What About My Privacy?
Helping Users Understand Online Privacy Policies

Wasja Brunotte
Leibniz University Hannover
Software Engineering Group
Cluster of Excellence PhoenixD
Hannover, Germany
wasja.brunotte@inf.uni-hannover.de

Larissa Chazette
Leibniz University Hannover
Software Engineering Group
Hannover, Germany
larissa.chazette@inf.uni-hannover.de

Lukas Köhler
Leibniz University Hannover
Software Engineering Group
Hannover, Germany
lkoehler@se.uni-hannover.de

Jil Klünder
Leibniz University Hannover
Software Engineering Group
Cluster of Excellence PhoenixD
Hannover, Germany
jil.kluender@inf.uni-hannover.de

Kurt Schneider
Leibniz University Hannover
Software Engineering Group
Cluster of Excellence PhoenixD
Hannover, Germany
kurt.schneider@inf.uni-hannover.de

ABSTRACT
Software systems have become an integral part of our daily lives. However, users tend to forget that they are not only consuming information, but also delivering personal information to service providers. This data collection means that users' privacy sphere is increasingly at stake. Informing users about what and how data is collected is pivotal for reaching transparency, trustworthiness, and ethics in modern systems. The main purpose of privacy policies is to inform users about what happens to their personal data. But instead they are extensive and purposefully obfuscating. Information about data practices are hidden in long and ambiguous text passages. To mitigate this, in this paper, we present a concept implemented as a web extension to support the end-user in dealing with privacy policies by providing easier access and visual explanations to privacy-related information. We evaluated the usefulness of our tool in a user study with 65 participants. The results show that our approach helps users to find a privacy policy faster and also supports users to better comprehend the relevant information. Our tool is a first step towards facilitating to deal with privacy policies from the end-user perspective. The results of the study and the positive feedback from the participants show a high degree of acceptance and potential for the tool to increase users' privacy awareness.

CCS CONCEPTS
• **Software and its engineering**; • **Information systems** → Collaborative and social computing systems and tools; • **Security and privacy** → **Social aspects of security and privacy**; **Usability in security and privacy**;

ICSSP'22, May 20–22, 2022, Pittsburgh, PA, USA

ACM ISBN 978-1-4503-9674-5/22/05.
https://doi.org/10.1145/3529320.3529327

KEYWORDS
Explainability, privacy, privacy policy, privacy awareness, visual explanations

ACM Reference Format:
Wasja Brunotte, Larissa Chazette, Lukas Köhler, Jil Klünder, and Kurt Schneider. 2022. What About My Privacy? Helping Users Understand Online Privacy Policies. In *Proceedings of the International Conference on Software and System Processes and International Conference on Global Software Engineering (ICSSP'22), May 20–22, 2022, Pittsburgh, PA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3529320.3529327

1 INTRODUCTION
Our daily lives are heavily interwoven with the use of software systems and the associated information processing. Systems support users in the efficient purchasing and distribution of goods, enable them to obtain information quickly and easily, and in many cases simplify their everyday lives, thus offering great convenience. However, users often forget that they are not only consuming information, but also delivering personal information to service providers. Data is a byproduct of computing, and as a result, numerous digital traces are "left behind" while utilizing software systems, and much of this data reveals a lot about individuals. For instance, purchasing patterns (online and offline) are recorded [19] and they can reveal if someone is pregnant, sick, or unemployed [41]. This information has not only a personal value, but also great monetary value. Indeed, this information is sufficiently valuable that it has spawned an entire industry that buys and sells information, the so-called data brokers [3, 41].

During the use of software systems, users automatically and usually unconsciously make a trade-off between the benefits a system offers and the data they reveal about themselves. Normally, this is not a privacy threat, as long as a company gathers and uses this data with "good intention" and to provide value and improved service to users. Here, it is crucial that a company's data practices are fair and transparent. End-users have the right to greater transparency and accountability in respect to their data. In practical terms, users should have the right to know what, how, and by whom data is gathered, processed, stored, and used. This kind of information is pivotal for reaching transparency, trustworthiness, and ethics in modern systems. This information disclosure is essential, so we can be sure, as users, that our data is not sold and used for purposes beyond our knowledge and consent.

Problem Statement. Privacy policies (PPs) are intended to inform users about what happens to their personal data, but instead they are extensive and purposefully obfuscating. Information about data practices are hidden in long, vague, and ambiguous text passages [30, 34, 36]. End-users cannot easily access the relevant information on the processing of their data, decreasing the transparency of modern systems and violateing the users' rights about getting to know what happens to their data.

Objective. To increase the transparency and to ease the access to the information presented in the privacy policy, in this paper, we want to support the users understanding privacy policies.

Contribution. We present a concept that consists of three steps: (1) checking a website for the presence of a privacy policy, (2) analyzing the privacy policy, and (3) providing visual explanations for the user to support the understanding of the policy's content. These three core ideas are implemented in a tool called Online Privacy Policy eXplainer (PriX) that explains PPs to end-users. In line with our research agenda [12], the purpose of this tool is to progress towards a more practical and useful solution that assists end-users in dealing with PPs by providing contextual information about PPs in a friendlier manner and assisting them in improving their privacy awareness. We conducted a usability study in combination with an online questionnaire to understand users' behavior regarding their online privacy, and to assess the tool's support in locating and interpreting information in online PPs. 65 people participated in our study. The results confirm that three out of four participants care about data privacy but rarely or never read the available PPs (87.7%). In addition, 88% of the participants agreed that the visual explanations on PriX helped them better understand the presented privacy information. As a result, we may infer that offering explanations about privacy policies is an useful way to make users more aware of a system's privacy conditions.

Outline. The rest of the paper is structured as follows: Section 2 presents background and related work. Section 3 summarizes our concept and the resulting prototype. Section 4 provides information on the used study design to evaluate the concept. In Section 5, we present the results of the study which we discuss in Section 6. We conclude the paper in Section 7.

2 BACKGROUND AND RELATED WORK

The current data protection law in the European Union (EU) is regulated by the European General Data Protection Regulation (GDPR). The GDPR is characterized by the principle of prohibition of data processing without permission. This implies that personal data processing is typically prohibited unless there is a legal permission or the data subject's agreement. It preserves the fundamental right to personal data protection, allowing each individual to make their own decisions about the disclosure and use of their personal data [22]. Yet, to protect and enforce this fundamental right, the GDPR requires that users are fully informed about the data practices of a given service provider. It also stipulates that data processing must be fair and transparent. However, PPs comprise rather extensive texts with a lot of legal technical terms and often contain ambiguous wording. As a consequence, such extensive texts are neither read nor understood, and finally ignored [14, 30, 34–36, 42].

