

SkillVet: Automated Traceability Analysis of Amazon Alexa Skills

Jide S Edu, Xavier Ferrer-Aran, Jose Such, Guillermo Suarez-Tangil

Abstract—Skills, are essential components in Smart Personal Assistants (SPA). The number of skills has grown rapidly, dominated by a changing environment that has no clear business model. Skills can access personal information and this may pose a risk to users. However, there is little information about how this ecosystem works, let alone the tools that can facilitate its study. In this paper, we present the largest systematic measurement of the Amazon Alexa skill ecosystem to date. We study developers' practices in this ecosystem, including how they collect and justify the need for sensitive information, by designing a methodology to identify over-privileged skills with broken privacy policies. We collect 199,295 Alexa skills and uncover that around 43% of the skills (and 50% of the developers) that request these permissions follow bad privacy practices, including (partially) broken data permissions traceability. In order to perform this kind of analysis at scale, we present SkillVet that leverages machine learning and natural language processing techniques, and generates high-accuracy prediction sets. We report several concerning practices, including how developers can bypass Alexa's permission system through account linking and conversational skills, and offer recommendations on how to improve transparency, privacy and security. Resulting from the responsible disclosure we did, 13% of the reported issues no longer pose a threat at submission time.

Index Terms—Voice Assistants, Smart Personal Assistants, Alexa, Smart Home, Privacy, Permission, Personal Data, Access Control.



1 INTRODUCTION

Smart Personal Assistants (SPA) are changing how users interact with technology and the way they consume services [1], driving new experiences and expectations [2], [3]. Advances in natural language processing enable SPA to take voice commands, which they process using machine learning to deduce users' intentions and fulfill users' requests. SPA are becoming very popular. According to recent reports, nearly 90 million adults in the US use at least one SPA device [4]. The adoption is expected to rocket even further as SPA integrate into smartphones [5] or as a chatbot [6].

Despite their growing popularity, intrinsic features of SPA expose users to various risks [7], [8], including 1) the open nature of the voice channel, 2) the complexity of their architecture, and 3) the use of a wide range of underlying technologies. Although considerable research has been recently conducted to understand and address these risks, many of them remain unresolved [7]. In particular, much of the research has focused on architectural elements of SPA devices (e.g., Amazon Echo) [9], [10], [11] and the speech recognition capabilities of SPA [8], [12]. However, SPA incorporate voice-driven applications generally developed by third parties, so-called *skills*. The skills ecosystem widens the attack surface of SPA [7], as malicious third-party actors may develop potentially harmful or unwanted software [13], [14]. Unfortunately, it is unclear what are the risks third-party skills may pose to users, beyond the potential for skill impersonation [14], [15] and other potentially suspicious

behaviors [16]. Authors in [17] have very recently shown that the Amazon skill certification process can be subverted by crafting mock policy-violating skills. However, there is a critical knowledge gap in understanding what are the *actual* data collection practices of skill developers in the wild. This gap prompts four core research questions: 1) *What is the current state of affairs in the third-party ecosystem of skills?* 2) *Is the collection of personal information explained?* 3) *Can we pinpoint fine-grained privacy issues (i.e., at the permission level)?* 4) *How can we do this in an automated way at scale?*

In this paper, we perform the first systematic measurement study of Amazon Alexa skills across the entire marketplace of 11 different countries to shed light on the third-party skills ecosystem and developers' practices. We characterize skills using their market categories, names, and developers, among others. We also characterize the data permissions they request during runtime and their traceability with data practice statements in the skill's privacy policy. Since it is time-consuming to analyze privacy policies manually, we also propose an automated machine-learning-based system, SkillVet, to help with this at scale. SkillVet identifies relevant policy statements and assesses the traceability as complete, partial, or broken. Our system overcomes challenges in modeling privacy policies, such as contradictions like negations and statements related to multiple types of personal information.

We then analyze suspicious skills in the look for concerning privacy practices. We do this by systematically analyzing the traceability between the permissions requested by the skill and the data practices stated by the skill developer in the privacy policy. This is done following a black-box approach since the skills' code or executable is not available — skills run in the cloud instead of the users' device

Our study makes the following key contributions:

- J. Edu, X.Ferrer Aran, J.Such were with King's College (KCL) London. J.Such was also with Universitat Politècnica de València. E-mail: {jide.edu, xavier.ferrer_aran,jose.such}@kcl.ac.uk.
- G.Suarez-Tangil was with IMDEA Networks Institute and KCL. E-mail: guillermo.suarez-tangil@imdea.org

- 1) We perform the largest systematic study of the Alexa skill ecosystem across 11 Amazon marketplaces (§2). We take a snapshot of all markets and measure the distribution of skills, developers’ diversity, and the invocation names used for 199k skills.
- 2) We study the permission model Amazon offers to developers, and investigate how skills make use of permissions to collect personal data (§3). We then design a methodology that systematically identifies potential privacy issues by analyzing the traceability between the permissions and the data practices stated by developers. We find a wealth of over-privileged skills with *broken* traceability.
- 3) We design a system to automate the traceability analysis at scale, *SkillVet*, based on machine learning and natural language processing (§4). Our system can be deployed by SPA providers to automatically conduct traceability analysis across marketplaces regularly and by developers to check the traceability of their skills. We demonstrate *SkillVet*’s accurate performance.
- 4) We provide insights regarding the developers’ business models and concerning practices we observe (§5). These practices include skills bypassing the permission system of Alexa to collect personal data from users by: a) account linking, and b) conversational skills. We provide evidence of skills (and their traceability) conducting these practices and discuss the implications of our findings (§6).

To foster research in the area, we release our datasets together with the ground-truth we have produced and our code in <https://github.com/jideedu/SkillVet>

2 THE SKILL ECOSYSTEM

Skills are, in a way, similar to mobile apps, as they enhance the capability of the ecosystem they are part of. They extend SPA functionality, allowing them to offer extensive services tailored to users’ interests and experiences. The SPA skill ecosystem provides an environment that enables the user to run functions such as music playback, online purchase, appointment booking, and different innovative home automation tasks. For instance, a user can use a “Uber” skill to order a ride through the Amazon Echo.

A skill can simply be activated with an invocation utterance, i.e.: the trigger phrase and the skill invocation name. For example, to invoke the “My Chef” skill, a user can say ‘*Alexa, open My Chef*’, or ‘*Alexa, start My Chef*’. Note that “Alexa, open” and “Alexa, start” are trigger phrases to put the SPA into recording mode. When a skill is successfully invoked, specific phrases can then be uttered to call intents (functions) within the skill. However, in some cases, a skill needs to be enabled before it can be invoked, for instance, when the skill requires permissions. Skills can be enabled via compatible apps or online at the skill store.

2.1 Threat Model

The attack surface of SPA is considerable large and can lead to several security and privacy issues — for an extensive review of the literature refer to [7]. One of the attack entry points in SPA, and the focus of this article, is SPA skills.

Unlike mobile apps, skills do not run on any user-controlled device. Instead, skills either run in the Amazon cloud (e.g., as an AWS Lambda function) or in a server controlled by the developer [18]. This could allow a malicious developer to sneak malicious code into their software via the application backend [13]. They could manipulate skills to covertly introduce misperception about reported events [19] and also use skills with duplicate invocation names, similar phonemes, or paraphrased invocation names to impersonate another skill [14], [15]. Recently, there has been report of how malicious actors, through a phishing attack, could steal a token that can let them add suspicious skills to the users’ account and access voice history [20].

In this paper, we also consider the threat from skills developed by third-party developers, but focus on the data practices of such developers, whether malicious or just negligent. Skills allow users to interact with their services and require the exchange of data. This could be data that are part of the user’s account with Amazon, such as users *location, mobile number, email address, name, address* [21], sensitive data inferred from user actions with the skills, or personal data collected by the skills during their conversations with the users. By considering skill developers as adversaries, we look into potentially over-privileged skills to understand what personal information they collect. This issue has been explored by several previous studies on permission gaps in other domains [22], [23] but not in SPA and their skills. While declaring more permissions than required seems to have no impact on a skill’s functionality, it could, however, be leveraged by a malicious developer to achieve malicious goals [22], which could impact the users’ privacy.

This research presents *SkillVet*, a system that could detect skills with inappropriate usage of data permissions at large and help protect users’ privacy. For instance, “Mock Interview” skill by *Graylogic Technologies* as detailed later. In addition, our system could help users understand skills data practices to make a better decision about their data.

2.2 Collection and Characterization

Amazon uses different online marketplaces to offer targeted services to several market segments. Amazon has 14 online marketplaces, out of which 11 have an online store with third-party skills. We crawl these 11 markets: US, UK, India, Australia, Canada, Germany, Japan, Italy, Spain, France, and Mexico. Skills are only available to registered users and a user can only be linked with one marketplace at a time.

For all the skills we collect, we make a characterization of the market category they belong to, the utterances that activate the skill, and the developer that has created it. Two further key elements we extract are i) the permissions that the skill requires to access personal information from the Amazon user’s account **via the Alexa API at runtime**, which Amazon makes publicly available in the marketplace; and ii) the privacy policy declared by the developers. We then use these two elements, together with an interactive analysis, to study when there is a privacy violation or a concerning practice that may harm users.

Unlike other platforms like Android, Amazon Alexa runs the skills in the cloud and *the code of the skills is not publicly available*. To analyze the ecosystem of skills, we built

a Web scrapper that recursively crawls all available marketplaces and extracts the metadata published by Amazon about the skills. Our crawler exhaustively collects all skills available in the Amazon Alexa marketplaces (for details refer to [24]). Note that one skill can belong to different categories and be hosted in multiple marketplaces. The URL may have different parameters in these cases, but the path displays a unique identifier per skill. We identify unique skills posted across markets in a process we call de-duplication.

We use the following attributes (which we scrape from the marketplace) to characterize every skill: invocation name, permissions, the category, developer’s information, privacy policy, terms of use, skill description, rating information and reviews. While new skills are added daily, we base our study on those skills accessible to users between the 15th of July 2020 and the 29th of July 2020. This is the time-frame that took our crawler to visit all marketplaces.

