

Understanding iOS Privacy Nutrition Labels: An Exploratory Large-Scale Analysis of App Store Data

Yucheng Li
yuchengl@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Deyuan Chen
deyuanc@andrew.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Tianshi Li
tianshil@cs.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Yuvraj Agarwal
yuvraj@cs.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Lorrie Cranor
lorrie@cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

Jason I. Hong
jasonh@cs.cmu.edu
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA

ABSTRACT

Since December 2020, the Apple App Store has required all developers to create a privacy label when submitting new apps or app updates. However, there has not been a comprehensive study on how developers responded to this requirement. We present the first measurement study of Apple privacy nutrition labels to understand how apps on the U.S. App Store create and update privacy labels. We collected weekly snapshots of the privacy label and other meta-data for all the 1.4 million apps on the U.S. App Store from April 2 to November 5, 2021. Our analysis showed that 51.6% of apps still do not have a privacy label as of November 5, 2021. Although 35.3% of old apps have created a privacy label, only 2.7% of old apps created a privacy label without app updates (i.e., voluntary adoption). Our findings suggest that inactive apps have little incentive to create privacy labels.

CCS CONCEPTS

• **Security and privacy** → **Human and societal aspects of security and privacy**; • **Software and its engineering** → *Software creation and management*.

KEYWORDS

Privacy, Privacy Nutrition Label, iOS development

ACM Reference Format:

Yucheng Li, Deyuan Chen, Tianshi Li, Yuvraj Agarwal, Lorrie Cranor, and Jason I. Hong. 2022. Understanding iOS Privacy Nutrition Labels: An Exploratory Large-Scale Analysis of App Store Data. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '22 Extended Abstracts)*, April 29-May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3491101.3519739>

1 INTRODUCTION

About a decade ago, researchers first introduced the concept of a “Privacy Nutrition Label” to provide users with a more concise

and easy-to-understand summary of how their sensitive data might be accessed and used. The concept was inspired by the “Nutrition Facts” panel, which displays nutrition information about food items in an easy-to-read format. Similarly, privacy nutrition labels offer an alternative to traditional privacy policies written in natural language, giving users more transparency about data usage [5]. Recently, this concept has gained more adoption by the industry. For example, since December 8, 2020, Apple has required all new and updated apps on their App Store to create a privacy label that discloses the app’s data collection practices. Google subsequently announced a similar requirement for the Google Play Store with a targeted date of July 2022 for all apps to have their Data Safety Section approved.

The responsibility of keeping privacy labels accurate largely falls on developers since they are supposed to self-report their apps’ privacy practices. However, based on previous observations of iOS developers creating privacy labels, Li et al. [12] identified patterns of misreporting and uncovered numerous challenges that developers face for creating accurate privacy labels. By quantitatively analyzing the privacy labels of apps on the App Store, we aim to contribute further understanding of how well iOS developers comply with the new privacy requirement after it had been enacted for a year. Our findings can also offer guidance for other app stores to adopt privacy nutrition labels (e.g., Google Play).

We took an exploratory data analysis approach by collecting and analyzing a large-scale dataset of the privacy labels and other metadata of 1.4 million apps on the U.S. Apple App Store. We started collecting the first weekly snapshot on April 2, 2021, and continue to do so even now. For this paper, we perform our data analysis on seven months of data from April 2, 2021 to Nov 5, 2021. Our analysis aims to gain a better understanding of the main drivers for developers to create privacy labels, as well as how promptly they create and update their privacy labels. So we formalize them into three research questions:

RQ1 How promptly do developers react to the call of creating a privacy label?

RQ2 How often do developers update privacy labels after the initial version?

RQ3 How do apps collect and use sensitive data according to their privacy labels?

In this paper, we present the preliminary results of our study and outline important future work directions inspired by our findings.



This work is licensed under a Creative Commons Attribution International 4.0 License.

CHI '22 Extended Abstracts, April 29-May 5, 2022, New Orleans, LA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9156-6/22/04.

