 female (50.2%) respondents. The minimum respondent age was 18 and the maximum was 69. We found no statistical differences in the distribution of age and gender between experimental condition groups. Half of respondents (52%) indicated that they read the privacy policies of websites at least occasionally. In addition, a majority of respondents indicated they would not use English websites or apps (81%), or read privacy policies in English (88%).

### 4.2 Interest in the Privacy Policy

We first compared initial interest for the privacy policies in the different language conditions. The results of a Mann-Whitney U test showed that interest in the contents of the Japanese privacy policy was significantly higher ( $p=0.024$ ) than interest in the English one.

Next, we used separate Wilcoxon Signed-Rank tests to evaluate interest in the privacy policy after viewing PrivacyGuide, for both risks levels in each language condition. Interest in the contents of the Japanese privacy policy significantly increased for both risk levels (low:  $p=0.0004$ ; high:  $p=0.04$ ). On the other hand, interest in the English privacy policy increased only after viewing the low risk results ( $p=0.0002$ ), but not for the high-risk result. The results provide evidence that PrivacyGuide can indeed promote interest in the privacy policy. In the case of the high-risk English privacy policy, the lack of interest may be due to respondents completely dismissing the possibility of using the website itself

and therefore considering that they no longer have to worry about the contents of its privacy policy.

### 4.3 Website Perception

We evaluated website perception to address whether PrivacyGuide communicated risk to respondents. First, we compared the respondents' initial perception of the English and Japanese privacy policy language versions of the website, using separate Mann-Whitney U tests for the statistical analysis. Respondents were less willing to register on a website with an English privacy policy ( $p=0.015$ ) than a website with a Japanese privacy policy, although we found no statistically significant differences in ( $p=0.07$ ) or privacy concern ( $p=0.34$ ).

We also compared the time spent viewing the experiment website, the PrivacyGuide result screen and the total time taken to finish the survey, between participants in different language conditions. Time data was not normally distributed, so we used separate one-tailed Mann-Whitney U tests for the analysis. There were no significant differences for any of the viewing times or for the total survey time ( $p>0.05$ ). Since the majority of participants indicated low self-perceived English ability, this suggests that participants' lower behavioral intention was influenced simply by seeing the English text.

We then evaluated the effect of PrivacyGuide on perception of the website, using separate Wilcoxon Signed-Rank tests for both risk levels in each language condition. There was a significant increase in behavioral intention ( $p=0.0004$ ) and a significant decrease in risk perception ( $p=0.004$ ) for the low-risk English privacy policy condition, but there were no significant differences for any of the other conditions.

Considering that behavioral intention towards the websites with an English privacy policy was initially lower, the results suggest that PrivacyGuide effectively communicated risk level information to respondents, whose perception of the website improved for the low-risk condition but not for the high-risk condition. This also might explain why interest in the contents of the high-risk English policy did not increase: users would not be interested in the privacy policy of a website that they are not considering using.

In the case of the Japanese privacy policy conditions, the initial intention and risk perception may not have been greatly influenced by the privacy policy itself, but rather by website unfamiliarity. Therefore, although the respondents were more interested about the contents of the Japanese privacy

policy, any additional information about the privacy policy did not have significant influence on the website itself. On the other hand, we did not find any statistical difference in privacy concern for any of the conditions.

### 4.4 Interest in Privacyguide

To address the question of interest in PrivacyGuide, we coded the responses to the open-ended question. A native Japanese speaker familiar with the goal and structure of the survey coded the responses to the open-ended question on participants' interest in trying PrivacyGuide according to whether they were positive, negative or neutral ("I don't know"). If the answer did not correspond to either of these types, it was coded as "other". Table 1 shows the comments by type (blank or "Other" responses are not included).

Table 1: Interest in PrivacyGuide.

	Positive		Negative		Don't know	
	N	%	N	%	N	%
EN-Low R	62	26%	123	51%	17	7%
EN-High R	46	19%	136	56%	14	6%
JP-Low R	67	27%	134	53%	8	3%
JP-High R	56	22%	138	55%	10	4%

The proportion of positive, negative and neutral responses was similar for the same risk level of different language groups. We validated this by quantitatively testing for the difference in interest in using PrivacyGuide between risk levels for each language, using separate one-tailed Mann-Whitney U tests. There was slightly higher interest in using PrivacyGuide for a low risk English privacy policy, but the difference was not statistically significant ( $p>0.05$ ). Similarly, the Japanese group showed no statistical difference in interest between the risk levels ( $p>0.05$ ).