2.3 Skills and Developers

Table 1 shows the breakdown of the total number of Alexa skills and developers across Amazon marketplaces. Overall, we see 199,295 skills published by 88,391 developers. Note that most of these skills have been developed in just a few years.¹ English-speaking marketplaces host the highest number of skills and developers. On the other end, the smallest market is Mexico with 1,972 skills and representing 1% of our dataset. However, there is overlap across markets as developers can publish the same skills in different markets [26]. This overlap is larger among English-speaking markets such as the US, UK, CA, and AU where Amazon tends to migrate existing skills [27]. In particular, we observe that close to 20,000 different skills in the UK market are also in the US market. When we de-duplicate skills, we observe that 53.76% (47,520) of the developers and 56.2% (112,029) of the skills we collect are unique. In what follows we look at unique skills unless we analyze specific markets.

TABLE 1

Number of skills & developers. English-speaking markets represent 82.74% (92k skills and 36k devs).

Marketplace	Skills		Developers	
	N	Percent	N	Percent
US	61,362	30.79%	27,030	30.58%
UK	32,822	16.47%	14,487	16.39%
India	29,344	14.72%	12,823	14.51%
Canada	26,027	13.06%	11,509	13.02%
Australia	23,909	12.00%	10,854	12.28%
Germany	9,096	4.56%	3,385	3.83%
Spain	4,759	2.39%	2,441	2.76%
Italy	4,203	2.11%	2,049	2.37%
Japan	3,513	1.76%	1,407	1.59%
France	2,288	1.15%	1,194	1.35%
Mexico	1,972	0.99%	1,212	1.37%
Total	199,295	100.00%	88,391	100.00%
Unique	112,029	56.2%	47,520	53.76%

Developers. We illustrate the relationship between developers and skills per marketplace in Figure 1. While most of the developers publish only one skill, a few of them have tens, and a handful have hundreds and even thousands of skills.

1. Alexa had only 135 skills in early 2016 [25].

In particular, 125 developers have published more than 50 skills, 69 developers more than 100 skills, and 5 developers more than 1K skills. We then study the description of all the skills published by prominent developers and see that many implicitly indicate that the skill has been developed with an automated platform. This is usually done by referring to the name or the URL of the framework that produced the skill. Popular among them are *VoiceApps.com*, *VoiceXP.com* and *VoiceFlow.com*. These platforms provide free sample skills that a developer can customize regardless of their technical background. In particular, they offer visual, drag-and-drop editors that lowers down any development barriers. We mine all descriptions in our dataset and see that at least 7% of the skills have been developed with these platforms. In particular, we observe a ratio of 1 developer to 30 skills with the most popular automated platform (*VoiceApps.com*).

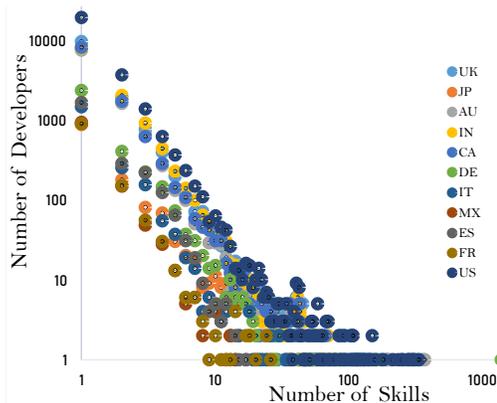


Fig. 1. Developers vs the number of skills they publish per marketplace.

Skill Invocation Names. Previous research showed that an attacker could impersonate skills by crafting specific malicious skill invocation names [14], [15], [17] by reusing the same as or (phonetically) similar invocation names to other skills. We measure the prevalence of same and similar skill invocation names across the Alexa ecosystem to understand the potential for this risk. We find 24,468 unique skills (21.8%), 6,876 (11.2%) in the US, with the same invocation name as another/other skill/s after cross-market de-duplication. We also find a considerable amount of very similar skill invocation names. For instance, we see that between 5% and 15% of all the skills in the US have a different but phonetically similar name. This is in line with recent measurements in less markets [28]. For detailed results, including the most popular names, the breakdown per country, the method to measure phonetic similarity, and the breakdown per similarity, refer to [24].

Take-aways. Our market characterization provides a high-level overview of the skill ecosystem and how Alexa marketplaces are structured. In summary, the main takeaways include:

- The skill ecosystem has grown considerably fast in recent years and it currently has few developers with thousand of skills each. English-speaking markets generally dominate and push skills to other markets.
- We see skills developed with automated platforms. The use of these platforms lower down the development barriers and enable bulk skill creation.
- We observe a high prevalence of skills with the same or

similar name, with the associated potential for impersonation and Sybil attacks as shown in previous works.

3 PERMISSIONS AND TRACEABILITY

In this section, we first explore the data permissions in all Alexa marketplaces, analyzing which permissions are more often requested by Alexa skills and developers among the different markets, and how they are distributed between skills. Then, we look at skill privacy policies to understand how developers disclose and justify the data permissions they ask for. To do so, we create the **largest traceability dataset for Alexa** known to date, namely the traceability-by-policy dataset (TBPD), which includes the traceability analysis of 1,758 skills requesting data permissions and their policies from the five English-speaking marketplaces (Australia, UK, US, Canada, and India), representing 67.4% of all skills requesting data. Using the dataset, we analyze the skill traceability at the skill category and developer dimensions and identify developers’ practices.

3.1 Data Permissions

Skills can request access to user information through the Alexa skill API. This information becomes available on a per-skill and per-data-type basis when the user consents. To manage the consent: i) Alexa declares the permissions of skill on their store and gives users control through the Web or the mobile app, and ii) skills have to inform users in their privacy policies how they will use the personal information they collect. Alexa currently supports 11 types of permissions [21]. These are: Device Address (15.0%), Location Services (4.1%), Email Address (15.7%), Device Country and Postal Code (13.7%), Reminder (9.1%), Customer Name (10.3%), List Read Access (5.3%), List write Access (4.9%), Mobile Number (4.3%), Amazon Pay (1.8%), and Skill Personalization (0.0%).² The most prevalent permissions are generally used to offer services based on the user’s location.

Overall, 2,608 (2.3%) skills across markets ask for data access permissions. This could, in principle, be seen as good news. Still, 2.6K skills is a sizable set of third-party software accessing personal information, and the practices of their developers require scrutiny. We also argue that skills may collect personal information using other means, i.e., through account linking or conversations, as discussed in §5.

Table 2 shows the permission distribution of skills per marketplace. As shown, most skills ask for one permission (typically, the device address as discussed above), with all marketplaces showing a similar pattern overall — albeit with slight differences between those markets with more or less number of skills. There is also an important minority of skills asking for two or three permissions. Finally, there are few exceptions where skills ask for over three permissions. The most noticeable one is a skill called “BILH Staff Chatbot” by *BIDMC* (the Beth Leahy hospital). The skill is available in the US, and it is requesting several unique permissions (i.e., Name, Email Address, Mobile Number,

2. At the time of our analysis, we do not see any skills with Skill Personalization yet. Also, our analysis reveals Notification (15.5%), which is now deprecated and replaced by Reminder and Timers (0.3%).

TABLE 2
Permissions distribution by skills and marketplace.

Skills	No of Unique Permissions				
	1 N (%)	2 N (%)	3 N (%)	4 N (%)	≥5 N (%)
US (1,698)	1,169 (68.85)	301 (17.73)	143 (8.42)	51 (3.00)	34 (2.00)
UK (581)	418 (71.94)	120 (20.65)	26 (4.48)	7 (1.20)	10 (1.72)
CA (366)	262 (71.58)	74 (20.22)	14 (3.83)	6 (1.64)	10 (2.73)
IN (358)	250 (69.83)	66 (18.44)	20 (5.59)	7 (1.96)	15 (4.19)
AU (348)	245 (70.40)	78 (22.41)	10 (2.87)	5 (1.44)	10 (2.87)
DE (341)	236 (69.21)	83 (24.34)	12 (3.52)	7 (2.05)	3 (0.88)
JP (126)	100 (79.37)	21 (16.67)	4 (3.17)	1 (0.79)	
ES (89)	56 (62.92)	24 (26.97)	6 (6.74)	3 (3.37)	
IT (85)	58 (68.24)	21 (24.71)	5 (5.88)	1 (1.18)	
FR (63)	44 (69.84)	15 (23.81)	4 (6.35)		
MX (46)	30 (65.22)	13 (28.26)	3 (6.52)		
Total 4,101	2,868 (69.93)	816 (19.90)	247 (6.02)	88 (2.15)	82 (2.00)
Uniq. 2,608	1,807 (69.29)	499 (19.13)	191 (7.32)	67 (2.57)	44 (1.69)

Reminders, Location Services, and Device Country and Postal Code). The skill is a chatbot for employees to help them fill out a COVID-19 staff daily health questionnaire. While the skill is not intended to diagnose users (according to the description), it asks for 13 different symptoms. The use of this skill poses an important privacy risk to employees, because the Web form version of the questionnaire does not ask for data such as postal code or location, so it is not clear why these are needed in Amazon.

3.2 Traceability Analysis

We look at privacy policies to understand how developers disclose the permissions they request. Specifically, we study the traceability between data operations obvious to users (recall users need to enable skills requesting permissions) and the data actions defined by developers in the skill policies. Note that Amazon’s privacy requirements for skill developers mandate that a skill must come with an adequate privacy policy if it collects personal information. In particular, any collection and use of personal information need to comply with what is stated in the privacy policy [29].

To evaluate traceability between permissions and policies, we collect and tag privacy policies of 1,758 Alexa skills requesting data permissions in the five English-speaking marketplaces. We develop a Selenium module in Python that automatically visits and downloads every skill’s privacy policy page, cleaning the HTML code to remove unnecessary markup code, normalizing punctuation, and extra white-spaces discarding non-ASCII characters, and finally converting text to lowercase (this is later referred as the pre-processing phase). Afterwards, each skill is analyzed and evaluated as having *broken*, *partial* or *complete* traceability, following previous studies of traceability analysis in other domains [30], [31], [32], [33].