Rudolph et al. [37] investigate why users are not willing to take care about appropriate privacy actions, such as reading PPs as well as managing available privacy settings. Their study reveals that users perceive these texts as well as taking appropriate privacy actions as too complicated and time consuming. Based on these results, they provide an intention model to better understand why users neglect to take appropriate privacy actions. They conducted a case study and found out that their model is able to explain correlations between users' privacy needs, their intentions, and behavior. Lederer et al. [27] point out that there should not be a separation between privacy information and the user's activities, but it should be integrated into the user's activities, according to [40]. Ortloff et al. [32] confirmed this in a user study comparing context-aware with traditional PPs.

Over time, several efforts have been made to support end-users in dealing with online PPs. The World Wide Web Consortium (W3C), through the Privacy Preferences Project (P3P), strived to create a standardized computer-readable format for PPs that would allow user agents (e.g., web browsers) to read and process them. This way, user agents would be able to automate decisions based on a user's privacy settings, so they would not have to read privacy policies on every website they visit [17, 18]. Several prototypes based on the P3P have been developed, including the *Privacy Bird*: a user agent that compares P3P-compliant PPs to the end-user's privacy settings and makes data exchange decisions based on these settings [20]. However, the P3P project was suspended since there was insufficient support from current browser implementers.

Harkous et al. [23] proposed *Polisis.* It is an automated framework for PP analysis. An application that builds on Polisis is the freeform question-answering system *PriBot.* It is some sort of chatbot, which enables a user to actively ask questions regarding a PPsuch as *"Do you share my data with third parties?"* and PriBot gives the corresponding answer.

Both systems, Privacy Bird and PriBot require the active involvement of the end-user. If the user wants information about a specific PP, he or she must first access the service or, in the case of PriBot, enter a question into the system before receiving information about the PP in question. We believe that access to information regarding privacy should be made as easy and simple as possible. This is the only way we can get the user to actually use such an information service. Another important aspect is that the user should be able to access the information in his usual working environment, in this case the web browser, so that he or she does not have to interrupt his or her workflow.

Considering the efforts and results of previous research, there is a lack of a system that takes into account and implements the findings of this research. It is important to integrate privacy information into end-users' activities as well as to contextualize it. With PriX, we aim to close this gap.

3 CONCEPTS

To reach our goal of supporting a user understanding online privacy policies, we developed a three-step approach that we prototypically implemented to evaluate its applicability: PriX[1] is implemented as a browser extension. Figure 1 outlines its underlying concept. When a user visits a website, the extension 1) checks if the website provides a PP. If this is the case, 2) the PP is automatically analyzed by the extension. Next, 3) the result of the analysis is visually presented to the user. The three steps are described in details below.

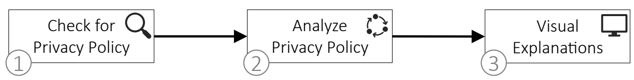

Figure 1: Concept of the Online Privacy Policy eXplainer

3.1 Check for a Privacy Policy

There is no uniform standard or format for the structure of a PP. Furthermore, it is also not specified how a website that provides a PP must link to it. To meet this challenge, we manually inspected 100 of the most visited websites on the Internet[2]. We chose this list to ensure that we inspect websites that have a high level of relevance in the daily browsing behavior of end-users. The goal of this inspection was to identify commonalities in the way a PP is structured, how a website embeds the link to it, and what PPs have in common regarding their content. The result regarding the linking of PPs was that 100% of the inspected websites contain the word "*privacy*" in the directory path of the URL or in the subdomain. In rare cases, websites may have more than one privacy policy, for instance, if a company sells more than one product and these refer to different privacy statements. In these cases, the shortest link has usually led to a general privacy policy. This was taken into account when designing the algorithm that automatically searches for the privacy policy. However, these facts complicate finding a privacy policy "by accident" to get information if the user does not explicitly ask for this kind of information.

Regarding the content of a website, our inspection has revealed that PPs usually contain a specific set of words which we refer to as *privacy indicators*. Typical privacy indicators are *"privacy notice"*, *"privacy statement"*, *"privacy policy"*, *"GDPR"*, *"your privacy"*, or *"data protection"*. Based on these findings, we are able to automatically find the PP of a website with a very high degree of accuracy.

3.2 Analysis of a Privacy Policy

Even if the user has found a privacy policy, it still requires a lot of effort to analyze it with respect to a specific information (e.g., about the data storage). Hundreds of pages have to be read in order to find the respective paragraph in the privacy policy. Hence, our second concept strives to automatically analyze a privacy policy with respect to typically requested information.

The analysis process is based on classification algorithms. The algorithms were trained by using the *OPP-115 Corpus* (Online Privacy Policies, set of 115) [45]. It is a set of real website PPs that contains data annotations. To ensure its legal validity, three law graduate students read and annotated real PPs with respect to data practices that were found in the text for the aforementioned project [45]. This resulted in ten categories into which the individual sections of a privacy policy can be divided, the so-called *data categories*, as shown in Figure 2.

A separate classification algorithm was trained for each category to achieve high accuracy. We chose *Naive Bayes* and *Random Forest* as classifiers. The two classifiers perform differently in the respective data categories, so we use the best of the two classifiers for each category. Currently, PriX only processes English-language PPs.

3.3 Visual Explanations

As a last part of our concept, the found information resulting from the second step need to be presented in an adequate way. As the user needs to understand the privacy policy, an explanation is requested in order to inform the user. According to the definition of explainable systems [16], a system is explainable when it presents an explanation I about a given system aspect X to an addressee A, so that A can understand this X of the system. Based on this definition, we define a *visual explanation* as follows:

- the aspect X to be explained is a *privacy policy*
- the addressee A is the *end-user* who needs a privacy explanation
- the explanation I is a *visual representation* of a subset of X

The first component of such a visual representation is a category that describes a passage of a PP, e.g., "third-party collection" if a website shares the data with other parties (cf. Section 3.2). This category summarizes a passage of text and discloses relevant information for understanding the policy and supporting decision-making, which can be considered an explanation. To enrich the explanation and to facilitate its comprehension, icons are the second component. Icons are an essential part of user interface (UI) design. They make a UI usable, help the user understand it, and even speed up the understanding process [2, 4, 7, 28]. We employed icons from the *privacy icons association*[3].