In addition, the results of cross tabulation showed a significant relationship between higher privacy policy reading frequency and positive interest in PrivacyGuide (chi square=26.52,  $df=2$ ,  $p<0.001$ ). Respondents who read privacy policies at least some times are the ones more likely to be interested in PrivacyGuide.

### 4.5 Use Cases and Barriers

In addition to the quantitative analysis, we also qualitatively analyzed the content of responses. Regarding which websites they would be interested in analyzing, respondents in both language groups mentioned a variety of use cases. They gave as an

example types of websites ranging from online shopping and SNS websites to financial and government websites.

In addition, they indicated interest in trying PrivacyGuide on the privacy policies of websites they frequently used—with mentions of Google, Yahoo and Instagram, among other well-known international websites—but they also mentioned wanting to use it on unfamiliar websites. In particular, they mentioned wanting to use PrivacyGuide when registering on a new website, if they felt the website was asking for too much personal information. Respondents mentioned personal information in general, only specifying address, phone number and credit card as examples. Interestingly, respondents from the English language groups mentioned an interest on trying PrivacyGuide on Japanese websites; conversely, a respondent from the Japanese language groups mentioned the potential usefulness for analyzing foreign websites' privacy policies.

With regards to respondents who indicated no interest in using PrivacyGuide, for the most part they did not specify a reason for their answer. We consider that non-interactive PrivacyGuide result screen could have limited further the interest shown in the application. Next in frequency were respondents who mentioned that they had concerns regarding the trustworthiness and reputation PrivacyGuide, and therefore would not use it. Similarly, respondents who answered neutrally mentioned that they would consider using PrivacyGuide if it could be trusted, if it was provided by a well-known company or "if everybody used it". Negative and neutral respondents also mentioned that they did not know the accuracy of PrivacyGuide, and therefore did not know whether they could rely on its results. Other respondents indicated that they did not need to use PrivacyGuide because they would not use risky websites in the first place.

## 5 DISCUSSION

The results indicate that privacy policy summarization, as the one provided by PrivacyGuide, has potential to be beneficial for Japanese users, in particular for foreign language privacy policies. PrivacyGuide effectively communicated risk in case of the English language privacy policy, reflected in their perception of the website, although there were no changes for the Japanese language privacy policy. In the case of

increase of interest in the privacy policy contents and interest in using PrivacyGuide, results were similar for the Japanese and English conditions.

In general, results are consistent with previous research in other countries. In the case of privacy concern, on which Privacy had no effect, the results are similar to those found in (Gluck et al., 2016). We consider the possibility that understanding of the privacy practices of a website by itself cannot ease users' feelings of concern, in particular for an unknown website. The results also indicate challenges and areas for improvement. Addressing lack of reputation is one way that could help improve adoption among Japanese users. This could be realized if well-known organizations or companies provide or support these applications, although what constitutes "well-known" has to consider the local context.

An important challenge to consider is Japanese users' concern about the trustworthiness of PrivacyGuide. We could not determine whether there were specific aspects that would influence trustworthiness perception, but we consider that Japanese users' comments on the accuracy of results may be a factor. Research on automation and machine learning-based applications suggests that providing explanation of results can increase trust (Lee and See, 2004), but there is need to consider how much information to provide according to the context and user expectations (Kizilcec, 2016). In the case of the Japanese privacy policy, presenting users the fragment of text used to decide on a privacy aspect risk level would have provided some additional information that would help them evaluate accuracy. Future research is planned to evaluate how to communicate trustworthiness and accuracy of the application, in particular for users who have to rely on it to understand foreign language privacy policies.

## 6 CONCLUSIONS

We conducted an experimental survey among Japanese users to evaluate of the effects of a privacy policy summarization application, PrivacyGuide. We considered two scenarios, native language (Japanese) and foreign language (English) privacy policies. The results showed that PrivacyGuide can achieve its goal of increasing interest in the content of privacy policy, for both languages. And in the case of the English privacy policy, it can effectively communicate risk and affect perception of a website. In addition, we found that Japanese users would

want analyze the privacy policies of different types of website —familiar and unfamiliar, domestic and foreign. On the other hand, users indicate that lack of trustworthiness, reputation and explanation about the results are barriers for use of the application. In future research, we will address the barriers identified by users, in particular regarding trust and how to provide explanations for automated analysis.