The type of traceability is identified by comparing the permissions requested by the skill through the Amazon Alexa API with the data practices covered in the skill policy. The different traceability types are explained next:

Complete: A skill is said to offer complete traceability if it provides adequate information in its privacy policy document about its data practices, i.e., the data action defined in the privacy policy document can be completely mapped to the access of data permissions. For instance, a skill like the “Aircraft Radar” in the UK market developed by *Chris Dzombak* offers complete traceability since it provides adequate information about its data practices in its privacy

policy. The statement “aircraft radar uses your devices address to find your location and search for aircraft around you” can be mapped with the Device Address permission the skill collects.

Partial: This is when the transparency of data action defined in the privacy policy document maps partially to the access permissions. For example, when the privacy document states that a skill collects personally identifiable information without explicitly stating what this information is. A statement such as “we may require you to provide us with certain personally identifiable information” offers partial traceability since it does not explicitly state what information is collected. A skill is also said to offer partial traceability if not all its data permissions are covered in its privacy document. Partial disclosure of data practices may also occur when data practices in a skill privacy document are not well mapped with the skill’s data permission. For instance, the statement “we may also collect your zip code” is partially traceable as the skills collect Device Address, and this refers to the user’s full address, including the zip code and the street number. Examples that provide partial information in their privacy policy include: 1) “Vote Sam Feldt” by *Fanspoke* that collects Mobile Number 2) “Flight Level - An Aircraft Radar” by *O. Schafer* that collects Device Country and Postal Code.

Broken: If a skill requests for permissions but does not have a privacy policy or the link to the privacy policy given is not working, we consider the traceability broken. Even when providing a privacy policy, A skill is said to have broken traceability if it provides no data implication in its privacy policy document. For instance, when a skill collects the device address, and its privacy document states nothing related to device address, we mark such skill data practices disclosure as broken. We only mark a broken policy when permissions are not traceable to the data practices defined in the privacy document. In the UK marketplace, skills like “Daily Yoga” by *Siva Pandeti* collects the Device Country and Postal Code, “Bike Weather Man” by *Marc Easen* collect Device Country & Postal Code, and “Smoggy Alerts” by *Teal Dreams Software, Inc.* collects the Device Country and Postal Code, and they all exhibit broken traceability.

When analyzing traceability, a skill requesting “Device Country and Postal Code” permission is tagged to have adequately disclosed its data practices if “Device Address” is mentioned in the privacy policy. This is because “Device Address” refers to the user’s full address, including the device country and postal code and the street number. Also, we did not find any skills using Skill Personalization, and grouped “List Write Access” and “List Read Access” under the “Personal Information” category as acknowledging personal information collection should be sufficient to disclose these permissions. Finally, reminders/notifications are not explicitly considered in the privacy requirements for skills by Amazon, as they may not have a personally identifying implication. Therefore, the traceability of all skills was evaluated considering: “Location Services”, “Amazon Pay”, “Mobile Number”, “Email Address”, “Name”, “Device Address”, “Device Country and Postal Code”, “Personal Information”.

3.3 Traceability Results

A total of 1,758 skills request data permissions in English-speaking marketplaces. Out of the 1,758 skills: 442 (25%) have broken traceability, 306 (18%) have partial traceability, and 1,010 (57%) have complete traceability. When we exclude the complete ones, we see 43% of the skills displaying bad privacy practices. Figure 2 shows the results by market, where we can see a very similar trend across marketplaces, with UK appearing to have slightly more broken and less complete traceability results, but it also has the least number of skills and developers.

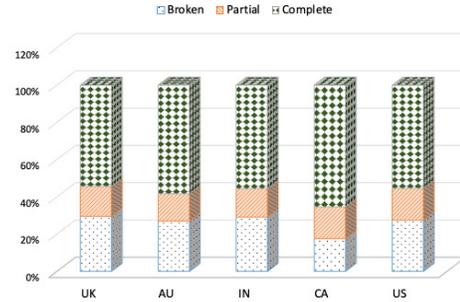


Fig. 2. Traceability for skills requesting data permissions in English-speaking markets.

Traceability by Categories. To understand how traceability affects different type of skills, we first compare our results considering the 1,758 unique skills collected across different skill subcategories and English-speaking marketplaces, as shown in Table 3. We show the breakdown using subcategories to give a more fine-grained perspective of particular types of (or sector-specific) skills and associated traceability. We rank the different subcategories based on the number of issues (broken and partial) normalized by the number of well defined policies (complete). Note that we do not list subcategories with less than 3 skills for the sake of clarity.

TABLE 3
Traceability per subcategory in English-speaking markets.

Categories	R	B	P	C
Sports	1	43 (91.49%)	1 (2.13%)	3 (6.38%)
Social	2	98 (85.96%)	7 (6.14%)	9 (7.89%)
(Uncategorised)	3	56 (90.32%)	1 (1.61%)	5 (8.06%)
Novelty & Humour	4	5 (71.43%)	2 (28.57%)	0 (0.00%)
Kids	5	10 (83.33%)	1 (8.33%)	1 (8.33%)
Games & Trivia	6	89 (57.42%)	34 (21.94%)	32 (20.65%)
Schools	7	3 (100.00%)	0 (0.00%)	0 (0.00%)
Utilities	8	17 (56.67%)	5 (16.67%)	8 (26.67%)
News	9	30 (63.83%)	4 (8.51%)	13 (27.66%)
Health & Fitness	10	60 (48.78%)	25 (20.33%)	38 (30.89%)
Organizers & Assistants	11	7 (63.64%)	1 (9.09%)	3 (27.27%)
Lifestyle	12	74 (46.54%)	32 (20.13%)	53 (33.33%)
Weather	13	22 (41.51%)	13 (24.53%)	18 (33.96%)
Food & Drink	14	39 (42.86%)	20 (21.98%)	32 (35.16%)
Streaming Services	15	4 (57.14%)	1 (14.29%)	2 (28.57%)
Productivity	16	42 (51.85%)	9 (11.11%)	30 (37.04%)
Business & Finance	17	82 (52.23%)	16 (10.19%)	59 (37.58%)
Smart Home	18	44 (57.14%)	4 (5.19%)	29 (37.66%)
Travel & Transportation	19	28 (37.33%)	18 (24.00%)	29 (38.67%)
Education & Reference	20	59 (43.70%)	16 (11.85%)	60 (44.44%)
Movies & TV	21	5 (50.00%)	1 (10.00%)	4 (40.00%)
Navigation & Trip	22	2 (40.00%)	1 (20.00%)	2 (40.00%)
Connected Car	23	9 (39.13%)	3 (13.04%)	11 (47.83%)
Public Transportation	24	2 (16.67%)	4 (33.33%)	6 (50.00%)
Novelty & Humor	25	3 (25.00%)	3 (25.00%)	6 (50.00%)
Home Services	26	5 (19.23%)	6 (23.08%)	15 (57.69%)
Self Improvement	27	2 (50.00%)	0 (0.00%)	2 (50.00%)
Wine & Beverages	28	1 (20.00%)	1 (20.00%)	3 (60.00%)
Music & Audio	29	94 (29.38%)	5 (1.56%)	221 (69.06%)
Shopping	30	26 (20.63%)	10 (7.94%)	90 (71.43%)
Calendars & Reminders	31	1 (14.29%)	1 (14.29%)	5 (71.43%)
Pets & Animals	32	1 (33.33%)	0 (0.00%)	2 (66.67%)
Event Finders	33	1 (25.00%)	0 (0.00%)	3 (75.00%)
Local	34	0 (0.00%)	1 (25.00%)	3 (75.00%)
Knowledge & Trivia	35	1 (8.33%)	0 (0.00%)	11 (91.67%)

B = Broken, P = Partial, C = Complete, R (Rank) ~ (B+P)/(C+1)

Sports has the largest proportion of skills with broken traceability, totaling 91% of all the skills in the subcategory. For instance, the “Western States” skill by *de Peck, Inc* collects the Device Address without any privacy document to disclose the data practices. Like many other subcategories, there are no skills with complete traceability in Novelty & Humour. When looking at the combination of broken and partially broken traceability, Games & Trivia is also one of the most problematic subcategories, particularly considering the number of skills it has. For instance, we find “Pixated Salat”, a prayer skill “to find Salah or Prayer time”. The skill is developed by *Pixated Ltd*, a strategic advertising group, and asks for the device’s location. However, the link to the policy provided is broken, and it links to the main site of the advertising group instead (<http://pixated.agency>).

It is also important to note that even in categories that do not rank high, there are also skills with concerning data collection practices, like the Education & Reference category. This subcategory has a high proportion of complete traceability (44%), but it also has a very sizable proportion of broken traceability skills (43%). For instance, in this category, we find a very interesting case of “A Sales Guy” by *VoiceXP*. This skill is supposed to provide information about Keenan, that according to his website, is a person who has been “selling something to someone for his entire life”. Users that want to know about Keenan would have to give up their mobile number, email address, full name, or device address. The developer, *VoiceXP*, is a company that does so-called *voice domain registration services*. We refer the reader to §6 for further discussions on the type of concerning practices used by skill developers.

Regarding complete traceability, it is worth highlighting the bottom part of Table 3. Particularly, the *Music & Audio* subcategory has the largest number of complete traceability skills (221), which is also a sizable proportion of skills within the subcategory (69%). This is closely followed by Shopping, with many complete traceability skills (90) that make up 71% of skills within that subcategory. Note that these two categories relate to industries with a larger tradition of offering services in the Web, where privacy has been in scrutiny for longer. Still, adding up the skills with broken and partial traceability across the two subcategories gives a total of 135 skills. Therefore, even the best ranking categories have a sizable number of broken/partial traceability skills.

Finally, and rather interestingly, the Schools and Kids subcategories are ranked in positions #7 and #5, and have 100% and 0%, and 83% and 8% of their skills broken or partially broken, respectively, for a total of 11 skills. This is especially concerning as they may collect children’s personal information, as we discuss more thoroughly in §6.

