Problematic Privacy Policies of Voice Assistant Applications

Song Liao, Christin Wilson, and Cheng Long | Clemson University
Hongxin Hu | University at Buffalo
Huixing Deng | Clemson University

This work conducts a comprehensive empirical analysis on the privacy policies of 64,720 Amazon Alexa skills and 16,002 Google Assistant actions. We show that a substantial number of problematic privacy policies exist in the Amazon Alexa and Google Assistant platforms.

Voice personal assistants (VPAs), such as Amazon Alexa and Google Assistant, have been seamlessly integrated into our daily lives. An estimated 4.2 billion voice assistants were being used around the world in 2020, and that number is forecasted to reach 8.4 billion by 2024. Despite these assistants having many convenient features, there is an increasing concern about the privacy risks of their users. Both the Amazon Alexa and Google Assistant platforms allow third-party developers to build new voice apps (which are called *skills* on the Amazon Alexa platform and *actions* on the Google Assistant platform) and publish them to app stores. To comply with privacy regulations [such as the Children's Online Privacy Protection Act (COPPA)[8] and the Health Insurance Portability and Accountability Act (HIPAA)[13]] and to protect consumers' privacy, developers are required to provide privacy policies and to notify users about their voice apps' data practices.

Typically, a proper privacy policy should, at minimum, have answers to three important questions: 1) What information is being collected? 2) How is this information being used? and 3) What information is being shared? Privacy policies provided by third-party developers can be diverse and poorly written, which results in users choosing to not read them. This also leads to users using privacy-sensitive services without having a proper understanding of the data that is being collected and what developers will do with it. In this work, we seek to understand whether voice app developers provide informative and meaningful privacy policies in two mainstream VPA platforms.

Privacy Policy Requirements

Voice apps, like smartphone apps, are mostly created by third-party developers and are available through a website known as the *app store*. Each voice app has a unique webpage that displays the developer's information, description, and sample utterances and the privacy policy link, user reviews, and so on. Google and Amazon have taken different approaches when it comes to the requirement for a privacy policy of each voice app available to users. While Google has made it mandatory for developers to provide a privacy policy along with

Digital Object Identifier 10.1109/MSEC.2021.3082474
Date of current version: 21 June 2021

 Copublished by the IEEE Computer and Reliability Societies

each action, Amazon is more lenient and only makes it a requirement for skills that declare they collect personal information. Both Amazon and Google prevent the submission of a voice app for certification if their respective requirements are not met.[1,6]

Google has a "Privacy Policy Guidance" page[4] in their documentation for action developers, which explains what Google's minimum expectation is for a privacy policy document. According to the guide, the privacy disclosures included in a policy should be comprehensive, accurate, and easy for users to understand and should disclose all of the information that an action collects through all of the interfaces, including the data that is automatically collected. How the information is collected, to whom it is shared, and when it is shared should also be specified. Google rejects an action if developers do not provide (or even misspell) the action name, company name, or email in the privacy policy. The link should be valid and should also be a public document viewable by everyone. The Amazon Alexa platform doesn't provide a guideline for privacy policy content.

Our Measurement Study

In this work, we conducted an empirical analysis to measure the effectiveness of privacy policies provided by voice app developers on both Amazon Alexa and Google Assistant platforms.[12] We built a crawler to collect each voice app's ID, name, developer information, description, permission (only for Alexa skills), and privacy policy link from Amazon and Google's voice app stores. As of March 2020, we were able to collect 64,720 unique skills under 21 categories from Amazon Alexa's skill store, 17,952 (28%) of which provided privacy policy links. Among the 16,002 Google actions that we collected, 9,955 (62%) had privacy policy links.

