SKILLPOV: Towards Accessible and Effective Privacy Notice for Amazon Alexa Skills

Jingwen Yan*†, Song Liao*‡, Mohammed Aldeen†, Luyi Xing§, Danfeng (Daphne) Yao¶, Long Cheng†
†Clemson University, {jingwey, mshujaa, lcheng2}@clemson.edu
‡Texas Tech University, song.liao@ttu.edu
§Indiana University Bloomington, luyixing@iu.edu
¶Virginia Tech, danfeng@vt.edu

Abstract*—Despite the popularity and many convenient features of Amazon Alexa, concerns about privacy risks to users are rising since many Alexa voice-apps (called skills) may collect user data during the interaction with Alexa devices. Informing users about data collection in skills is essential for addressing their privacy concerns. However, the constrained interfaces of Alexa pose a challenge to effective privacy notices, where currently Alexa users can only access privacy policies of skills over the Web or smartphone apps. This in particular creates a challenge for visually impaired users to make informed privacy decisions. In this work, we propose the concept of *Privacy Notice over Voice*, an accessible and inclusive mechanism to make users aware of the data practices of Alexa skills through the conversational interface: for each skill, we will generate a short and easily understandable privacy notice and play it to users at the beginning of the skill in voice. We first conduct a user study involving 52 smart speaker users and 21 Alexa skill developers to understand their attitudes toward data collection and the *Privacy Notice over Voice* mechanism. 92.3% of participants liked the design of *Privacy Notice over Voice* and 70.2% of participants agreed that such mechanism provides better accessibility and readability than traditional privacy policies for Alexa users. Informed by our user study results, we design and develop a tool named SKILLPOV (Skill's Privacy Notice over Voice) to automatically generate a reference implementation of *Privacy Notice over Voice* through static code analysis and instrumentation. With comprehensive evaluation, we demonstrate the effectiveness of SKILLPOV in capturing data collection (91.3% accuracy and 96.4% completeness) from skill code, generating concise and accurate privacy notice content using ChatGPT, and instrumenting skill code with the new privacy notice mechanism without altering the original functionality. In particular, SKILLPOV receives positive and encouraging feedback after real-world testing conducted by skill developers.

I. INTRODUCTION

In recent years, Voice Personal Assistants (VPAs) such as Amazon Alexa and Google Assistant have gained substantial popularity. Predictions indicate that by 2024, the number of VPA devices could surpass 8.4 billion, exceeding the global population [16]. Amazon Alexa has a "dominant share" of the U.S. market for smart speakers and 64% of Americans own an Amazon Echo as of 2023 [21]. A notable feature of Amazon Alexa is that it allows third-party developers to build their own voice-apps (called skills) to the skills store, which is the largest voice-app marketplace and currently boasts over 130,000 published skills [26]. This greatly expands Alexa's capabilities, offering functions from weather forecasts, restaurant reservations to controlling smart home devices. Despite many advantages, concerns have been raised about the privacy risks to users as Alexa skills may collect users' personal data [39], [12]. Lau *et al.* conducted interviews with both Alexa users and non-users, and revealed that privacy concerns could be the main deterring factor for new users [44].

Privacy-aware interaction with Alexa by notifying users about the data collection of skills is of significant importance to address the user's privacy concerns. On the other hand, law and policy makers around the world are increasingly recognizing the importance of privacy notices and enacting comprehensive laws and regulations, *e.g.*, the European General Data Protection Regulation (GDPR) [10], the California Consumer Privacy Act (CCPA) [6], and the Children's Online Privacy Protection Act (COPPA) [7]. Specifically, Article 13 of the GDPR mandates that organizations provide clear information to individuals when collecting their personal data, emphasizing the necessity for companies to inform users about data collection practices. The transparent privacy notice also serves the goal of ensuring Alexa services' legal and regulatory compliance.