As a result, we define *visual explanation* as the interaction between a PP's explanation (data category) and the visual element (icon). An overview of the privacy icons used and the corresponding data categories from Figure 2 is shown in Figure 3. An overview of PriX's visual representation of a PP (using the example of mozilla.org) can be found at [10].

4 EVALUATION

The overall goal of the research presented in this paper is to support a user understanding a privacy policy. To achieve this goal, we developed the concept consisting of three sub-concepts presented in Section 3. In PriX, we prototypically implemented the core ideas of the concept. In a next step, we evaluate the prototype for its usefulness.

[1] https://prix.se.uni-hannover.de
[2] https://moz.com/top500
[3] https://privacy-icons.ch

First-Party Collection Privacy practice describing data collection or data use by the company/organization owning the website.	***Third-Party Sharing/Collection*** Privacy practice describing data sharing with third parties or data collection by third parties. A third party is a company/organization other than the first party company/organization that owns the website.
Data Security Practice that describes how users' information is secured and protected, e.g., from confidentiality, integrity, or availability breaches. Common practices include the encryption of stored data and online communications.	***User Access, Edit and Deletion*** Privacy practice that allows users to access, edit or delete the data that the company/organization has about them.
User Choice/Control Practice that describes general choices and control options available to users.	***Data Retention*** Privacy practice specifying the retention period for collected user information.
Policy Change The company/organization's practices concerning if and how users will be informed of changes to its privacy policy, including any choices offered to users.	***Do Not Track*** Practices that explain if and how Do Not Track signals (DNT) for online tracking and advertising are honored.
International and Specific Audiences Specific audiences mentioned in the company/organization's privacy policy, such as children or international users, for which the company/organization may provide special provisions.	***Other*** Other information, such as possibilities to contact the operator of the website.

Figure 2: Data categories of PriX

Figure 3: The privacy icons used by PriX

4.1 Study Goal

We applied the goal definition template by Wohlin et al. [46] to formulate the goal of our study.

> **Goal Definition:** Analyze the developed concept of PriX to evaluate the impact and perception of software-side support for PPs with respect to usability, effectiveness, and understandability from the perspective of end-users in the context of an online study using the think-aloud method.

In particular, we pay attention to the following four aspects that are at the core of our concept (cf. Figure 4):

(1) Finding the privacy policy (Q1)
(2) Finding specific information in the privacy policy (Q2)
(3) Understanding a privacy policy (Q3)
(4) Perceived ease of use (Q4)

(1) Many websites offer a PP. However, there is no standardized specification of how it must be embedded and linked to in a website (cf. Sec. 3.1). Therefore, as a first aspect, we investigate whether PriX supports end-users in finding a website's PP more efficiently and, most importantly, more quickly.

(2) In cases where users only want to learn about certain aspects of data processing and do not have the time to read the entire PP, finding a specific information may pose a challenge. Many PPs do not offer a table of contents that users could use to navigate directly to the section they are interested in. A table of contents may help users to find specific information more quickly. Therefore, the second aspect is to examine whether our tool helps end-users finding a specific information more quickly.

(3) Without a legal background, PPs are often difficult to understand and the information is "hidden" in long passages of text. This may overwhelm users and result in PPs not being read at all. For this reason, we analyze whether PriX assists the end-user in understanding a PP.

(4) Software usability is an important point for the acceptance and use of a system. Poor usability can have far-reaching consequences, ranging from annoying users to endangering their lives [26]. Therefore, we also investigate the usability of PriX and the acceptance of the tool, identifying possible usability drawbacks.

4.2 Study Design

To achieve the aforementioned goal and to analyze the aspects, we combined a *synchronous remote usability test* with an online questionnaire. To ensure the quality of our survey, we followed established guidelines for survey design [25, 43]. Due to the pandemic, it was not possible to conduct a conventional *in loco* usability study. However, Brush et al. [13] state that there is no quantitative difference between a remote and local study. We decided for a *between-groups* design, since we wanted to avoid familiarization as well as learning effects that could influence the study results. Therefore, the participants were randomly divided into two groups. The *experimental group*, which used PriX to complete given tasks, and the *control group*, which had to complete the tasks without using PriX. The same tasks were given to both groups without prior interactions with any of the study's PPs, since this could have introduced a bias. During the study, the *Think-Aloud* [8] method was employed and the conversations (audio) as well as the screen (video) of the participants were recorded. The study was conducted with a total of 65 participants. The experimental group had 33 participants, whereas the control group had 32. There were no prerequisites for participating in the study.

Both groups were given the same tasks based on two real PPs. One PP stemmed from *mozilla.org* and the other one from *netflix.com*. We picked these PPs because Mozilla's PP is short and the provider collects less information. In contrast, Netflix's PP is rather long and contains a lot of information about the provider's data practices. The experimental group obtained no introduction to the use of PriX. They were only told, at the beginning of the experiment, which icon in the toolbar of the browser could be used to open PriX. The task t_1 was to find the PP of Mozilla.org. Analogously, the task t_2 was to find the PP of Netflix.com. In task t_3, Mozilla's privacy policy should be looked up to see if data is collected and if so, in what way. In task t_4, the subjects were asked to check whether Mozilla protects its users' data and, if so, in what way. Tasks t_5 and t_6 were related to Netflix's PP. Task t_5 asked whether Netflix uses cookies. If so, the respondent should find out what the cookies are used for. Task t_6 dealt with finding out whether Netflix collects data Netflix about its users. If Netflix does, subjects were asked to answer what data Netflix collects.

4.3 Data Collection

We conducted the study between April and May 2021. We collected data from 65 participants. Due to the nature of an online study, we used tools such as TeamViewer or AnyDesk to grant the participants the right to access our computer. This way, participants did not need to install any software or plugin on their computer. In addition, we used BigBlueButton hosted by the Leibniz University Hannover for communication. OBS was used to record the screen of the experimenters computer to analyze the experiment afterwards. This screen-recording also allowed us to exactly measure the task times.

4.4 Data Analysis

During our study, we measured specific variables that help evaluating the four aspects mentioned above. Figure 4 summarizes the four aspects and the metrics used to measure the influence of PriX on these aspects.

