Implementation and In Situ Assessment of Contextual Privacy Policies

Anna-Marie Ortloff
University of Regensburg
Regensburg, Germany
Anna-Marie.Ortloff@student.ur.de

Maximiliane Windl
University of Regensburg
Regensburg, Germany
Maximiliane.Windl@student.ur.de

Valentin Schwind
Frankfurt University of Applied Sciences
Frankfurt, Germany
valentin.schwind@fb2.fra-uas.de

Niels Henze
University of Regensburg
Regensburg, Germany
Niels.Henze@ur.de

ABSTRACT

Online services collect an increasing amount of data about their users. Privacy policies are currently the only common way to inform users about the kinds of data collected, stored and processed by online services. Previous work showed that users do not read and understand privacy policies, due to their length, difficult language, and often non-prominent location. Embedding privacy-relevant information directly in the context of use could help users understand the privacy implications of using online services. We implemented Contextual Privacy Policies (CPPs) as a browser extension and provide it to the community to make privacy information accessible for end-users. We evaluated CPPs through a one-week deployment and in situ questionnaires as well as pre- and post-study interviews. We found that CPPs were well received by participants. The analysis revealed that provided information should be as compact as possible, be adjusted to user groups and enable users to take action.

Author Keywords

privacy; privacy policies; online services; contextual privacy

CCS Concepts

•**Security and privacy** → **Human and societal aspects of security and privacy;** •**Human-centered computing** → *Web-based interaction; Human computer interaction (HCI);*

INTRODUCTION

Most modern online services collect data about users accessing them for various purposes, such as personalization of advertisements or content. In recent years, this created increasing concerns as the collected data can also be used for malicious actions [13, 27, 37]. Not understanding the way data is used can lead to loss of the users' trust [23] and in the case of online businesses, also to financial losses [69]. Privacy policies are the most common method to communicate information about data practices. They consist of a text which describes the accessed data and how it is processed. This is referred to as the principle of notice and choice [44, 70] or notice and consent [7]. Most users are concerned about their online privacy [67], however, they do not read privacy policies and often do not understand them [47]. Reasons for not reading privacy policies are their length, the difficult legal language in which they are written, and their location, which is often not very prominent [42, 40, 54, 47, 22]. It can be concluded that privacy policies, when they are presented as they are today, do not fulfill the purpose of informing the users.

One way to make privacy-related information accessible is displaying relevant information directly in the context of use [35, 22, 58]. This can decrease the amount of information which has to be displayed and avoids hiding information behind an inconspicuous link. A first assessment of Contextual Privacy Policies (CPPs) using a mockup of an online service suggests that users might prefer the contextual presentation of privacy information compared to the presentation as continuous text [49]. Unfortunately, CPPs are not available for current online services. Online services could implement them on their websites but this is unlikely as users' privacy awareness does not directly increase their revenue. CPPs could also be integrated into web browsers, but this is challenging as it requires determining the context where information should be displayed and to provide the information that should be displayed to users. Thus, CPPs might seem beneficial but we do not know how to implement them, if they are accepted by users, and if they actually help users understand privacy implications of their online behavior.

In this paper, we present the implementation and in situ evaluation of CPPs over one week. We implemented the CPPs as a browser extension for the most common web browsers. Similar to ad-blockers, the extension can be configured for different websites to show privacy information within the con-

DIS '20, July 6–10, 2020, Eindhoven, Netherlands.

ACM ISBN 978-1-4503-6974-9/20/07 ...$15.00.
http://dx.doi.org/10.1145/3357236.3395549

text the information is relevant. We exemplarily implemented CPPs for some of the most frequently used online services, such as Google.com, Facebook.com, and Amazon.com. Using the developed browser extension, we collected qualitative and quantitative data. We identified four main themes and provide design recommendations for the implementation of CPPs. Our results suggest that privacy policies should – most importantly – be provided within the users' context of use. They should be presented in compact and readable ways, potentially using visual metaphors, and be adaptable to different user groups. Users should get actionable advice for changing their behavior. They should not be overloaded by repeatedly showing the same information. We provide the extensions' source code to the research and open source community to enable future research and development[1].

RELATED WORK

A simple categorization of privacy distinguishes privacy of the personal sphere, which refers to the right to be left alone and information privacy, from privacy of personal data, which refers to control over one's personal data [53]. This work focuses on the second kind, enabling users to be in control of their personal data. Privacy awareness in this context means that a user has an accurate understanding of how, by whom, and to what extent their personal data is processed [53]. Today, privacy policies are virtually the only way to communicate privacy information. They consist of a text which is supposed to inform the users about the data that is collected, what happens with it, and their choices. This approach of providing privacy information and asking users for confirmation is called the principle of notice and choice [44, 70] or notice and consent [7].

While notice and choice is necessary, it is not sufficient in its current textual form [62, 14, 18]. Users do not understand privacy policies the way they are currently implemented, so they do not fulfill their aim of informing [56]. Several reasons for this are given in the literature. One is simply the length of the texts [42]. It takes a long time to read them, but users usually do not take that time. Obar found that the average time for reading privacy policies is 73 seconds [47], which is not enough to read and comprehend them. The difficulty of the legal language in which they are written is another reason for users' struggle to understand privacy policies [40]. Their wording is suitable for fulfilling legal requirements, but not for informing users. Proctor et al. found that even college students have poor comprehension of privacy policies' content [54]. This contributes to the generally abstract nature of privacy policies in their current state. It is hard for users to connect the general information given in privacy policies to their current context of using the software [22]. This is aggravated when different websites, such as Google and Youtube [16], share the same privacy policy. Based on an analysis of privacy policies, Jensen and Potts conclude that significant changes need to be made to meet regulatory and usability requirements [29].

The users' understanding of data practices is important because online consumers' trust is influenced by a feeling of security concerning private data [23]. When information on privacy is displayed clearly and concisely to users, it can become a relevant factor for decision making [64]. Some users might even be willing to pay more for offered products if their privacy is protected [64]. It should thus be desirable for companies and users alike to have understandable and accessible privacy policies. The increasing adoption of privacy by design [34, 15], user-tailored privacy by design [66] or previously proposed design guidelines [52, 59] can help assure privacy awareness. Unfortunately, it seems as if the commercial value of data makes companies refrain from embracing previous research. Therefore, previous work developed tools to supplement privacy policies and increase privacy awareness. Cranor et al. developed a user interface [19] for the Platform for Privacy Preferences (P3P) [55] which allowed websites to specify which data they collect and how the data is used. Unfortunately, P3P was suspended in 2006 as it received insufficient support from browser implementers and was not adopted by popular online services. In an iterative process, Kelley et al. developed what the authors call "Nutrition Labels" for privacy which visualize privacy policies in a compact tabular format [32, 33]. Tsai et al. developed Privacy Finder, a prototypical shopping search engine which clearly displays privacy policy information [64]. Results from a controlled experiment suggest that participants would prefer to purchase from online retailers who better protect their privacy. However, previous tools often require the user to state their preferences on privacy first [20, 64], which can pose an initial hurdle to their adoption and in the case of the privacy notices, it is not clear where and when they would be shown to users [32]. Furthermore, the service that wants to access information is the most important factor for users to decide if they want to provide data [51] which makes automated negotiation challenging.

Many problems of privacy policies as they are currently used could be avoided by displaying privacy information immediately when they become relevant [50] and directly in the user's current context of use [22, 58, 35]. It reduces the amount of text necessary to be displayed at once. Information is more accessible and concrete, because it is connected with a context. It has been shown that presenting only relevant information from privacy policies can make users more privacy-aware [8]. This is in line with Nissenbaum's theory of contextual integrity, which views context as essential to determine whether a specific action is a violation of privacy [45]. The principle of contextuality has also been applied to some extent in practice. Current versions of iOS [48] and Android [21] follow the paradigm of *ask-on-first-use*. However, this process still lacks contextuality, as the same permission may be required in different contexts of use, but only the first one is known to the users [65]. Furthermore, it was shown that users need support to engage with application permissions and to make choices which are good for their privacy [4]. Therefore, Feth proposed CPPs which show brief snippets of privacy information in the user's context of use [22]. Ortloff et al. evaluated CPPs by integrating them into a mockup of a social media website. The results suggest that users prefer the contextual presentation of privacy information compared to the representation as continuous text [49].