Traceability by Permissions. To understand how traceability varies across the different type of permissions, we also look at the traceability of skills per permission requested. Table 4 shows the distribution of traceability across different permissions for the 1,758 analyzed skills. The permissions are first grouped into Broken, Partial, Complete, with respect to the policies of the skills where these permissions are requested. A total of 2,616 permissions are requested (622 by skills with broken traceability, 485 by skills with partial traceability (75 of these are traceable), and 1,509 by skills that exhibit complete traceability). The most requested

TABLE 4
Distribution of traceability across different permissions for the 1,758 analyzed skills in the 5 English-speaking marketplaces.

Permission	D	R	B	P	C
Device Address	464	648	188 (29%)	130 (20%)	330 (51%)
Device Country	330	569	106 (19%)	77 (14%)	386 (68%)
Email Address	251	428	99 (23%)	77 (18%)	252 (59%)
Personal Info.	144	324	67 (21%)	41 (13%)	216 (67%)
Name	173	350	86 (25%)	82 (23%)	182 (52%)
Mobile Number	97	139	36 (26%)	37 (27%)	66 (47%)
Location Services	90	102	35 (34%)	31 (30%)	36 (35%)
Amazon Pay	33	56	5 (9%)	10 (18%)	41 (73%)
Total	1,582	2,616	622 (24%)	485 (19%)	1,509 (58%)
Unique	1,123	1,758	442 (25%)	306 (18%)	1,010 (57%)

D = Developer, R = Requested, B = Broken, P = Partial, C = Complete,

TABLE 5
Developers’ disclosure practices (Broken, Partial or Complete) in all 5 English-speaking marketplaces.

Markets	Developers	B	P	C	P+C
US	731	219	145	330	15
UK	402	197	41	146	7
CA	224	88	27	102	2
IN	215	107	23	79	5
AU	202	90	22	82	3
Total	1,730	701	258	739	32
Unique	1,123	350	207	566	21

D = Developer, B = Broken, P = Partial, C = Complete

permission is the *Device Address* which is requested 648 times by 464 developers. While *Amazon Pay* is the least asked permission requested only 56 times by 33 developers, it tends to be requested more by skills that have complete traceability. In contrast, *Location Services* permission requested by 90 unique developers is found more in skills that exhibit broken traceability.

The Good, the Bad and the Ugly Developers. We next study how many “good”, “bad” and “average” developers there are. Table 5 shows the number of developers per type of traceability considering the 5 English marketplaces. Overall, we see a total aggregate across markets of 1,730 cases where developers request permissions in their skills, out of which 1,123 are unique developers that post skills in several markets. When looking at unique developers, we see:

The Good. There are 566 developers with all their skills showing complete traceability. All the skills developed by these developers have statements in their privacy policies clearly stating and justifying the permissions they request. This accounts for about 50% of the developers. For instance, developers such as *GoVocal.AI*, *Blutag Inc*, and *Ixartz* in the US marketplace have complete traceability between the permissions their skills ask and their privacy policies.

The Bad. There are 350 developers with all their skills broken. This accounts for about 31% of the developers. Skills do not offer an adequate explanation in general when we analyze the partial skills, their privacy statements, and we look at their reviews. One example is *Tagrem Corp* in the US market, with all the skills they developed exhibiting broken traceability, as it does not acknowledge the collection of any personal information in the skills’ privacy policy while the skills do request for permissions.

The Ugly. We see 207 developers with all their skills with partial traceability. This accounts for about 18% of the developers. They appear to have a sloppy attitude when

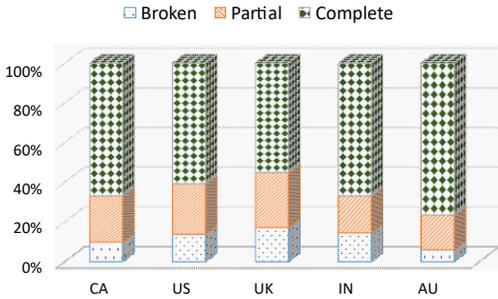


Fig. 3. Traceability in skills reusing privacy policies (English-speaking markets). Total: 1498 skills, broken: 196, partial: 358, complete: 944.

writing privacy policies and informing users of how the personal information they request is used. For example, the developer *Geekycoders* with over 30 skills in the US market have partial traceability in all of them. While the developer requests for only one type of permission (First Name), the statement “we do not collect your personal information when you use any part of our products, unless the app specifically lists in any pre-download description” only offers partial traceability since it does not explicitly state what information is collected.

This shows that both average and bad developing practices are commonplace and widespread in Alexa across marketplaces, particularly in the US. Interventions for developers, instead of just for individual skills, may also be an effective way of vetting marketplaces.

Reused Policies. We observe that some skills are reusing the privacy policies of others. Thus, we measure the prevalence of reused policies and study the impact this practice has on traceability by systematically mapping policy links to skills in the same English-speaking marketplace. Figure 3 shows an overview of our results, where 175 privacy policy links are reused across marketplaces by 1,498 skills (with the following breakdown³: CA-172, US-825, UK-180, IN-172 and AU-149) from 235 developers (CA-37, AU-26, IN-9, UK-39, US-124). Out of all the skills where we see reusing policies, only 29% offer a comprehensive policy that adequately describes their practices, 44% show broken traceability between policy documentation and data permissions, while 27% show partial traceability. For example, the skills “Curious City” and “Hits 94” developed by *Voxmatic.io* both with data permission Device Address have a privacy policy hosted in a domain that is down (<http://voxmatic.io/privacy>), though both are fully functional when spoken to.

Interestingly, although many of the skills with the same policy are published by the same developer and collect the same permissions, there are cases where the skill collects a different set of permissions (CA-111, AU-92, IN-109, UK-63, US-548). For example, in the Canada marketplace, the skills “Big Sky” and “I’m Driving” (both from developer *Philosophical Creations* and with partial traceabilities) use the privacy policy link <https://driving.big-sky-alexa.com/privacy>. However, the first asks from Alexa Notifications and Device Address permissions while the second asks for Email Address and Mobile Number. There are also some

3. The breakdown is read as follows: CA-172 means 172 skills using one of the policies in the Canadian market.

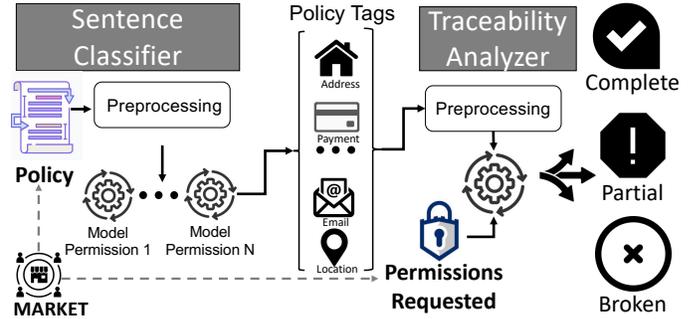


Fig. 4. Overview of SkillVet.

cases where completely different developers still reused the same privacy policy. We see 24 unique instances involving 54 different developers across all of the English speaking markets. For instance, in the US market, 4 developers, *Evezilla Ltd*, *Kavson Ltd*, *Rai Integration Ltd*, and *Mikk London*, use the same policy link <https://www.starfishmint.com/policy/privacy.html>.

Finally, another interesting observation is that while some skills reuse broken policies and, therefore, automatically inherit the broken traceability, other skills reusing policies have different traceability due to the different permissions the skill request. For instance, the “Mastering Python Networking Facts” skill by *Network Automation Nerds LLC*, which asks for Device Country and Postal Code, exhibits complete traceability with the privacy policy at “<https://www.alexa.com/help/privacy>”. However, the “Stock Price” skill by *JJ* exhibits partial traceability when the same policy is used while asking for the Lists Read Access permission, which differs from the other skill’s permissions.

4 AUTOMATED TRACEABILITY ANALYSIS

To facilitate the traceability analysis of skills, we create a system, *SkillVet*, based on machine learning and natural language processing. Given a skill, it first identifies and classifies all statements in its privacy policy that relate to data practices over personal information. We map each data statement with one or more Alexa permissions. These are the permissions the skill justifies in the privacy policy. We then compare these permissions with the ones the skill is authorized to request through the Amazon API during runtime. Depending on if the permissions requested match those found in the policy, the skill is then classified as having a *complete*, *partial*, or *broken* privacy policy. Figure 4 presents an overview of the system. The automated traceability analysis system consists of two parts, the *Sentence Classifier* and the *Traceability Analyzer*. The *Sentence Classifier* (§4.1) spawns several models that perform the mapping mentioned above. The *Traceability Analyzer* (§4.2) collects all data permissions found in the privacy policy and compares them with the original set of data permissions requested by the skill. This is done to determine the traceability between the use of a permissions and its justification in the policy.

4.1 Sentence Classifier

To identify the privacy permissions requested in each sentence automatically, we manually created the *permission-by-sentence dataset (PBSD)*. The PBSD compiles 10,409 annotated sentences from 532 original Alexa policies randomly

chosen from the TBPB described in §3 (the rest of TBPB is used as *unseen* data for the evaluation as detailed in §4.3). The annotations are tags that associate each statement to one (or more) of the Alexa permission categories of interest. That is, PBSD contains sentences with permissions needed for the traceability analysis as discussed in detail in §3.2 (*Amazon Pay, Device Address, Device country and postal code, Email Address, Location Services, Mobile Number, Name, and Personal Information*) and the negative label, named *None*, which includes sentences that do not belong to any of the Alexa permission classes considered. The distribution of tagged sentences per permission category in the PBSD is presented in Table 6. Note the uneven distribution of sentences among classes. The permissions with the largest number of statements are “Personal Information” and “Name”. Although we strive to have representative statements for all categories in Alexa, some less actionable permissions such as “Device Address” or other recently incorporated permissions like “Amazon Pay” are still not commonly used among skills, and they are therefore not easy to find and tag. To mitigate this issue, we extended the PBSD with sentences from the 350 Android privacy policies (APP-350) [33] only for permissions that map with Alexa’s data permission categories, i.e.: the *Device Address, Device country and postal code, Email Address, Location Services, and Mobile Number*.