The Amazon Alexa platform requires that skills collecting personal information must have a privacy policy [1], which is a legal document for notifying users about a skill's data practices, and the user's rights related to the data. It is crucial for app developers to provide a privacy policy that aligns with privacy regulations, as failure to do so can result in substantial financial penalties. A notable instance occurred in 2019 when Google faced a €50 million fine from the French government due to insufficient privacy policies that did not fully comply with GDPR [11]. However, privacy policies are frequently lengthy, complex, and challenging for users to comprehend [55]. The intricate legal jargon and dense presentation may deter users from thoroughly understanding, leading many to skim or entirely skip reading these policies [35]. These challenges are also present in the Amazon Alexa platform [49]. What's more concerning is that the constrained interfaces on

*The first two authors contributed equally to this work.

Network and Distributed System Security (NDSS) Symposium 2025
24-28 February 2025, San Diego, CA, USA
ISBN 979-8-9894372-8-3
https://dx.doi.org/10.14722/ndss.2025.240677
www.ndss-symposium.org

Alexa devices pose a unique challenge to effective privacy notice. Privacy policies of skills are available on the store's webpages in text form, where users have to access them over the Web or through smartphone apps. Thus, these privacy policies are inaccessible to users who mainly use Alexa services through the conversational interface. This in particular creates a challenge for people with special needs (*e.g.*, the visually impaired) to make informed privacy decisions [53].

In this work, we introduce *Privacy Notice over Voice*, an accessible and inclusive privacy notice mechanism to enable users for privacy-aware interaction with Alexa. Our key idea is to enhance Alexa skills with the new privacy notice functionality by adding extra logic in the code to notify users about the skill's data practices through the voice interface (*i.e.*, playing a short privacy notice at the beginning of skills in voice). We identify several challenges to achieve *Privacy Notice over Voice*. 1) There is a lack of understanding regarding users' and developers' attitudes toward VPA's privacy notice, such as their preferred format, length, and placement. 2) The lack of efficient methods to automatically generate privacy notice content in a concise, accurate, and easily understandable manner from skill code poses a challenge to ensuring effective privacy notification. 3) The unique code structure of Alexa skills also presents a challenge in integrating privacy notices into the skill code, as it requires careful consideration of the interaction between the front-end and back-end code. A new tool for instrumenting skills with the *Privacy Notice over Voice* functionality at the source code level is needed.

To address these challenges, we propose and develop a tool, named SKILLPOV (Skill's Privacy Notice over Voice), to automatically generate an easy-to-digest privacy notice and instrument skills with a customized mechanism to make users aware of the data practices of a skill through the conversational interface. SKILLPOV significantly reduces the burden for developers on revising their legacy code to add the *Privacy Notice over Voice* functionality for more accessible and inclusive privacy notice. In summary, we make the following contributions:

- **New insights from user study on privacy notice.** We conducted a user study to understand users' and developers' attitudes towards data collection and their expectations on different aspects of *Privacy Notice over Voice*. 92.3% of users expressed that they like the design of *Privacy Notice over Voice*. 70.2% of users acknowledged that *Privacy Notice over Voice* provides better readability and accessibility than traditional privacy policies, potentially benefiting individuals with visual impairments. 81% of Alexa skill developers showed interest in using a tool that can automatically provide the *Privacy Notice over Voice*. Our user study provides new insights regarding the attitudes of both users and developers towards VPA's privacy notice, which guides our design of the SKILLPOV.
- **A new tool development.** We designed and developed the tool SKILLPOV for developers to automatically generate *Privacy Notice over Voice* based on skill code. SKILLPOV captures data practice (*i.e.*, data collection and data sharing/storage) in skill codes, generates a concise, understandable, and accurate privacy notice using LLMs (large language models), and seamlessly inserts the generated privacy notices into the skill codes without altering the original functionality. We shared the SKILLPOV and related results to facilitate future research [1].
- **Evaluation.** We conducted a comprehensive evaluation of SKILLPOV's performance, including functionality, reliability, readability, accuracy, and completeness. By comparing privacy notices generated by SKILLPOV with traditional privacy policies, we demonstrated the effectiveness and advantages of SKILLPOV. It performs effectively during real-world testing conducted by skill developers. SKILLPOV achieves an accuracy of 91.3% and completeness of 96.4% in data collection, which are much higher than that of the current privacy policies of skills (35.3% and 50%, respectively).