[1] https://github.com/Maxikilliane/CPP-browser-extension

Overall, a substantial body of work shows that the current presentation of privacy policies is not suitable to provide users with an accurate understanding of data practices. Industry provides very limited support or even prohibits efforts, such as privacy by design [34, 15], user-tailored privacy by design [66] or P3P [55], that increase users' understanding of data practices. Proposed alternative approaches to textual notice, like visceral notice [11], would also require industry support to be implemented. This may be the reason why previous attempts to provide better ways to communicate privacy policies have not been realized for popular online services. Previous work, such as interfaces for P3P [19], the "Nutrition Labels" for privacy by Kelley et al. [32], the Privacy Finder [64], or CPPs [22], could only be evaluated by showing participants mockups of actual online services. Unfortunately, as highlighted by the privacy paradox, this limits the validity of the results as users' intentions and actions concerning privacy do not always match [46, 17]. In conclusion, we do not know how to successfully implement better communication strategies for privacy policies without the support by the online services themselves. We also do not know how previous proposals, such as CPPs, are adopted by users after longer periods of time and when used for natural browsing.

SYSTEM

We developed a browser extension to provide privacy policies in the users' context of use, similar to browser-based blocking extensions that prevent online tracking [38]. Instead of blocking parts of online services, the extension embeds contextual privacy-related information in websites. Considering the dimensions as discussed by Schaub et al. [59], CPPs are just in time notices and can be compared to location permissions on smartphones, even though they do not provide control over blocking. In the following, we provide an overview of the extension's implementation and the currently supported websites.

Implementation of In Situ CPPs

Current web browsers provide an (almost) unified API [41] for in situ placement of context-sensitive content on existing websites. Using stylesheets and the execution of runtime code those extensions allow consistent rendering of content even on web pages with different designs or layouts. As websites differ in their structure, it is important to choose an adequate anchor to place the aggregated policies into the corresponding context. As there is currently no systematic approach to crawl through legal terms and to render them in an automated way, we developed a system presuming two kinds of manually maintained data: (1) A context-sensitive anchor for CPP placement on website and (2) aggregated and easy to read information extracted from the websites' privacy policy. Both requirements were implemented using a predefined "white list" of websites.

The main functionality of our system is displaying privacy-related information to raise privacy awareness in the context of use. Specific situations on visited websites trigger a CPP container in a simple and unified style close to the context. An outlined container with an icon as a title is used to communicate that the CPP belongs to the website and the current context but is provided by the extension. Thereby, the content is easily distinguishable from the design of the surrounding website. The provided information includes a short summary of the websites' privacy policies (data collected and purpose) in a readable way and directly related to the current context of use (see Table 1). We present the information using a simple language and bullet-point lists. We iteratively refined the design with three people, which were different from those in the main study, who used the CPPs for multiple days, before providing feedback, which we incorporated.

Three examples of websites using in situ CPPs are shown in Figure 1. The system was developed using the browser extension API for Firefox, Chrome, and Edge. Opera, Safari, and Internet Explorer were not supported as they either do not share the same extension API [43], make it challenging to embed content into websites, or lack support for features the other browsers provide.

Online applications like our system benefit from continuous feedback and the preparation of content from an active user community. Automated procedures (e.g. [73, 68]) could help generate policies for more platforms and potentially place them. Unfortunately, reliable procedures are currently not available. Thus, the plug-in contains an option to manually provide feedback by clicking on the icon of the extension in the browser toolbar. We also used this form (see Figure 2) to collect feedback during our system evaluation.

Supported Websites

We integrated seven commonly used websites into the extension. On these websites, CPPs are shown to users in their relevant context. We selected the websites from the most visited websites in Germany [3], to ensure that they would receive traffic from our participants during the study. However, some websites were excluded for different reasons. Websites with pornographic content, such as Livejasmin.com and Xhamster.com, were excluded because participants in our study might find it embarrassing that the researchers conducting the study know they visited these sites. Second, Russian websites, namely Mail.ru, Vk.com, and Yandex.ru, were excluded as they can only be used with knowledge of the Russian language. We excluded Instagram because the main privacy related functionality of Instagram, sharing photos and videos, is not available from a normal browser, but only from the designated apps [26]. Wikipedia was also excluded because its privacy policy did not yield scenarios in which displaying CPPs would be useful. The remaining web pages were Amazon, Facebook, eBay, eBay Kleinanzeigen (a service similar to eBay Classifieds in the United States and Gumtree in the United Kingdom), Google, Twitter, and YouTube. An overview of the implemented CPP scenarios is shown in Table 1.

The privacy policies of the seven websites were manually assessed for situations in which information from the privacy policy is relevant. Two authors read the PPs of the chosen websites separately and highlighted passages with concrete examples of collecting and using user data. We compared them, e.g. by their ability to fit into a triggered situation, and the final scenarios were selected together. We chose the triggers based on educated guesses, as to when a specific section of

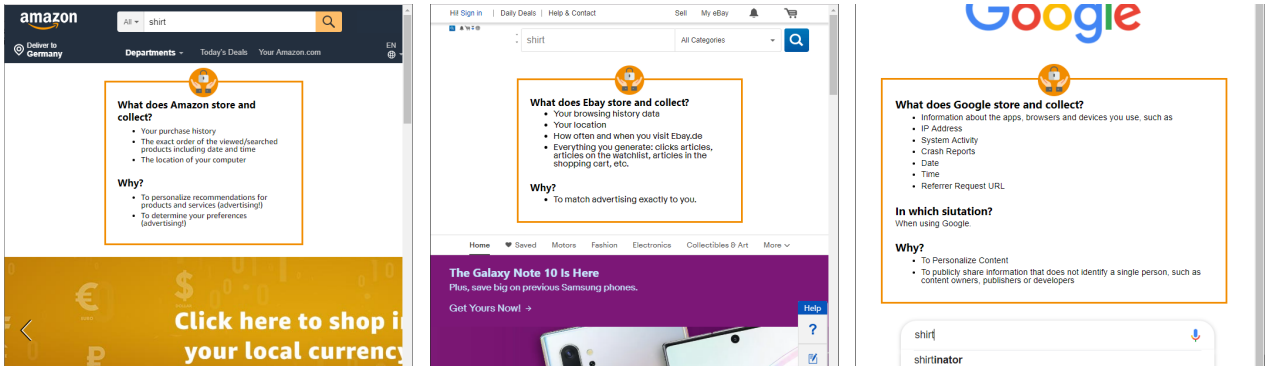

Figure 1. Three screenshots of the in situ CPPs presenting context sensitive information about the corrsponding web pages: Amazon, eBay, and Google. CPPs were provided using the context sensitive content (e.g. search policies after clicking on a search text input).

the privacy policy becomes relevant, e.g. location information in Google is displayed at the bottom of the result page, so we displayed a CPP when the user scrolled there. A total of 16 scenarios were extracted from the privacy policies. They consist of privacy information and a corresponding trigger to display the information.

METHOD

To learn how CPPs are adopted by users after longer periods of time and when used for natural browsing, we conducted an in situ study with 15 participants. The study took one week per participant. During this week, participants were asked to use their computer as usual. We used a mixed-method approach by combining in situ questionnaires, qualitative interviews, and recordings of log data.

Measures

We used a range of qualitative and quantitative measures to develop an understanding of participants' interaction with and opinion about CPPs. At the start and the end of the study, we assessed participants' general privacy awareness using the 7-point Likert item proposed by Schaub et al. [61]. We also asked them to rate how comfortable they are with how the 7 websites for which the CPPs were implemented, handle their data.