TABLE 6
Sentence count per permission in PBSD and the extended PBSD+APP350.

Alexa Data Permission	#PBSD	#PBSD+APP350
Amazon Pay	135	135
Device Address	254	776
Device Country and Postal Code	244	518
Email Address	205	2181
Location Services	186	533
Mobile Number	97	941
Name	416	416
Personal Information	2,482	2,482
None	6,390	6,390
Total	10,409	14,372

The problem we tackle is a multi-label classification problem, as sentences can be referring to more than one permission. For instance, the sentence “We collect your name and postal address” indicates the collection of permissions *Name* and *Device country and postal code*. One well-known way to tackle multi-label classification problems is to transform the problem into a number of binary classification problems.⁴ In particular, the sentence classifier is formed by 9 binary models, one per each of the 8 Alexa data permission categories we model, and one for the sentences that do not belong to any permission category. All binary classifiers are trained following a one-vs-all strategy. For every classifier modeling a particular permission, we label all sentences relating to that permission as the positive class and all other sentences as the negative one. Given a sentence from a privacy policy, each of the classifiers then evaluates whether that sentence belongs to the permission category being tested or not. Therefore, we obtain a multi-label classification of the sentence. That is, every classifier makes

4. Note that other transformations may be possible [34], e.g., transforming the multi-label classification problem into a multi-class classification problem but that would require 2^8 classes.

an independent assessment and as a result we are able to model sentences describing more than one permission.

4.2 Traceability Analyzer

The traceability analyzer collects the data permissions output by the Sentence Classifier, and compares them with the permissions requested by the skill in two steps:

Policy Preprocessing. Given an Alexa policy, the Traceability Analyzer first splits the policy into sentences and preprocesses each of the sentences in the same way as presented in §3.2. After this, it removes sentences non-related to data permissions, such as those in which the authors provide some contact details, e.g.: “if you have any inquiries about this skill, please contact us by email”. These sentences are often structured similarly in all policies, use similar terms, and are usually wrongly assessed by classifiers, which are unable to identify that the sentence does not describe a data permission request. To detect this special set of sentences, we use a keyword blacklist containing words such as ‘contact us’ or ‘call us’. The same is done with negation terms such as “does not” or “doesn’t”. This step filters out negative sentences such as “this skill does not need your email address”, which are quite common among skill policies. If a sentence contains any of these blacklisted keywords, the sentence is ignored and not classified, as modeling negative statements is important to determine the correct meaning of a privacy policy [35].

Traceability. After the policy preprocessing, the Traceability Analyzer runs each of the 9 classifiers through all the remaining sentences, collecting all data permissions in the policy. Note that the *None* classifier prevails over the others to handle contradictions in data policies [35]. That is, we consider that a policy does not have a (proper) data statement if there are permissions associated to a policy in addition to a positive classification of the *None* classifier. Next, through the comparison between permissions requested by the skill and those automatically found in its policy, the Traceability Analyzer classifies the traceability of the skill as *broken, partial* or *complete*, in the same way as in §3.

4.3 Implementation and Evaluation

Sentence Classifier Training and Validation. Each of the 9 models is a binary SVM⁵ classifier built on top of an n-gram binary vectorizer and a tf-idf layer⁶. We test various parameters over a set of 5-fold cross-validation runs for each of the 9 models, including the size of the n-grams, the SVM loss method, the SVM alpha value, and different oversampling and undersampling strategies to balance the classes of the *permission-by-sentence* dataset (PBSD). We make sure that each data permission is represented consistently among the instances selected for training and testing in order to increase the robustness of the classifiers against sentences of all types. The parameters that consistently returned better F1 and accuracy scores over the different

5. SVM is suited for binary classification problems [36], as we do not have enough data to use deep learning techniques, and it usually performs better for NLP tasks with n-grams than other classifiers [37].

6. Note that approaches such as sentence embeddings [38] could be used, but the repetition of similar constructs among policies indicates that a simpler, n-gram classifier is also appropriate

TABLE 7
K-fold validation for the Sentence Classifier.

Alexa Data Permission	F1	Acc
Amazon Pay	.987	.991
Device Address	.919	.908
Device Country + P.C.	.921	.909
Email Address	.967	.965
Location Services	.960	.973
Mobile Number	.978	.977
Name	.977	.983
Personal Information	.986	.982
None	.993	.991

5-fold cross-validation are then used to train with all data. In a nutshell, the best parameters are: we use n-grams of size 1, 2 and 3, stratified random undersampling and oversampling as balancing strategy, *modified huber loss* as the SVM loss function, and an *svm alpha* of 10^{-5} for the SVM classifier. After undersampling and oversampling certain classes, on average, each classifier was trained and tested using between 2K and 8K of the total sentences found in the *permission-by-sentence* dataset. Interestingly, we observe that the classifiers performs best when we use *at the same time* n-grams of different sizes (1,2,3).

The final version of all binary classifiers that are deployed in *SkillVet*, are trained using 100% of the the PBSO dataset, and then evaluated using a set of extra *unseen* 523 sentences to check their performance at a sentence-permission level (see below for the evaluation of *SkillVet* as a whole). The average F1-score and accuracy metrics obtained are reported in Table 7. Note how the sentence classifiers are able to differentiate and classify each of the permissions correctly, obtaining F1-scores and an accuracy of over .9 for all data permissions.

SkillVet Evaluation. To evaluate the performance of *SkillVet* as a whole, we compare the traceability results we obtained with the *unseen* remaining subset of 972 skills and their policies from the *traceability-by-policy dataset (TBPD)*⁷. That is, the policies of these skills are not used to create the PBSO as detailed in §4.1, so they are not used at any point for training, validation, fine-tuning, or evaluation of the sentence classifiers. It takes *SkillVet* 7 mins 23 secs to process all 972 skills (an average of 0.45 secs per skill). As shown in Table 8, *SkillVet* is able to correctly classify 93.1% (905) of the previously unseen 972 policies. The highest accuracy measures are obtained when identifying broken and complete policies (98.5% and 93.9% accuracy), while partial traceability seems to work comparatively worse, with a few miss-classified as complete.

We further investigate the reasons behind the 67 errors in total (out of 972 skills), which we organize into the following three main types: i) *Class error* (32 cases, 48%) occurs when the policy is wrongly analyzed due to an error in one of the sentence classifiers producing the wrong class; ii) *Fair error* (15 cases, 22%) occurs when even a human struggles to identify the correct class because the sentence can be interpreted in different ways, e.g., the most common

7. Note that we did not consider a further 254 skills in TBPD, because their traceability is trivial: they have missing policies, dead links, or empty policies. A simple check without sentence classification is enough to assess them broken. Including them would potentially bias the evaluation away from the more complex, error-prone cases.

TABLE 8
Confusion matrix comparing *SkillVet* with a human analysis for the 972 *unseen* skills from TBPD.

<i>SkillVet</i>	Actual Traceability on 972 unseen policies			
	Broken (B)	Partial (P)	Complete (C)	Total
B	264 (98.5%)	8 (4.9%)	11 (2.0%)	283
P	0 (0.0%)	134 (81.7%)	22 (4.1%)	156
C	4 (1.5%)	22 (13.4%)	507 (93.9%)	533
Total	268	164	540	972

cases are sentences that refer to *address* but it is not clear whether it is the device address or the email address, in fact, the manual analysis for traceability would consider this as partial precisely because it is not exactly clear what address the sentence is referring to; iii) *Filter error* (20 cases, 30%) is when the error is attributed to the preprocessing stage of *SkillVet*, e.g., most of the cases are due to sentences where there is a complex contradiction, in the sense that there is a negation but associated to a positive non-collection action, actually suggesting a good practice (e.g., not shared with, not disclosed to) in an overall positive sentence.

We next discuss how one might be able to improve the performance of *SkillVet* further in light of the errors. For Class errors, one way to improve would be to have more tagged sentences in the PBSO. In general, however, *SkillVet* shows a very good performance, considering that PBSO has less than 250 instances for most of the permissions (cf. Table 6). For Fair errors, this is more challenging, but one possible improvement could consider a separate sentence classifier trained with cases where it is not clear what address is being considered, where a human looking at the wider context can make the right assessment (though in most cases as stated above, the human tagger may not be able to ascertain this either). Finally, for Filter errors, there may be room for improvement by considering a more sophisticated preprocessing approach perhaps aided with an ontology, similarly to [35]. However, this may also introduce errors in turn, and the way *SkillVet* accounts for negations and contradictions seem effective in general for this domain.

5 BEYOND API PERMISSIONS

Amazon enforces access to personal data through a permission model embedded in their APIs as discussed in §3.1. However, skills could also request personal data directly from the user without Amazon’s API. This can be done via the Web through account linking, or via conversations.

5.1 Account Linking

Account linking allows users to connect their identity with the one they use in a different system like Google, Amazon, or Facebook [39]. This is implemented in Alexa using OAuth 2.0. A total of 5,230 skills (about 4.7% of the skills in our dataset) are using the account linking feature across marketplaces (with the breakdown: FR-23%, DE-10%, US-6%, UK-6%, AU-4%, CA-4%). This is over twice the 2,608 skills (2.3%) that request for at least one of the Alexa permissions (see beginning of §3). By systematically enabling skills with account linking using a script developed with Selenium [40], we report where they connect to. To do this, we check

whether the skill name or developer’s name is in the domain used on the account linking URL, and/or whether a name of an OAuth service provider is present (for curation process and full list of OAuth providers refer to [24]).

Table 9 shows that most of the skills connect to the developer’s system, although an important number of skills rely on third-party services such as Google or Facebook for authentication. In particular, 52% of the skills with account linking connect to the developer’s site, and 29% to services like Facebook, Twitter, Amazon, or Google, among others. Some skills (14% of all 5,230) allow both, authenticating through a third-party using OAuth before being redirected to the developer’s site. Finally, the rest 5% (Unresolved) could not be labelled as developer or third-party.

TABLE 9

How skills connect user’s Alexa identity with their identity in another system.