<https://doi.org/10.1145/3491101.3519739>

This is the first large-scale analysis of Apple privacy labels to the best of our knowledge. Our analysis showed that 51.6% of apps in the U.S. App Store still do not have a privacy label as of November 5, 2021. Specifically, although 35.3% of old apps (i.e., apps published before Dec. 8, 2020) have created a privacy label, only 2.7% of old apps created a privacy label without simultaneous app updates (i.e., voluntary adoption). Furthermore, privacy label creation was highly associated with app updates, which suggests that there is little incentive for developers of inactive apps to create a privacy label. For apps that created a privacy label in April, 2021, only 5.8% of them ever updated the privacy label while 43.4% of them updated the app. This large gap suggests apps may not update the privacy labels in time.

2 RELATED WORK

In this section, we summarize two lines of research that are closely related to this work.

2.1 Privacy Nutrition Label Research and iOS Privacy Nutrition Label

Privacy researchers have designed privacy nutrition labels for various platforms (e.g., websites [5], mobile apps [7], IoT apps [2]) to provide a clear, uniform, and concise summary of app data practices. Prior research has shown multiple benefits of privacy nutrition labels for users, such as increased speed of finding privacy information and better comprehension of the app's privacy practices [6]. These benefits make privacy nutrition labels a compelling alternative of privacy policies, which are widely acknowledged to be lengthy, ambiguous, and hard to understand [4, 14].

Apple introduced *app privacy details* to their App Store in December 2020, marking the first ever large-scale adoption of the concept of privacy nutrition labels.¹ With this feature, users can learn at a glance what data will be collected by an app, whether the data is linked to users or used to track users, and the purposes for which data may be used. Accuracy is a key requirement for privacy nutrition labels, while the fact that Apple privacy labels are self-reported by developers without a systematic review process leads to potential inaccuracy issues. By observing twelve iOS developers creating a privacy label for their app and interviewing them, Li et al. [12] identified recurring errors in privacy nutrition labels due to developers' knowledge blindspots and the significant overhead and ambiguity involved in this process. Specifically, they found that many developers had not heard about the privacy label requirement before the study or had misunderstanding about when they could create the privacy label for their apps. Their findings also suggest that the challenges of creating a privacy label may make developers reluctant to update their privacy labels in the long run.

In this work, we present preliminary findings of the first large-scale analysis study of Apple's privacy nutrition labels. Cranor et al. [1] analyzed standardized bank privacy notices as a form of privacy nutrition label, while their sample size ($N = 6,000$) is much smaller than ours ($N = 1,437,605$). By quantitatively measuring how developers created and updated privacy nutrition labels, we can better understand the challenges in promoting the adoption

of privacy nutrition labels. We found that an app update appeared to be the key driver of the creation of the first privacy label for an app. We also confirmed that developers rarely updated the privacy labels after creating the first version.

2.2 Large-Scale Privacy Analysis of Mobile Apps

Another line of related work is large-scale analysis studies of mobile apps regarding privacy. Some work examined app privacy behaviors using static or dynamic program analysis to identify data leaks [3, 8, 16]. Specifically, there is growing interest in automatically identifying inconsistencies between privacy policies and the app data practices using program analysis to analyze the app and using NLP to analyze the policies [16]. Although the consistency requirement is also crucial for privacy nutrition labels, Apple's definitions of certain terms make it impossible to conduct this analysis without access to the backend data storage [12]. Therefore, we chose to not involve program analysis in this work.

Other research took a similar approach, focusing on analyzing various types of privacy notices such as privacy policies [15] and permission request rationales [13], which requires a great amount of work simply for parsing the content of the privacy notices. Thanks to the standardized nature, privacy labels are less ambiguous than privacy policies and provide more clear purposes than the rationale messages created by developers. We can directly obtain a privacy label in a machine-readable format, making the analysis a lot easier. In Section 4.3, we demonstrate the potential benefits of analyzing privacy nutrition labels on a large scale for the app store, researchers, developers and users.

3 METHOD

We summarize the data collection methods and the data preprocessing methods in this section.

3.1 Data Collection Methods

We have been collecting data from the U.S. App Store every week since April 2, 2021.




Figure 1: Number of Apps Captured Each Week

Every Thursday, we start the data collection process by updating the app list to account for the removal and addition of apps. Then on

¹<https://developer.apple.com/app-store/app-privacy-details/>

Friday, we collect privacy labels and app metadata by querying the two types of data for all the apps in the updated app ID list. Finally, we run error-checking scripts to identify empty or incomplete data instances and run the crawling scripts again to fix them. Each privacy label instance contains the app ID and the privacy details, namely the data types collected and whether the data is linked to users or used to track users.