> **While Google has made it mandatory for developers to provide a privacy policy along with each action, Amazon is more lenient and only makes it a requirement for skills that declare they collect personal information.**

We aimed to identify problematic privacy policies, such as those without any data practices or with misleading/inconsistent information. Unlike smartphone platforms (for example, Android or iOS), the source code of voice apps in the Amazon Alexa and Google Assistant platforms are not publicly available. This limits the extent of our privacy analysis since we do not have the actual code of voice apps to find inconsistencies with the provided privacy policies. Instead, we used each voice app's description and permission information in its introduction page for our analysis to detect inconsistent privacy policies. For this reason, our results on the inconsistencies of privacy policies are not focused on the exact number of mismatches and errors but on the existence of problems potentially affecting the overall user experience. We first processed all of the privacy policy links to identify high-level issues, such as broken and duplicate uniform resource locators (URLs). Then, we obtained all of the privacy policy content and conducted a natural language processing (NLP)-based analysis to identify inconsistencies.

Capturing Data Practices

We used PolicyLint,[7] a privacy policy-analysis tool, to capture the data practices in each privacy policy document. In this step, we also identified privacy policies without data practices. PolicyLint can detect both positive and negative data practices. For example, it captures the sentence, "We never collect data with our skills," as a negative data practice. In this case, if a skill mentions data collection in its description, we conclude that it has an inconsistent privacy policy.

Since PolicyLint is mainly focused on privacy policy analysis, it does not perform well in capturing data practices from app descriptions due to the diversity of descriptions written by different developers. For this reason, we developed an NLP-based approach to capture the data practices in a voice app's description. First, we defined 40 verbs commonly used in data practices,[14] such as "collect," "use," "retain," and "disclose." We also collected a dictionary of 16 nouns related to data practices from Amazon's skill-permission list[3] and the Amazon Developer Services Agreement.[2] Next, we parsed a description into sentences and searched for sentences containing both of these verbs and nouns. We used the SpaCy library[5] to obtain the attribute for each word and the straight correlation between a noun and a verb and then generated phrases related to data practices. Then, the library generated phrases related to data practices. The phrase would follow the patterns of "subject|verb|object" or "subject|passive-verb" (for example, "Alexa asks your location" or "user name required"). Our NLP-based analysis achieved an accuracy of 80%, and we identified 117 skills (confirmed by our manual analysis) that had data practices in their descriptions.

Inconsistency Checking

We checked for potential inconsistencies between the data practices from the descriptions and those from the policies. We mainly aimed to detect three types of inconsistencies in the privacy policies. First, if a voice app didn't provide a privacy policy although one was required, it was defined as lacking one. For example, the "Heritage Flag Color" skill mentioned that "the device location is required" in its description. However, the developer didn't provide a privacy policy. Second, if a skill had a privacy policy but did not fully disclose its data-collection practices in it, we considered it as incomplete. For example, the description of the "Thought Leaders" skill mentioned "Permission required: Customer's Full name, Customer's Email Address, Customer's Phone number," but none of these were mentioned in its privacy policy. Third, if a skill explicitly mentioned it did not collect data in its description but claimed the opposite in the privacy policy (or vice versa), it was considered as inconsistent.

High-Level Issues

Not all voice apps have correct privacy policy URLs. Table 1 lists the statistics of privacy policies and high-level issues. In the case of Alexa, only 17,952 out of the 64,720 skills we collected had privacy policies. This is partially because of the lenient skill certification on the Amazon Alexa platform. After conducting further experiments on the skill certification,[9] we have come to understand that even if a skill collects personal information, the developer can choose not to declare it during the certification stage and bypass the privacy policy requirement. This is achieved by collecting personal information through the conversational interface (for example, asking users' names). As a result, developers may choose to not provide a privacy policy. Amazon only requires skills that collect personal data to provide privacy policies, thus, not all of these 46,768 skills required privacy policies.

Out of the 16,002 actions we collected from the Google action store, 6,047 did not provide privacy policies. Among these actions lacking privacy policies, only seven provided the developer information, and only 32 were rated by at least one user, which indicates that most of these actions are rarely used by VPA users. Since it is not possible to submit an action for certification without including a privacy policy URL, it is puzzling how these actions were available in the store.