Each time a CPP was displayed, we recorded a timestamp as well as the site and which CPP was triggered. After a CPP was displayed, we asked participants to fill out a short questionnaire. The questionnaire asked for benefits of the provided privacy information, if and how this will affect their behavior, as well as to rate the relevance of the information. We ensured that the questionnaire does not appear more than once per hour to avoid fatigue effects. Participants could also provide feedback at any time by clicking extension's icon in the browser toolbar (see Figure 2).

At the end of the study, we showed images of all situations in which a CPP could have appeared, alongside with a short explanation of the action which triggers the appearance of the CPP. We asked participants to rate the relevance for their privacy, whether and how this information would influence their behavior, and for further feedback regarding the CPPs in this situation. We also conducted semi-structured interviews which took about one hour per participant.

Procedure

At the start of the study, we met with the participants to explain the purpose and the procedure of the study. We asked them for informed consent and to fill out the questionnaire about privacy awareness. We helped them install the CPP extension on their personal computers. In addition, participants received a short introduction to the CPP system and were reminded to use their computers as usual, but keeping the extension active and answering questions when prompted. Afterward, the participants used the CPP extension for one week while completing the questionnaires when prompted by the extension.

After one week, we invited participants back to our lab to collect usage data and feedback on the CPPs through a semi-structured interview. We administered the questionnaire about privacy awareness once more, to measure changes in the participants' privacy awareness due to the usage of the system. Finally, we thanked the participants and compensated them with course credits or sweets.

Participants

The study took place in Germany with participants, who were fluent German speakers, as the information in the CPPs was provided in German. The participants were recruited through our university's online forums and snowball sampling starting with the authors' acquaintances. Participants were required to use their personal computer, as well as a browser, regularly and have Chrome, Firefox or Edge as their preferred browser.

The browser extension was installed on 20 participants' personal computers in Spring 2019 and remained on the computers for at least 7 days. We chose this duration based on previous work with similar study designs [71, 36, 39, 4]. Five participants were excluded from the final analysis, for different reasons. Three persons did not use their browser often enough and were shown less than 5 CPPs throughout the study. One person accidentally deleted the data recorded by the extension when using a virus scanner and another person canceled the appointment for the final interview. Of the 15 remaining participants, 7 were female and 8 male. Their ages ranged from 19 to 38 ($M = 25.0$, $SD = 4.96$). 13 participants were

Website	Action Trigger	Data Collected	Purpose
Amazon	Typing into the product search field	• Your purchase history • The exact order of the viewed/searched products • Date and time • The location of your computer	• Personalize recommendations for products and services (advertising) • Determine your preferences (advertising)
	Purchasing a product (checkout)	• Information about your credit history	• Using scoring methods to ensure your solvency
eBay	Clicking on the log in with Facebook-button	• Facebook automatically gives access to personally identifiable information, such as content you've viewed or liked, information about the ads you've been shown or you've clicked on, and more.	• So that Facebook advertising can be better personalized to you.
	Typing in product search field	• Your browsing history data • Your location • How often and when you visit ebay.(de|com) • Everything you generate: clicks on articles, articles on the watchlist, articles in the shopping cart, etc.	• To personalize advertisement
eBay Kleinanzeigen	Typing in product search field	• Your browsing history data • Your location • How often and when you visit ebay.(de|com) • Everything you generate: clicks on articles, articles on the watchlist, articles in the shopping cart, etc.	• To personalize advertisement
Facebook	Typing into the *new post* field	• Among other things, the content that you post.	• Share data with other companies in the Facebook group, such as Instagram, WhatsApp, and Oculus • Personalize advertising • Re-use, modify, translate, or copy them • Show or display them publicly • Create derivate works
	Uploading pictures, videos, or files	• Meta information about uploaded files, such as the location of a photo and the date the file was created	• Share data with other companies in the Facebook group, such as Instagram, WhatsApp, and Oculus • Use for personalizing content • Use for personalized advertising
	Being on the event section or on a single event an signing up	• Location-related information: • Your current location if access is allowed • Your place of residence • Places you like to visit • Companies/people in your area	• Personalize advertisement
Google	Typing in the search field on the Google start page	• IP Address • System Activity • Crash Reports • Date • Time • Referrer Request URL	• Personalize content • Publicly share information that does not identify a single person with content owners, publishers or developers
	Typing a search query into Google search field	• Your search queries • Third-party sites using Google services, including a Google Captcha for example	• Match content to you • Publicly share information that does not identify a single person with content owners, publishers or developers
	Seeing a Google Ad	• Information about viewing advertising • Information about interaction with advertising • Purchases that you make	• Personalize advertisement • Publicly share information that does not lead to the personal identification of an individual, such as rights owners, publishers or developers
	Scrolling to the bottom of the Google results page	Location-specific data collected by various methods such as • GPS • Your IP address • Sensors on the device (motion sensors on mobile phones) • Information near the device like Wi-Fi access points, mobile towers or Bluetooth equipment	For personalization like an adaption of content and advertisement to you
Twitter	Being on a Twitter page	Individual personal information like • The type of device you use • Your IP address	• Personalize content such as displayed tweets • Match ads to you
	Typing something into the *new post* field	• The content of your tweet	• Publish your tweet (in the default settings, the tweet is completely public and the entire internet has access to it) • Show you more relevant content
Youtube	Playing a video	Information about activities you do on the internet such as • Videos you watch on YouTube • Interaction with content such as likes/dislikes and comments	• Personalization of content
	YouTube ad appearing	• Information about viewing advertising • Information about interaction with advertising • Purchases that you do	• Personalize ads • Publicly share information that does not identify an individual

Table 1. All CPP scenarios implemented with their action trigger and information about data collection and purpose for data collection.

students, one was employed as a software developer and another one as an editor. Eight participants used a computer with Microsoft Windows, five used MacOS and one participant used Linux. Seven participants used Mozilla's Firefox and eight participants used Google Chrome.

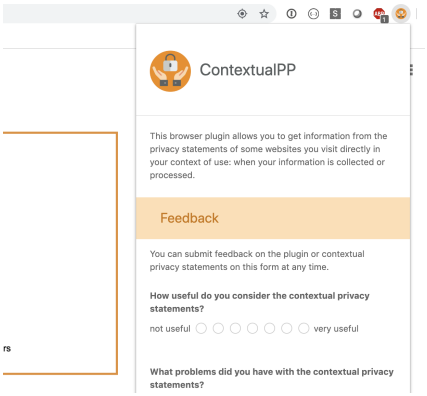

Figure 2. Form to provide feedback about the extension we also used in the in situ study.

Data Analysis

We used thematic analysis, as described by Braun and Clarke [10], to make sense of the qualitative data from the interviews and the questionnaires collected throughout the in situ deployment. Braun and Clarke describe a series of phases which make up the process of thematic analysis [10]. We apply a realist approach [10], and capture and report the experience of our participants with the CPPs.

The in situ questionnaires were already in text form. The interviews were transcribed using InqScribe[2]. Dialect was resolved to standard language, except in some cases, when it was hard to find an equivalent. Interjections, such as "ähm", were left out in most cases, except when omitting them would make the transcript harder to understand. Some nonverbal utterances, such as laughter, sighs, and pauses were transcribed as well. We also transcribed the interviewer in case of deviations from the protocol.

To generate the initial codes [10], two interviews from the sample were randomly selected. The two interviews were separately coded by two authors using the *RQDA*-Package [57] and the freely available software R. Following this process of open coding, the coded interviews were compared and differences in coding were resolved through discussion. The codebooks were merged and the refined codebook was used to code the remaining interviews. We wrote the codes on slips of paper to find themes in the data and formed categories of related codes. During this process, we repeatedly reviewed the themes by looking at the extracts of coded data and comparing them to extracts from other themes. We identified relationships and connections between themes and gave each theme a representative name. During this process, two codes (*miscellaneous* and *feedback_on_study*) were excluded from further analysis, because they did not add value to our analysis.

