Measuring the Effectiveness of Privacy Policies for Voice Assistant Applications

Song Liao*, Christin Wilson*, Long Cheng, Hongxin Hu, Huixing Deng
School of Computing, Clemson University, USA

ABSTRACT

Voice Assistants (VA) such as Amazon Alexa and Google Assistant are quickly and seamlessly integrating into people's daily lives. The increased reliance on VA services raises privacy concerns such as the leakage of private conversations and sensitive information. Privacy policies play an important role in addressing users' privacy concerns and informing them about the data collection, storage, and sharing practices. VA platforms (both Amazon Alexa and Google Assistant) allow third-party developers to build new voice-apps and publish them to app stores. Voice-app developers are required to provide privacy policies to disclose their apps' data practices. However, little is known whether these privacy policies are informative and trustworthy or not on emerging VA platforms. On the other hand, many users invoke voice-apps through voice and thus there exists a usability challenge for users to access these privacy policies.

In this paper, we conduct the first large-scale data analytics to systematically measure the effectiveness of privacy policies provided by voice-app developers on two mainstream VA platforms. We seek to understand the quality and usability issues of privacy policies provided by developers in the current app stores. We analyzed 64,720 Amazon Alexa skills and 16,002 Google Assistant actions. Our work also includes a user study to understand users' perspectives on privacy policies of voice-apps. Our findings reveal a worrisome reality of privacy policies in two mainstream voice-app stores. For the 17,952 skills and 9,955 actions that have privacy policies, there are many voice-apps with incorrect privacy policy URLs or broken links. We found that 1,755 Alexa skills and 192 Google actions provide a broken privacy policy URL. Amazon Alexa has more than 56% of skills with duplicate privacy policy URLs. While the Google Assistant platform has 9.0% of actions with duplicate privacy policy URLs. There are also skills/actions with inconsistency between the privacy policy and description. 6,047 Google actions do not have a privacy policy although they are required to provide one. Google and Amazon even have official voice-apps violating their own requirements regarding the privacy policy. We have reported our findings to both Amazon Alexa and Google Assistant teams, and received acknowledgments from both vendors.

*The first two authors contributed equally to this work.

ACSAC 2020, December 7–11, 2020, Austin, USA

ACM ISBN 978-1-4503-8858-0/20/12…$15.00
https://doi.org/10.1145/3427228.3427250

ACM Reference Format:
Song Liao, Christin Wilson, Long Cheng, Hongxin Hu, Huixing Deng. 2020. Measuring the Effectiveness of Privacy Policies for Voice Assistant Applications. In *Annual Computer Security Applications Conference (ACSAC 2020), December 7–11, 2020, Austin, USA.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3427228.3427250

1 INTRODUCTION

Voice Assistants (VA)[1] such as Amazon Alexa and Google Assistant have been seamlessly integrated into our daily life. An estimated 4.2 billion voice assistants are being used around the world in 2020, and the number is forecasted to reach 8.4 billion units by 2024 which is higher than the current world population [9]. VA handles a wide range of queries that humans are posing, *e.g.*, from ordering everyday items, managing bank accounts, controlling smart home devices to recommending clothing stores and new fashions. Despite the many convenient features, there is an increasing concern on privacy risks of VA users [16, 19, 24, 27, 33, 34, 37, 39].

Privacy and data protection laws are in place in most of the countries around the world to protect end users online. These compliance requirements are mostly satisfied by providing a transparent privacy policy by developers. Google was fined €50 million by a French data protection regulator after its privacy policy failed to comply with General Data Protection Regulation (EU GDPR) [6]. This fine was not for failing to provide a privacy policy but for not having a one that was good enough and failing to provide enough information to users. Researchers have shown that there are many discrepancies between mobile apps (*e.g.*, Android apps) and their privacy policies [17, 38, 41, 49], which may be either because of careless preparation by benign developers or an intentional deception by unscrupulous developers [43]. Such inconsistencies could lead to public enforcement actions by the Federal Trade Commission (FTC) or other regulatory agencies [50]. For example, FTC fined $800,000 against Path (a mobile app operator) because of an incomplete data practice disclosure in its privacy policy [13]. In another case, Snapchat transmitted geolocation information from users of its Android app, despite the privacy policy states that it did not track such information. In 2014, FTC launched a formal investigation requesting Snapchat to implement a comprehensive privacy program [14].

VA 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, respectively) and publish them to app stores. In order to comply with privacy regulations (such as COPPA [18]) and protect consumers' privacy, voice-app developers are required to provide privacy policies and notify users of their apps' data practices. Typically, a proper privacy policy

[1]Also known as voice personal assistants, smart home personal assistants or smart speakers.

is a document that should have answers to a minimum of three important questions [7]: 1) What information is being collected? 2) How this information is being used? and 3) What information is being shared? Third-party skills and actions are in very high number in the respective stores. Privacy policies provided by third-party developers could be diverse and poorly written, which results in more users ignoring the privacy policy and choosing to not read it. This also leads to users using a privacy-sensitive service without having a proper understanding of the data that is being collected from them and what the developer will do with it. On the other hand, the feature that makes VA devices like Amazon Echo and Google Assistant interesting is the ability to control them over the voice without the need of physically accessing them the convenience, it poses challenges on effective privac to enable users to make informed privacy decisions. The policy may be missing completely in the conversational unless users read it over VA's companion app on smart through the web.

In this work, we mainly investigate the following three questions (RQs):

- RQ1: What is the overall quality of privacy policies pro voice-app developers in different VA platforms? Do the informative and meaningful privacy policies as requir platforms from a user's perspective?
- RQ2: For a seemingly well-written privacy policy that vital information regarding the service provided to u we trust it or not? Can we detect inconsistent privacy p voice-apps?
- RQ3: What are VA users' perspectives on privacy p voice-apps? What is possibly a better usability choic users to make informed privacy decisions?

We conduct the first empirical analysis to measure the ness of privacy policies provided by voice-app developers Amazon Alexa and Google Assistant platforms. Such an not previously been reported. The major contributions and are summarized as follow[2].

- We analyze 64,720 Amazon Alexa skills and 16,002 Goo tant actions. We first check whether they have a priva For the 17,952 skills and 9,955 actions that have one, unfortunately, we find there are many voice-apps in app stores with incorrect privacy policy URLs or broken links. Surprisingly, Google and Amazon even have official voice-apps violating their own requirements regarding the privacy policy.
- We further analyze the privacy policy content to identify potential inconsistencies between policies and voice-apps. We develop a Natural Language Processing (NLP)-based approach to capture data practices from privacy policies and apps' descriptions. We then compare the data practices of a privacy policy against the app's description. We find there are privacy policies that are inconsistent with the corresponding voice-app descriptions. We also find voice-apps which are supposed to have a privacy policy but do not provide one.
- We conduct a user study with 91 participants to understand how users engage with privacy policies and their perspectives on VA's privacy policies, using the Amazon Mechanical Turk crowdsourcing platform. Our survey results suggest the need of VA platforms to take measures to improve the quality of privacy policies and provide effective privacy notices to make informed privacy decisions for VA users. We also briefly discuss possible solutions to improve the usability of privacy notices to VA users.

Responsible disclosure. We have reported our findings to both Amazon Alexa and Google Assistant teams. We have received acknowledgments from both vendors. It is worth mentioning that

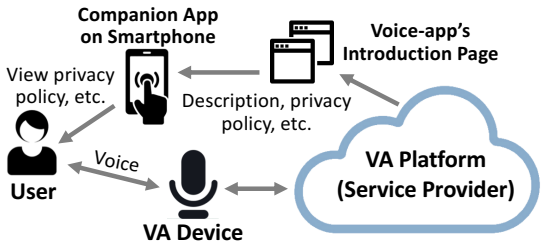

Figure 1: The privacy policy URL is provided in a voice-app's introduction page. A privacy policy can be accessed either over the VA's companion app or through the web.

Voice-app listing on the store. We mainly focus on two mainstream VA platforms, *i.e.*, Amazon Alexa and Google Assistant, both with conceptually similar architectures. These platforms allow third-party developers to publish their own voice-apps on VA stores. As shown in Fig. 1, a voice-app's introduction page that is shared by the developer on the store contains the app name, a detailed description, the category it belongs to, developer information, user rating and reviews, privacy policy link, and example voice commands which can be viewed by end users. The source code is not included in the submission and therefore is not available either to the certification teams of VA platforms or to end users. Users who enable a skill/action through the voice-app store may make their decisions based on the description. It explains the functionality and behavior of the voice-app and what users can expect from it. Some developers also mention the data that is required from users (*i.e.*, data practices) in the description.

VA platform's requirements on privacy policy. Application developers are often required to provide a privacy policy and notify users of their apps' privacy practices. VA platforms have different requirements regarding the privacy policies of voice-apps. Google Assistant requires every action to have a privacy policy provided on submission. Amazon Alexa requires only skills that collect personal information to mandatorily have a privacy policy. Both Amazon and Google prevent the submission of a voice-app for certification if their respective requirements are not met [1, 10]. In addition to the privacy policy URL, both platforms offer an option for developers to

[2] Accompanying materials of this work including the dataset, empirical evidences for inconsistent privacy policies, and tools are available at https://github.com/voice-assistant-research/voice-assistant.

provide a URL for the terms of use as well. These URLs, if provided by the developers, are made available along with the voice-app's listing on the store.