For those skills and actions that have provided privacy policy links, not every URL leads to a page containing the privacy policy. We found 3,131 Alexa skills (17%) and 169 Google actions (2%) that provided broken privacy policy URLs, as displayed in Table 1. There are also URLs which lead to other developers' privacy policies. An example of this is the "NORAD Tracks Santa" skill, which provides a privacy policy URL that links to Amazon's privacy policy page instead of to a privacy policy written by the developer. The privacy policy URL of "Rubetek SmartHome," which is both an Alexa skill and a Google action, leads to the company's home page, which promotes its products rather than linking to the privacy policy page.

Duplicate URLs

We found that a substantial portion of privacy policies shared the same URLs. In particular, more than 56% of Amazon Alexa skills had duplicate privacy policy URLs. Out of the 17,952 Amazon skills with privacy policies, 7,828 had unique URLs. The other 10,124 skills (56.4%) shared 1,206 different privacy policy URLs. 1,783 skills (9.9%) provided the same link (https://getstoryline.com/public/privacy.html). Note that these 1,783 skills were not from the same developer, which indicated that the privacy policy was irrelevant to these skills. Here, "irrelevance" means that the privacy policy provided in a URL was not written specifically for the developer or the voice app (for example, without including the skill name, company name, or developer's email).

The issue of duplicate URLs was more serious on the Amazon Alexa platform. Table 2 lists the top five duplicate privacy policy URLs of Alexa skills. These URLs were shared by 3,765 skills, constituting 21% of the total skills that had privacy policies. Nine percent of the actions on the Google Assistant platform had duplicate privacy policy URLs. Out of the 9,955 actions, 9,056 had unique privacy policy links. The other 899 actions shared 204 different privacy policy URLs.

To understand why there was such a large number of voice apps with duplicate privacy policy URLs,

Table 1. High-level issues of privacy policies in two mainstream VPA platforms.

	Alexa skills		Google actions	
	Total number	Percentage	Total number	Percentage
With privacy policy	17,952	—	9,955	—
Broken privacy policy URL	3,131	17%	169	2%
Duplicate privacy policy URL	10,124	56%	899	9%
Privacy policy without any data practices	793	11%	305	3%

especially on the Amazon Alexa platform, we further examined the developer information of these voice apps. Our intuition was that developers who published multiple voice apps may have used the same privacy policy URLs for them. We found that, for developers who developed more than one voice app, 77% of their skills used duplicate privacy policy URLs. For the top five developers who published the most skills with privacy policies on the Amazon Alexa platform, 2,064 out of their 2,069 skills (99.8%) used duplicate privacy policy URLs. A serious problem happens if such a privacy policy link is broken, which results in thousands of skills being affected. For example, we found two broken links in the top five duplicate privacy policies URLs, as listed in Table 2. There were 1,275 skills using these two links; thus, all of their privacy policies were inaccessible. As for Google actions, we observed a similar issue.

For the developers who published more than one action, 31% of their actions had duplicate privacy policy URLs. For the top 10 developers who published the most actions, 86% of their actions used duplicate privacy policy links. The content of these privacy policy URLs was not specific to voice apps, and users may have skipped reading the privacy policies even though they were provided.

Some official Google and Amazon voice apps violated their own requirements. We collected 309 official Amazon Alexa skills (that is, developed by Amazon, Amazon Alexa Prime Video, Amazon.com, and so on), out of which 222 skills came with privacy policy URLs, but 211 of these were duplicates. Among these privacy policy links, 181 pointed to the general Amazon privacy notice, and 27 were to Alexa's "Terms of Use" or the Amazon Web Services' privacy notice. Surprisingly, seven privacy policy links were totally unrelated to privacy notices: three linked to the Amazon home page, two were developer documents, and two were pages about insurance or an Amazon product. We found two official weather skills on Amazon Alexa's skill store, and one of them asked for the user's location information according to its description, but it didn't provide a privacy policy. Figure 1 presents the weather skill developed by Amazon with the ID "B071Z29JLY." This skill may be automatically enabled and available on all Alexa devices since it is built in. This example demonstrates that Amazon Alexa violates its own requirement by publishing skills capable of collecting personal information without providing a privacy policy.