The focus of this study was to gain explorative insights in CPP's in situ usage, however to bestow credibility on our findings, we conducted an exploratory quantitative analysis using R. In this analysis, we considered p-values of below .05 to be significant. Because the data was not normally distributed, we utilized Wilcoxon signed-rank tests to compare the data of pre- and post-study levels of privacy concern. Additionally, we calculated descriptive statistics for usage data of the CPPs, and the relevance ratings of the CPP information given in different situations displayed to the participants during the post-study interviews.

[2] https://www.inqscribe.com/

RESULTS

Through the thematic analysis, we identified four main themes in the data, which are essential for understanding how participants perceived CPPs and the information that they contained. Examples of codes and quotes connected to these themes are shown in Table 2. Participants' quotes have been translated from German into English. In our quantitative analysis, we highlight links to the results of the thematic analysis.

Qualitative results

We identified the four main themes: *conceptions about data collection*, *handling of personal information*, *feedback on CPPs* and *considerations for future use of CPPs*, see Table 2. The former of these two themes are both linked to the concept of handling of data, while the latter two themes both relate to the CPPs. Naturally, several aspects of the participants' impressions are connected to more than one theme. This should be considered as inherent to the data because a single statement of a participant cannot be viewed as an isolated concept, as it is connected to this participant's other utterances.

Conceptions about Data Collection

This theme consists of knowledge and opinions the participants had about what happens to the data they generate by using certain websites. Participants have different opinions about the data which is collected about them and how this data is used. In some cases open disapproval of the way data is handled by the websites was voiced, such as "that's so bad that they always just want to tailor advertising to you." (P11). In other cases, the participants were okay with it: "[...] that advertisements are matched to me that, [Pause] I don't know, somehow you also benefit from it." (P16). It became apparent that participants opinions on the severity of the collected information differ. These opinions are formed based on different aspects of information. In some cases, the participants focused on what is collected, "something like system activity, crash reports that's none of their business. For what do they need that?" (P2). In other situations it is more important why data is collected, "as long as that is only being used to tailor advertisement to me, it's okay for me that data is collected." (P14).

There are several reasons why participants approved that data is collected. One is when they perceive it as the company's right to get that data, e.g. because they are interacting with the company's site, "yes, I use their website, I look at things, of course, they track it. [...] In my opinion, they are free to collect usage statistics so to speak, of their things." (P2). Additionally, participants approved data collection, when they see a benefit for themselves: "Date and time can be helpful when searching for the right result. They support me." (P2). In addition to these clear cut opinions, participants speculated about how the data collected by websites is used, imagining scenarios beyond the information given in the CPPs: "I think

Theme	Subtheme	Example Statements
Conceptions about data collection	Negative conception	*"That this whole collection of data only serves to manipulate me that I buy more stuff and that's against my interest of course." [P10]*
	Neutral conception	*"I think it can happen for example [...] that one's image is reused, in some kind of advertisement for example." [P10]*
	Positive conception	*"Date and time can be helpful when searching for the right result, they support me there." [P2]*
Handling of personal information	Behavior change	*"I plan on using DuckDuckGo." [P2]*
	No behavior change	*"No I don't think, I would change my behavior. I use Google, because it works best for me. Also YouTube. There are no real alternatives." [P5]*
	Current handling of data	*"Because I already use an adblocker, I almost receive no advertisement at all." [P20]*
Feedback on text and CPPs	General feedback	*"It was pleasant. It wasn't really disturbing or anything." [P2]*
	Feedback on text	*"So, for me it is clear, what is meant by the referrer url. But not everyone knows that, it's a bit too cryptic. I think, if it is planned for the masses... my mother doesn't know what the referrer url is." [P4]*
	Feedback on look	*"Besides, the layout is very pretty, also the icon is pretty and the choice of color is fitting." [P9]*
Considerations for future use of CPPs	Novelty of information	*"After I saw the same privacy notice at Google for the 7th time, it has only a limited value for me." [P10]*
	Feature requests	*"Maybe that you put the information into categories and mark them by color. Just like a traffic light." [P3]*

Table 2. Four themes identified by thematic analysis, with examples of associated codes and participants' statements

it can happen for example [...] that one's image is reused, in some kind of advertisement for example" (P10).

Handling of Personal Information

Participants exhibit different kinds of behavior in reaction to being exposed to information related to data practices. While the previous theme focused on participants' opinions about privacy, this theme focuses on planned or performed actions. Efforts are made to avoid the collection of data, for example by using a different search engine which collects less data, "I would probably still change to a less data-hogging search engine" (P9), with some being even more concrete, "I plan to use DuckDuckGo" (P2). This intended behavioral change is due to the disapproval of data collection practices.

Conversely, the reasons why participants do not change or do not want to change their data-related behavior vary. Sometimes participants do not see a necessity to change their behavior, because they endorse the websites' data handling practices, "in this context it makes sense to save these data" (P2), or at least they do not care about them, "My God, so what if they saw what I googled" (P16). This mirrors the relationship between the disapproval of data practices and behavioral changes. In other cases, there was simply a lack of knowledge on how to prevent data collection: "With the other things I don't really know how I can prevent it." (P11). Other participants weighed their options, but the ease, the comfort or other features of the web service they were currently using, outweighed their privacy concerns: "You won't get around Amazon, because there are other criteria, like payment comfort or ordering being too convenient that you would look for something else." (P14). Finally, the lack of alternatives can be the reason why participants keep on using web services, even though they do not completely agree with the way their data is handled, "no I don't think, I would change my behavior. I use Google because it works best for me. [...] There are no real alternatives." (P5).

When thinking about their personal behavior concerning privacy-related information, participants also described their difficulties to understand privacy policies the way they are currently used. In many ways, participants' thoughts about their interaction with privacy policies mirrored the literature. Participants criticized their length, "[...] it [the information] is normally wrapped in as much text as possible, so that it's too bothersome for the user to deal with it" (P2), legal language, "they are normally in some [pause] legal language, which is, I'd say not totally super difficult, but you do have to deal with it for some time" (P10), and lack of context and concreteness, "very very long rambling, from which you can't discern what's actually happening." (P1). In contrast to this common means of accessing knowledge on data processing, participants recognized the merit of CPPs, which leads to the next theme.

Feedback on CPPs

Participants considered it beneficial to see the CPPs directly in their context of use, "That's really well-positioned because it's near the relevant place." (P10). In contrast to current presentations of privacy policies, participants were also content with the brevity of the CPPs: "[...] it was good that it was very concise. And it was little to read through." (P14). In general, participants were not disrupted much from their daily workflow, "it was pleasant. It wasn't really disruptive or anything" (P2).

Participants also commented on the visual features of the CPPs. One participant specifically criticized the CPPs' positioning on Amazon: "Because on Amazon, you first have to enter something into the search field, then this thing pops up. But you don't really see it, because Amazon makes search sugges-

tions which cover the information." (P4). This participant also suggested placing the CPPs consistently, "I generally would prefer that this information always appears in the same place, so that I can prepare myself for it, so that I don't have the feeling that this is the 50,000th pop-up and I think: What's going on here? So that I know: This is the plugin I've installed previously and it's okay like that." (P4).

Feedback concerning the layout of the CPPs was diverse. Generally, participants positively commented on the design, "besides, the layout is very pretty, also the icon is pretty and the choice of color is fitting." (P9) and "the layout is great. Very clear." (P5). However, the feedback on the size of the CPP was ambiguous. While one participant proposed to reduce the size CPPs, "maybe you could have made this a bit more compact. A bit smaller." (P9), another participant suggested to increase the size: "maybe that you stretch it a bit. So that it is as wide as the screen. And maybe a bit higher." (P7).

Regarding the phrasing of the CPPs' text, the participants appreciated the use of bullet points: "I really liked that, this threefold division with the bullet points." (P2) and "I like the phrasing. Also that there are always bullet points." (P11). Sometimes the phrasing of the CPPs caused difficulties to understand the information. "In this case, I would like a bit of an explanation, [...] because I don't fully understand what is meant. And also what are scoring procedures [referring to Amazon's scoring procedure while checking out a product]. That could be explained as well." (P11). Another participant mentioned that even though they understand what is meant, other people might have problems to understand the information: "So, for me it is clear, what is meant by the referrer url. But not everyone knows that, it's a bit too cryptic. [...] my mother doesn't know what the referrer url is." (P4).