Requirements on specific content in privacy policies. Google has a "Privacy Policy Guidance" page [7] in their documentation for action developers. The guide explains what Google's n expectation is for a privacy policy document. Accordin guide, the privacy disclosures included in the policy should prehensive, accurate and easy to understand for the us privacy policy should disclose all the information that a collects through all the interfaces including the data th lected automatically. How the collected information is u who and when the collected information is shared with sl specified. Google rejects an action if developers do not pr even misspell) the action name, company name, or develop in the privacy policy. The link should be valid and should a public document viewable by everyone. Amazon Alexa provide a guideline for the privacy policy content in the documentation.

Voice-app enablement. VA users enable official (*i.e.*, d by VA platforms) or third-party voice-apps to expand the f ality of their devices. For the Amazon Alexa platform, skil enabled by saying a simple command through voice or b it from the VA's companion app on smartphone (*i.e.*, the Amazon Alexa app on Android/iOS used for managing VA devices). A skill for which the developer has requested permission to access user's data sends a permission request to VA's companion app on smartphone during enablement. The other voice-apps are directly enabled. Google Assistant does not require users to enable an action before using it, where users can directly say the invocation command to invoke an action. As illustrated in Fig. 1, a privacy policy can be accessed either over the VA companion app or through the web. However, it is not accessible through the VA devices through the conversational interface. VA platforms do not require end users to accept a privacy policy or the terms of use of a voice-app before enabling it on their devices. It is left for the users to decide whether to go through the privacy policy of the voice-app they use or not.

2.2 Challenges on privacy policy analysis

Existing privacy policy analysis on smartphone platforms [17, 38, 41, 43, 49, 50] typically conduct static code analysis to analyze potential inconsistencies between an app's privacy policy and its runtime behavior. Unlike smartphone app (*e.g.*, Android or iOS) platforms, the source code of voice-apps in the Amazon Alexa and Google Assistant platforms are not publicly available. A voice-app is hosted in a server selected by its developer and only the developer has access to it. As far as we know, the source code is not available even to the VA platform's certification teams. This limits the extent of our privacy analysis since we do not have the actual code of voice-apps to find more inconsistencies with the privacy policies provided. The only useful information that we have about a voice-app is the description that is provided by the developer. Descriptions do not have a minimum character count and developers can add a single line description or a longer description explaining all functionalities and other relevant information. Regardless, due to the unavailability of other options, we use the voice-app descriptions for our analysis to detect problematic privacy policies. For this reason, our results on the inconsistency checking of privacy policies (in Sec. 3.3) are not focused on the exact number of mismatches and errors but on the existence of problems potentially affecting the overall user experience.

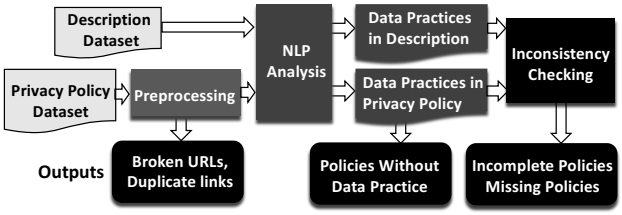

Figure 2: Processing pipeline of our privacy policy analysis.

In this section, we first present an overview of our approach, and then detail the major modules including data collection process (Sec. 3.1), capturing data practices based on the NLP analysis (Sec. 3.2), and inconsistency checking (Sec. 3.3). We seek to understand whether developers provide informative and meaningful privacy policies as required by VA platforms. Fig. 2 illustrates the processing pipeline of our privacy policy analysis. As previously mentioned, each skill/action's listing page on the store contains a description and a privacy policy link (URL). We first collect all these webpages, and pre-process them to identify high-level issues such as broken URLs and duplicate URLs. Then, we conduct an NLP based analysis to capture data practices provided in privacy policies and descriptions. We seek to identify three types of problematic privacy policies: i) without any data practice; ii) incomplete policies (*e.g.*, a skill's privacy policy lacks data collection information but it has been mentioned in the skill's description); and iii) missing policies (*e.g.*, a skill without a privacy policy but requires one due to its data collection practices).

3.1 Data collection

We built a crawler to collect a voice-app's id, name, developer information, description and privacy policy link from the Amazon Alexa's skills store and Google Assistant's actions store. There were several issues for crawling introduction pages of voice-apps. First, for the skills store, 23 categories of skills are listed but these are not mutually exclusive. For example, the category "communication" is a subcategory in the "social" category and the category "home services" is a subcategory in the "lifestyle" category. Some skills are classified and listed in multiple categories. We need to remove duplicates during the data collection. Second, Alexa's skills store only provides up to 400 pages per category, and each page contains 16 skills. Though the Amazon Alexa claimed there are over 100,000 skills on its skills store, we were able to crawl only 64,720 unique skills as of March 2020. Third, the Google Assistant's actions store lists actions in pages that dynamically load more actions when users reach the end of the page. We were unable to directly use the crawler to automatically get information about all the actions. To address this issue, we used the Selenium WebDriver (SWD) [12] to load the dynamic content of a webpage. Finally, we crawled 16,002 actions belonging to 18 categories from the Google actions store.

The total numbers of skills and actions by category we collected are listed in Table 10 and Table 11 in Appendix.

Another issue was to obtain the privacy policy content. Given the privacy policy links, we observed that there are five types of policy pages: i) normal html pages; ii) pdf pages; iii) Google Doc and Google Drive documents; iv) txt files; and v) other types of files (*e.g.*, doc, docx or rtf). For normal html pages, we used the webdriver [12] tool to collect the webpage content when they are opened. For the other types of pages, we downloaded these files and then extracted the content from them. Finally, we converted all the privacy policies in different formats to the txt format.

Privacy Policy Dataset. We collected 64,720 unique skills under 21 categories from the Amazon Alexa's skills store and 17,952 of these skills provide privacy policy links. Among the 16,002 Google actions that we collected, 9,955 have privacy policy links.

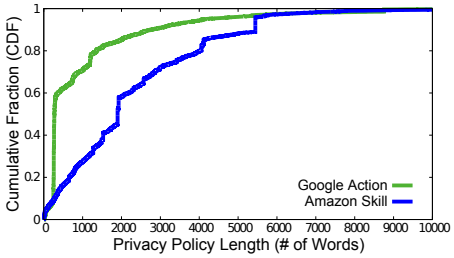

Figure 3: Length of a privacy policy.

For each skill/action with a valid policy link, we calculated the number of words in the document. Fig. 3 shows the cumulative distribution function of the privacy policy length. The average length is 2,336 words for Alexa skills and 1,043 words for Google actions, respectively. We noticed that a large number (4,572) of actions use the same template, and thus their privacy policies have a similar length (around 260 words). We also observed many very short privacy policies which are not informative. An example is the Google action "Mister Baavlo" which says "*We do not store any of your data*" but does not mention what data it collects. Examples of short privacy policies are listed in Table 1.

Voice-app name	Skill/Action	Privacy policy content
Passive income tips	Skill	"This is just a sample privacy policy link, You can use this url, If you do not have it."
Activity Book	Skill	"This skill does not collect or save any personal information."
BestDateCalendar	Skill	It directs to Google home page
Story Time	Skill (Kids)	" No information is collected during the use of Story Time"
KidsBrushYourTeethSong	Skill (Kids)	"Privacy Policy" (no content)
Mister baavlo	Action	"We do not store any of your data"
Sanskrit names for yoga poses	Action	"Google Docs: You need permission to access this published document."

Table 1: Examples of short privacy policies.

Description Dataset. Description of a voice-app is intended to introduce the voice-app to end users with information regarding its functionality and other relevant information. It may also contain data practices (*e.g.*, the data required to be collected to achieve a functionality) of the voice-app. We collected voice-app descriptions and used them as baselines to detect potentially inconsistent privacy policies. In our dataset, all skills/actions come with descriptions.

3.2 Capturing data practices

In order to automatically capture data practices in privacy policies and descriptions of voice-apps, we develop a keyword-based approach using NLP. However, we want to emphasize that we do not claim to resolve challenges for comprehensively extracting data practices (*i.e.*, data collection, sharing, and storing) from natural language policies. Instead, we mainly focus on obtaining empirical evidences of problematic privacy policies using a simple and accurate (*i.e.*, in terms of the true positive) approach. We discuss the limitation of our approach in Sec. 6.

Verb set related to data practices. Researchers in [20, 43] have summarized four types of verbs commonly used in privacy policies: `Collect`, `Use`, `Retain` and `Disclose`. Each type contains semantically similar verbs in terms of the functionality. `Collect` means an app would collect, gather, or acquire data from users; `Use` indicates an app would use or process data; `Retain` means storing user data; and `Disclose` indicates an app would share or transfer data to another party.

Verb Set	Access, Ask, Assign, Collect, Create, Enter, Gather, Import, Obtain, Observe, Organize, Provide, Receive, Request, Share, Use, Include, Integrate, Monitor, Process, See, Utilize, Retain, Cache, Delete, Erase, Keep, Remove, Store, Transfer, Communicate, Disclose, Reveal, Sell, Send, Update, View, Need, Require, Save
Noun Set	Address, Name, Email, Phone, Birthday, Age, Gender, Location, Data, Contact, Phonebook, SMS, Call, Profession, Income, Information

Table 2: Keyword dictionary related to data practices.

Noun set related to data practices. From Amazon's skill permission list [5] and Amazon Developer Services Agreement [3], we manually collected a dictionary of 16 nouns related to data practices. Table 2 lists a dictionary with 40 verbs and 16 nouns that we used in our privacy policy analysis.