In the Google Assistant's action store, we found 92 official actions developed by Google. All of the 92 actions provided privacy policy links, but they pointed to two different Google privacy policy pages—both of which were general privacy policies. Google requires that every action should have an app-specific privacy policy provided by developers upon submission (including the action name, company name, or developer's email in the privacy policy). However, our analysis revealed that this requirement had not been enforced in a proper manner at the submission time of these 92 actions. To confirm whether such a requirement was enforced for third-party developers, we submitted multiple actions purposely violating this requirement (for example, without providing the action name or providing a wrong name). Our submissions were rejected due to this reason.

No Data Practices

Using the PolicyLint tool,[7] we captured data practices for each privacy policy. For these privacy policies with data practices, the average number of practices was 25.9 for Alexa skills and 15.5 for Google actions. The maximum number of data practices in a privacy policy was

Table 2. The five most common duplicate privacy policy URLs shared by Alexa skills.

Duplicate privacy policy URLs	Shared by skills	Percentage	Is it live?
https://getstoryline.com/public/privacy.html	1,783	9.9%	Live
https://corp.patch.com/privacy	1,012	5.6%	Broken
https://cir.st/privacy-policy	410	2.3%	Live
http://spokenlayer.com/privacy	297	1.7%	Live
http://www.lottostrategies.com/script/ showpage/1001029/b/privacy_policy.html	263	1.5%	Broken

Figure 1. An official skill lacks a privacy policy. Even though it collects the user's location according to the description, no privacy policy is provided.

433, which is likely to be a general privacy policy rather than an app-specific one.

We detected 942 Alexa skills and 389 Google actions having privacy policies but with no data practices. After manually checking these privacy policies, we saw that we achieved an accuracy of 82%, with 149 false positives for the 942 Alexa skills and 84 false positives for the 389 Google actions. We found that most of these cases occurred because of the crawler failing to correctly obtain privacy policies. For example, when privacy policies were embedded in the Web framework, we could not get the correct content while crawling the privacy policy webpage.

In particular, 793 privacy policies provided with Alexa skills did not have data practices (verified by our manual analysis). Figure 2 displays the breakdown of the different issues of these privacy policies. Privacy policy URLs, 393, led to totally unrelated pages, which had advertisements and shopping options. URLs, 191, led to actual website domains, but the links were not found. These could also be considered as broken links. 181 URLs led to actual privacy policy pages but did not mention any data practices. Seventeen URLs led to a page where the actual link to the privacy policy did exist but also redirected to some other pages. Another 11 skills needed logins to access the documents.

After the manual analysis, we found 305 Google actions having privacy policies without any data practices, as shown in Figure 2. URLs, a total of 143, led to pages that were not found. An amount of 48 URLs led to unrelated links with shopping options and product advertisements. Five URLs led to pages containing links to actual privacy policies. Three URLs were privacy policies but did not have any data practices. Different from the Alexa skills, 106 actions provided their privacy policies as Google Docs, which do not have the correct permissions, resulting in users not being able to access them. Obviously, these violate Google's restriction of "the link should be a public document viewable by everyone."

Inconsistency in Privacy Policies

Inconsistency Between the Privacy Policy and Permission Information

Both the Amazon Alexa and Google Assistant platforms provide permission-requesting application programming interfaces (APIs) for voice apps collecting specific types of data from users, and they require developers to request permissions through the provided APIs. The Amazon Alexa platform allows developers to request permissions for collecting the device address and the customer's name, email address, and phone number. If a skill collects these data through permission APIs, the permission information will be displayed in the skill's introduction page. We checked whether the privacy policies disclosed such data practices or not. Google states that an action must "Request all sensitive user data (location and name) via the Permissions API." However, Google doesn't display the permission information in each action's webpage. Thus, we only analyzed the inconsistency between privacy policies and the permission information for Alexa skills.