Considerations for Future Use of CPPs

While the previous theme focused on concrete feedback on the CPPs, participants also provided ideas to further improve the use of CPPs. This theme can also be considered as a collection of possibilities for the interaction with privacy-related information. The interviews and the in situ questionnaire revealed that participants became less interested when the same information was provided repeatedly. The qualitative analysis suggests that the information presented through CPPs was perceived less useful or relevant over time, "after I saw the same privacy notice at Google for the 7th time, it has only a limited value for me" (P10). This was, however, only limitedly supported by our quantitative analysis which showed that the aggregated rated relevance remained almost constant on a high level.

Participants proposed multiple approaches to avoid overexposure with privacy-related information. One participant highlighted the novelty of information as a factor which contributes to the usefulness, "If it was adjusted a bit, I'd consider it a really useful tool. Especially when I visit sites, which I didn't use before." (P2). Most suggestions include methods to avoid showing information all the time but, for example, only once a week (P19) or with certain time intervals between the presentation (P11). Participants also proposed manual hiding, "I would do it like that that you see the information the first time you visit the website. And then maybe on the second day or after some time. And then maybe not at all. So that there is somehow a possibility to say, 'Don't show again!'" (P10).

Further suggestions include adjusting or customizing information for the person using the extension, "and it would be quite nice if it would adjust itself. I mean for example, if I disabled my location information that this information doesn't appear every time." (P2). As seen in the *Handling of Personal Information* theme, participants stated that they would indeed like to change their behavior, after being confronted with information on data practices of websites, but they do not know how to do so. To overcome this, participants suggested integrating tips or advice on how data collection can be limited, "I would like it, if there were tips on how to prevent this. That the location is tracked for example. Or that, when browser history is being tracked that there is a tip on how to prevent it." (P11). A further suggestion to augment the content of CPPs was connected to trust issues concerning the information given in the CPPs: "If there would have been citations from the privacy policies... So that it is proven." (P10).

Participants also proposed approaches to make the CPPs even easier to understand, such as icons or colors, thus further reducing the amount of text: "Maybe that you use symbols. Because it's always text. For example that you do it with icons, so that you have a quick overview of what is being collected." (P3). This participant also suggested using a color system similar to traffic lights: "Maybe that you put the information into categories and mark them by color. Just like a traffic light. [...] So that you capture it faster and don't always have to read the whole thing."

It was also emphasized that the CPPs' usefulness increases the more sites are supported: "I would definitely prefer it to work on every site." (P7) With more sites being supported and accessible through CPPs, it becomes possible to further support users in making informed decisions and choices, e.g. by directly comparing websites by their privacy policies: "it would be handy if there was a way to compare them." (P19).

Quantitative results

To augment our findings from the qualitative analysis, we first examined the usage data collected through the extension. The number of CPPs displayed to participants ranged between 11 and 332 ($M = 80.7$, $SD = 86.5$). Users saw CPPs on all websites we implemented them for, except Twitter. Even though some of the participants were Twitter users, they did not use their laptop or computer to access the site during the duration of the study. However, the frequency of CPPs displayed on websites differed by a large margin. On four of the websites (Facebook, Amazon, eBay and eBay Kleinanzeigen) they were each displayed 30 times or less, while on Google and YouTube, CPPs appeared more than 400 times. This shows clearly that these two websites received the most traffic from our participants.

A Wilcoxon signed-rank test was conducted to examine whether privacy concern differs before and after being exposed to CPPs. There was no significant difference in privacy concern before ($M = 4.00$, $SD = 1.31$) and after the study ($M = 4.27$, $SD = 1.16$), $V = 21$, $p = .28$. When examining

the data by participant, we found changes in both directions. For 2 participants, the exposure to CPPs made them less concerned about privacy, while 6 became more concerned. Quantitatively, the privacy concern of 7 participants did not change during the week. Means (not illustrated) show a slight trend that having seen more CPPs during the week makes it more likely for a change in privacy concern to occur. In general, the exposure to CPPs may have an effect on individual privacy concerns, but this effect differed from one participant to another. This is further supported by the qualitative data, where even though the reasons for approval or disapproval of data collection were similar across participants, their application to specific situations differed.

We also examined the perceived relevance of the information given in the CPPs. This was measured for each of the 16 different situations presented to the participants during the post-study interview, as well as in the in situ questionnaires collected during their interaction with the prototype. We explored measures of spread and central tendency for different variables.

Comparing the relevance between the different sites, the difference between the ratings is small, with relevance ratings ranging from 4.23 for Twitter to 5.55 for Google ($M = 5.12$, $SD = 0.46$). Taking the unique situations into account ($M = 5.18$, $SD = 0.53$), larger average differences arise, with the minimum average relevance being 3.87 for the CPP displayed when writing a tweet on Twitter and the maximum average relevance 5.93 for typing a query on the Google start page. The largest difference between average relevance becomes apparent when comparing the participants' individual ratings which range from 2.94 to 7.00 ($M = 5.18$, $SD = 1.07$). We conclude that the relevance of CPP-related information is highly subjective and depends more on the individual than on the site or the exact information, which is supported by our qualitative data.

Finally, we examined how participants' relevance ratings of the CPPs developed over time. In the interviews, participants often mentioned the usefulness of the information declining, depending on how often they had already seen it. However, when aggregating data by calculating an average rating on each day after the beginning of the study the means (not illustrated) do not indicate a clear tendency. As not every participant filled out a diary on every day of the study, and average relevance ratings differed substantially for the different participants, the aggregated data may not allow visualizing the complete trend. We found several instances in the data, where a specific CPP was repeatedly shown to a participant and relevance ratings decreased over time. However, often only two or three repeated measurements of the same CPP for the same participant were available, and so the trend is not quite as obvious.

DISCUSSION AND LIMITATIONS

In general, the CPPs received positive feedback from the participants who used them in their natural environment. They especially appreciated the CPPs presentation within their context of use and also positively highlighted the CPPs concise nature and clear structure. While some participants commented on the repeated exposure to CPPs, the aggregated rated relevance remained almost constant throughout the study.

Participants provided a large number of constructive comments to further improve the CPPs' implementation. They appreciated the condensed presentation in lists of bullet points. Other visualization techniques, such as the traffic light metaphor proposed by P3, have also been discussed and evaluated in a nutritional context [6]. This ties in with the work of Kelley et al. [32], whose privacy labels were also designed similarly to nutrition labels. Participants provided suggestions to reduce habituation effects that could occur when users are frequently exposed to the same CPPs. Future work could compare these approaches to find the best combination to keep the user aware, but not overwhelmed with privacy information. There was also feedback on how the phrasing of the CPPs could be improved.

While participants generally appreciated the current length and granularity of the CPPs, users who are especially concerned for their privacy might want to receive further details that could be provided through links to more detailed information, Furthermore, it might be necessary to adapt the use of technical terms to the needs of the person using the extension.

The results could be limited by the study's duration, the limited effects on behavior change, the recruited sample, and the information provided by the browser extension. The participants experienced CPPs in situ, during their normal browsing, but only for one week. Nonetheless, we assume that results are transferable beyond one week as the aggregated relevance ratings remained almost stable throughout the study. Furthermore, participants could also imagine continuing to use CPPs after finishing the study and provided important directions for further improvement. Participants expressed the intent to change their behavior but only a few implemented the change. Thus, the results could still be affected by the privacy paradox [46]. Participants may act differently than their stated intentions. At least, we found anecdotal evidence of participants who changed their online behavior during or after study. P20 reported changing the behavior during the study in the post-study interview by switching the search engine, and P1 reported a behavior change after the study. More specifically, P1 started to use a dedicated browser just for eBay. Furthermore, CPPs are not designed to necessarily change users' behavior but to increase their understanding of websites' data practices.