3.2 Data Preprocessing Methods

We examined the apps in our dataset along several dimensions related to our three research questions. Below, we introduce these dimensions and how we pre-processed the data to derive corresponding attributes.

3.2.1 New vs. Old Apps. All new apps published on the App Store after Dec. 8, 2020 must provide a privacy label. Apps published before that date can voluntarily add a privacy label at any time, but they will only be forced to add a privacy label on their next app update. We refer to apps published before and after Dec. 8, 2020 as *old apps* and *new apps* respectively. As of Nov. 5, 2021, there are 1,162,748 old apps (80.9% apps on the U.S. App Store) and 274,857 new apps (19.1%).

3.2.2 Old apps without a label on April 2, 2021. For old apps, we further made a distinction based on whether we captured the time that the first privacy label was created. Since we started data collection in April, a few months after the enactment of this new policy, 266,740 apps already had a privacy label in our first data snapshot collected on April 2, 2021 (22.9% old apps). That is to say, we could only identify the time of the first compliance for the rest of the apps (77.1%), which are *old apps without a label on April 2, 2021*. Most of the RQ1 and RQ2 analysis only considered this part of the old apps.

3.2.3 Use of "Ratio" in Figures. As reflected in Figure 1. The number of apps we capture each week is not under a perfectly uniform distribution. With the total count fluctuating around 1.2 million and the number of old apps constantly decreasing, raw counts may not reveal the actual trend in some of our analyses. To minimize this factor, we chose to use a weekly ratio over the raw number to describe the trends in most of the subsequent plots. That is, plots with "ratio" as y-axis do not have a constant denominator. Instead, the ratios are calculated dynamically with numbers captured each week as denominators to account for individual situations and smooth out the fluctuation. Particularly, we frequently use "old apps without a label on April 2, 2021" as the denominator. Note that this number is not a constant, as some apps might be deleted by their developers and some might be purged by the app store. This number has a similar weekly trend as the "Old Apps Published Before Dec. 8, 2020" line in Figure 1.

4 PRELIMINARY FINDINGS

We present findings based on data for 1,437,605 apps collected from April 2 to November 5, 2021 (32 weeks).




Figure 2: Ratios of Apps That Have Labels (Cumulative). The denominator is the number of apps in the U.S. App Store in each week as reflected in Figure 1. Among all the apps, 51.6% of them still do not have a privacy label as of Nov. 5, 2021.




Figure 3: Ratios of old apps that added the first privacy label in each week. The denominator is the number of old apps without a label on April 2, 2021 captured each week. This chart shows an overall decreasing trend in the compliance speed of old apps.

4.1 How Promptly Do Developers React to the Call of Creating a Privacy Label? (RQ1)

4.1.1 The majority of apps on the U.S. App Store are still missing a privacy label. Our analysis showed that 51.6% of apps in the U.S. App Store still do not have a privacy label as of Nov. 5, 2021, eleven months after the app store enacted the new requirement. Due to Apple's requirement, all new apps published after Dec. 8, 2020 must have a privacy label to enter the App Store, resulting in an 100% compliance rate. On the other hand, only 35.3% of old apps created a privacy label as of Nov. 5, 2021. Although the percentage of apps that have labels is increasing steadily overtime (see Figure 2), it largely owes to the newly published apps every week who are forced to have a privacy label. The reality is that the number of old apps adding first label each week is stepping downward, as suggested in Figure 3 by the decreasing trend of the ratio of old apps that created the first label in each week.

4.1.2 Privacy label creation is associated with app updates. Although 35.3% of old apps have created a privacy label, only 2.7% of old apps created a privacy label without simultaneous app updates (i.e., voluntary adoption). Therefore, we further looked into the relationship between privacy label creation time and app update time. Figure 4 shows that the number of old apps that created the first privacy label in the same week of an app update are consistently higher than privacy labels added before an app update, both using the number of old apps without a label on April 2, 2021 as the denominator. Figure 5 further shows that among apps that created the first privacy label over the seven months, 64.2% made an app update at the same time of creating the first label and 14.4% made an app update within one week after the first label.