Account Type	Number of Skills	Percent
Developer	2,720	52%
Third-Party	1,517	29%
Third-Party or Developer	732	14%
Unresolved	262	5%
Total	5,230	100%

We then look at the traceability of the skills with account linking and at least one permission. Table 10 shows that 47% of them have broken or partial traceability. While these broken and partial results are conclusive, when it comes to *complete* traceability, results may not be conclusive — recall that account linking could enable the collection of further personal data without the need to use data permissions (i.e., taking the user out of Alexa). Accounting for the exact data collected by developers through account linking is challenging because: 1) third-party sites used for account linking have all different formats and are thus difficult to scrape; 2) developers may ask for any type of personal information (at any point, and not just during the registration).

TABLE 10

Traceability results for English-speaking Markets skills with account linking and at least one permission.

Market	Skills	Broken		Partial		Complete	
		N	%	N	%	N	%
US	133	30	22.6%	35	26.3%	68	51.1%
UK	46	6	13.0%	13	28.3%	27	58.7%
CA	30	1	3.3%	6	20.0%	23	76.7%
AU	27	2	7.4%	4	14.8%	21	77.8%
IN	26	5	19.2%	5	19.2%	16	61.5%
Total	262	44	16.8%	63	24.0%	155	59.2%
Unique	164	39	23.8%	38	23.2%	87	53.1%

To study further the process of account linking and data involved, we randomly choose 70 skills and perform the account linking process manually. We see 65 skills (92.85%) requesting personal data without using the Alexa API. We next discuss some case studies we find as a result of this analysis. First, we look at “GCSE Revision”, a skill developed by *Palm UK Ltd* that helps users to prepare for exam-board specific GCSE (General Certificate of Secondary Education) science questions. While the skill asks for permissions Full Name and Email Address, we see that during the linking process the skill asks, in addition, for the user’s *address*. This is particularly problematic as the skill is effectively bypassing Alexa’s permission system, using account

linking as a proxy to obtain additional data from users. This also shows a prevalent issue we have discussed earlier (see §3.3): this skill targets students between 12 and 16 year old students. Since GCSE has access to the Full Name, the skill could have been collecting the family name of children in the UK that are preparing for the test, together with how well they perform. Note that children’s last names can be usually inferred even if the Amazon account was under their parent’s names. Therefore, the developers of GCSE could send targeted postal advertisement to the household. Finally, we also look at “DaddySays” skill by *Tagrem Corp* in the Kids category. The skill collects the email address of the users during the account linking process in addition to the information collected through the permissions Lists Read Access and Lists Write Access.

5.2 Collection via Conversation

We interact with 100 randomly-chosen skills from those that do not request permissions using the Amazon API and that do not use the account linking feature either. We see that 35 are single-interaction skills (they just provide an answer to the user without further interaction), 45 are conversational skills, and the remaining 20 return no response or responded with error (e.g., “Sorry; I am not sure about that”). Out of the 45 conversational skills, we find 3 *skills* asking for personal information directly when conversing with users, effectively bypassing the Amazon model for data collection practices. These are ‘Would You Rather for Family’ by *Voice Games*, “The Bartender” by *Pylon ai*, and “Phone Tracker” by *S4 Technology, Inc*. For instance, the latter asks for a phone number as part of the conversation.

To study conversational skills more systematically, we used our system, *SkillVet*, in tandem with recent work, *SkillExplorer* [16], which has provided an initial method to discover these skills. In particular, we studied the dataset of 100 skills (developed by 89 distinct developers) collecting personal information via conversation provided by the authors of *SkillExplorer* [16]. Out of all these skills, 82 collect personal information only through conversation, 11 collect personal information via Alexa API in addition to the one collected during a conversation, and 7 also use account linking. Table 11 summarizes the collection method used and shows the results of the traceability analysis. Overall, we see that 35% the skills exhibit broken traceability. In this case, all issues are attributed to personal data collected alone during a conversation. We also see that 20% partially disclose their data practices in their privacy policy, and the remaining 45% exhibit complete traceability. One good example of a broken skill is the “Praise Me” skill by *Jackson Jacob* with skill id B07G5B4P7R and a customer rating of 4.8 out of 5 in the US market. This skill is available across all five English speaking countries. According to its description, this skill is supposed to praise the user after asking for their name. However, the skill not only request the name of the user during the conversation without declaring the permission for it, but it also fails to disclose its data practices in a privacy policy. Furthermore, the skill has a privacy policy link that takes the user to a porn site. There are also two developers with skills having different traceability outputs. *Voice first tech* has skills that exhibit partial and complete traceability (1 skill in

TABLE 11
Traceability Result with SkillVet on Subset of skills found in [16]
collecting personal information via conversation.

How PII were collected	Traceability			Total	
	B	P	C	Skills	Dev
Conversation + Account Linking		1	6	7	5
Conversation + Alexa API		5	6	11	9
Conversation Only	35	14	33	82	79
Total number of Skills	35	20	45	100	–
Total number of Developers	35	16	40	–	89

TABLE 12
Confusion matrix comparing SkillVet with a human-centered
analysis for the set of the skills given by [16].

SkillVet	Actual Traceability			Total
	Broken (B)	Partial (P)	Complete (C)	
B	33 (94.3%)	2 (8.7%)	0 (0%)	35
P	1 (2.9%)	17 (73.9%)	2 (4.8%)	20
C	1 (2.9%)	4 (17.4%)	40 (95.2%)	45
Total	35	23	42	100

each category), while *Volley inc.* has broken and complete traceability skills (1 broken, 2 complete skills).

We also compare the results given by SkillVet with a human-based traceability analysis made by one of the authors for the 100 skills. As shown in Table 12, SkillVet achieves a very good accuracy, being able to correctly identify the traceability for 90 of the 100 skills. It is particularly very good at spotting complete (95.2%) and broken (94.3%) traceability, and good for partial (73.9%). Of the 10 errors in total, 8 are Class errors, 1 Filter error and 1 Fair error. These results align with what we report for the unseen 972 skills that request permissions through Alexa API in § 4.3.

Finally, note that SkillExplorer is restricted by design to those skills that have unique invocation names [16]. Future work should look into methods to discover this type of skills considering skills that share invocation names, which are very prevalent, as shown in this paper (§ 2). Importantly, however, SkillVet would easily integrate with any such tool when provided with the data the skills ask for and any accompanying privacy policies.

6 DISCUSSION

6.1 Key Findings and Limitations

Potential privacy violations. We see bad privacy practices in about 748 skills (43%) of those that request permissions in English-speaking marketplaces⁸, involving 557 (50%) of the developers with skills that request permissions. Although it could be some developers have multiple Amazon developer accounts and the real number of entities behind those accounts are less, this still paints a very worrying picture. Particularly worrying cases are those in categories like Kids, Schools, and Education, which exhibit broken (72 skills in total) and partial (17 skills in total) traceability. Beyond the traceability itself, this may have additional implications, as skills in those categories might be subject to even tighter privacy regulations for children, such as the US Children’s Online Privacy Protection Act (COPPA) of 1998 [41]. Other categories, such as the Utilities category, also have countless

8. Recall English-speaking skills are 82.7% of the total

examples of unjustified collection of personal information – i.e., with broken traceability. Skills, such as “Mock Interview” by *Graylogic Technologies*, ask for permissions — like Device Address — without acknowledging such collections in its privacy document. Furthermore, we even see a case of a skill collecting information, while stating that it is not currently used. This is the case of the “Bin reminder” skill by *Shane* that collects the user’s device address.

Bulk skill creation. There are a few developers that publish hundreds and even thousands of skills (cf. §2). We see some developers using automated systems to develop and/or deploy skills. This democratizes the creation of skills, but also lowers down the entry barriers for criminals and can foster the commoditization of unwanted skills. Previous works have recently shown that miscreants leverage online app generators to publish unwanted apps in Android [42]. We believe that Amazon Alexa will suffer from the same challenges than Google or Apple with the vetting of apps in Android or iOS markets [43]. In fact, recent research seems to point towards issues with the current vetting process in Alexa [17]. We posit that the research community needs to develop detection mechanisms tailored to the SPA ecosystem. However, one important limitation is the lack of access to the skill software (e.g. deployed in the cloud), which also constrains our analysis.

Market migration. Amazon allows developers to publish the same skill in different markets. Our study shows that there is an important overlap of skills across markets, which includes the use of the exact same privacy policy. This has important regulatory implications derived from the different local regulations in a globalized ecosystem, similar to what we have seen for cookies in the wake of GDPR [44], [45], and the different languages across marketplaces. Note, however, that skills that are pushed from English-speaking markets to other markets now go through a language migration process, and privacy policies may differ for the same skill [27]. This process was not in place at the time we started our study. As future work, we plan to re-crawl all policies and study how many have changed with the skill migration feature. Another limitation is that we only conduct the traceability analysis for English-speaking marketplaces, as this is the source language of our ground-truth. We could overcome this limitation using automated translation, but this could introduce errors in the classification. Evaluating how our model performs with translated policies is part of our future work.

6.2 Developers’ Business Models

Our key findings warrant a further discussion about the motivations developers may have to develop skills. In this regard, our study identifies over 47K developers that have contributed to the Alexa marketplace with a rich ecosystem of skills (cf. Table 1 in §2). Since Amazon forbids advertisements in skills, an important open question is what do these developers gain and what are their motivations. First, companies may be interested in offering a voice-over interaction with the user through Alexa. Examples are skills in the Travel & Transportation category (e.g. to order a ride) or in the Food & Drink category (e.g. to order a pizza). In these situations, developers need to bind the user’s Amazon

account with the external service via account linking, as mentioned before. Interestingly, we only observe 5% of the skills requesting account linking, which is circa 5K unique skills (cf. §5.1). We can, thus, conclude that account linking is not yet the primary motivation for most developers.