We recruited a specific sample, which mainly consisted of participants with a technical background. They might have more background knowledge about online privacy and might also be more interested in it. As technically adept users are more likely to be early adopters of new technology and more likely to install a browser plugin, we considered this group a good target group for the study. As even these users had problems deciding how they could change their behavior to influence the collection of their data, we assume that less technical users are even more in need of advice on how to take action. Additionally, previous work in the domain of usable privacy suggests that students and older adults do not differ enough to generally assume that findings based on sampling from a student population are not generalizable [63]. Finally, the extension provided specific information for 16 scenarios. As stated by participants, CPPs for more websites should be provided in the future. While we derived the information and

the relevant context from websites privacy policies, the provided information could be inaccurate. We believe that the number of scenarios and potentially inaccurate information are not necessarily limitations. CPPs do not aim to replace traditional privacy policies, but even augmenting most websites with CPPs will require significant effort. The study shows that supporting some websites already is beneficial, which is encouraging. There will always be CPPs which are inaccurate just as ad-blockers will never block all ads and occasionally block content. We assume that providing inaccurate information will encourage companies to contribute CPPs for the extension.

DESIGN SUGGESTIONS

The analysis allows us to derive a number of implications that are important for the further development of CPPs and similar systems.

Novelty is important for the relevance of privacy policies. The analysis revealed that participants perceived the privacy information less relevant when they had already seen it often. This trend can be found in the quantitative data and more prominently in the statements in the post-study interviews. As also discussed by Anderson et al., attention towards privacy policies can diminish through frequent exposure [5]. Therefore, developers have to find ways to overcome this challenge. Future implementations should enable users to select showing CPPs only the first time they are relevant or when they change. Furthermore, it should be possible to disable CPPs for selected sites and to only show them on request.

CPPs should be quickly and easily comprehensible. While participants stated that they liked the presentation through bullet points, they also suggested approaches that do not always require reading to understand the core of the information. Suggestions included incorporating a color system as it is used in traffic lights, where red stands for very sensitive data, yellow for somewhat sensitive data and green for not sensitive data. Participants also suggested using icons for conveying information. Developers should take this into account and use icons or color schemes to make the information easier to comprehend for certain user groups.

The used phrasing should be adjustable to the user. The qualitative feedback clearly shows that not only the amount of information but also how the information is conveyed depends on the person reading it. As phrased by Schaub et al., there should be different notices for different audiences [59]. While some users who are especially concerned about privacy want extensive details provided, e.g. through a link to a detailed explanation, some users appreciate the current length of the CPPs. The use of technical terms also needs to be adaptable to different user groups. Developers could include a setting to make it possible to adjust the phrasing to the users' needs. Different texts could be implemented similar to adapting software to different languages.

Users should be supported in changing their behavior. Several participants stated that they were quite shocked about the kinds and amount of data being collected and wanted to change their behavior. Nevertheless, they often did not know how to accomplish the change, because they did not know about other web pages or how to adjust their settings. In line with the context-adaptive privacy model proposed by Schaub et al., CPPs have to go beyond providing awareness to also enable users to make decisions and stay in control [60]. Consequently, CPPs should explain to users how to configure the online services they use and should also provide alternative services on request.

CPPs should be adaptable to users' concerns. Some participants saw it as a fair price to pay with their data for an otherwise free service. It was even considered beneficial by some to receive tailored recommendations due to the collected data. This is similar to previous work [72, 31]. In other cases participants did not see an alternative or enjoyed the comfort the service is offering too much, to be wanting to switch to a different service. While for some participants information about their location was most privacy relevant, other participants enjoyed the comfort of a web page automatically adjusting to their current location. Developers should, therefore, enable incorporating users' preferences when implementing color schemes or icons representing different privacy levels.

CONCLUSION

In this paper, we presented contextual privacy policies that embed privacy policies directly in the context where they are relevant. We implemented CPPs as a browser extension and exemplarily integrated support for 7 commonly used websites and 16 scenarios. In an in situ study, participants used the extension for one week. Through qualitative analysis we derived four themes that are either linked to handling of data or CPPs. The analysis revealed that participants prefer CPPs over regular privacy policies. They highlighted CPPs' shortness, clear structure, and their comprehensibility. Participants also provided suggestions for further improvement which we synthesized to five design recommendation. We release the CPP browser extension as open-source to foster future research and development.

Our system relies on the manual integration of privacy notices. This is surprisingly similar to other privacy-related extensions. Ad- and tracking-blockers also partially rely on manual integration of new ad-services or tracking approaches. Nonetheless, future work should continue looking into semi-automated [73, 68] and automated approaches [25] to make CPPs available for more websites. Future work should also investigate how to support devices which raise even more privacy challenges. Mobile devices have limited screen space to provide privacy-related information [24], IoT devices might not have screens [72, 28], and pervasive environments extend beyond a single device [12, 30, 59]. CPPs might already be helpful for users who are visually impaired due to the much shorter and more comprehensible presentation. Supporting additional devices and modalities could also benefit users with special needs who are especially vulnerable when it comes to privacy [1, 2]. Finally, we focused on the individual's privacy. In a networked world, communicating privacy implications becomes even more challenging and novel communication strategies must be the focus of future work [9].

REFERENCES

[1] Tousif Ahmed, Roberto Hoyle, Kay Connelly, David Crandall, and Apu Kapadia. 2015. Privacy Concerns and Behaviors of People with Visual Impairments. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 3523–3532. DOI: http://dx.doi.org/10.1145/2702123.2702334

[2] Tousif Ahmed, Roberto Hoyle, Patrick Shaffer, Kay Connelly, David Crandall, and Apu Kapadia. 2017. Understanding Physical Safety, Security, and Privacy Concerns of People with Visual Impairments. *IEEE Internet Computing* (2017), 1–1. DOI: http://dx.doi.org/10.1109/MIC.2017.265103316

[3] Alexa Internet. 2019. Top Sites in Germany. https://www.alexa.com/topsites/countries/DE. (2019). Last accessed: 2019-05-30.

[4] Hazim Almuhimedi, Florian Schaub, Norman Sadeh, Idris Adjerid, Alessandro Acquisti, Joshua Gluck, Lorrie Faith Cranor, and Yuvraj Agarwal. 2015. Your Location Has Been Shared 5,398 Times! A Field Study on Mobile App Privacy Nudging. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 787–796. DOI:http://dx.doi.org/10.1145/2702123.2702210

[5] Bonnie Anderson, Tony Vance, Brock Kirwan, David Eargle, and Seth Howard. 2014. Users Aren't (Necessarily) Lazy: Using NeuroIS to Explain Habituation to Security Warnings. In *Proceedings of the 35th International Conference on Information Systems*. http://aisel.aisnet.org/icis2014/proceedings/ISSecurity/13

[6] Kelvin Balcombe, Iain Fraser, and Salvatore Di Falco. 2010. Traffic lights and food choice: A choice experiment examining the relationship between nutritional food labels and price. *Food Policy* 35, 3 (2010), 211 – 220. DOI:http://dx.doi.org/https://doi.org/10.1016/j.foodpol.2009.12.005

[7] Solon Barocas and Helen Nissenbaum. 2009. On notice: The trouble with Notice and Consent. In *Proceedings of the engaging data forum: The first international forum on the application and management of personal electronic information*. 7. https://ssrn.com/abstract=2567409

[8] Mike Bergmann. 2009. Testing Privacy Awareness. In *The Future of Identity in the Information Society*, Vashek Matyáš, Simone Fischer-Hübner, Daniel Cvrček, and Petr Švenda (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 237–253.

[9] Danah Boyd. 2012. Networked privacy. *Surveillance & Society* 10, 3/4 (2012), 348.

[10] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. DOI: http://dx.doi.org/10.1191/1478088706qp063oa

[11] Ryan Calo. 2011. Against notice skepticism in privacy (and elsewhere). *Notre Dame Law Review* 87 (2011), 1027–1072.