Variables. We collected data on the variables presented in Figure 4. In particular, we measured the times required to solve specific tasks, the task accuracy, and the usability score. These metrics are defined and calculated as follows:

- **Task Time** (in seconds) defined as the time needed to solve a specific task, e.g., finding the privacy policy or an information.
- **Task Accuracy** defined as an ordinal scale ranging from *A* to *C*:
  - *"A"* means the respondent was able to solve the task correctly
  - *"B"* means that the respondent was in the correct paragraph of the PP, but in the end could not provide the required information correctly
  - *"C"* means that the task could not be answered or that the provided answer was wrong
- **Usability Score** defined as the System Usability Scale (SUS) [9], which is regarded as an established metric for assessing software usability [5, 33].

In addition to these quantifiable metrics, we also considered the perceptions of the participants reported on Likert scales. This way, we obtained additional information on the participants' agreement with specific items, including

- The privacy policy was easy to find
- The required information from the privacy policy was easy to find
- The content of the privacy policy was easy to understand

Qualitative Analysis. Open-ended and closed-ended questions were asked in the questionnaires. For example, open-ended questions were asked about task $t6$ as well as questions for the experimental group about the visual explanations (e.g., what a particular data category might mean). We followed an open-coding approach [38] to analyze the open-ended questions. It consists of two consecutive coding cycles. We applied *In Vivo Coding* [15] during the first coding cycle to reflect the perspectives of the participants. The second coding cycle consists of *Pattern Coding* [31] to break down the initial codes into a smaller number of themes or constructs. The coding process was conducted independently by two authors of this work. In cases of discrepancies, the authors discussed them until they found a common consensus.

Quantitative Analysis. In order to measure the influence of our prototype with respect to the four aspects mentioned before, we analyzed the data from both treatment groups with respect to differences. In particular, we tested the null hypotheses summarized in Table 1.

These two hypotheses H1 and H2 are related to the first two aspects (i.e., finding the privacy policy and finding a specific information in the PP). The subjects of both groups were asked to open the given website (t_1 : mozilla.org, t_2 : netflix.com) and then to search for and open the PP. Since task times have a strong positive skew, we applied a log transformation to the raw times to improve the accuracy from the results [39]. According to the central limit theorem, since there is a sufficient sample size ($n > 30$), a normal distribution can be assumed [24]. Thus, we used the z-test to analyze $H1.1_0$ and $H1.2_0$. However, as we performed two tests to test the main hypothesis H1, we apply the Bonferroni correction leading to an adjusted p-value of $p_{corr} = p/2 = 0.025$. That is, as soon as one of the two hypotheses H1.1 and H1.2 is significant with a p-value smaller than p_{corr}, we can assume that there are differences comparing the two groups and, thus, we can reject H1.

To compare the times required to find a specific information in the PP, we applied a similar procedure: Both groups were given the same tasks ($t_3 - t_6$). Tasks t_3 and t_4 were related to Mozilla's PP and t_5 and t_6 were related to Netflix's PP. Task t_3 requested that participants responded whether Mozilla collects data about the user and, if so, what data. In task t_4, participants were asked whether

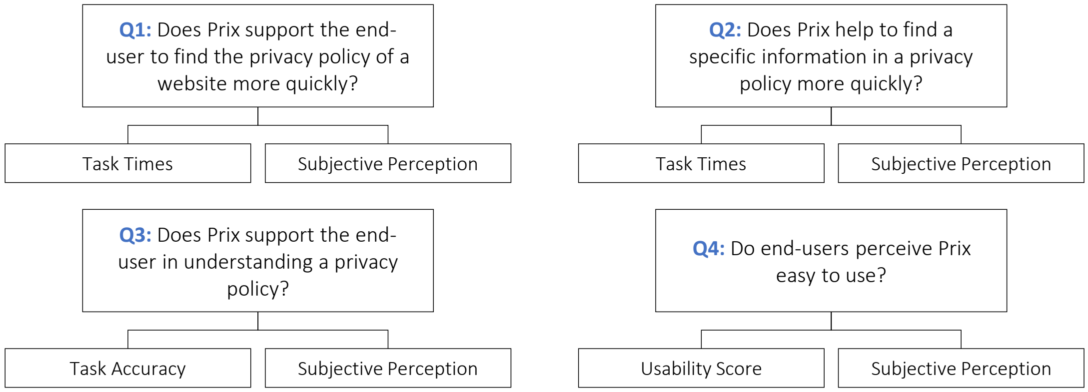

Figure 4: Questions and Related Metrics

Table 1: Null hypotheses to test the data for differences between the two treatment groups

ID	Hypothesis
$H1_0$	There is no difference between the two treatment groups with regard to the time needed to find the privacy policy.
$H1.1_0$	There is no difference between the two treatment groups with regard to the time needed to find Mozilla's privacy policy.
$H1.2_0$	There is no difference between the two treatment groups with regard to the time needed to find Netflix's privacy policy.
$H2_0$	There is no difference between the two treatment groups with regard to the time needed to find a specific information.
$H2.1_0$	There is no difference between the two treatment groups with regard to the time needed to find information on the data collection in Mozilla's privacy policy.
$H2.2_0$	There is no difference between the two treatment groups with regard to the time needed to find information on the data protection in Mozilla's privacy policy.
$H2.3_0$	There is no difference between the two treatment groups with regard to the time needed to find information on the use of cookies in Netflix's privacy policy.
$H2.4_0$	There is no difference between the two treatment groups with regard to the time needed to find information on the data collection in Netflix's privacy policy.

they think that Mozilla protects user's data and, if so, in what way. In t_5, we asked participants whether Netflix uses cookies and, if so, why. For the last task (t_6), we asked the participants whether Netflix collects data about a user, if so, in what way.

According to the central limit theorem, normal distribution can be assumed since there is a sufficient sample size ($n > 30$). We also applied a log transformation to the task times and used the z-test to test $H2.1_0$ to $H2.4_0$ for statistical significance. Again, as we performed four tests to test the main hypothesis H2, we apply the Bonferroni correction leading to an adjusted p-value of $p_{corr} = p/4 = 0.0125$. That is, as soon as one of the two tests is significant with a p-value smaller than p_{corr}, we can assume that there are differences comparing the two groups.

For the third aspect, the understanding of a privacy policy, we analyzed the accuracy of the participants' results when performing a specific task using the scheme described in Section 4.4.

For the fourth aspect, we considered the system's usability score. With this measure we wanted to investigate the perceived usability of PriX.

5 RESULTS

We conducted the online experiment as described in Section 4 and analyzed the data accordingly. In the following, we present the results.