II. BACKGROUND

In this section, we first present how skills can collect user data and their code structure. Then, we discuss the threat model of this work.

A. Data Collection and Privacy Disclosure in Alexa Skills

Alexa skills can collect user data through conversational interactions and permission requests. For example, Figure 1 depicts a skill named "Quotify", which contains both types of data collection. Upon activation with the phrase "Alexa, open Quotify", a welcome message is presented, followed by the phrase "What's your name?" When users respond, the skill is able to gather users' names during the conversation. Additionally, developers can request permission to obtain specific types of data directly from users' Alexa accounts with their consent [8]. In Figure 1, "Quotify" requests permissions to access the user's device address, full name, mobile number and email address information. Users have the option to grant these permissions to the skill when it is first invoked.

The Amazon Alexa platform requires each skill that collects user data to provide a privacy policy link [1]. However, to visit the privacy policy of a skill and understand what data it will collect, users need to first find the skill from a list of similar skills on the skills store and then visit the privacy policy on the skill's listing webpage or the Amazon Alexa companion app on their smartphones. This process makes privacy policies difficult for users to access. Meanwhile, as revealed by previous works [60], [31], [59], [49], the quality of privacy policies written by third-party skill developers is unsatisfactory and usually does not reflect the actual data collection behaviors in skills. Moreover, most privacy policies are lengthy and require a considerable time to read out such documents, leading users to skip reading them.

B. Skill Code Structure

A skill's code comprises both a front-end interaction model (*i.e.*, front-end code) and back-end code. The front-end code

[1]The details of our tool and user study with results are available at https://github.com/CUSecLab/SkillPoV.

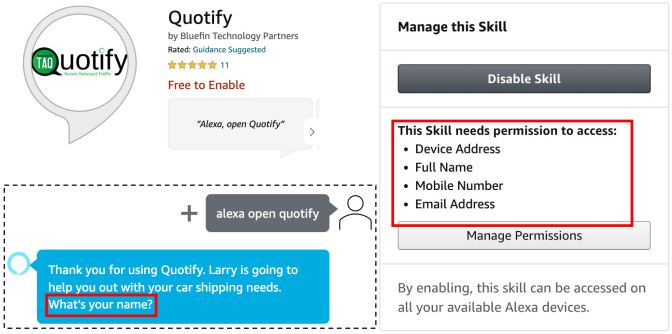

Figure 1: A skill that collects user data through both conversation channel and permission request.

defines all users' possible spoken input (*i.e.*, words, phrases) when they interact with a skill. The back-end code processes user input, executes skill functions, and generates appropriate skill responses. The front-end and back-end code are tightly coupled during skill running. Additionally, the skill's code includes a manifest file, which contains various details such as the skill's name, category, description, privacy policy, and permissions information.

1) Front-end code: The front-end code is written in JSON format and it consists of the following key components:

Intents: An intent corresponds to an action that satisfies a user's verbal request. Each intent is associated with a specific functionality or operation within the skill. For example, when a user provides a username in a sentence, the `CaptureUserNameIntent` in Listing 1 will be triggered to capture the username. When a skill is created, several default Amazon Alexa built-in intents will be created automatically, such as `Amazon.StopIntent` for stopping a skill.