[12] Jean Camp and Kay Connelly. 2007. Beyond consent: privacy in ubiquitous computing (Ubicomp). In *Digital Privacy – Theory, Technologies, and Practices* (1 ed.), Alessandro Acquisti, Stefanos Gritzalis, Costos Lambrinoudakis, and Sabrina di Vimercati (Eds.). Taylor and Francis, New York, Chapter 16. DOI:http://dx.doi.org/https://doi.org/10.1201/9781420052183

[13] John Edward Campbell and Matt Carlson. 2002. Panopticon.com: Online Surveillance and the Commodification of Privacy. *Journal of Broadcasting & Electronic Media* 46, 4 (2002), 586–606. DOI: http://dx.doi.org/10.1207/s15506878jobem4604_6

[14] F. H. Cate. 2010. The Limits of Notice and Choice. *IEEE Security Privacy* 8, 2 (March 2010), 59–62. DOI: http://dx.doi.org/10.1109/MSP.2010.84

[15] Ann Cavoukian. 2012. Privacy by Design. *IEEE Technology and Society Magazine* 31, 4 (2012), 18–19.

[16] Sushain K. Cherivirala. 2018. UsablePrivacy.org Explore Google.com. https://explore.usableprivacy.org/youtube.com/?view=machine. (2018). Last accessed: 2019-06-02.

[17] Kay Connelly, Ashraf Khalil, and Yong Liu. 2007. Do I Do What I Say?: Observed Versus Stated Privacy Preferences. In *Human-Computer Interaction – INTERACT 2007*, Cécilia Baranauskas, Philippe Palanque, Julio Abascal, and Simone Diniz Junqueira Barbosa (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 620–623.

[18] Lorrie Faith Cranor. 2012. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *Journal on Telecommunications and High Technology Law* 10 (2012), 273.

[19] Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. 2006a. User Interfaces for Privacy Agents. *ACM Trans. Comput.-Hum. Interact.* 13, 2 (June 2006), 135–178. DOI: http://dx.doi.org/10.1145/1165734.1165735

[20] Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. 2006b. User Interfaces for Privacy Agents. *ACM Transactions on Computer-Human Interaction (TOCHI)* 13, 2 (June 2006), 135–178. DOI: http://dx.doi.org/10.1145/1165734.1165735

[21] Android Developers. n.d. Request App Permissions. https://developer.android.com/training/permissions/requesting.html. (n.d.). Last accessed: 2019-05-10.

[22] Denis Feth. 2017. Transparency through Contextual Privacy Statements. In *Mensch und Computer 2017-Workshopband*, Manuel Burghardt, Raphael Wimmer, Christian Wolff, and Christa Womser-Hacker (Eds.). DOI: http://dx.doi.org/10.18420/muc2017-ws05-0406

[23] Carlos Flavián and Miguel Guinalíu. 2006. Consumer trust, perceived security and privacy policy. *Industrial Management and Data Systems* 106, 5 (2006), 601–620. DOI:http://dx.doi.org/10.1108/02635570610666403

[24] Martin Gisch, Alexander De Luca, and Markus Blanchebarbe. 2007. The Privacy Badge: A Privacy-awareness User Interface for Small Devices. In *Proceedings of the 4th International Conference on Mobile Technology, Applications, and Systems and the 1st International Symposium on Computer Human Interaction in Mobile Technology (Mobility '07)*. ACM, New York, NY, USA, 583–586. DOI: http://dx.doi.org/10.1145/1378063.1378159

[25] 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. https://www.usenix.org/conference/usenixsecurity18/presentation/harkous

[26] Instagram Inc. n.d. Help Center - How do I post a photo? https://help.instagram.com/442418472487929. (n.d.). Last accessed: 2019-06-02.

[27] Jim Isaak and Mina J. Hanna. 2018. User Data Privacy: Facebook, Cambridge Analytica, and Privacy Protection. *Computer* 51, 08 (aug 2018), 56–59. DOI: http://dx.doi.org/10.1109/MC.2018.3191268

[28] Timo Jakobi, Sameer Patil, Dave Randall, Gunnar Stevens, and Volker Wulf. 2019. It Is About What They Could Do with the Data: A User Perspective on Privacy in Smart Metering. *ACM Trans. Comput.-Hum. Interact.* 26, 1, Article 2 (Jan. 2019), 44 pages. DOI: http://dx.doi.org/10.1145/3281444

[29] Carlos Jensen and Colin Potts. 2004. Privacy Policies As Decision-making Tools: An Evaluation of Online Privacy Notices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 471–478. DOI: http://dx.doi.org/10.1145/985692.985752

[30] Apu Kapadia, Tristan Henderson, Jeffrey J. Fielding, and David Kotz. 2007. Virtual Walls: Protecting Digital Privacy in Pervasive Environments. In *Pervasive Computing*, Anthony LaMarca, Marc Langheinrich, and Khai N. Truong (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 162–179.

[31] David J. Kaufman, Juli Murphy-Bollinger, Joan Scott, and Kathy L. Hudson. 2009. Public Opinion about the Importance of Privacy in Biobank Research. *The American Journal of Human Genetics* 85, 5 (2009), 643 – 654. DOI:http://dx.doi.org/https://doi.org/10.1016/j.ajhg.2009.10.002

[32] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. 2009. A "Nutrition Label" for Privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security (SOUPS '09)*. ACM, New York, NY, USA, Article 4, 12 pages. DOI: http://dx.doi.org/10.1145/1572532.1572538

[33] Patrick Gage Kelley, Lucian Cesca, Joanna Bresee, and Lorrie Faith Cranor. 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, 1573–1582. DOI: http://dx.doi.org/10.1145/1753326.1753561

[34] Marc Langheinrich. 2001. Privacy by Design — Principles of Privacy-Aware Ubiquitous Systems. In *Ubicomp 2001: Ubiquitous Computing*, Gregory D. Abowd, Barry Brumitt, and Steven Shafer (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 273–291.

[35] 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 (2004), 440–454.

[36] Shaun Alexander Macdonald and Stephen Brewster. 2019. Gamification of a To-Do List with Emotional Reinforcement. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19)*. Association for Computing Machinery, New York, NY, USA, Article Paper LBW1621, 6 pages. DOI:http://dx.doi.org/10.1145/3290607.3313060

[37] Elena Maris, Timothy Libert, and Jennifer Henrichsen. 2019. Tracking sex: The implications of widespread sexual data leakage and tracking on porn websites. (2019). https://arxiv.org/abs/1907.06520

[38] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. 2018. Characterizing the Use of Browser-Based Blocking Extensions To Prevent Online Tracking. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. USENIX Association, Baltimore, MD, 103–116. https://www.usenix.org/conference/soups2018/presentation/mathur

[39] Matthew Louis Mauriello, Brenna McNally, and Jon E. Froehlich. 2019. Thermporal: An Easy-To-Deploy Temporal Thermographic Sensor System to Support Residential Energy Audits. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, Article Paper 113, 14 pages. DOI:http://dx.doi.org/10.1145/3290605.3300343

[40] Aleecia M McDonald and Lorrie Faith Cranor. 2008. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society* 4 (2008), 543–568. http://hdl.handle.net/1811/72839

[41] MDN Web Docs. 2019. Browser Extensions. (Jun 2019). https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions

[42] George R. Milne, Mary J. Culnan, and Henry Greene. 2006. A Longitudinal Assessment of Online Privacy Notice Readability. *Journal of Public Policy & Marketing* 25, 2 (2006), 238–249. DOI: http://dx.doi.org/10.1509/jppm.25.2.238

[43] Mozilla. 2019. Building a cross-browser extension. https://developer.mozilla.org/en-US/docs/Mozilla/Add-ons/WebExtensions/Build_a_cross_browser_extension. (2019). Last accessed: 2019-05-30.