5.1 Demographics

The participants ranged in age from 18 to 59 years old (M=23.8, SD=5.4). The majority of participants (80%) were students. A total of 64 participants (98.46%) stated that they used the Internet at least once a day. 49 subjects (75.38%) agreed with the statement *"privacy is important to me"*. This contrasts with the subjects' statements on how often they pay attention to whether a website site has an PP. In this case, 72.31% of the subjects said they rarely to never do so. In addition, only six participants (9.23%) stated that they often read the PP of the website they are currently visiting. Three participants (4.62%) said that they always read it, and nine (13.84%) stated that they occasionally do it.

5.2 Finding a Privacy Policy

First, we compared the time required to find a privacy policy with or without PriX. The subjects of both groups were asked to open the given website (t_1 : mozilla.org, t_2 : netflix.com) and then to search for and open the PP. To test hypothesis H1 (cf. Table 1), we tested the hypotheses H1.1 and H1.2 using the z-Test. The results are presented in Table 2.

The results show a significant difference between the two groups in favor of the experimental group, so we can reject the null hypotheses $H1.1_0$ and $H1.2_0$. In particular, as the p-value for Mozilla's

Table 2: Results of the z-test to test H1

ID	Results	Sign.?
$H1_0$		yes
$H1.1_0$	$z \approx -7.59,\ p < 0.01$	yes
$H1.2_0$	$z \approx -2.13,\ p \approx 0.03$	yes

privacy policy is below the corrected p-value of 0.025, we can also reject the main hypothesis $H1_0$ and we can assume that **the use of PriX significantly reduces the time to find a privacy policy**.

When asked about the level of agreement to the statement *"The privacy policy was easy to find"*, all of the participants of the experimental group agreed (or strongly agreed) for both PPs (Mozilla and Netflix). In the control group 12.49% strongly agreed, 37.5% agreed, 18.75% neither agreed nor disagreed, 21.88% disagreed and 9.38% strongly disagreed with the statement for Mozilla's PP. In the case of Netflix, 53.11% strongly agreed, 37.5% agreed, and 3.13% each neither agreed nor disagreed, disagreed, or strongly disagreed (cf. Figure 5).

	Strongly Agreed	Agreed	Neither Nor	Disagreed	Strongly Disagreed
Mozilla's PP	12.49%	37.5%	18.75%	21.88%	9.38%
Netflix's PP	53.11%	37.5%	3.13%	3.13%	3.13%

Figure 5: Contronl Group - PP was easy to find

Further analysis of the recorded data also showed that some subjects from the control group, despite having problems searching for the privacy policy (both Mozilla and Netflix) during our experiment, later claimed to have had no problems. A reason for this could be that the participants might been have uncomfortable admitting the possible problems they had or admitting their unawareness with respect to the PPs.

Summarizing, we can say that the z-Test has proven significant differences in task times and the comparison of the answers of the participants also confirm that PriX may support end-users in finding a PP of a website more quickly.

5.3 Finding specific Information in a Privacy Policy

In a next step, we compared the times required to find a specific information in a privacy policy. Both groups were given the same tasks ($t_3 - t_6$). Tasks t_3 and t_4 were related to Mozilla's PP and t_5 and t_6 were related to Netflix's PP. Task t_3 requested that participants responded whether Mozilla collects data about the user and, if so, what data. In task t_4, participants were asked whether they think that Mozilla protects user's data and, if so, in what way. In t_5, we asked participants whether Netflix uses cookies and, if so, why. For the last task (t_6), we asked the participants whether Netflix collects data about a user, if so, in what way.

To test the main hypothesis H2, we tested the hypotheses $H2.1_0$ to $H2.4_0$ using the z-test. The results are summarized in Table 3.

For the tasks related to Mozilla's PP, a significant difference for task times was found in favor of the experimental group. Thus, using

Table 3: Results of the z-test to test H2

ID	Results	Sign.?
$H2_0$		yes
$H2.1_0$	$z \approx -2.10, p \approx 0.04$	yes
$H2.2_0$	$z \approx -3.86, p < 0.01$	yes
$H2.3_0$	$z \approx -0.26, p \approx 0.80$	no
$H2.4_0$	$z \approx -1.19, p \approx 0.24$	no

the PriX significantly reduces the time needed to find a specific information in Mozilla's privacy policy.

In contrast, no significant difference was found regarding Netflix's PP. Even though we could not prove a significant difference, the experimental group was quicker on average here as well (see Table 4).

In total, as the p-value for H2.2 is below the corrected p-value of 0.0125, we can nevertheless reject H2 and, thus, we can assume that **there is a significant difference in the required time when searching for a specific information in a privacy policy**. Only about half of the subjects from the control group used the search function to find the specific information. After completing the tasks, we asked the participants if the task information was easy to find. For Mozilla's PP, 72.73% of the participants in the experimental group strongly agreed, 24.24% agreed, and one participant neither agreed nor disagreed (cf. Figure 6). 46.88% of the subjects in the control group agreed, 40.62% strongly agreed, and 12.5% neither agreed nor disagreed (cf. Figure 7).

	Strongly Agreed	Agreed	Neither Nor	Disagreed	Strongly Disagreed
Mozilla's PP	72.73%	24.24%	3.03%	-	-
Netflix's PP	42.43%	48.48%	6.06%	-	-

Figure 6: Exp. Group - PP information was easy to find

When asked about Netflix's PP , 42.43% of the participants in the experimental group strongly agreed, 48.48% agreed, 6.06% neither agreed nor disagreed, and one subject disagreed (cf. Figure 6). In the control group, on the other hand, 18.75% disagreed, 31.25% neither agreed nor disagreed, 21.87% agreed, and 28.13% strongly agreed (cf. Figure 7).

	Strongly Agreed	Agreed	Neither Nor	Disagreed	Strongly Disagreed
Mozilla's PP	40.62%	48.88%	12.5%	-	-
Netflix's PP	28.13	21.87%	31.25%	18.75%	-

Figure 7: Contronl Group - PP information was easy to find

PriX highlights the specific passage of text in a PP that corresponds to the respective data category. 93.9% of the participants agreed that this feature is very useful in finding specific information in a PP. PriX's visual explanations have also been shown to provide a performance advantage for more extensive PPs, which, while not significant in this case, was present on average.