Phrases extraction. We first parsed a privacy policy into sentences. We used the SpaCy library [8] to analyze each sentence, and obtained the attribute for each word. SpaCy can effectively find the straight correlation between a noun and a verb and ignore other words in a sentence. We identified three types of basic phrases:

- noun (subject) + verb, *e.g.*, "Alexa (will) tell" or "email (is) required"
- verb + noun (object), *e.g.*, "send (a) message"
- verb + noun (object) + noun + noun, *e.g.*, "tell (you) (the) name (of) meeting (on) (your) calendar"

Next, we combined two basic phrases to generate a longer phrase if they share the same verb. The combined phrase would follow patterns: "subject+verb+object" or "subject+is+passive verb". For example, for a sentence "Alexa skill will quickly tell you the name and time of the next meeting on your Outlook calendar", we obtained the phrase "Alexa skill tell name, meeting, calendar".

Identifying data practices. Given all phrases extracted from the privacy policy and description, we used the verb and noun sets in Table 2 to identify data practice phrases. For each phrase, we obtained the noun with the related verb and checked whether they are in our keyword dictionary. For example, our tool identified privacy policies of 680 skills in the Amazon Alexa platform and 403 actions in the Google Assistant platform having zero data practice. We manually analyzed the results and it shows that our analysis

tool achieves an accuracy of 85% on average for these skills and actions (details are in Sec. 4.2.2).

3.3 Inconsistency checking

With description phrases and privacy policy phrases for each voice-app, we checked any potential inconsistency between them. First, if the data practice phrases in a description are not semantically similar to any data practice phrase in the corresponding privacy policy, we consider this privacy policy to be incomplete. For example, the description of skill "Thought Leaders" mentions "Permission required: Customer's Full name, Customer's Email Address, Customer's Phone number", but none of them are mentioned in its privacy policy. We consider it an incomplete privacy policy. To measure the semantic similarity of two data practice phrases, we used the similarity measurement based on the word2vec model (a technique for natural language processing) provided by SpaCy [8]. We set the similarity threshold to 0.9 in our analysis. A higher threshold value means stricter rules will be applied in the semantic similarity measurement. As a result, we obtained a relatively large set of the potentially incomplete privacy policies. Considering the limitation of NLP techniques, we then conducted a manual analysis to identify the true incomplete privacy policies.

Second, since the Amazon Alexa platform only requires skills that collect personal information to provide a privacy policy, we detected whether a privacy policy of an Alexa skill is missing although it is required. If the description mentions that a skill collects some data but the skill has no privacy policy, we consider that the skill lacks a privacy policy. For example, a skill "Heritage Flag Color" mentions "The device location is required" in its description. But the developer doesn't provide a privacy policy. Note that it only reflects an inconsistency between the privacy policy and description. To validate whether the skill really collects the location information or not, we need to conduct a dynamic testing to explore the skill's runtime behavior (details are in Sec. 4.2.3).

4 MAJOR FINDINGS

In this section, we discuss the major findings from our analysis of the privacy policies available in the stores of both Amazon Alexa and Google Assistant. We first present high-level issues such as broken and incorrect privacy policy URLs, duplicate privacy policy links, and issues in Google and Amazon's official voice-apps. Then, we conduct a content analysis of privacy policies, and discuss the issues such as zero data practice and inconsistency in privacy policies. In addition, we discuss usability issues of privacy policies for voice-apps. We back our findings with representative examples that we found from the app stores during our analysis.

4.1 High-level issues

4.1.1 Not all voice-apps have a privacy policy URL. Both Google and Amazon have taken different approaches when it comes to the requirement of a privacy policy for each voice-app available to users. While Google has made it mandatory for developers to provide a privacy policy along with each action, Amazon is more lenient and makes it a requirement only for skills that declare that they collect personal information through the skill. On analyzing the stores, we have noticed irregularities concerning this, as illustrated in Table 3. Out of the 16,002 actions we collected from the Google action directory, 9,955 have privacy policies provided which means that 38% of the actions do not have a privacy policy provided. While it is not possible to submit an action for certification without including a privacy policy URL, it is puzzling how these actions are available in the store without providing one. Out of these 6,047 actions that do not have privacy policies, only 7 actions provide the developer information, and only 32 actions were rated by at least one user. Interestingly, these actions are all from three categories, "Food & Drink", "Music & Audio", and "News & Magazines", where only 831 actions (12%) out of 6,877 actions provide a privacy policy. 1,949 actions in the "Food & Drink" category are named like "Recipe Results from {Store Name}", *e.g.*, "Recipe Results from Wendy's". 1,791 actions from the "News & Magazines" category are named as "News Results from {Media Name}", such as "News Results from CNN".

	Alexa skills		Google actions	
	Total #	Percentage	Total #	Percentage
Without privacy policy	46,768	72%	6,047	38%
Valid privacy policy URL	16,197	25%	9,763	61%
Broken privacy policy URL	1,755	3%	192	1%

Table 3: Statistics of privacy policies on two VA platforms.

In the case of Alexa skills, as shown in Table 3, only 17,952 (28%) skills have a privacy policy out of the 64,720 skills we collected (*i.e.*, 46,768 skills without a privacy policy). It is partially because of the lenient skill certification on the Amazon Alexa platform. After conducting further experiments on the skill certification, we have understood that even if a skill collects personal information, the developer can choose to not declare it during the certification stage and bypass the privacy policy requirement [23]. This is achieved by collecting personal information through the conversational interface (*e.g.*, asking users' names). Even though this data collection is prohibited, the certification system of Amazon Alexa doesn't reject such skills. As a result, developers may choose to not provide a privacy policy. Amazon only requires skills that collect personal data to provide a privacy policy, and thus not all these 46,768 skills require a privacy policy. In Sec. 4.2.4, we identify skills that potentially lack a required privacy policy.

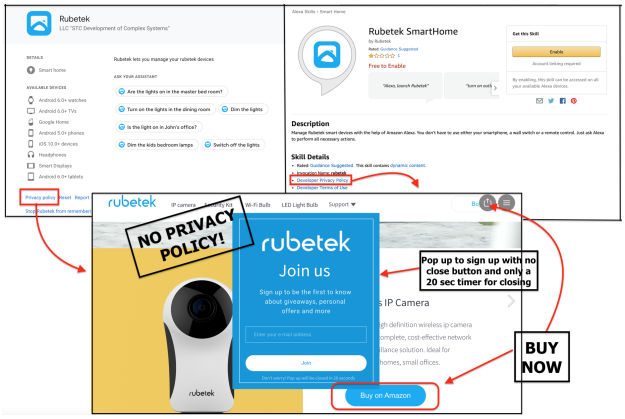

Figure 4: Landing page of the privacy policy provided with the Google action and Alexa skill developed by Rubetek.

4.1.2 Broken links and incorrect URLs. For those actions and skills that have provided a privacy policy URL, not every URL leads to

the page containing a privacy policy. Through our experiments, we found 192 Google actions and 1,755 Alexa skills that have provided broken privacy policy URLs, as shown in Table 3. There are also URLs which lead to other developer's privacy policies. An example for this is the skill "NORAD Tracks Santa" by NORAD which provides a privacy policy URL that links to Amazon's privacy policy page instead of 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 homepage which promotes its products, as shown in Fig. 4, rather than linking to the privacy policy page. Sec. 4.2 presents our content analysis of privacy policies, which provides more details about the voice-apps with incorrect privacy policy URLs.

4.1.3 Duplicate URLs. We found a substantial portion of privacy policies share same URLs. In particular, Amazon Alexa has more than 56% of skills with duplicate privacy policy URLs. Fig. 5 shows the prevalence of duplicate privacy policy URLs in both platforms. Out of the 17,952 Amazon skills with privacy policies, 7,828 skills have a unique privacy policy URL. The other 10,124 skills (56.4%) share 1,206 different privacy policy URLs. Out of these, 1,783 skills (9.9%) have provided the same link (https://getstoryline.com/public/privacy.html) as their privacy policy URLs. Note that these 1,783 skills are not from the same developer which indicates that the privacy policy is irrelevant to these skills. Here the irrelevance means that the privacy policy provided in the URL was not written specifically for the developer or the voice-app (*e.g.*, without including the voice-app name, company name, or developer email).

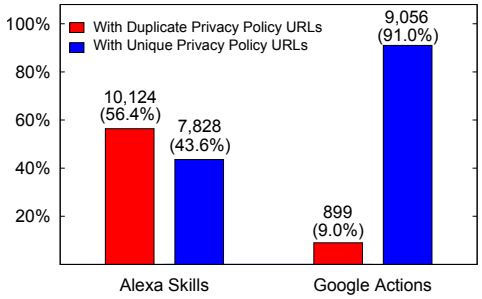

Figure 5: Duplicate privacy policy URLs in two platforms.

Table 4 lists the most common privacy policy URLs in the Amazon Alexa and Google Assistant platforms. The issue of duplicate URLs is more serious on the Amazon Alexa platform. The top three duplicate URLs are shared by 3,205 skills, constituting 17.8% of the total skills that have a privacy policy. As shown in Fig. 5, the Google Assistant platform has 9.0% of actions with duplicate privacy policy URLs. 9,056 out of 9,955 actions have a unique privacy policy. The other 899 actions share 204 different privacy policy URLs.