[44] Kanthashree Mysore Sathyendra, Shomir Wilson, Florian Schaub, Sebastian Zimmeck, and Norman Sadeh. 2017. Identifying the Provision of Choices in Privacy Policy Text. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, 2774–2779. DOI: http://dx.doi.org/10.18653/v1/D17-1294

[45] Helen Nissenbaum. 2004. Privacy as contextual integrity. *Washington Law Review* 79 (2004), 119–158.

[46] Patricia A. Norberg, Dan Horne, and David Horne. 2007. The Privacy Paradox: Personal Information Disclosure Intentions Versus Behaviors. *Journal of Consumer Affairs* 41, 1 (2007), 100 – 126. DOI: http://dx.doi.org/10.1111/j.1745-6606.2006.00070.x

[47] Jonathan A. Obar. 2016. The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services. *SSRN Electronic Journal* (2016). DOI: http://dx.doi.org/10.2139/ssrn.2757465

[48] Jason D. O'Grady. 2014. New privacy enhancements coming to iOS 8 in the fall. https://www.zdnet.com/article/new-privacy-enhancements-coming-to-ios-8-in-the-fall/. (2014). Last accessed: 2019-05-10.

[49] 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.

[50] Sameer Patil, Roberto Hoyle, Roman Schlegel, Apu Kapadia, and Adam J. Lee. 2015. Interrupt Now or Inform Later?: Comparing Immediate and Delayed Privacy Feedback. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1415–1418. DOI: http://dx.doi.org/10.1145/2702123.2702165

[51] Sameer Patil, Yann Le Gall, Adam J. Lee, and Apu Kapadia. 2012. My Privacy Policy: Exploring End-user Specification of Free-form Location Access Rules. In *Financial Cryptography and Data Security*, Jim Blyth, Sven Dietrich, and L. Jean Camp (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 86–97.

[52] Andrew S. Patrick and Steve Kenny. 2003. From Privacy Legislation to Interface Design: Implementing Information Privacy in Human-Computer Interactions. In *Privacy Enhancing Technologies*, Roger Dingledine (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 107–124.

[53] Sabine Pötzsch. 2009. Privacy Awareness: A Means to Solve the Privacy Paradox?. In *The Future of Identity*, V. Matyáš, S. Fischer-Hübner, D. Cvrček, and P. Švenda (Eds.). IFIP International Federation for Information Processing, 226–236. DOI: http://dx.doi.org/10.1007/978-3-642-03315-5_17

[54] Robert W. Proctor, M. Athar Ali, and Kim-Phuong L. Vu. 2008. Examining Usability of Web Privacy Policies. *International Journal of Human-Computer Interaction* 24, 3 (2008), 307–328. DOI: http://dx.doi.org/10.1080/10447310801937999

[55] Joseph Reagle and Lorrie Faith Cranor. 1999. The Platform for Privacy Preferences. *Commun. ACM* 42, 2 (Feb. 1999), 48–55. DOI: http://dx.doi.org/10.1145/293411.293455

[56] Joel R. Reidenberg, Travis Breaux, Lorrie Faith Carnor, Brian French, Amanda Grannis, James T. Graves, Fei Liu, Aleecia McDonald, Thomas B. Norton, Rohan Ramanath, N. Cameron Russell, Norman Sadeh, and Florian Schaub. 2014. Disagreeable Privacy Policies: Mismatches Between Meaning and Users Understanding. *Berkeley Technology Law Journal* 30, 1 (2014). DOI:http://dx.doi.org/10.15779/Z384K33

[57] Huang Ronggui. 2016. RQDA: R-based Qualitative Data Analysis. R package version 0.2-8. http://rqda.r-forge.r-project.org/. (2016). Last accessed: 2019-07-14.

[58] Florian Schaub, Rebecca Balebako, and Lorrie Faith Cranor. 2017. Designing Effective Privacy Notices and Controls. *IEEE Internet Computing* 21, 3 (2017), 70–77. DOI:http://dx.doi.org/10.1109/MIC.2017.75

[59] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015a. A Design Space for Effective Privacy Notices. In *Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*. USENIX Association, Ottawa, 1–17. https://www.usenix.org/conference/soups2015/proceedings/presentation/schaub

[60] F. Schaub, B. Konings, and M. Weber. 2015b. Context-Adaptive Privacy: Leveraging Context Awareness to Support Privacy Decision Making. *IEEE Pervasive Computing* 14, 01 (jan 2015), 34–43. DOI: http://dx.doi.org/10.1109/MPRV.2015.5

[61] Florian Schaub, Aditya Marella, Pranshu Kalvani, Blase Ur, Chao Pan, Emily Forney, and Lorrie Faith Cranor. 2016. Watching them watching me: Browser extensions impact on user privacy awareness and concern. In *NDSS workshop on usable security*.

[62] Paul M Schwartz and Daniel Solove. 2009. Notice & Choice. In *The Second NPLAN/BMSG Meeting on Digital Media and Marketing to Children*. 7.

[63] Andreas Sotirakopoulos, Kirstie Hawkey, and Konstantin Beznosov. 2011. On the Challenges in Usable Security Lab Studies: Lessons Learned from Replicating a Study on SSL Warnings. In *Proceedings of the Seventh Symposium on Usable Privacy and Security (SOUPS '11)*. Association for Computing Machinery, New York, NY, USA, Article Article 3, 18 pages. DOI: http://dx.doi.org/10.1145/2078827.2078831

[64] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research* 22, 2 (2011), 254–268. DOI: http://dx.doi.org/10.1287/isre.1090.0260

[65] Primal Wijesekera, Arjun Baokar, Ashkan Hosseini, Serge Egelman, David Wagner, and Konstantin Beznosov. 2015. Android Permissions Remystified: A Field Study on Contextual Integrity. In *24th USENIX Security Symposium (USENIX Security 15)*. USENIX Association, Washington, D.C., 499–514. https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/wijesekera

[66] Daricia Wilkinson, Saadhika Sivakumar, David Cherry, Bart P Knijnenburg, Elaine M Raybourn, Pamela Wisniewski, and Henry Sloan. 2017. User-tailored privacy by design. In *Proceedings of the Usable Security Mini Conference*. https://www.ndss-symposium.org/wp-content/uploads/2017/09/usec2017_03_4_Wilkinson_paper.pdf

[67] Craig E Wills and Mihajlo Zeljkovic. 2011. A personalized approach to web privacy: awareness, attitudes and actions. *Information Management & Computer Security* 19, 1 (2011), 53–73.

[68] Shomir Wilson, Florian Schaub, Frederick Liu, Kanthashree Mysore Sathyendra, Daniel Smullen, Sebastian Zimmeck, Rohan Ramanath, Peter Story, Fei Liu, Norman Sadeh, and others. 2018. Analyzing privacy policies at scale: From crowdsourcing to automated annotations. *ACM Transactions on the Web (TWEB)* 13, 1 (2018), 1.

[69] Jochen Wirtz, May O. Lwin, and Jerome D. Williams. 2007. Causes and consequences of consumer online privacy concern. *International Journal of Service Industry Management* 18, 4 (2007), 326–348.

[70] Kuang-Wen Wu, Shaio Yan Huang, David C. Yen, and Irina Popova. 2012. The effect of online privacy policy on consumer privacy concern and trust. *Computers in Human Behavior* 28, 3 (2012), 889–897. DOI: http://dx.doi.org/10.1016/j.chb.2011.12.008

[71] Shaomei Wu, Lindsay Reynolds, Xian Li, and Francisco Guzmán. 2019. Design and Evaluation of a Social Media Writing Support Tool for People with Dyslexia. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, Article Paper 516, 14 pages. DOI: http://dx.doi.org/10.1145/3290605.3300746

[72] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article Article 200 (18 2018), 20 pages. DOI: http://dx.doi.org/10.1145/3274469

[73] Sebastian Zimmeck and Steven M. Bellovin. 2014. Privee: An Architecture for Automatically Analyzing Web Privacy Policies. In *23rd USENIX Security Symposium (USENIX Security 14)*. USENIX Association, San Diego, CA, 1–16. https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/zimmeck