Slots: Slots are the variables that can capture a specific type of user's verbal reply, such as username or user address. The slots have two components: name and type. Developers can use built-in slot types by Amazon, *e.g.*, Amazon.FirstName, Amazon.PhoneNumber, which are trained with thousands of popular first names commonly used or phone numbers, respectively. Lines 3-5 in Listing 1 show a slot `firstname` in `CaptureUserNameIntent` designed for capturing user name and its type is `Amazon.FirstName`.

```
1 {
2     "name": "CaptureUserNameIntent",
3     "slots": [{
4         "name": "firstname",
5         "type": "AMAZON.FirstName"}],
6     "samples": [
7         "{firstname}", "I'm {firstname}"]
8 }
```

Listing 1: Front-end code example.

Sample utterances: Sample utterances comprise a collection of probably spoken phrases that are linked to intents. Developers should provide representative phrases so that the interaction model can better learn the sentence pattern. User requests will be compared with each sample utterance and find the most suitable one. For instance, in the `CaptureNameIntent` shown in Listing 1, "{firstname}" and "I'm {firstname}" serve as sample utterances that utilize the slot "firstname". When a user says either "Bob" or "I'm Bob", the `CaptureNameIntent` will be triggered and the name "Bob" will be extracted as the value of slot `firstname`.

2) Back-end code: For each intent in the front-end code, there is a corresponding intent handler in the back-end code. After an intent captures a specific type of user input, the intent handler will process data, execute skill functions, and provide a response to users. For example, as shown in Figure 2, the `CaptureUserNameIntentHandler` is the intent handler corresponding to the `CaptureUserNameIntent`. When users reply to the previous question "What is your name?" with their names, such as "Bob", the `CaptureUserNameIntentHandler` will capture "Bob" as slot value (`firstname`) and send it to the back-end code. The `CaptureUserNameIntentHandler` in the back-end code will generate an output using the value of the `firstname` slot ("Bob") after performing other necessary tasks. After that, the intent handler will provide the output, "Hello Bob! What is your phone number?", to users to continue the conversation.

In addition to the intent handlers, a special handler named `LaunchRequestHandler` will always be triggered first when a skill is invoked. As defined by Alexa, it does not have a corresponding front-end intent but will be triggered by skill invocation. For example, in Figure 2, when users invoke the skill by speaking "Alexa, open {skill_name}", the `LaunchRequestHandler` will be triggered and provide users a greeting "Hello! What is your name?".

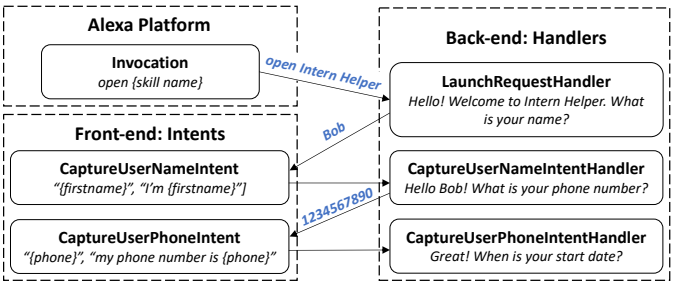

Figure 2: Relation between front-end and back-end code.

C. Threat Model

Considering that users may not be aware of data practices during conversations and often grant permissions during the installation of skills without checking the permission details, users could be totally unaware of their data being collected, shared and stored. Such unawareness about data practices may result in serious discrepancies in privacy expectations and pose potential privacy risks to users.

In addition, existing research shows that there is a notable gap between VPA's privacy requirements and skill developers' practices [48]. It is reasonable to assume that inexperienced developers may inadvertently develop privacy non-compliant skills and are unaware of the privacy requirements, such as providing necessary privacy policies to disclose data collection behaviors in skills. This motivates us to propose the SKILLPOV

tool to generate a reference instantiation of *Privacy Notice over Voice* through static code analysis and instrumentation. Since SKILLPOV is designed for benign developers who wish to enhance the privacy compliance of their skills, we assume the source code of skills is available. For malicious developers who intend to hide their data collection behaviors from users, addressing such cases is beyond the scope of this work; however, the Alexa platform may detect privacy issues in skills using existing tools such as SkillExplore [38], SkillDetective [61], and Vitas [47].