Table 4: Values of task times - RQ2

Task	**M** (experimental group)	**M** (control group)	**SD** (experimental group)	**SD** (control group)
t_5	138.27	145.38	81.69	85.55
t_6	54.09	67.69	60.36	51.99

5.4 Understanding a Privacy Policy

In the next step, we analyzed if the PriX supports the understanding of a privacy policy. In Figure 8, an overview of the accuracy of the results for tasks t_3 - t_6 is depicted. We have aggregated the results for a better overview.

	Experimental Group			Control Group		
	A	*B*	*C*	*A*	*B*	*C*
$t_{3,4}$	87.8%	4.6%	7.6%	70.31%	18.75%	10.94%
$t_{5,6}$	65.2%	18.1%	16.7%	62.4%	14.1%	23.5%

Figure 8: Accuracy of task results - RQ3

A chi-square test of independence was performed to examine the relation between the two groups and the accuracy of the results. The relation between these variables was significant, χ^2 (2, N = 65) = 7.3451, p = .0029. The subjects in the experimental group were more likely to solve the tasks correctly.

After participants have completed their tasks, we asked them if they knew what data the provider collects and what happens to it. The subjects in the experimental group reported being somewhat better informed about the provider's data practices. Here, 31.28% answered with yes, 68.13% answered with no. In the control group 21.87% answered with yes and 78.13% answered with no.

Considering the data in the Figure 8, the significant result of the chi-square test, and the fact that more participants in the experimental group said they were more informed about the provider's data practices, it is possible to conclude that **PriX not only improves understanding, but also can give the user a greater sense of being informed**. However, we cannot establish this conclusion based only on the provided data because it might be also attributable to the respondents' differing viewpoints. Because of our random assignment, a greater number of more distrustful users may have ended up in the control group.

We asked the participants of the control group whether they understood Mozilla's PP. 43.75% (28.12% Netflix) of the subjects strongly agreed, 43.75% (31.25% Netflix) agreed, 9.37% (31.25% Netflix) neither agreed nor disagreed, and only one subject disagreed (9.38% for Netflix's PP). The experimental group was not asked this specific question on subjective impressions of understanding since they mostly engaged with a PP through PriX, and we wanted to know how participants perceived the visual explanations. The experimental group perceived the data categories as supportive for their understanding which is evidenced by an aggregated agreement score of 87.88%. 6.1% disagreed that the data categories supports them in better understanding the PP. With respect to the icons, 64.52% perceived them as supportive, while 9.68% did not. Two subjects did not even notice them. This impact, we believe, was produced by users seeing the icons for the first time. With more frequent use, users might be able to recognize and correlate the symbol with the appropriate data category. Some participants' comments support this hypothesis. PriX also gives detailed text information about a specific data category. 78.8% of the participants perceived the information as useful, 18.17% neither agreed nor disagreed, and only one participant (3.03%) did not perceive it useful. We were also able to demonstrate that PriX assists users in dealing with privacy policies in a more efficient manner. It also helps users to better understand a PP, as supported by the objective metrics (accuracy of task results) that we analyzed and the subjective results regarding how users perceived the tool support.

5.5 Perceived Ease of Use

We calculated the System Usability Scale (SUS) [9] for PriX. The calculated SUS score was of 87.95 on average, which is 19.95 points higher than that of an average web application [6]. According to the central limit theorem and a since we had a sample size of 33, normal distribution can be assumed. The calculated 95% confidence interval [84.49, 91.41] reveals that PriX is a system with at least *above-average* or *almost excellent* usability [6]. 96.7% of the participants of the experimental group stated that PriX has facilitated them in gaining an overview of the PPs. Only one participant neither agreed nor disagreed.

As the findings reveal, **PriX gives effective assistance in coping with privacy policies.** Subjects perceived the visual explanations as helpful, both in terms of finding information and of understanding a privacy policy.

6 DISCUSSION

We end this paper with the interpretation of our results, the discussion of threats to validity, and the presentation of future work.

6.1 Interpretation

Despite the constant data collection that takes place in today's systems, many users remain unaware of the importance of protecting their privacy. There's no question that today's systems support and help us in many areas. As users, we use fitness applications, smart home gadgets, and medical systems, but we often forget that they are invading and compromising our privacy more and more.

PPs are the primary channel that service providers use to inform end-users about data collection and data practices. Unfortunately, the mere existence of a PP does not enable end-users to make conscious choices regarding the use of a particular service in matters of privacy. One reason for this is that the text is often long and difficult to understand. Many users neglect their privacy due to a lack of knowledge about the privacy subject, mental overload, or lack of time. However, if PPs are neither read, nor understood, on what basis can users make decisions regarding their privacy? Nevertheless, users want to be informed and have an interest in their own privacy. 75.38% of the subjects of our study claimed that they care about data privacy. Other studies also confirm the interest in data privacy among a majority of users [21, 44]. This underpins users' right and need to be able to make informed and conscious decisions about their privacy. Our privacy heavily and increasingly depends on the design choices of such systems. Therefore, as software engineers, we should already make our processes transparent, accountable and, above all, equip them with the necessary fairness. Because the software development process should take ethical aspects and equality into account as well as following a human-centered approach. If we align our processes in this way, we will also be able to develop systems that are privacy-aware and respect the privacy of their users. We can also adapt and improve existing systems with processes of this type as well as also harness already existing approaches and adapt them [1, 29].

To this end, we proposed PriX, an *Online Privacy Policy eXplainer*. PriX makes PPs more accessible and comprehensible by providing *visual explanations*, helping users to improve their privacy awareness. Our study results clearly show that visual explanations are an appropriate way to foster and increase users' understanding of a PP. Our results also reveal that, if end-users are able to easily access and understand privacy information, they are willing to consume this information. This is underpinned by the fact that, at the end of the study, more than half of the subjects from the experimental group asked if we will make PriX available, since they would like to use it on a regular basis.

75.6% of participants from the experimental group claimed that privacy is important to them. Meanwhile, only 24.2% of participants occasionally or often pay attention to whether a website offers a PP. Furthermore, only 15% of participants affirmed that they, occasionally or often, actually read a PP. This may also indicate that PriX could be a facilitator in the contact with PPs , offering more accessible privacy information and thus having a positive impact on end-users' privacy awareness.

6.2 Limitations and Threats to Validity

In this section, we discuss the threats to validity according to Wohlin et al. [46].

Construct Validity: The answers of the participants heavily rely on self reporting. Some participants may have claimed to understand a PP better because they were in a test situation. The experimenter was present throughout the entire experiment to explain the task and assist with questions. This might also have had an influence on the results. We tried to address this threat by instructing the experimenter to be as unobtrusive as possible.