To understand why there exists such a large number of voice-apps with duplicate privacy policy URLs especially on the Amazon Alexa platform, we further examined the developer information of these voice-apps. Our intuition is that developers who published multiple voice-apps may use the same privacy policy URLs. We found that for the developers who developed more than one skill, 77% of their skills use duplicate privacy policy URLs. Table 5 lists the top 5 developers who published the most skills with a privacy policy on the Amazon Alexa platform. As illustrated in the table, 2,064 out of 2,069 skills (99.8%) use duplicate privacy policy URLs. Obviously, the content of these privacy policy URLs are not skill-specific, and users may skip reading the privacy policy although it is provided. A serious problem happens if such a privacy policy link is broken, which results in hundreds of skills being affected. For example, we found a broken link "https://www.freshdigitalgroup.com/privacy-policy-for-bots" (shown in Table 5). There are 217 skills using this link, and thus all their privacy policies become inaccessible. As to the Google actions, we also observed the similar issue. Although Google requires that a privacy policy must include one of the following: action name, company name or developer email, there are developers using a general privacy policy with the company name or email for all their actions. For the developers who published more than one action, 31% of actions have duplicate privacy policy URLs. For the top 10 developers who published the most actions, 86% of their actions use a duplicate privacy policy link.

Platform	Duplicate privacy policy URLs	Total #	Percentage
Amazon	https://getstoryline.com/public/privacy.html	1,783	9.9%
	https://corp.patch.com/privacy	1,012	5.6%
	https://cir.st/privacy-policy	410	2.3%
Google	https://policies.google.com/privacy	97	1.0%
	https://xappmedia.com/privacy-policy/	55	0.6%
	https://docs.google.com/document/d/1yHGyixM2n6n32VoefxZIk8fxEMg0Lb-ELFm_tSqHPF0/edit?usp=sharing	40	0.4%

Table 4: The most common duplicate privacy policy URLs.

Developer	# of skills developed	Skills with duplicate URLs	Top duplicate URLs used by the developer
Patch.com	1,012	1,012	http://corp.patch.com/privacy
Radio.co	295	292	http://www.lottostrategies.com/script/showpage/1001029/b/privacy_policy.html
Tinbu LLC	264	263	http://spokenlayer.com/privacy
FreshDigitalGroup	259	258	https://www.freshdigitalgroup.com/privacy-policy-for-bots
Witlingo	239	239	http://www.witlingo.com/privacy-policy

Table 5: Top 5 developers that published the most skills with a privacy policy on Amazon Alexa platform.

4.1.4 There are Google and Amazon's official voice-apps violating their own requirements. We found two official "Weather" skills on Amazon Alexa's skills store, and one of them asks for user's location according to the description but it doesn't provide a privacy policy. Fig. 6 shows the "Weather" skill developed by Amazon with the product ID "B071Z29JLY". This skill may be automatically enabled and available on all Alexa devices since it is a built-in skill. This example demonstrates that Amazon Alexa violates its own requirement by publishing voice-apps capable of collecting personal information without providing a privacy policy.

We collected 98 Amazon Alexa official skills (*i.e.*, developed by Amazon, Amazon Alexa Devs, and Amazon Education Consumer Team), out of which 59 skills come with privacy policy URLs (but all are duplicate URLs). Among these privacy policy links, 30 links point to the general Amazon privacy notice and 6 links are the AWS (Amazon Web Services) privacy notice, Amazon payment privacy or Alexa term of use. Surprisingly, 23 privacy policy links are

Figure 6: An official skill lacks a privacy policy. Even though it collects the user's location according to the description, no privacy policy is provided.

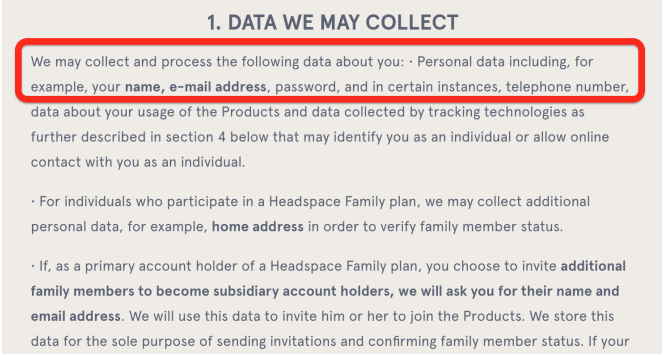

Figure 7: Privacy policy URL provided with a kids skill "Headspace Bedtime Story" disclosing the collection of personal data which is prohibited according to Amazon Alexa's privacy requirements [2].

totally unrelated to privacy notice, in which 17 links are Amazon homepage and 6 links are pages about insurance. In the Google Assistant's actions store, we found 92 official actions developed by Google. All the 92 actions provide a privacy policy link, but they point to two different Google Privacy Policy pages and both of them are general privacy policies. Google requires that every action should have an app-specific privacy policy provided by developers on submission (including the action name, company name, or developer email in the privacy policy). However, our analysis reveals that this requirement had not been enforced in a proper manner at the submission time of these 92 actions. Note that in our testing, we have submitted multiple actions purposely violating this requirement (*e.g.*, without providing the action name or providing a wrong name). Our submissions got rejected due to the above reason.

4.2 Content analysis of privacy policies

4.2.1 Irrelevance to specific voice-app. It is important to cover all aspects of a service's data practices in the privacy policy. The contradiction is providing these data practices for a service that is not capable of doing any of the data collections mentioned in the privacy policy (we acknowledge that there is a legal side to this problem so that developers may mention all possible data practices in a privacy policy). This is especially evident in the Amazon Alexa's skills store where most skills have a privacy policy that is common across all services that the developers provide. These policies do not clearly define what data practices the skill is capable of. Some of these privacy policies do not even mention the Alexa skill or Google action as a service and state that it is the privacy policy of a specific service such as the website domain. We analyzed whether a voice-app mentions the app name in its privacy policy. There are only 3,233 skills out of 17,952 skills (around 18%) mentioning skills' names in their privacy policies. For Google actions, 5,297 out of 9,955 actions (around 53%) mention action names in their privacy policies.

There were also privacy policies provided for kids skills which mention that the service is not intended to be used by children and also that the service can collect some form of personal information, which is not allowed for skills in the kids category according to Amazon Alexa's privacy requirements [2]. Fig. 7 shows an example where the privacy policy URL provided with a kids skill disclosing the collection of personal data. In addition, we found 137 skills in the Amazon Alexa's kids category whose privacy policies mention data collection is involved. But they just provide a general privacy policy. All these skills potentially violate Amazon Alexa's privacy requirements on kids skills, which state that any personal information is not supposed to be collected from kids.

4.2.2 Zero data practice. We applied our method described in Sec. 3.2 to capture data practices in each privacy policy. Fig. 8 illustrates the cumulative distribution function of data practices we identified using our privacy policy dataset. For these privacy policies with data practices, the average amount is 26.1 in Amazon Alexa and 13.4 in Google Assistant, respectively. The maximum number of data practices in a privacy policy is 428, which is likely a general privacy policy rather than an app-specific one.

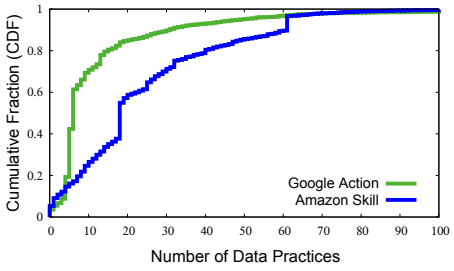

Figure 8: Number of data practices in a privacy policy.

Our tool detected 680 Alexa skills and 403 Google actions having privacy policies but with zero data practice. After manually checking these privacy policies, our method achieved an accuracy of 85% with 87 false positives in the 680 Alexa skills and 76 false positives in the 403 Google actions, respectively. We found that most cases were because of failures of the crawler to correctly obtain privacy policies. For example, when privacy policies are embedded in the web framework, we could not get the correct content while crawling the privacy policy webpage.

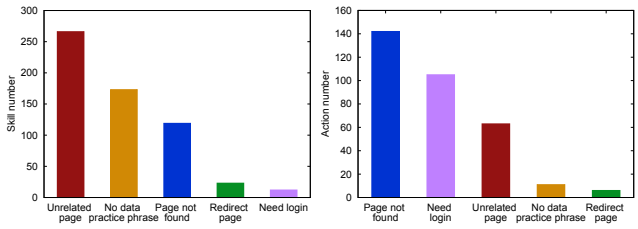

(a) Alexa skills with no data practice (b) Google actions with no data practice

Figure 9: Different issues of privacy policies that have zero data practice in two VA platforms.

In particular, 593 privacy policies provided with Alexa skills have zero data practices (confirmed by our manual analysis). Fig. 9 shows the breakdown of different issues of these privacy policies. 266 privacy policy URLs lead to totally unrelated pages which have advertisements and shopping options. 173 URLs lead to an actual privacy policy page but has no data practices mentioned. 119 URLs lead to an actual website domain but the link is not found. These too can be considered as broken links. 23 URLs lead to a page where the actual link to the privacy policy does exist but will be redirected to some other pages. Another 12 skills need logins to access to the documents.

After the manual analysis, we found 327 Google actions having privacy policies with zero data practice, as shown in Fig. 9. 142 URLs lead to a page that is not found. 63 URLs lead to unrelated links with shopping options and product advertisements. 11 URLs are privacy policies but with no data practice. 6 URLs lead to a page containing the link to the actual privacy policy. In addition, 105 actions provide their privacy policy as a Google Doc which does not have the correct permissions set resulting in users not being able to access it. Obviously, they violate Google's restriction "the link should be a public document viewable by everyone".