Table 3 summarizes the results of our inconsistency checking. We obtained 1,369 skills using permission APIs to collect user data in the Amazon Alexa platform. First, we checked whether the skills had privacy policies. Surprisingly, a skill, "Adopt A Pet," used permission APIs to access "Device Country and Postal Code" data but didn't

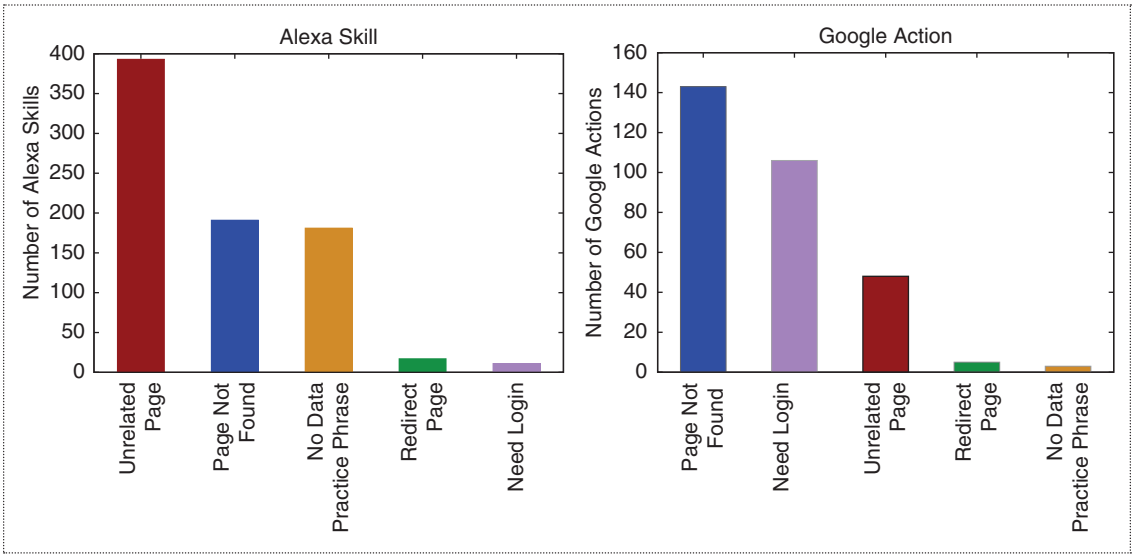

Figure 2. Some different issues of privacy policies that do not have data practices in two VPA platforms.

provide a privacy policy. Next, we used the PolicyLint tool to analyze privacy policies, and we found that 375 skills requested permissions but provided incomplete privacy policies. We also observed that 48 skills provided inconsistent privacy policies (they claimed to not collect data in the privacy policies, but, in fact, they did). For example, the "Cursed Painting" skill collected the email address with permission APIs but claimed that "We do not collect any personal information" in its privacy policy.

Inconsistency Between the Privacy Policy and Description

For the Amazon Alexa platform, we identified 117 skills mentioning data collection in their descriptions. We only found two Google actions mentioning data collection in the description. After checking their privacy policies, we detected whether voice apps lacked privacy policies or had incomplete or inconsistent ones. The results are presented in Table 3.

There were 24 skills collecting data without providing privacy policies. For example, the "Name My Grandkids" skill included in its description that it asks users for personal information and stores it for future use. In another case, the Lapel Athletics skill required the device location according to its description. However, neither of these skills provided privacy policies. We identified 61 Alexa skills that described the collection of personal data in the description but provided an incomplete privacy policy. Among the 61 skills, 20 asked for the address or location, 14 requested the user name, 10 asked for the email account and password, and the others asked for birthday, phone number, contact, and gender or health-related data. For example, the "Running Outfit Advisor" skill mentioned collecting gender information in its description but did not mention this data practice in its privacy policy. We found five skills providing inconsistent privacy policies. Three skills mentioned that they would collect data in their descriptions but claimed that they would not do so in their privacy policies. The other two skills claimed they would not collect data in their description but disclosed data-collection practices in their privacy policies. For example, the "meet talk" skill (displayed in Figure 3) described "This skill will ask your name." But, in the privacy policy, it claimed "We never collect or share personal data with our skills." Another skill, "Caren," mentioned "This skill does not collect personal information" in its description, but in the privacy policy it said "We collect your cell phone number."