Conclusion Validity: We performed significance tests to help us ensure the significance of the results. Despite the fact that the overall sample size was 65 participants, some of our results may be threatened due to the small sample size. A replication of our study with a larger sample size and a more heterogeneous distribution is needed to confirm our results.

Internal Validity: The experiment results might be affected by the quality of the questions. We attempted to reduce this threat by piloting our study; also, the questions were prepared in the subjects' native language to avoid misunderstandings.

External Validity: The participant selection might have threatened the validity of our results. 80% of our participants reported being students. Although our study did not require specific background knowledge, the results might reflect the needs and perceptions of the whole population. This represents a limitation to the generalization of the results.

6.3 Future Work

Building up on the work of Ortloff et al. [32], we plan to design PriX to adapt according to the context. Our goal is to provide users more context-sensitive explanations regarding the actions they are currently performing on a website, which in turn, are linked with the provided information of a provider's PP. For instance, if a user is currently seeing a website's contact form, PriX should be able to analyze the PP with respect to the data practices and explain to the user how the entered data will be processed and used. The technological difficulty regarding explanations about privacy, in general, is to ensure that the texts continue to be legally compliant. According to [11], and to address this challenges, we need to explore how explanations can be adapted to the relevant context and how existing processes can be harnessed [1, 29]. As a further point for future research, a long-term study with users of PriX is planned. In this study, we want to examine whether PriX has an influence on users' privacy awareness and if PriX is seen as a useful tool for everyday life.

7 CONCLUSION

In a world where data is treated as a commodity, data privacy has become an increasingly important and prominent topic. Legal regulations are being established that require corporations to disclose their data-related policies to users. Such PPs are frequently lengthy and complex texts that overwhelm the user with information, prompting many users to ignore them. However, ignoring these PPs can hurt users' privacy and even lead to adverse consequences for them.

To address this issue, we created a tool (PriX) to assist users in understanding online PPs when browsing websites. PriX's goal is to provide visual explanations that seek to clarify the information included in PPs, making the information more accessible. We have analyzed our tool with respect to metrics such as: task completion times, time needed to find a specific information, how the tool supported users in understanding PPs, and the overall tool's acceptance. Therefore, we conducted a user study (sRUT and online questionnaire) with 65 participants. The results were satisfactory, since our study showed that visual explanations are an appropriate way to foster privacy awareness and can help users understand PPs. Besides that, PriX was well accepted, and more than half of the participants were interested in using the tool outside the experiment. We see this study as an important step in understanding how explanations regarding privacy can be presented and, on a broader level, how awareness about data privacy can be increased. We plan to build on this research and investigate other aspects of PPs' accessibility, as well as how explanations might be an useful strategy in this area.

ACKNOWLEDGMENTS

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy within the Cluster of Excellence PhoenixD (EXC 2122, Project ID 390833453). This work was also supported by the research initiative Mobilise between the Technical University of Braunschweig and Leibniz University Hannover, funded by the Ministry for Science and Culture of Lower Saxony.