Figure 4: Trends of the ratios of old apps that created the first privacy label. Denominator is old apps without a label on April 2, 2021 captured that week.




Figure 5: Distribution of the length of time between first privacy label creation and its following app update. The denominator is the number of old Apps that have at least published one privacy label. $N = 104914$. It implies a strong correlation between the version update and first privacy label creation. Over 64.1% of the apps released version updates at the same time as they first published their privacy label.

4.2 How Often Do Developers Update Privacy Labels? (RQ2)

4.2.1 Additional privacy label updates were rare. We further investigated how often developers updated their privacy labels after adding the first one. In this analysis, we only considered the 137,088 apps that created the first privacy nutrition labels in April 2021, including both new and old apps, so we can observe their app updates and privacy label updates for a fixed length of time (28 weeks). As indicated in Figure 6, out of the 137,088 apps, we found that only 7,884 of them ever updated the privacy label later (5.8%). On the other hand, 59,555 of them at least released one app version update (43.4%). Although app updates do not always cause changes in data practices, the large gap between the app updates and additional privacy label updates suggest that developers may not update privacy labels frequently enough to reflect data practice changes in time.

4.3 How Do Apps Collect and Use Sensitive Data According to Their Privacy Labels? (RQ3)

Privacy labels provide information about apps' data practices that may be hard to learn from conventional privacy notices such as privacy policies and permissions. This potentially offers an efficient way to gain a holistic understanding of app data use on the App Store. For example, Figure 7 shows a stable trend of apps reporting "Data not collected" on the U.S. App Store. The denominator for each point is the number of apps that have had a privacy label in that week. For apps that reported some data collection behaviors in the privacy label, we analyzed the ratios of apps that reported different types of data practices over the seven months, including "Data Not Linked to You", "Data Linked to You", and "Data Used to Track You".²

Note that these three categories are not mutually exclusive and we count an app as reporting a certain type of data practice if they at least mentioned one data type associated with this data practice. Figure 8 shows that among apps that reported data collection, both "Data Not Linked to You" and "Data Linked to You" remain stable, while the ratios of "Data Used to Track You" started decreasing in late April, when iOS 14.5 was released. Note that the denominator for each point is the number of apps that have created a privacy label showing data collection practices in that week. These results suggest that the App Tracking Transparency framework introduced in iOS 14.5, which aims to give users control over tracking behaviors,³ may have caused a positive effect on reducing app tracking practices.

We further break down the apps into different categories. There are 26 app categories in total on the App Store. From Figure 9 and 10, we can learn that apps from different categories show very different patterns of data use. This result shows the potential to benefit app developers and users using the privacy label data. Developers can obtain insights into how apps in the same category of their own app tend to collect and use data. Users can see whether a specific

²"Data Not Linked to You" means data not identifiable on its own and not stored with other identifiable data; "Data Linked to You" means data either identifiable on its own or stored with other identifiable data; "Data Used to Track You" means data shared with third parties for advertising purposes.

³<https://developer.apple.com/app-store/user-privacy-and-data-use/>




Figure 6: Distribution of the app version update count and the privacy label update count. The denominator for both charts is the number of apps that created the first privacy label in April ($N = 137,088$).




Figure 7: Trends of ratios of apps claiming “Data not collected”. The denominator for each point is the number of apps that are equipped with a privacy label in that week.

app they want to use has disclosed data practices that are common to other apps in that category.

5 DISCUSSION AND FUTURE WORK

We discuss the challenges and opportunities of promoting privacy labels and future directions based on our findings.