As for the Google Assistant platform, 6,047 Google actions did not provide privacy policies, which violates its own restriction of "Google requires all actions to post a link to their privacy policy in the directory." Since there were only two Google actions mentioning data collection in their descriptions, we did not find any inconsistencies in their privacy policies. In addition, for these skills and actions with data collection, we found that 28 skills and one action should have asked for permissions through APIs (since they collect the restricted data explicitly defined by VPA platforms), but they did not.

Table 3. A summary of the inconsistency-checking results.

	Data collection through permission APIs	Data collection mentioned in the description
Number of skills with data collection	1,369	117
Lacking a privacy policy	1	24
Having an incomplete privacy policy	375	61
Having an inconsistent privacy policy	48	5

Potential Noncompliance With Legal Regulations

In our analysis, there were multiple skills that collected personal information to be published on the Amazon Alexa skill store under the kids category without providing privacy policies. This is not compliant with COPPA[8] regulations, which require every developer collecting personal information from children to follow certain rules. Providing a privacy policy with accurate information about what data are collected and what they are used for is one of the main requirements. The objective is to clearly let parents know about what personal information can be collected by the skill from their children. Health-related information can also be collected by a voice app through conversational interface without a privacy policy being provided, even though only the user can decide whether to provide it or not. However, voice apps having this capability might be a violation of the HIPAA regulation.[13] The California Online Privacy Protection Act requires developers to provide a privacy

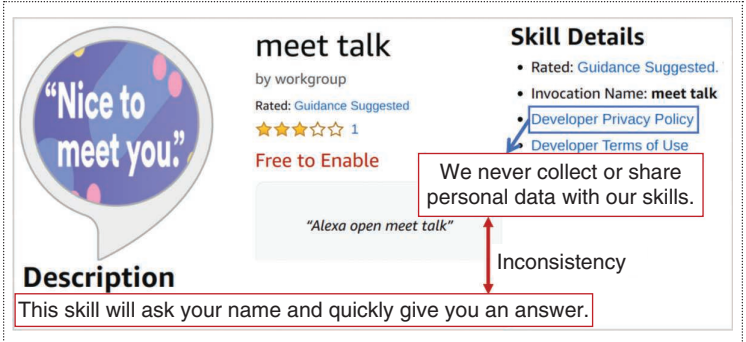

Figure 3. A skill with an inconsistent privacy policy.

policy that states exactly what data can be collected from users. We found that 137 kids' skills provide general information in their privacy policies without providing specifics on what personal data they actually collect. These voice apps and their privacy policies may not be in compliance with legal regulations.

Discussion

Why Poor-Quality Privacy Policies?

The Amazon Alexa platform not explicitly requiring app-specific privacy policies may result in developers providing the same document to explain the data practices of all of their services. This leads to uncertainty and confusion among end users. There are skills with privacy policies containing up to 433 data practices, and most of these data practices are not relevant to the skill. Thus, these documents do not give a proper understanding of the capabilities of a skill to end users. The poor quality of privacy policies provided with voice apps is partially due to the lack of app-specific privacy policies and to the lenient certification process at the skill-submission phase. The content of a privacy policy is not thoroughly checked when a skill is submitted for certification, which has resulted in a large number of inactive and broken links and also privacy policies not related to their skills. Some privacy policies mentioned data practices that were in violation of the privacy requirements that Amazon and Google set, but these voice apps were still certified.