4.2.3 Inconsistency between the privacy policy and description. As mentioned in Sec. 3.3, we set a large threshold value (*i.e.*, 0.9) in our semantic similarity measurement so as to obtain more potentially incomplete privacy policies (although there exist false positives). Using the method presented in Sec. 3.3, we found 102 skills with possible incomplete privacy policies. Through a manual analysis, we finally identified 44 Alexa skills that have a privacy policy which is inconsistent with the corresponding description. These skills describe the collection of personal data in the description but these data practices are not mentioned in the privacy policy provided. The consequence of such occurrences is that the users are not informed about what happens to their information and who it is shared with. Among the 44 skills, 17 skills ask for address or location; 7 skills request email/ account/ password; name is asked by 7 skills and 4 skills require the birthday information; and the other skills ask for phone number, contact, gender or health related data. Fig. 10 shows an example, where the skill "Running Outfit Advisor" mentions collecting the gender information in the description, but does not mention this data practice in its privacy policy. In another case, the description of the skill "Record - Journal - Things to Do Calendar" describes the collection of personal information like the address of the user. The description has the following line: "*Device Address. Your device address will be used to provide responses with events local to your area.*" In the skill's privacy policy, the data practices are not disclosed clearly enough but only says "we will collect personal information by lawful". We treated this kind of privacy policy as inconsistent (incomplete) privacy policy since it fails to give a clear idea about its data practices. Table 6 shows the list of these skills with inconsistency between the privacy policy and description. We also identified 2 Google actions: "Money manager" asks for income data and "Joke Generator" asks for user names. But their privacy policies do not mention such data practices.

To validate the above results, we manually checked whether these skills really collect data as they claimed in their descriptions but with an incomplete privacy policy. We could confirm that 32 skills involve data collection. We found 25 skills (highlighted with the light gray background in Table 6) using the skill's built-in feature to ask for username, email and address when users invoke skills at the first time (permissions will be taken from users when skills are first enabled). There are also 7 skills (highlighted with the dark gray background in Table 6) asking for user data through the voice channel. We also found that 6 skills didn't work properly. For the other 6 skills, we didn't find any data collection during our testing.

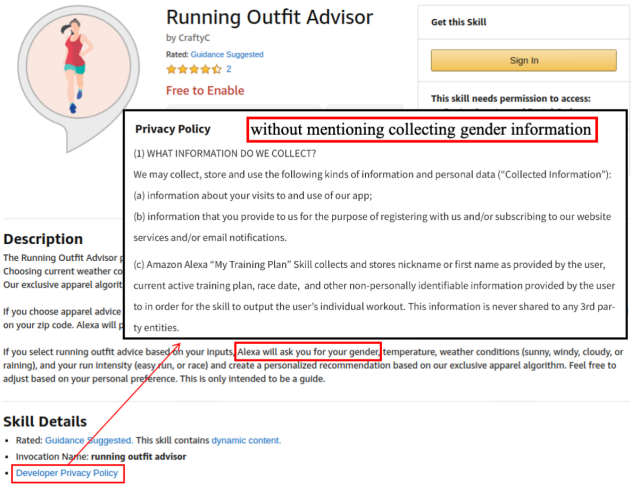

Figure 10: "Running Outfit Advisor" skill mentions collecting the gender information in the description, but does not mention this data practice in its privacy policy.

4.2.4 Missing required privacy policies. In Sec. 4.1.1, we have shown 6,047 Google actions do not have a privacy policy provided, which violates its own restriction "Google require all actions to post a link to their privacy policy in the directory". Here we focus on Amazon Alexa skills and identify cases with missing required privacy policies using our tool.

To collect user's personal data for use within the voice-apps, developers can use the built-in feature of collecting the personal information directly from their Amazon account after taking permissions from the user. This permission is taken from the user when the skill is first enabled. While this is appropriate and respect the users privacy, there is another channel that can be misused for collecting personal information. A developer can develop a skill to

Arbonne My Office , Best Roomies , CitySpark Events , Conway Daily Sun Calendar , FortiRecorder , K5 , WP6 , garage control , ISS: Distance From Me? , Kotipizza , Laconia Daily Sun Calendar , Mailbox Assistant , Maui Time Calendar , My Air Quality , Natural Hazards , Novant Health , Portland Phoenix Calendar , Record-Journal - Things to Do Calendar , SkyHome , SkyView Academy , Thought Leaders , Trivia Quest , The Transit Oracle (Bus Predictions for SF Muni) , What Should I Wear , what's nearby , Crush Calculator , Find My Phone , Flu Season , Hal9000 , Running Clothes , Running Outfit Advisor , walk cake , Arm My Guardzilla, Ash Timber Flooring, Cake Walk, GINA Talk, group messenger, Happy birthday, Home Workout Exercise Video Fitness 7 Day Videos, hugOne, Kamakshi Cloud's GPS Finder, Neighbor Knocker, OMS Customer Care, Trip Tracker,

Table 6: Skills with incomplete privacy policies as of May 2020. We manually tested these skills, and confirmed 32 skills with incomplete privacy policies (highlighted with the light/dark gray background).

ask for the personal information from the user through the conversational interface. Both Amazon and Google prohibit the use of conversational interface to collect personal data. But, in the case of Amazon, this is not strictly enforced in the vetting process. By collecting personal information in this manner, the developer can avoid adding a privacy policy URL to the skill's distribution requirements. This is possible because Amazon requires only skills that publicly declare that they collect personal information to mandatorily have a privacy policy. The developer can easily bypass this requirement by lying about not collecting personal information [23].

Data	# of Skills	Skills names
Name	10	First Name Analysis , Mr. Tongue Twister (Kids) , My daily task , Name My Grandkids , Social Network , Uncle Tony (Kids) , who's right , Haircut Scheduler, insurance service, LOVE CALCULATOR
Location	6	Doctor Locator , Heritage Flag Color , World Time , Lapel Athletics, OC Transpo, Weather
Gender	1	Interactive Bed Time Story (Kids)
Age	1	cadmiumgreen , bright smile
Birthday	1	Cake Walk
Ip Address	1	Network-Assistant

Table 7: Skills with no privacy policies despite mentioning the collection of users data in their descriptions.

Fig. 11 illustrates an example where the skill "Name My Grandkids" includes in its description that it asks the users for personal information and stores it for future use. In another case, the skill "Lapel Athletics" requires the device location according to its description. But both these skills do not provide a privacy policy. Table 7 lists skills which are supposed to have a privacy policy but do not provide one. To validate the results, we manually tested these skills. We found that none of them use the VA's built-in feature to collect data when skills are invoked at the first time, and 12 skills (highlighted with the gray background in Table 7) ask for user data

Figure 11: Although the skill description mentions collection of personal information, no privacy policy is provided.

through the voice channel. There were 4 skills not working. For the other 4 skills, we didn't find any data collection during our testing.

4.2.5 Cross-platform inconsistency. For a few voice-apps that are present on both Alexa and Google platforms, we found that the privacy policies provided with each are not the same. Comparing Google actions and Alexa skills, we found that 82 voice-apps which are present on both the platforms have differences in the privacy policy links provided despite the name, the descriptions and the developer name being the same. 40 of these pairs have different privacy policies links all together. For example, the skill "Did Thanos Kill Me" uses a duplicate privacy policy link "https://getstoryline.com/public/privacy.html" (shown in Table 4), but the corresponding Google action version provides a specific privacy policy. Since Google requires every action to provide a privacy policy link, developers provide one with the Google action but may choose to not provide one along with the Alexa skill since the Alexa platform doesn't have this requirement such as skill and action "Website Reader". We found 42 such pairs of skills/actions, where a skill doesn't have a privacy policy while the Google action version has one. The detailed voice-app names are listed in Table 12 in Appendix.

4.2.6 Potential noncompliance with legal regulations. We observed skills that collect personal information being published on the Amazon Alexa skills store under the kids category without providing a privacy policy. For example, Table 7 lists 3 skills (which are marked with "Kids" in the table) in the kids category lacking a privacy policy. This is not compliant with the COPPA regulations which require every developer collecting personal information from children to follow certain rules. Providing a privacy policy with accurate information about the data being collected and what it is used for is one of the main requirements. The objective is to clearly let the parents know about what personal information can be collected by the skill from their children. Health related information can also be collected by a skill through the conversational interface without providing a privacy policy even though only the user can decide whether to provide it or not. But skills having the capability to do so might be a violation of the HIPAA (Health Insurance Portability and Accountability Act) regulation. CalOPPA (California Online Privacy Protection Act) requires developers to provide a privacy policy that states exactly what data can be collected from users. In

Sec. 4.2.1, we found that 137 kids skills provide general information without providing specifics on what personal data they actually collect. These voice-apps and their privacy policies may not be in compliance with the legal regulations.

4.3 Usability issues

4.3.1 Lengthy privacy policies. One of the main problems associated with privacy policies regardless of the type of service it is provided with is the length of the privacy policy document. Most developers write long policies that decreases the interest that a user has in reading it. From the analysis of the privacy policies in our datasets, as shown in Fig. 3, we observed that 45% of the privacy policies have more than 1,500 words. Being a legal document, it takes an average of 12 mins to read 1,500 words. This makes the privacy policy hard to read for the users and almost impossible to be read out through voice. The privacy policies of Google and Amazon themselves have more than 4,300 words each. The average number of words in a privacy policy provided along with Alexa skills is 2,336 and that of Google actions is 1,043. This results in users frequently skipping reading the privacy policy even if it is provided. The participants of the user study we conducted as shown in Sec. 8 complained about the length of the privacy policy being the major reason for them not reading a privacy policy.