5.1 Challenges in Having Inactive Apps Create Privacy Labels

As shown in Section 4.1, apps published before the new requirement comprise 80.9% of apps on the U.S. App Store, and only 35.3% of these apps have created a privacy label. Although this number is increasing, the rate of change appears to be slowing over time. We noticed that privacy label updates often happened together with app updates, which suggests that there is not much incentive for developers to create privacy labels for apps that are not being updated at this point. This observation echoes prior work’s finding

Figure 8: Trends of ratios of apps reporting different types of data practices. The denominator for each point is the number of apps that have a privacy label showing data collection practices in that week. The decreasing trend of “Data Used to Track You” suggests the positive effect of the App Tracking Transparency framework introduced in iOS 14.5 at late April.

that developers often hold a passive attitude towards privacy [10]. Although it appears to be somewhat remarkable that so many developers had created labels for their old apps and 2.7% old apps created a privacy label without app updates, which suggests voluntary adoption. It would be interesting to understand better what motivated the voluntary adoption. The overall low label adoption rate for old apps makes the label system less useful for users, since they can only view labels for about half the apps they might be interested in.

To address this problem, there are two important directions for future research. The first direction is to gain more in-depth understanding on what factors affect how promptly developers create a privacy label. For example, Li et al. [12] interviewed developers and learned that many of them had not heard about privacy label




Figure 9: Ratios of apps reporting no data collection per app category. The denominator of each bar is the total number of apps in that category that have a privacy label. Note that some categories are abbreviated (Dev: Developer Tools, Edu: Education, Life: Lifestyle, Med: Medical, Ref: Reference)




Figure 10: Ratios of apps reporting different types of data practices per app category. The denominator of each bar is the total number of apps in that category that have a privacy label reporting data collection practices. Note that some categories are abbreviated (Dev: Developer Tools, Edu: Education, Life: Lifestyle, Med: Medical, Ref: Reference)

or had misperceptions about how to create one, which potentially explains the problem with inactive apps. Future research should further investigate this problem using different methods and more diverse and representative samples. A second direction is to design more effective techniques to increase developers' knowledge about privacy labels and help them create the label. Specifically, there should be more methods to inform developers about this requirement in addition to the app store. For example, the IDE may be enhanced to prompt developers about this requirement when they are coding [9] and it may even further help developers create the privacy label by analyzing the source code [11].

5.2 Challenges in Having Developers Update Privacy Labels over Time

As shown in Section 4.2, privacy labels seem to be rarely updated after they are created. Among apps that created the first label in April, only 5.8% of them ever updated the label within 28 weeks after the first privacy label was created, while 43.4% of them released app updates. Although a privacy label update is not needed if the data practices are not changed, this large gap suggests that it may be challenging for apps to keep their privacy labels up-to-date. Currently, the App Store does not show further alerts as long as there is a privacy label already in place.

Hence, future research needs to develop methods to automatically detect privacy labels that fail to reflect changes in data practices. A potential idea is to leverage the information in the release

notes of app updates (part of the app metadata in our dataset). For example, keywords like “cloud storage” in release notes can indicate potential changes to data practices on an app update. We tested this idea and found 982 apps that contained the keyword “cloud storage” and we did find apps that stored sensitive user data on cloud but did not report it on their privacy label, such as one location logging app that stores the location data on the cloud but had a privacy label claiming “Data Not Collected”.

5.3 Opportunities in Improving Understanding of App Privacy

On the other hand, we have shown in Section 4.3 that because privacy labels can be described in a machine-readable format, it provides a simple way to analyze app data practices on a large scale. Our analysis demonstrates that the App Transparency Framework introduced in iOS 14.5 appears to result in less data used for tracking, which is beneficial to user privacy. We also demonstrate the different patterns of data use across app categories which can serve as references for developers and users to develop and select apps. Overall, privacy labels have the potential to help users, developers and researchers in different aspects. Although we want to note that these labels may contain errors and whether there are systematic errors that can cause misunderstanding in the overall trend remains to be investigated by future work.

6 LIMITATIONS

There are several limitations in our study methodology. First, as mentioned in Section 3.2.1, we used the “Date published” information to categorize apps into old and new apps. However, a closer examination of the dataset revealed that this information may not always accurately reflect the first day the app was published on the app store. Specifically, we noticed that 1.3% of new apps had a first version release date earlier than the published date and 0.52% of new apps do not have a privacy label provided, which suggests that there may be errors in the release date and published date information that we crawled from the app store. Second, because we started crawling four months after the requirement of privacy labels, we could not figure out the exact time that the 266,740 apps that created a privacy label before we started crawling (22.9% old apps) and we had to exclude this portion of apps from some analysis about RQ1 and RQ2. This means the trends we identified about old apps without a label on April 2, 2021 may not generalize to other old apps.