## ACKNOWLEDGEMENTS

We would like to thank Naonori Kato for his help during the preparation of this study.

## REFERENCES

- Benjamini, Y., Hochberg, Y., 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society. Series B (Methodological)* 57, 289–300.
- Curran, P.J., West, S.G., Finch, J.F., 1996. The Robustness of Test Statistics to Nonnormality and Specification Error in Confirmatory Factor Analysis. *Psychological methods* 1, 16–29.
- EU, 2016. REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation).
- Gideon, J., Cranor, L., Egelman, S., Acquisti, A., 2006. Power Strips, Prophylactics, and Privacy, Oh My! Presented at the *Symposium On Usable Privacy and Security (SOUPS) 2006*, ACM Press, p. 12.
- Gluck, J., Schaub, F., Friedman, A., Habib, H., Sadeh, N., Cranor, L.F., Agarwal, Y., 2016. How Short Is Too Short? Implications of Length and Framing on the Effectiveness of Privacy Notices, in: *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. USENIX Association, Denver, CO, pp. 321–340.
- Harbach, M., Hettig, M., Weber, S., Smith, M., 2014. Using Personal Examples to Improve Risk Communication for Security & Privacy Decisions, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '14*. ACM, New York, NY, USA, pp. 2647–2656.
- Harkous, H., Fawaz, K., Lebret, R., Schaub, F., Shin, K., Aberer, K., 2018. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning [WWW Document].
- Kelley, P.G., Cesca, L., Bresee, J., Cranor, L.F., 2010. Standardizing Privacy Notices: An Online Study of the Nutrition Label Approach, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10*. ACM, New York, NY, USA, pp. 1573–1582.
- Kim, D.J., Ferrin, D.L., Rao, H.R., 2008. A Trust-Based Consumer Decision-Making Model in Electronic Commerce: The Role of Trust, Perceived Risk, and Their Antecedents. *Decision Support Systems* 44, 544–564.
- Kizilcec, R.F., 2016. How Much Information?: Effects of Transparency on Trust in an Algorithmic Interface, in: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, CHI'16*. pp. 2390–2395.
- Kline, R.B., 2005. Principles and practice of structural equation modeling, 2nd ed, Principles and practice of structural equation modeling, 2nd ed. Guilford Press, New York, NY, US.
- Lee, J.D., See, K.A., 2004. Trust in Automation: Designing for Appropriate Reliance. *Hum Factors* 46, 50–80.
- McDonald, A.M., Cranor, L.F., 2008. The Cost of Reading Privacy Policies. *ISJLP* 4, 543.
- Proctor, R.W., Ali, M.A., Vu, K.-P.L., 2008. Examining Usability of Web Privacy Policies. *International Journal of Human-Computer Interaction* 24, 307–328.
- Rosseel, Y., 2012. lavaan: An R Package for Structural Equation Modeling. *Journal of Statistical Software*, Articles 48, 1–36.
- Steinfeld, N., 2016. “I agree to the terms and conditions”: (How) do users read privacy policies online? An eye-tracking experiment. *Computers in Human Behavior* 55, 992–1000.
- Sunyaev, A., Dehling, T., Taylor, P.L., Mandl, K.D., 2015. Availability and quality of mobile health app privacy policies. *J Am Med Inform Assoc* 22, e28–e33.
- Tesfay, W.B., Hofmann, P., Nakamura, T., Kiyomoto, S., Serna, J., 2018. PrivacyGuide: Towards an Implementation of the EU GDPR on Internet Privacy Policy Evaluation, in: *Proceedings of the Fourth ACM International Workshop on Security and Privacy Analytics, IWSPA '18*. ACM, New York, NY, USA, pp. 15–21.
- Wilson, S., Schaub, F., Dara, A.A., Liu, F., Cherivirala, S., Giovanni Leon, P., Schaarup Andersen, M., Zimmeck, S., Sathyendra, K.M., Russell, N.C., B. Norton, T., Hovy, E., Reidenberg, J., Sadeh, N., 2016. 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)*. pp. 1330–1340.
- Zimmeck, S., Bellovin, S.M., 2014. Privee: An Architecture for Automatically Analyzing Web Privacy Policies, in: *Proceedings of the 23rd USENIX Conference on Security Symposium, SEC'14*. USENIX Association, Berkeley, CA, USA, pp. 1–16.