In some cases, even if the developer wrote the privacy policy with proper intention and care, there could be some discrepancies between the policy and the actual code since updates made to a voice app might not be reflected in the privacy policy. This is especially possible with the current VPA architecture because the back-end code of a voice app can be updated at any time by the developer and does not require any recertification to be made available to end users. The outdated policy may lead to developers unintentionally collecting personal information without informing the users.

Usability of Privacy Notice for Voice Apps

The constrained interfaces on VPA devices pose a challenge to effective privacy notices. Privacy policies are only available on a voice app store's webpages. The unavailability of privacy policies through the voice channel requires users to access them over the Web or through VPAs' companion apps on their smartphones. One possible reason for this could be the large size of privacy policies and the time required to read the long documents. Users who only use voice assistant services through their VPA devices may not necessarily be aware of the existence of the privacy policies in the respective stores.

To understand how users engage with privacy policies and their perspectives on VPAs' privacy policies, we conducted a user study with 91 participants (VPA users) using the Amazon Mechanical Turk platform.[12] 48% of the participants claimed that they were aware of the privacy policies of the voice apps they used. However, when asked about how often they actually read the privacy policies provided by the developers, 73% responded with "rarely." Of the participants, 66% said that they never read the privacy policy, and 47% were not aware of which data are being collected by the skill. When asked about the issues they face with privacy policies,

- 20% of the participants responded by saying they are hard to access.
- 44% of the participants felt that the documents were too long.
- 24% claimed that they felt inconsistencies between a privacy policy and a skill's actual functionality and description.

Our survey results suggest the need for VPA platforms to take measures to improve the quality of privacy policies and to provide effective privacy notices for VPA users to make informed privacy decisions. As our future work, we plan to improve the usability of privacy notices to VPA users, such as providing privacy notices through voice responses.

Limitation

A main limitation of this work is that we largely depend on the description and permission information provided with the voice apps for detecting inconsistencies of privacy policies, without exploring the voice apps' runtime behaviors. This leaves our findings on the inconsistency checking incomplete. The availability of source code can significantly increase the knowledge of what personal data a voice app is able to collect and where it is stored. Without a baseline, a future research effort could be to dynamically test voice apps by enabling them and checking their data-collection practices. Recently, SkillExplorer[10] has been proposed to dynamically explore voice apps' runtime behaviors and to detect privacy violations in voice apps. In our future work, we will leave the dynamic analysis of voice apps to identify more inconsistent privacy policies. Lentzsch et al. also studied whether privacy policies of skills consistently disclose the data accessed using similar methodologies.[11] They mainly focused on identifying skills that request permissions (that is, collect specific types of data through permission APIs) but do not have accompanying privacy policies. However, differently from this,[11] we aim to identify inconsistencies between a skill's description and its privacy policy. In addition, the authors of this article[11] only considered the Amazon Alexa platform while we discussed the quality and usability issues of privacy policies on both the Amazon Alexa and Google Assistant platforms.

In this work, we conducted a comprehensive empirical analysis on the privacy policies of 64,720 Amazon Alexa skills and 16,002 Google Assistant actions. We identified 4,004 Alexa skills and 331 Google actions having problematic privacy policies. We found that 6,047 Google actions didn't have privacy policies, which violates Google's policy requirement related to privacy policies. The results reveal a worrisome reality about privacy policies in two mainstream voice app stores. ■

Acknowledgments

This work is supported in part by the National Science Foundation under grants 2114920, 2031002, 1846291, 1700499, and 1642143.