7 CONCLUSION

In this paper, we present the first work to analyze the privacy nutrition labels of apps on the U.S. Apple App Store. From our analysis of 32 weeks of data about 1.4 million apps, we identified various challenges in the adoption of privacy nutrition labels, including inactive apps lacking incentives to create privacy labels and developers may not update privacy labels to reflect data practice changes in time. We discuss future research directions based on our findings.

ACKNOWLEDGMENTS

This research was supported in part by the National Science Foundation under Grant No. CNS-1801472, Air Force Research Laboratory

under agreement number FA8750-15-2-0281, and Innovators Network Foundation. Tianshi Li was supported in part by the CMU CyLab Presidential Fellowship. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Research Laboratory or the U.S. Government. We thank the anonymous reviewers for their constructive feedback.

REFERENCES

- [1] Lorrie Faith Cranor, Pedro Giovanni Leon, and Blase Ur. 2016. A large-scale evaluation of US financial institutions’ standardized privacy notices. *ACM Transactions on the Web (TWEB)* 10, 3 (2016), 1–33.
- [2] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. 2020. Ask the experts: What should be on an IoT privacy and security label?. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 447–464.
- [3] Catherine Han, Irwin Reyes, Amit Elazari Bar On, Joel Reardon, Álvaro Feal, Serge Egelman, Narseo Vallina-Rodriguez, et al. 2019. Do you get what you pay for? Comparing the privacy behaviors of free vs. paid apps. In *Workshop on Technology and Consumer Protection (ConPro 2019)*, in conjunction with the 39th IEEE Symposium on Security and Privacy, 23 May 2019, San Francisco, CA, USA.
- [4] Carlos Jensen and Colin Potts. 2004. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *Proceedings of the 2004 conference on Human factors in computing systems - CHI '04*. ACM Press. <https://doi.org/10.1145/985692.985752>
- [5] 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*. 1–12.
- [6] 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*. 1573–1582.
- [7] Patrick Gage Kelley, Lorrie Faith Cranor, and Norman Sadeh. 2013. Privacy as part of the app decision-making process. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 3393–3402.
- [8] Konrad Kollnig, Anastasia Shuba, Reuben Binns, Max Van Kleek, and Nigel Shadbolt. 2021. Are iPhones Really Better for Privacy? Comparative Study of iOS and Android Apps. *arXiv preprint arXiv:2109.13722* (2021).
- [9] Tianshi Li, Yuvraj Agarwal, and Jason I Hong. 2018. Coconut: An IDE plugin for developing privacy-friendly apps. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 4 (2018), 1–35.
- [10] Tianshi Li, Elizabeth Louie, Laura Dabbish, and Jason I. Hong. 2021. How Developers Talk About Personal Data and What It Means for User Privacy: A Case Study of a Developer Forum on Reddit. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW3 (jan 2021), 1–28. <https://doi.org/10.1145/3432919>
- [11] Tianshi Li, Elijah B. Neundorfer, Yuvraj Agarwal, and Jason I. Hong. 2021. Hon-ey-suckle: Annotation-Guided Code Generation of In-App Privacy Notices. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3 (sep 2021), 1–27. <https://doi.org/10.1145/3478097>
- [12] Tianshi Li, Kayla Reiman, Yuvraj Agarwal, Lorrie Faith Cranor, and Jason I Hong. 2022. Understanding Challenges for Developers to Create Accurate Privacy Nutrition Labels. *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (2022).
- [13] Xueqing Liu, Yue Leng, Wei Yang, Wenyu Wang, Chengxiang Zhai, and Tao Xie. 2018. A large-scale empirical study on android runtime-permission rationale messages. In *2018 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 137–146.
- [14] 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.
- [15] Shomir Wilson, Florian Schaub, Aswath 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.
- [16] Sebastian Zimmeck, Ziqi Wang, Lieyong Zou, Roger Iyengar, Bin Liu, Florian Schaub, Shomir Wilson, Norman Sadeh, Steven Bellovin, and Joel Reidenberg. 2016. Automated analysis of privacy requirements for mobile apps. In *2016 AAAI Fall Symposium Series*.