4.3.2 Hard to access. The constrained interfaces on VA devices pose challenges on effective privacy notices. According to the current architecture of Amazon Alexa and the Google Assistant, the privacy policy is not available directly through VA devices used at home like the Amazon Echo and the Google Home. No prompt is delivered either during any part of a user's interaction with the voice-app that requests the user to take a look at the privacy policy. If at all the user wants to view the privacy policy, he/she has to either find the voice-app listing on the store webpage or check the VA's companion app on smartphone, and find a voice-app's privacy policy URL provided in the listing. The permissions set by the developer to collect personal information from users is shown as a prompt in the smartphone companion app to the user while enabling the voice-app. But as mentioned in Sec. 4.2.4, developers do not necessarily have to take permission from user and can instead collect it during the conversation. We discuss solutions to improve the usability of privacy notice for voice-apps in Sec. 6.3.

5 USER STUDY

We conducted a preliminary user study using the Amazon Mechanical Turk crowdsourcing platform [4], and our study has received an IRB approval. Different from prior user studies that focus on understanding security and privacy concerns of VA devices [15, 26, 33, 39, 45], we aimed to understand how users engage with privacy policies and their perspectives on them. We looked for the frequency of checking the privacy policies and any issues the users might have encountered with them. Our participants were MTurk workers who reside in USA, having a HIT (Human Intelligence Tasks) acceptance rate greater than 98 and have at least 500 HITs approved prior to this study. These filters were added to reduce the amount of junk data that we may have collected. All participants were initially presented with a consent form approved by the IRB office. Participants who did not consent to the form were denied to proceed with the study. We rewarded $0.2 to each participant who completed the study.

Question	Response	% of users
Are you aware of the privacy policies of your skills/actions?	Yes	48
	No	52
How often do you read the privacy policy of a skill/action?	Rarely	73
	Half the time	11
	Most of the time	16
Do you read the privacy policy from the skill/action's webpage/Alexa app?	No	66
	Yes	34
Do you know what personal data the skills/actions you use are capable of collecting from you?	No	47
	Maybe	21
	Yes	32
Do you read the privacy policy before using a new skill / action?	No	79
	Maybe	7
	Yes	14
Do you read the privacy policy before enabling a kid's skill / action?	No	75
	Maybe	7
	Yes	18

Table 8: Survey responses.

We had a total of 98 participants who took part in our study. We had included a question to ensure that the user is answering the survey authentically. Based on the responses to this question, we rejected the answers of 7 participants. Our results for the user study were thus based on responses from 91 users. The participants are either Amazon Alexa users or Google Assistant users. We didn't include participants who use other assistants like Siri and Cortana in our study. We had 66 participants who are Alexa users and 25 participants who use Google assistant at home.

Table 8 shows the survey responses. When asked about whether they are aware of the privacy policies of the voice-apps they use, about 48% of the participants claimed that they are aware of it. But when asked about how often they actually read the privacy policy provided by the developer, 73% responded with "rarely". 11% responded that they read it half the time. 34% of our participants said that they use the VA's companion app on smartphone or the web browser to read the privacy policy while the rest 66% said that they never read it. 47% were not aware of what data is being collected by the skill from them and another 21% were not entirely sure either. This shows a major usability issue where the users

Question	Do you think all voice-apps should have a privacy policy?
Responses	A privacy policy is always necessary to give users a piece of mind.
	Users should be able to know the risks involved such as if others could be listening in illegally.
	Privacy policy is definitely required so it can assure consumers that it is unlikely that malicious actions will occur with their data.
	They should be made easily accessible too.
	Required if in fact there are things that customers should be warned about prior to using it.
	A privacy policy would be completely necessary. I feel like the skills need to disclose everything being done with a user's data.
	But it should be easily explained and controls easily learned.
	If data is being collected, this is personal information that the user should have some control over.
	It needs to be more digestible so people will actually read it.
	I do think it is necessary to have a privacy policy, but I do think it should be short and easy to understand.

Table 9: User's view on the necessity of privacy policies. We present a few selected responses received from the participants in our user study when asked the question "Do you think all voice-apps should have a privacy policy?"

ignore the privacy policy even when it is provided by the developer. When asked about the issues they face with privacy policies, 20% of the participants responded by saying it is hard to access. 44% of participants felt that the document was too long. 24% claimed that they felt inconsistencies between the privacy policy and the skill's actual functionality and description. Users also had problem with developers not providing a privacy policy at all and the ones provided being not informative. The document being too legal and hard to comprehend was a concern for the users. Only 14% of participants felt that they always check the privacy policy before enabling a skill. 79% of our participants did not check the privacy policy before enabling a general skill and 75% did not check it before enabling a kids skill. The lack of usage of the privacy policy by the users shows the need of the VA platforms to address the concerns and take measures to improve the quality as well as the usability of the privacy policies provided by developers. We have included a few responses from the participants about their perspectives on whether privacy policies should be required for every voice-app in Table 9.

6 DISCUSSION

In this section, we discuss the limitation of this work and further research that can help in addressing the user frustration over privacy policies in VA platforms.

6.1 Limitation

We were unable to examine the actual source code of voice-apps. The availability of the source code can significantly increase the knowledge of what personal data a voice-app is able to collect and where it is stored. This can be compared with the privacy policy to ensure the developer is not performing any malicious activity or misusing the user's trust. With having no baseline, a future research effort that can be done on this regard is to dynamically test voice-apps by enabling them and check their data collection practices. Recently, SkillExplorer [28] has been proposed to dynamically explore skills' runtime behaviors and detect privacy violations in skills. As our future work, we plan to extend SkillExplorer to identify more inconsistent privacy policies of voice-apps.

Most developers provide short descriptions which will introduce the skill/action to end users, but data practices are not frequently defined in the descriptions. Since the data related to voice-apps is very limited, we largely depend on the descriptions provided with the voice-apps. This makes our findings on the inconsistency checking not complete. As mentioned in Sec. 3.2, our focus is on revealing the existence of problematic privacy policies, rather than identifying all the inconsistent privacy policies. For capturing data practices, we use a keyword-based method, and check whether keywords exist in the phrases. However, the keyword set can be incomplete. In our future work, we plan to use machine learning techniques to train a model to identify data practices from natural language documents. Nevertheless, we have collected strong evidence in revealing issues over privacy policies on VA platforms.

We only conducted a preliminary user study to understand how users engage with privacy policies and their concerns about them. The user survey motivates us to propose the privacy policy over voice mechanism (details are in Sec. 6.3). As our future work, we will expand our user study to involve more participants and add more questions. In particular, we will investigate which voice-apps users usually install/invoke, and how this may have influenced the survey results.

6.2 Why poor-quality privacy policies?

The Amazon Alexa platform not explicitly requiring app-specific privacy policies results in developers providing the same document that explains data practices of all their services. This leads to uncertainties and confusion among end users. There are skills with privacy policies containing up to 428 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 the skill to end users. The poor quality of privacy policies provided with voice-apps is partially due to the lack of an app-specific privacy policy and due to the lenient certification system. During the certification process, the content of a privacy policy is not checked thoroughly when the skill is submitted for certification, which has resulted in a large amount of inactive and broken links and also privacy policies not related to the skill. Some privacy policies mention data practices that are in violation of the privacy requirements that Amazon and Google have set but these voice-apps are still certified.

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

6.3 Privacy policy through voice

The unavailability of privacy policies through the voice channel requires users to access them over the web or through VA's companion apps on their smartphones. One possible reason for this can be due to the large size of the privacy policies and the time required to read out the long document. Users who only use voice assistant services through their VA devices, may not necessarily be aware of the existence of the privacy policies in the respective stores. Also, it is completely left to the user to decide whether to view the privacy policy or not. There is no approval asked prior to enabling the voice-app for the user. In order to address these issues, we propose to introduce a built-in intent (*i.e.*, functionality) for a voice-app that gives information to users about the privacy policy of the voice-app through a voice response. The major challenge for this is that the privacy policies are usually too long to be read out to users. Thus, the response provided by the built-in intent has to be marginally short.

Prior work has been done to summarize the privacy policies to make it more readable to the user. Tools like Polisis [29] and Privacycheck [11] conduct privacy policy analysis and represent the data practices mentioned in the document in a simpler form to users. But from our analysis of the skills/actions available in the stores, we have noticed that most privacy policies are general policies and do not necessarily define what the behavior of the voice-app in

particular is. Since personal information can be collected through the conversational interface, our approach aims in understanding this capability from the voice-app's source code, automatically generating an easy-to-digest privacy notice, and letting the user know about it through the voice channel.