REFERENCES

[1] Simone Agostinelli, Fabrizio Maria Maggi, Andrea Marrella, and Francesco Sapio. 2019. Achieving GDPR Compliance of BPMN Process Models. In *Information Systems Engineering in Responsible Information Systems*, Cinzia Cappiello and Marcela Ruiz (Eds.). Springer International Publishing, Cham, 10–22.
[2] Peter Bøgh Andersen. 2001. What Semiotics can and cannot do for HCI. *Knowledge-Based Systems* 14, 8 (2001), 419–424. Semiotic Approaches to User Interface Design.
[3] Gary Anthes. 2014. Data Brokers Are Watching You. *Commun. ACM* 58, 1 (Dec. 2014), 28–30.
[4] Inc. Apple Computer. 1992. *Macintosh Human Interface Guidelines.* Addison-Wesley Publishing Company, USA.
[5] Aaron Bangor, Philip T. Kortum, and James T. Miller. 2008. An Empirical Evaluation of the System Usability Scale. *International Journal of Human–Computer Interaction* 24, 6 (2008), 574–594.
[6] Aaron Bangor, Philip T. Kortum, and James T. Miller. 2009. Determining what individual SUS scores mean: adding an adjective rating scale. *Journal of Usability Studies archive* 4 (2009), 114–123.
[7] Pippin Barr, James Noble, and Robert Biddle. 2003. Icons R Icons. In *Proceedings of the Fourth Australasian User Interface Conference on User Interfaces 2003 - Volume 18* (Adelaide, Australia) *(AUIC '03)*. Australian Computer Society, Inc., AUS, 25–32.
[8] T. Boren and J. Ramey. 2000. Thinking aloud: reconciling theory and practice. *IEEE Transactions on Professional Communication* 43, 3 (2000), 261–278.
[9] John Brooke. 1996. SUS: A quick and dirty usability scale. In *Usability Evaluation In Industry (1st ed.)* (1 ed.), Patrick W. Jordan, B. Thomas, Ian Lyall McClelland, and Bernard Weerdmeester (Eds.). Taylor & Francis, London, GB, 189–194.
[10] Wasja Brunotte. 2022. https://prix.se.uni-hannover.de/ui/vp.html
[11] Wasja Brunotte, Larissa Chazette, Verena Klös, and Timo Speith. 2022. Quo Vadis, Explainability? – A Research Roadmap for Explainability Engineering. In *Requirements Engineering: Foundation for Software Quality*, Vincenzo Gervasi and Andreas Vogelsang (Eds.). Springer International Publishing, Cham, 26–32. https://doi.org/10.1007/978-3-030-98464-9_3
[12] Wasja Brunotte, Larissa Chazette, and Kai Korte. 2021. Can Explanations Support Privacy Awareness? A Research Roadmap. In *2021 IEEE First International Workshop on Requirements Engineering for Explainable Systems (RE4ES)*. 176–180.
[13] A.J. Bernheim Brush, Morgan Ames, and Janet Davis. 2004. A Comparison of Synchronous Remote and Local Usability Studies for an Expert Interface. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems* (Vienna, Austria) *(CHI EA '04)*. Association for Computing Machinery, New York, NY, USA, 1179–1182.
[14] Fred H. Cate. 2010. The Limits of Notice and Choice. *IEEE Security Privacy* 8, 2 (2010), 59–62.
[15] Kathy Charmaz. 2006. *Constructing grounded theory: A practical guide through qualitative analysis.* sage.
[16] Larissa Chazette, Wasja Brunotte, and Timo Speith. 2021. Exploring Explainability: A Definition, a Model, and a Knowledge Catalogue. In *2021 IEEE 29th International Requirements Engineering Conference (RE)*. 197–208.
[17] L.F. Cranor. 2003. P3P: making privacy policies more useful. *IEEE Security Privacy* 1, 6 (2003), 50–55.
[18] Lorrie Cranor, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, and Joseph Reagle. 2002. The platform for privacy preferences 1.0 (P3P1. 0) specification. (2002).
[19] Lorrie Faith Cranor. 2004. I Didn't buy It for Myself. In *Designing Personalized User Experiences in eCommerce* (1 ed.). Springer, 57–73.
[20] Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. 2006. User Interfaces for Privacy Agents. *ACM Trans. Comput.-Hum. Interact.* 13, 2 (June 2006), 135–178.
[21] Lorrie F Cranor, Joseph Reagle, and Mark S Ackerman. 1999. Beyond concern: Understanding Net users attitudes about online privacy. AT&T Labs-Research Technical Report TR 99.4. 3. *Retrieved April* 14 (1999), 1999.
[22] Gloria González Fuster and Raphaël Gellert. 2012. The fundamental right of data protection in the European Union: in search of an uncharted right. *International Review of Law, Computers & Technology* 26, 1 (2012), 73–82.
[23] Hamza Harkous, Kassem Fawaz, Rémi Lebret, Florian Schaub, Kang G. Shin, and Karl Aberer. 2018. Polisis: Automated Analysis and Presentation of Privacy Policies Using Deep Learning. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, Baltimore, MD, 531–548.
[24] Mohammad Rafiqul Islam. 2018. Sample size and its role in Central Limit Theorem (CLT). *Computational and Applied Mathematics Journal* 4, 1 (2018), 1–7.
[25] Rüdiger Jacob, Andreas Heinz, and Jean Philippe Décieux. 2014. *Umfrage: Einführung in die Methoden der Umfrageforschung.* Oldenbourg Wissenschaftsverlag.
[26] Patrick W Jordan. 2020. *An introduction to usability* (1 ed.). CRC Press.
[27] Scott Lederer, Jason I. Hong, Anind K. Dey, and James A. Landay. 2004. Personal privacy through understanding and action: five pitfalls for designers. *Personal and Ubiquitous Computing* 8, 6 (01 Nov 2004), 440–454.
[28] Joseph Lindley, Haider Ali Akmal, Franziska Pilling, and Paul Coulton. 2020. *Researching AI Legibility through Design.* Association for Computing Machinery, New York, NY, USA, 1–13.
[29] Raimundas Matulevičius, Jake Tom, Kaspar Kala, and Eduard Sing. 2020. A Method for Managing GDPR Compliance in Business Processes. In *Advanced Information Systems Engineering*, Nicolas Herbaut and Marcello La Rosa (Eds.). Springer International Publishing, Cham, 100–112.
[30] Aleecia M. McDonald and Lorrie Faith Cranor. 2008. The Cost of Reading Privacy Policies 2008 Privacy Year in Review. *I/S: A Journal of Law and Policy for the Information Society* 4, 3 (2009 2008), 543–568.
[31] Matthew B. Miles and A. Michael Huberman. 1994. *Qualitative Data Analysis: An Expanded Sourcebook.* SAGE Publications.
[32] Anna-Marie Ortloff, Lydia Güntner, Maximiliane Windl, Denis Feth, and Svenja Polst. 2018. Evaluation kontextueller Datenschutzerklärungen. In *Mensch und Computer 2018 - Workshopband*, Raimund Dachselt and Gerhard Weber (Eds.). Gesellschaft für Informatik e.V., Bonn.
[33] S. Camille Peres, Tri Pham, and Ronald Phillips. 2013. Validation of the System Usability Scale (SUS): SUS in the Wild. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 57, 1 (2013), 192–196.
[34] Irene Pollach. 2007. What's Wrong with Online Privacy Policies? *Commun. ACM* 50, 9 (Sept. 2007), 103–108.
[35] President's Concil of Advisors on Science and Technology. 2014. *Big data and privacy: A technological perspective.* Technical Report. Executive Office of the President. Report to the president.
[36] Joel R. Reidenberg, Travis Breaux, Lorrie Faith Cranor, Brian French, Amanda Grannis, James T. Graves, Fei Liu, Aleecia McDonald, Thomas B. Norton, and Rohan Ramanath. 2015. Disagreeable Privacy Policies: Mismatches between Meaning and Users' Understanding. *Berkeley Technology Law Journal* 30, 1 (2015), 1–88.
[37] Manuel Rudolph, Denis Feth, and Svenja Polst. 2018. Why Users Ignore Privacy Policies – A Survey and Intention Model for Explaining User Privacy Behavior. In *Human-Computer Interaction. Theories, Methods, and Human Issues*, Masaaki Kurosu (Ed.). Springer International Publishing, Cham, 587–598.
[38] Johnny Saldaña. 2021. *The Coding Manual for Qualitative Researchers.* SAGE Publications.
[39] Jeff Sauro and James R. Lewis. 2012. *Quantifying the User Experience: Practical Statistics for User Research* (1st ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 54 pages.
[40] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. 2017. Designing Effective Privacy Notices and Controls. *IEEE Internet Computing* (2017), 1–12.
[41] Bruce Schneier. 2015. *Data and Goliath: The hidden battles to collect your data and control your world.* WW Norton & Company.
[42] Nili Steinfeld. 2016. "I agree to the terms and conditions": (How) do users read privacy policies online? An eye-tracking experiment. *Computers in Human Behavior* 55 (2016), 992–1000.
[43] Seymour Sudman and Norman M Bradburn. 1982. *Asking questions: A practical guide to questionnaire design.* Jossey-Bass.
[44] Humphrey Taylor. 2003. Most people are "privacy pragmatists" who, while concerned about privacy, will sometimes trade it off for other benefits. *The Harris Poll* 17, 19 (2003), 44.
[45] Shomir Wilson, Florian Schaub, Aswarth Abhilash Dara, Frederick Liu, Sushain Cherivirala, Pedro Giovanni Leon, Mads Schaarup Andersen, Sebastian Zimmeck, Kanthashree Mysore Sathyendra, N Cameron Russell, et al. 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)*. 1330–1340.
[46] Claes Wohlin, Per Runeson, Martin Höst, Magnus C Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in software engineering.* Springer Science & Business Media.