References

1. "Alexa skills policy testing." Amazon. https://developer.amazon.com/fr/docs/custom-skills/policy-testing-for-an-alexa-skill.html (accessed June 1, 2021).
2. "Amazon developer services agreement." Amazon. https://developer.amazon.com/support/legal/da (accessed June 1, 2021).
3. "Configure permissions for customer information in your skill." Amazon. https://developer.amazon.com/en-US/docs/alexa/custom-skills/configure-permissions-for-customer-information-in-your-skill.html (accessed June 1, 2021).
4. "Google privacy policy guidance." Google. https://developers.google.com/assistant/console/policies/privacy-policy-guide (accessed June 1, 2021).
5. "Industrial-strength natural language processing." SpaCy. https://spacy.io (accessed June 1, 2021).
6. "Policies for actions on Google." Google. https://developers.google.com/actions/policies/general-policies (accessed June 1, 2021).
7. B. Andow et al., "Policylint: Investigating internal privacy policy contradictions on google play," in *Proc. 28th USENIX Security Symp. (USENIX Security 19)*, 2019, pp. 585–602.
8. N. Apthorpe, S. Varghese, and N. Feamster, "Evaluating the contextual integrity of privacy regulation: Parents' IoT toy privacy norms versus COPPA," in *Proc. 28th USENIX Security Symp. (USENIX Security)*, 2019, pp 123–140.
9. L. Cheng, C. Wilson, S. Liao, J. Young, D. Dong, and H. Hu, "Dangerous skills got certified: Measuring the trustworthiness of skill certification in voice personal assistant platforms," in *Proc. ACM SIGSAC Conf. Comput. Commun. Security (CCS)*, 2020, pp. 1699–1716. doi: 10.1145/3372297.3423339.
10. Z. Guo, Z. Lin, P. Li, and K. Chen, "Skillexplorer: Understanding the behavior of skills in large scale," in *Proc. 29th USENIX Security Symp. (USENIX Security 20)*, 2020, pp. 2649–2666.
11. C. Lentzsch, S. J. Shah, B. Andow, M. Degeling, A. Das, and W. Enck, "Hey Alexa, is this skill safe?: Taking a closer look at the Alexa skill ecosystem," in *Proc. Network Distributed Systems Security (NDSS) Symp.*, 2021, pp. 1–18.
12. S. Liao, C. Wilson, L. Cheng, H. Hu, and H. Deng, "Measuring the effectiveness of privacy policies for voice assistant applications," in *Proc. Computer Security Appl. Conf.*, 2020, pp. 856–869. doi: 10.1145/3427228.3427250.
13. F. H. Shezan, H. Hu, G. Wang, and Y. Tian, "Verhealth: Vetting medical voice applications through policy enforcement," in *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2020, pp. 1–21.
14. L. Yu, X. Luo, X. Liu, and T. Zhang, "Can we trust the privacy policies of android apps?" in *Proc. 2016 46th Annu. IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN)*, pp. 538–549. doi: 10.1109/DSN.2016.55.

Song Liao is a Ph.D. student in computer science at Clemson University, Clemson, South Carolina, 29634, USA. His research interests include computer security and privacy. Liao received an M.S in software engineering from Xi'an Jiaotong University. Contact him at liao5@g.clemson.edu.

Christin Wilson is a software engineer at Global Lending Services LLC Greenville, South Carolina, 29603, USA. His research interests include computer security and privacy. Wilson received an M.S. from Clemson University. Contact him at cwils28@g.clemson.edu.

Cheng Long is an assistant professor in the School of Computing at Clemson University, Clemson, South Carolina, 29634, USA. His research interests include system and network security, the Internet of Things, and mobile computing. Long received a Ph.D. in computer science from Virginia Tech. Contact him at lcheng2@clemson.edu.

Hongxin Hu is an associate professor in the Department of Computer Science and Engineering at the State University of New York at Buffalo, Buffalo, New York, 14260, USA. His research interests include security in emerging networking technologies, security in the Internet of Things, machine learning for security and privacy, and security and privacy in social networks. Hu received a Ph.D. in computer science and engineering from Arizona State University. Contact him at hongxinh@buffalo.edu.

Huixing Deng is a software engineer at Amazon, Seattle, Washington, 98125, USA. His research interests include computer security and privacy. Deng received an M.S. from Clemson University. Contact him at huixind@g.clemson.edu.