We propose the privacy policy through voice mechanism. Such solution can be implemented as a plugin in the voice-app developer console when developers develop/deploy their voice-apps. Therefore, we assume the source code of a voice-app is available. We describe our preliminary approach based on the Amazon Alexa platform. We take the interaction model of a skill, which is a JSON (JavaScript Object Notation) file and scan for all the slots and their slot types specified. We categorise the built-in slot types based on what type of personal information they can collect. For custom slot types, we compare the values provided with the entries in datasets we assembled of possible values and check for a match. After we get all the types of information that can be collected by the skill, we create a response notifying the user that the skill has these capabilities and advise users to look at the detailed privacy policy provided by the developers. This functionality can be invoked when the skill is first enabled. On opening the skill for the first time, this brief privacy notice can be read out to the user. This will give the user a better understanding of what the skill he/she just enabled is capable of collecting and using. The users can also ask to invoke this functionality later to get a brief version of the privacy policy. As our future work, we plan to extend this approach to automatically generate easy-to-digest privacy policies for voice-apps at the development phase and make users aware of the data practices of a voice-app so that they are able to make informed privacy decisions before communicating with the voice-app.

7 RELATED WORK

7.1 Privacy concerns for voice assistants

Many research efforts have been undertaken to study user concerns (human factors) about the security/privacy of VA devices [15, 16, 19, 24, 26, 27, 33, 34, 39, 45]. Fruchter *et al.* [26] used natural language processing to identify privacy and security related reviews about VA devices from four major online retailers: Target, Walmart, Amazon, and Best Buy. The authors highlighted that users worried about the lack of clarity about the scope of data collection by their voice assistants. Through a semi-structured interviews with 17 VA users, Abdi *et al.* [15] uncovered the lack of trust users have with some of VA use cases such as shopping, and a very limited conception of VA ecosystem and related data activities. Malkin *et al.* [33] surveyed 116 VA owners and found that half did not know that their recordings were being stored by the device manufacturers. Similarly, authors in [45, 48] conducted interviews on smart home owners to examine user mental models and understand their privacy perceptions of IoT devices. Geeng *et al.* [27] investigated tensions and challenges that arise among multiple users in smart home environment. Lau *et al.* [31] conducted interviews with both VA users and non-users, and revealed that privacy concerns could be the main deterring factor for new users.

There has been an increasing amount of research on various attack vectors against VA systems and the corresponding defenses. One line of research is to exploit interpretation errors of user commands by speech recognition, such as voice squatting attack [30, 47], and generate hidden/inaudible voice commands [21, 35, 36, 40, 42, 44, 46]. Another line of research focuses on defense mechanisms, including, continuous authentication [25], canceling unwanted baseband signals [46], correlating magnetic changes with voice commands [22], and user presence-based access control [32]. Our work differs from these previous work in that we investigate the effectiveness of privacy policies provided by voice-app developers.

7.2 Privacy policy analysis for mobile apps

Privacy policies disclose an organization's or developer's data practices. Though researchers have conducted privacy policy analysis on Android platform [17, 38, 41, 43, 49, 50], there is an absence of privacy policy analysis on VA platforms. Zimmeck *et al.* [50] presented a privacy analysis system for Android to analyze apps' potential non-compliance with privacy requirements, and inconsistencies between privacy policies and apps. Results show that 71% of apps that lack a privacy policy should have one, and a substantial portion of apps exhibit potential privacy requirement inconsistencies. Wang *et al.* [41] developed a hierarchical mapping based approach for privacy policy analysis which is able to handle the data inputted by users in addition to the data accessed directly through the mobile device. The user input data is checked for possible privacy leaks and this is used to determine whether the app's privacy policy is in contradiction with this leakage. The consistency between the data collected by the app and the privacy policy provided is verified by using a data flow analysis. A major difference of our work from these works is that we rely on voice-app's description to detect inconsistency in privacy policies due to the unavailability of voice-app's source code. To the best of our knowledge, this is the first work to systematically measure the effectiveness of privacy policies for voice-apps.

8 CONCLUSION

In this work, we conducted a comprehensive empirical analysis on privacy policy of 64,720 Amazon Alexa skills and 16,002 Google Assistant actions. We designed an NLP-based approach to capture data practices in privacy policies and descriptions of voice-apps. Our results showed that a substantial number of problematic privacy policies exist in the Amazon Alexa and Google Assistant platforms, a worrisome reality of privacy policies on VA platforms. Google and Amazon even have official voice-apps violating their own requirements regarding the privacy policy. We conducted a user study to understand users' perspectives on voice-apps' privacy policies, which reflects real-world user frustrations on this issue. We also discussed possible approaches to improve the usability of privacy policies on VA platforms.

ACKNOWLEDGMENTS

We are grateful to anonymous reviewers for their constructive feedback. This work is supported in part by National Science Foundation (NSF) under the Grant No. 2031002, 1846291, and 1642143, and the NSF/VMware Partnership on SDI-CSCS program under the Grant No. 1700499.

REFERENCES

[1] 2020. Alexa Skills Policy Testing. https://developer.amazon.com/fr/docs/custom-skills/policy-testing-for-an-alexa-skill.html. (2020).

[2] 2020. Alexa Skills Security Requirements. https://developer.amazon.com/fr/docs/custom-skills/security-testing-for-an-alexa-skill.html. (2020).

[3] 2020. Amazon Developer Services Agreement. https://developer.amazon.com/support/ legal/da. (2020).

[4] 2020. Amazon Mechanical Turk. https://www.mturk.com/. (2020).

[5] 2020. Configure Permissions for Customer Information in Your Skill. https://developer.amazon.com/en-US/docs/alexa/custom-skills/configure-permissions-for-customer-information-in-your-skill.html. (2020).

[6] 2020. Google fined €50 million for GDPR violation in France. https://www.theverge.com/2019/1/21/18191591/google-gdpr-fine-50-million-euros-data-consent-cnil/. (2020).

[7] 2020. Google Privacy Policy Guidance. https://developers.google.com/assistant/console/ policies/privacy-policy-guide. (2020).

[8] 2020. Industrial-Strength Natural Language Processing. https://spacy.io. (2020).

[9] 2020. Number of digital voice assistants in use worldwide from 2019 to 2023. https://www.statista.com/statistics/973815/worldwide-digital-voice-assistant-in-use/. (2020).

[10] 2020. Policies for Actions on Google. https://developers.google.com/actions/policies/ general-policies. (2020).

[11] 2020. PrivacyCheck for Google Chrome. https://identity.utexas.edu/privacycheck-for-google-chrome. (2020).

[12] 2020. Selenium automates browsers. https://www.selenium.dev. (2020).

[13] 2020. Snapchat Transmitted Users' Location and Collected Their Address Books Without Notice Or Consent. https://www.orrick.com/Insights/2013/02/FTC-Assesses-800000-Fine-Against-Mobile-App-Operator-and-Issues-Mobile-Privacy-and-Security-Guidance . (2020).

[14] 2020. Snapchat Transmitted Users' Location and Collected Their Address Books Without Notice Or Consent. https://www.ftc.gov/news-events/press-releases/2014/05/snapchat-settles-ftc-charges-promises-disappearing-messages-were/. (2020).

[15] Noura Abdi, Kopo M. Ramokapane, and Jose M. Such. 2019. More than Smart Speakers: Security and Privacy Perceptions of Smart Home Personal Assistants. In *Symposium on Usable Privacy and Security (SOUPS)*.

[16] Tawfiq Ammari, Jofish Kaye, Janice Y. Tsai, and Frank Bentley. 2019. Music, Search, and IoT: How People (Really) Use Voice Assistants. *ACM Trans. Comput.-Hum. Interact.* 26, 3 (2019), 17:1–17:28.

[17] Benjamin Andow, Samin Yaseer Mahmud, Wenyu Wang, Justin Whitaker, William Enck, Bradley Reaves, Kapil Singh, and Tao Xie. 2019. Policylint: Investigating Internal Privacy Policy Contradictions on Google Play. In *USENIX Security Symposium (USENIX Security)*. 585–602.

[18] Noah Apthorpe, Sarah Varghese, and Nick Feamster. 2019. Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus COPPA. In *USENIX Security Symposium (USENIX Security)*.

[19] Alexander Benlian, Johannes Klumpe, and Oliver Hinz. 2019. Mitigating the intrusive effects of smart home assistants by using anthropomorphic design features: A multimethod investigation. *Information Systems Journal* (2019), 1–33. https://doi.org/10.1111/isj.12243

[20] TravisD. Breaux, Hanan Hibshi, and Ashwini Rao. 2014. Eddy, a formal language for specifying and analyzing data flow specifications for conflicting privacy requirements. *Requirements Engineering* (2014), 1–27.

[21] Nicholas Carlini, Pratyush Mishra, Tavish Vaidya, Yuankai Zhang, Micah Sherr, Clay Shields, David Wagner, and Wenchao Zhou. 2016. Hidden Voice Commands. In *25th USENIX Security Symposium (USENIX Security 16)*. 513–530.

[22] S. Chen, K. Ren, S. Piao, C. Wang, Q. Wang, J. Weng, L. Su, and A. Mohaisen. 2017. You Can Hear But You Cannot Steal: Defending Against Voice Impersonation Attacks on Smartphones. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. 183–195.

[23] Long Cheng, Christin Wilson, Song Liao, Jeffrey Young, Daniel Dong, and Hongxin Hu. 2020. Dangerous Skills Got Certified: Measuring the Trustworthiness of Skill Certification in Voice Personal Assistant Platforms. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*.

[24] H. Chung, M. Iorga, J. Voas, and S. Lee. 2017. "Alexa, Can I Trust You?". *IEEE Computer* 50, 9 (2017), 100–104.

[25] Huan Feng, Kassem Fawaz, and Kang G. Shin. 2017. Continuous Authentication for Voice Assistants. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking (MobiCom)*. 343–355.