Amazon promotes in-skill purchasing and paid subscriptions, which allows developers to sell premium content in skills such as extra features, in-game elements and interactive stories. This skill content can be offered as: i) subscriptions, where developers charge users recurrently to access premium features, ii) one-time purchases, where users pay once to have permanent access to the premium features and content, and iii) consumables, where content can be purchased, exhausted and purchased again (e.g., extra-lives in games [46]). In-skill purchasing is not yet supported by all markets. At crawling time, we only see 1,535 (1.37%) skills with in-skill purchasing or paid subscriptions in the US and UK. Prices for in-skill products range considerably, from 0.99 to 99 USD/GBP. “Sleep Sounds: Hair Dryer”, “Sleep Sounds: Harp Sounds”, and “Sleep Sounds: Heavy Rain” developed by *Voice Apps, LLC*. are examples supporting in-skill purchases. However, judging by the number of skills with in-skill purchasing, we can conclude that this is not the main source of income yet. Developers may also offer paid skills in the future, but this is not an option at writing time.

Overall, we see that most of the skills in our dataset are free (both to enable and to use fully). The total number of skills that are free and do not have in-skill purchasing is 110,494 (98.6%). Out of these, 106,289 (94.9%) do not offer account linking either. Therefore, though there might be small incentives, such as the Amazon incentive program [47], which rewards skill creation (e.g. with a smart plug), or the simple curiosity to develop in a new environment, we conclude that *there is an ample number of skill developers for which there is no clear business model*. Some of these skills might be in the data monetization business. An obvious example is “Pixated Salat”, the prayer skill developed by a large advertisement company (c.f., §3.3). Another example is the “Autochartist” skill, which states in their privacy policy that the data they collect may be used “to provide you with news, special offers and general information about other goods, services and events which we offer that are similar to those that you have already purchased or enquired about unless you have opted not to receive such information”. It is, however, unclear how users of skills can opt out.

6.3 Mitigation and Responsible Disclosure

Tackling bad practices from third-party developers is a challenge to systems security. Having a well-defined — fine-grained — permission model is an important step forward. However, we have seen that: 1) developers can bypass it, and 2) developers do not offer transparency about how they utilize users’ information. We next present a range of countermeasures that could mitigate the risk of enabling third-party developers access to users’ data. First, Amazon should not allow developers to reuse policies verbatim as this is error-prone, as we have seen. Second, both Amazon and the developers should perform a thorough review of the traceability themselves. *SkillVet* can help in the process of systematically reviewing permissions and privacy

policies. However, further research efforts are needed to ensure that policy notices are relevant, actionable and understandable by users [48]. Third, publishing the code of the skills or the binaries will considerably foster research in the area, similar to what we have witnessed in Android [49]. While approaches to detect unwanted skills can be adopted from Android, the Alexa ecosystem has unique features like voice recognition that present novel challenges (e.g., voice masquerading attack [15] — malicious skills that pretend to hand over the control to another skill). Finally, another important mitigation is to increase user awareness, as users are known to have incorrect or at best incomplete mental models of how assistants such as Alexa work [2], [50]. This may be even more challenging when users interact with skills aside from the policy enforcer, such as during the account linking process or with conversational skills. There have been works to defend users against malicious activities on online services (e.g., spam [51], or phishing detection [52], [53]). Solutions in this direction will have to consider that the assistants’ nature brings, again, unique challenges. We refer the reader to [7] for a security and privacy review of countermeasures for Alexa-like assistants. **Responsible Disclosure:** We perform a responsible disclosure process, starting from mid August 2020, as follows. First, we notify all skill developers who are not engaging in good data practices whenever we have their contact details. Second, we also report our findings to Amazon and have confirmed that the skill store team have taken action.

As a result, we have already seen that 150 skills — 13% of which we identified issues for — no longer pose a threat to users at the time of writing: 137 of these skills have been vetted and are no longer available on Alexa, while 13 of them now have complete traceability.

7 RELATED WORK

Skill Measurement, Privacy and Security. There have only been limited related work on SPA third-party skills and the mechanisms they use in the discovery process [54]. There has been work on measuring the availability of privacy policies for skills [55], and on whether the vetting process in Alexa can be compromised [17] (by crafting 234 mock policy-violating skills). However, previous works do not focus on data practices by actual skills developers and do not perform a traceability analysis between privacy policies and data permissions asked by the skills. The authors in [56] used skills description as a baseline to detect inconsistent privacy policies, but not the actual data permissions collected by the skills through Alexa API. This is a limitation as many developers don’t often mention what permissions they collect when describing their skills. For example, skills developed by *Vipology* such as “100.7 BIG”, “FM101.7”, “KISS FM 101.7”, and “The Wolf 103.3” request for Device Country and Postal Code through Alexa API. However, *Vipology* did not mention this permission in its voice-app descriptions. Besides, the permissions mentioned by the developers in the skill descriptions could be different from what the skills are requesting for. A good example is the “Interflora” by *Interflora British Unit*. This skill collects “Full Name”, “Email Address”, “Mobile Number”, and “Amazon

Pay” permissions through Alexa API but only mention the Amazon pay permission in its skill description.

Authors in [16] conducted a measurement study to understand skills’ behaviour by building a system called SkillExplorer. This system interacts with skills and identifies those that request personal information through conversations bypassing developer specifications. However, the system only works with a limited number of skills: those with a unique invocation name, so they missed thousands of skills with similar or same invocation name — cf. Section 2.3. Lentzsch et al. [28], in parallel to our work, provided an initial look at the effectiveness of skill privacy policies. We, however, offer a much more nuanced view of the data practices in third-party skills and their developers. This allows us to uncover many more partial traceability cases. Our crawling strategy also allows us to analyze the traceability of 1,758 unique skills (56% more skills). Together, this brings the number of skills with complete traceability to 57% (instead of the 77% reported in [28]). Besides, their automated method that uses PoliCheck [57] underneath has 83.3% precision when detecting complete policies, whereas SkillVet has a precision of 95%. This better precision could be attributed, among others, to SkillVet covering all data types in Alexa (incl. Lists and Amazon Pay), while in [28] the ontology they use for PoliCheck only considers a subset of the data types. Further, PoliCheck was trained *only* with android privacy policies, while SkillVet was trained with Alexa skill privacy policies, only complemented with a few android privacy policies (c.f 4.1).

Research on skills also focused on exploiting the speech recognition capabilities of SPA by crafting malicious skill invocation names, usually leveraging phonetic similarity [14], [15]. More recently, a study in [13] shown how malicious developers could potentially sneak malicious code into their software via the application backend after vetting, and another interesting study looked into how a malicious third-party skill could covertly reword responses from legitimate sources to intentionally introduce misperception about the reported events [19]. There is also a report of how Amazon bug could allow an attacker to access a user’s install Alexa skills and user’s Amazon profile information through a phishing attack [20]. Our threat model in this paper is different from these works (see Section 2.1).

Privacy Traceability Analysis. Our traceability analysis relates to previous work that focuses on ensuring software requirements compliance with governing legal texts and privacy policies [30]. In particular, previous work analyze privacy traceability in domains such as Online Social Networks [31]Social Media Aggregators [32], and Smartphone Applications [33]. We, however, focus in this work on analyzing the privacy traceability of skills. Also, in [33], the authors analyze Android apps characteristics for likely non-compliance with privacy requirements and some selected and applicable laws. In our case, we only analyze non-conformity with privacy requirements based on the requested data permissions by the skills. Finally, PolicyLint [35], a tool to identify policy contradictions in Android, is complementary to our work, and it could be used to understand if there are ill-defined policies in Alexa. While self-contradictions are orthogonal to the traceability analysis we do, SkillVet accounts for negations in a way

proven effective for traceability analysis. However, we plan to explore the use of ontologies as in [35] as future work.

8 CONCLUSION

The Amazon skill ecosystem has grown rapidly without a clear business model. We have presented a measurement study that provides a large-scale analysis of this ecosystem, analyzing over 199k third-party skills. By looking at the distribution of skills and the diversity of developers, we have shown current practices in the wild and offered a unique understanding of developers’ motivations. As a key novelty, we propose an automated system, SkillVet, that analyzes the traceability of permissions. While transparency is paramount to let users make informed decisions about disclosing their data, our result shows that 43% of skills do not comprehensively disclose their data practices. We have also seen how skills may bypass Alexa’s permission system by requesting personal information without using their APIs (e.g., on an external domain via account linking). This indicates that even skills with a complete privacy policy can pose a risk to users. We have discussed the most concerning practices and the implications of our findings, with our responsible disclosure helping address 13% of the issues found at the time of writing.

Acknowledgments

We thank the anonymous reviewers. This research was partially funded by EPSRC under grant EP/T026723/1 and the “Ramon y Cajal” Fellowship RYC-2020-029401-I. Edu was supported by the Petroleum Technology Development Fund (PTDF) for his PhD.

REFERENCES

- [1] T. Ammari, J. Kaye, J. Y. Tsai, and F. Bentley, “Music, search, and iot: How people (really) use voice assistants,” *ACM Trans Computer-Human Interaction (TOCHI)*, vol. 26, no. 3, pp. 1–28, 2019.
- [2] N. Abdi, K. M. Ramokapane, and J. Such, “More than smart speakers: Security and privacy perceptions of smart home personal assistants,” in *USENIX Symposium on Usable Privacy and Security (SOUIPS)*, 2019.
- [3] N. Abdi, X. Zhan, K. M. Ramokapane, and J. Such, “Privacy norms for smart home personal assistants,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–14.
- [4] B. Kinsella, “Nearly 90 million u.s. adults have smart speakers,” <https://voicebot.ai/2020/04/28/nearly-90-million-u-s-adults-have-smart-speakers>, 2020, [Online; last accessed May-2020].
- [5] Kinsella, Bret, “Amazon Launches New Alexa App Updates and Takes Its Mobile Strategy Another Step Forward,” <https://perma.cc/CYJ6-9JSH>, 2020.
- [6] B. Kinsella, “Alexa becomes a chatbot – you can now talk to alexa by typing,” <https://perma.cc/877Y-C3DV>, 2020, [Online; accessed Dec’20].
- [7] J. S. Edu, J. M. Such, and G. Suarez-Tangil, “Smart Home Personal Assistants: A Security and Privacy Review,” *ACM Computing Surveys*, 2020.
- [8] H. Abdullah, K. Warren, V. Bindschaedler, N. Papernot, and P. Traynor, “Sok: The faults in our asrs: An overview of attacks against automatic speech recognition and speaker identification systems,” 2020.
- [9] L. Xinyu, T. Guan Hua, A. X. Liu, L. Chi Yu, and T. Xie, “The insecurity of home digital voice assistants: Amazon alexa as a case study,” *CoRR*, 2017.
- [10] V. Kepuska and G. Bohouta, “Next-generation of virtual personal assistants (microsoft cortana, apple siri, amazon alexa and google home),” *CCWC*, pp. 99–103, 2018.

- [11] H. Feng, K. Fawaz, and K. G. Shin, "Continuous authentication for voice assistants," *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking - MobiCom '17*, 2017.
- [12] F. Brasser, T. Frassetto, K. Riedhammer, A.-R. Sadeghi, T. Schneider, and C. Weinert, "Voiceguard: Secure and private speech processing," *Interspeech 2018*, 2018.
- [13] SRLabs, "smart spies: alexa and google home expose users to vishing and eavesdropping," 2019. [Online]. Available: <https://srlabs.de/bites/smart-spies/>
- [14] D. Kumar, R. Paccagnella, P. Murley, E. Hennenfent, J. Mason, A. Bates, and M. Bailey, "Skill squatting attacks on amazon alexa," in *27th USENIX Security Symposium*, 2018, pp. 33–47.
- [15] N. Zhang, X. Mi, X. Feng, X. Wang, Y. Tian, and F. Qian, "Dangerous skills: Understanding and mitigating security risks of voice-controlled third-party functions on virtual personal assistant systems," in *IEEE Symposium on Security and Privacy (SP)*, 2019.
- [16] Z. Guo, Z. Lin, P. Li, and K. Chen, "Skilleplorer: Understanding the behavior of skills in large scale," in *29th USENIX Security Symposium*, 2020, pp. 2649–2666.
- [17] 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 *ACM Conference on Computer and Communications Security (CCS)*, 2020.
- [18] Amazon, "Host a custom skill as a web service," <https://developer.amazon.com/en-US/docs/alexa/custom-skills/host-a-custom-skill-as-a-web-service.html>, 2020, [Online; last accessed 29-May-2020].
- [19] F. Sharevski, P. Jachim, P. Treebridge, A. Li, A. Babin, and C. Adadevoh, "Meet malexa, alexa's malicious twin: Malware-induced misperception through intelligent voice assistants," *International Journal of Human-Computer Studies*, vol. 149, 2021.
- [20] BBC, "Amazon alexa security bug allowed access to voice history," <https://www.bbc.co.uk/news/technology-53770778>, 2020, [Online; last accessed 29-September-2021].
- [21] Amazon.com, "Configure permissions for customer information in your skill," <https://developer.amazon.com/en-US/docs/alexa/custom-skills/host-a-custom-skill-as-a-web-service.html>, 2019, [Online; last accessed 16-12-2019].
- [22] A. Bartel, J. Klein, Y. Le Traon, and M. Monperrus, "Automatically securing permission-based software by reducing the attack surface: an application to android," in *2012 Proceedings of the 27th IEEE/ACM International Conference on Automated Software Engineering*, 2012, pp. 274–277.
- [23] H. Wang, Y. Guo, Z. Tang, G. Bai, and X. Chen, "Reevaluating android permission gaps with static and dynamic analysis," in *2015 IEEE Global Communications Conference (GLOBECOM)*, 2015.
- [24] J. S. Edu, X. Ferrer-Aran, J. M. Such, and G. Suarez-Tangi, "Skillvet: Automated traceability analysis of amazon alexa skills (extended version)," 2021. [Online]. Available: <https://arxiv.org/abs/2103.02637>
- [25] S. Perez, "Amazon Alexa now has over 1,000 Skills, up from 135 in January," <https://techcrunch.com/2016/06/03/amazon-alexa-now-has-over-1000-skills-up-from-135>, 2016, [Online; last accessed May-2020].
- [26] Amazon, "Develop Skills in Multiple Languages," <https://developer.amazon.com/en-US/docs/alexa/custom-skills/develop-skills-in-multiple-languages.html>, n.d., [Online; last accessed May-2020].
- [27] —, "Alexa Skills English Variants Migration," <https://developer.amazon.com/es-ES/docs/alexa/faq/english-migration-faq.html>, 2019, [Online; last accessed 24-September-2019].
- [28] 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 *Network and Distributed Systems Security (NDSS) Symposium*, 2021.
- [29] Amazon, "Security Testing for an Alexa Skill," <https://developer.amazon.com/docs/custom-skills/security-testing-for-an-alexa-skill.html>, 2019.
- [30] J. D. Young and A. I. Anton, "A method for identifying software requirements based on policy commitments," in *IEEE International Requirements Engineering Conference*, 2010, pp. 47–56.
- [31] P. Anthonysamy, P. Greenwood, and A. Rashid, "Social networking privacy: Understanding the disconnect from policy to controls," *Computer*, no. 6, pp. 60–67, 2013.
- [32] G. Misra, J. M. Such, and L. Gill, "A privacy assessment of social media aggregators," in *IEEE/ACM Int Conf on Advances in Social Networks Analysis and Mining (ASONAM)*, 2017, pp. 561–568.
- [33] S. Zimmeck, P. Story, D. Smullen, A. Ravichander, Z. Wang, J. Reidenberg, N. C. Russell, and N. Sadeh, "Maps: Scaling privacy compliance analysis to a million apps," *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 3, pp. 66–86, 2019.
- [34] G. Madjarov, D. Kocev, D. Gjorgjevič, and S. Džeroski, "An extensive experimental comparison of methods for multi-label learning," *Pattern recognition*, vol. 45, no. 9, pp. 3084–3104, 2012.
- [35] B. Andow, S. Y. Mahmud, W. Wang, J. Whitaker, W. Enck, B. Reaves, K. Singh, and T. Xie, "Policylint: investigating internal privacy policy contradictions on google play," in *USENIX Security Symposium*, 2019, pp. 585–602.
- [36] C. Cortes and V. Vapnik, "Support-vector networks," *Machine Learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [37] S. Aiyar and N. P. Shetty, "N-gram assisted youtube spam comment detection," *Procedia computer science*, vol. 132, pp. 174–182, 2018.
- [38] M. Pagliardini, P. Gupta, and M. Jaggi, "Unsupervised learning of sentence embeddings using compositional n-gram features," *arXiv:1703.02507*, 2017.
- [39] Amazon, "Understand Account Linking," <https://developer.amazon.com/docs/account-linking/understand-account-linking.html>, 2020, [Online; last accessed 15-October-2020].
- [40] Selenium, "Selenium webdriver," <https://www.selenium.dev/>, 2020, [Online; last accessed 15-October-2020].
- [41] F. T. Commission., "Children's online privacy protection rule; final rule," 2013. [Online]. Available: <https://www.ftc.gov/system/files/2012-31341.pdf>
- [42] P. Kotzias, J. Caballero, and L. Bilge, "How did that get in my phone? unwanted app distribution on android devices," in *IEEE Symposium on Security and Privacy (SP)*, 2021.
- [43] H. Hu, L. Yang, S. Lin, and G. Wang, "Security vetting process of smart-home assistant applications: A first look and case studies." *ArXiv*, vol. abs/2001.04520, 2020.
- [44] X. Hu, G. Suarez-Tangil, and N. Sastry, "Multi-country study of third party trackers from real browser histories," in *IEEE European Symposium on Security and Privacy*, 2020.
- [45] C. Iordanou, G. Smaragdakis, I. Poese, and N. Laoutaris, "Tracing Cross Border Web Tracking," in *ACM IMC*, 2018.
- [46] Amazon, "Understand In-Skill Purchasing," <https://developer.amazon.com/docs/in-skill-purchase/isp-overview.html>, 2019, [Online; last accessed 24-September-2019].
- [47] —, "Build for alexa and earn fantastic perks," 2019. [Online]. Available: <https://developer.amazon.com/en-GB/alexa/alexa-skills-kit/alexa-developer-skill-promotion>
- [48] F. Schaub, R. Balebako, and L. F. Cranor, "Designing effective privacy notices and controls," *IEEE Internet Computing*, vol. 21, no. 3, pp. 70–77, 2017.
- [49] K. Chen, P. Wang, Y. Lee, X. Wang, N. Zhang, H. Huang, W. Zou, and P. Liu, "Finding unknown malice in 10 seconds: Mass vetting for new threats at the google-play scale," in *24th {USENIX} Security Symposium ({USENIX} Security 15)*, 2015, pp. 659–674.
- [50] M. Tabassum, T. Kosinski, and H. R. Lipford, "'i don't own the data': End user perceptions of smart home device data practices and risks," in *USENIX SOUPS*, 2019.
- [51] G. Stringhini, C. Kruegel, and G. Vigna, "Detecting spammers on social networks," in *ACSAC*, 2010, pp. 1–9.
- [52] I. Fette, N. Sadeh, and A. Tomasic, "Learning to detect phishing emails," in *WWW*, 2007, pp. 649–656.
- [53] Z. Dou, I. Khalil, A. Khreishah, A. Al-Fuqaha, and M. Guizani, "Systematization of knowledge (sok): A systematic review of software-based web phishing detection." *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2797–2819, 2017.
- [54] R. W. White, "Skill discovery in virtual assistants," *Communications of the ACM*, vol. 61, no. 11, pp. 106–113, 2018.
- [55] A. Alhadlaq, J. Tang, M. Almaymoni, and A. Korolova, "Privacy in the amazon alexa skills ecosystem," *Star*, vol. 217, no. 11, 1902.
- [56] S. Liao, C. Wilson, L. Cheng, H. Hu, and H. Deng, "Measuring the effectiveness of privacy policies for voice assistant applications," in *ACSAC*, 2020, pp. 856–869.
- [57] B. Andow, S. Y. Mahmud, J. Whitaker, W. Enck, B. Reaves, K. Singh, and S. Egelman, "Actions speak louder than words: Entity-sensitive privacy policy and data flow analysis with pollicheck," in *29th {USENIX} Security Symposium ({USENIX} Security 20)*, 2020, pp. 985–1002.