[26] Nathaniel Fruchter and Ilaria Liccardi. 2018. Consumer Attitudes Towards Privacy and Security in Home Assistants. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*.

[27] Christine Geeng and Franziska Roesner. 2019. Who's In Control?: Interactions In Multi-User Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI'19)*. 1–13.

[28] Zhixiu Guo, Zijin Lin, Pan Li, and Kai Chen. 2020. SkillExplorer: Understanding the Behavior of Skills in Large Scale. In *29th USENIX Security Symposium (USENIX Security 20)*.

[29] 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. *CoRR* abs/1802.02561 (2018). http://arxiv.org/abs/1802.02561

[30] Deepak Kumar, Riccardo Paccagnella, Paul Murley, Eric Hennenfent, Joshua Mason, Adam Bates, and Michael Bailey. 2018. Skill Squatting Attacks on Amazon Alexa. In *USENIX Security Symposium (USENIX Security)*. 33–47.

[31] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW (2018), 1–31.

[32] X. Lei, G. Tu, A. X. Liu, C. Li, and T. Xie. 2018. The Insecurity of Home Digital Voice Assistants - Vulnerabilities, Attacks and Countermeasures. In *2018 IEEE Conference on Communications and Network Security (CNS)*. 1–9.

[33] Nathan Malkin, Joe Deatrick, Allen Tong, Primal Wijesekera, Serge Egelman, and David Wagner. 2019. Privacy Attitudes of Smart Speaker Users. In *19th Privacy Enhancing Technologies Symposium (PETS)*.

[34] Graeme McLean and Kofi Osei-Frimpong. 2019. Hey Alexa: examine the variables influencing the use of artificial intelligent in-home voice assistants. *Computers in Human Behavior* 99 (2019), 28 – 37.

[35] Nirupam Roy, Sheng Shen, Haitham Hassanieh, and Romit Roy Choudhury. 2018. Inaudible Voice Commands: The Long-Range Attack and Defense. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 547–560.

[36] Lea Schönherr, Katharina Kohls, Steffen Zeiler, Thorsten Holz, and Dorothea Kolossa. 2019. Adversarial Attacks Against Automatic Speech Recognition Systems via Psychoacoustic Hiding. In *Network and Distributed System Security Symposium (NDSS)*.

[37] Faysal Shezan, Hang Hu, Jiamin Wang, Gang Wang, and Yuan Tian. 2020. Read Between the Lines: An Empirical Measurement of Sensitive Applications of Voice Personal Assistant Systems. In *Proceedings of The Web Conference (WWW)*.

[38] R. Slavin, X. Wang, M. B. Hosseini, J. Hester, R. Krishnan, J. Bhatia, T. D. Breaux, and J. Niu. 2016. Toward a Framework for Detecting Privacy Policy Violations in Android Application Code. In *2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE)*. 25–36.

[39] Maurice E. Stucke and Ariel Ezrachi. 2017. How Digital Assistants Can Harm our Economy, Privacy, and Democracy. *Berkeley Technology Law Journal* 32, 3 (2017), 1240–1299.

[40] Tavish Vaidya, Yuankai Zhang, Micah Sherr, and Clay Shields. 2015. Cocaine Noodles: Exploiting the Gap between Human and Machine Speech Recognition. In *9th USENIX Workshop on Offensive Technologies (WOOT 15)*.

[41] X. Wang, X. Qin, M. Bokaei Hosseini, R. Slavin, T. D. Breaux, and J. Niu. 2018. GUILeak: Tracing Privacy Policy Claims on User Input Data for Android Applications. In *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*. 37–47.

[42] Qiben Yan, Kehai Liu, Qin Zhou, Hanqing Guo, and Ning Zhang. 2020. SurfingAttack: Interactive Hidden Attack on Voice Assistants Using Ultrasonic Guided Wave. In *Network and Distributed Systems Security (NDSS) Symposium*.

[43] L. Yu, X. Luo, X. Liu, and T. Zhang. 2016. Can We Trust the Privacy Policies of Android Apps?. In *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 538–549.

[44] Xuejing Yuan, Yuxuan Chen, Yue Zhao, Yunhui Long, Xiaokang Liu, Kai Chen, Shengzhi Zhang, Heqing Huang, XiaoFeng Wang, and Carl A. Gunter. 2018. Commandersong: A Systematic Approach for Practical Adversarial Voice Recognition. In *USENIX Security Symposium (USENIX Security)*. 49–64.

[45] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End User Security and Privacy Concerns with Smart Homes. In *USENIX Conference on Usable Privacy and Security (SOUPS)*. 65–80.

[46] Guoming Zhang, Chen Yan, Xiaoyu Ji, Tianchen Zhang, Taimin Zhang, and Wenyuan Xu. 2017. DolphinAttack: Inaudible Voice Commands. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 103–117.

[47] Nan Zhang, Xianghang Mi, Xuan Feng, XiaoFeng Wang, Yuan Tian, and Feng Qian. 2019. Understanding and Mitigating the Security Risks of Voice-Controlled Third-Party Skills on Amazon Alexa and Google Home. In *IEEE Symposium on Security and Privacy (SP)*.

[48] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proc. ACM Hum.-Comput. Interact.* (2018).

[49] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel Reidenberg, N. Russell, and Norman Sadeh. 2019. MAPS: Scaling Privacy Compliance Analysis to a Million Apps. *Proceedings on Privacy Enhancing Technologies* 2019 (07 2019), 66–86.

[50] Sebastian Zimmeck, Ziqi Wang, Lieyong Zou, Roger Iyengar, Bin Liu, Florian Schaub, Shormir Wilson, Norman Sadeh, Steven M. Bellovin, and Joel Reidenberg. 2017. Automated Analysis of Privacy Requirements for Mobile Apps. In *24th Network & Distributed System Security Symposium (NDSS 2017)*.

Appendices

Category	Skills we crawled	Skills with a privacy policy
Business & Finance	3,599	1,420
Connected Car	140	100
Education & Reference	7,990	1,460
Food & Drink	1,379	407
Games & Trivia	11,822	1,461
Health & Fitness	2,026	844
Kids	3,252	461
Lifestyle	11,080	2,693
Local	1,283	377
Movies & TV	915	153
Music & Audio	9,216	3,155
News	6,810	2,907
Novelty& Humor	3,418	394
Productivity	4,263	1,434
Shopping	342	204
Smart Home	2,432	2,254
Social	1,479	531
Sports	1,592	343
Travel & Transportation	1,187	205
Utilities	1,025	191
Weather	853	150
Total skills	76,103	21,144
Total unique skills	64,720	17,952

Table 10: Alexa skills by category in our dataset as of March 2020. Some skills are classified and listed in multiple categories. After removing the cross-listed duplicates, we obtained 64,720 unique skills, and 17,952 of these skills provide privacy policy links.

Category	Actions we crawled	Actions with a privacy policy
Arts & Lifestyle	96	96
Business & Finance	532	532
Communication & Social	43	43
Education & Reference	1,600	1,600
Food & Drink	2,256	305
Games & Fun	4,043	4,043
Health & Fitness	227	227
Kids & Family	108	108
Local	67	67
Movies, Photos & TV	59	59
Music & Audio	2,339	427
News & Magazines	2,282	99
Productivity	65	65
Shopping	99	99
Smart Home	1,404	1,404
Sports	510	509
Travel & Transportation	209	209
Weather	63	63
Total actions	16,002	9,955

Table 11: Google actions by category in our dataset as of March 2020.

Issue	Skill name
Different privacy policies provided in a skill & action pair	AGL, Air Quality, Amdocs Connected Home, Because News Quiz, Burbank Town Center, Bustle, Central Mall Lawton, Debate Cruncher, Delmarva Power Smart Home Pilot, Desert Financial, Detective Mr Z, Did Thanos Kill Me, Eton, EV Car, FGLair Smart Home, Fox Chapel School Lunch, iMagic, iSmart Plus, Ithaca College Physical Therapy, KEEL Vodka, Legal Newswire, Lutron Connect, Mighty Mule, New York Daily News, Orbit B-Hyve, Orlando Sentinel, Rain Bird, Real Simple Tips, Robo Coach, Robonect lawn mower, Royal Credit Union, Sense, Smartenit, Symcon, The Daily Beast, The Hartford Small Business Insurance, The Morning Call, Ticketmaster, Turbo Tips
A skill doesn't have a privacy policy while the Google action version has one	A Precious Day, Ambient Woodstock Chimes, B96.5, Bob George Ministries, Celebration Rock, FIFA Ultimate Quiz, FM NEWS 101 KXL, Freedom 970, GodLife, GPS: God. People. Stories. from Billy Graham, Hive, Houston Baseball, Jingle Bells, Kurt Talk, Liverpool Football Quiz, LOVE Brentford, LOVE Spurs, Matt Lieber Bot, Michigan Insider, My Morning Prayer, Radio Chaser, Real Presence Radio, Really Untrue Facts, Sadguru Whispers, Sherlock Riddles, SimpliSpoken Voice Tester, Sleep by Nature Made, Stephen King Library, Super Over, The Andrew Klavan Show, The Danny Lakey Late Show, The Global Startup Movement, The Hot Breakfast, The Michael Knowles Show, The Ticket Top 10, This Week in Beatles History, Touch India, Triple M NRL, Voicebot Podcast, WE Hip Hop, Website Reader, wikiHow

Table 12: Same voice-apps with different privacy policies on two VA platforms as of May 2020.