III. UNDERSTANDING USERS' AND DEVELOPERS' ATTITUDES REGARDING PRIVACY NOTICES

To understand users' and developers' attitudes toward the data practices in skills and the concept of *Privacy Notice over Voice*, we conducted a user study via Amazon Mechanical Turk (MTurk) [3] and Qualtrics [19]. The user study was approved by our University's Institutional Review Board (IRB).

A. Pilot Study and Recruitment

With the initial concept of *Privacy Notice over Voice* (providing a short privacy notice at the beginning of skills in voice), we manually integrated a privacy notice sentence generated by us into a demo skill ("Intent Helper"), which has been shown in Figure 2. More details about how to automatically generate privacy notices for skills will be described in Section V. To understand users' attitudes toward the traditional privacy policies and the proof-of-concept *Privacy Notice over Voice* in the skill, we created a survey including comparisons between them using the Qualtrics platform [19]. The details of the survey and responses will be presented in Section III-B, and the complete survey is available in our GitHub repository (https://github.com/CUSecLab/SkillPoV). In addition, we conducted a short survey for developers to understand their attitudes toward *Privacy Notice over Voice* in Section III-C.

Pilot Study: In preparation for our formal study, we conducted a pre-survey involving 5 Ph.D. students from our lab. All participants had prior experience using voice assistants and none of them had been exposed to or had any prior knowledge of this project. The objective was to gather their feedback on the overall structure of the survey. Following their completion of all questions, we conducted face-to-face interviews to listen to their opinions and address any confusion they had. Based on this feedback, we refined both the questions and the format of our survey, including feedback on ambiguous sentences, difficulty in comprehension, and the perceived lack of logical sequence in question order.

Recruitment: After necessary adjustments, we recruited participants through Amazon MTurk and directed them to our survey hosted on Qualtrics. During this recruitment, participants were required to have prior experience using voice assistants, *e.g.*, Amazon Alexa or Google Assistant, and were familiar with their basic functionalities. To maintain the quality of our survey, we required users to be MTurk Masters and have approval ratings higher than 95%. We paid $3 for each participant who successfully completed our survey.

We received 110 responses from participants recruited through MTurk. We had to exclude 8 participants since they indicated that they do not use any voice assistant. To further ensure that participants were attentive and understood the necessary concepts being tested, we strategically placed three attention-check questions (*ACQ*) throughout the survey after we explained such concepts: *ACQ1:* Please choose the definition of a privacy policy. *ACQ2:* What is the name of Amazon Alexa's voice application? (skill) *ACQ3:* In this survey, we present a proof-of-concept method to inform users about data collection. What do we call it? (*Privacy Notice over Voice*) This approach effectively identified and excluded participants who failed to grasp the concepts of privacy policy or did not carefully read the questions and selected answers randomly. Finally, 52 participants passed the attention-check questions and we retained the responses from these participants for further analysis.

The average time for completing the survey was 15 minutes. Among the 52 participants, 34.6% of them were identified as female and 65.4% were identified as male. In terms of educational background, 57.7% of participants have a bachelor's degree, 13.5% have a master's degree and the remaining participants were high school graduates. 86.5% of participants reported using Amazon Alexa, 55.8% reported using Google Assistant, and the remaining participants chose others, *e.g.*, Apple Siri and Android voice assistants. 59.6% of the participants reported having 2 years or more of experience with voice assistants, and 88.5% of participants used voice assistants more than once a week.

B. VPA Users' Perspectives

In this section, we designed a survey to address the following questions regarding participants' attitudes toward data collection and their expectations regarding privacy notices. Differing from previous works [41], [46], [30], our new insight primarily focuses on users' understanding and preference regarding the proposed *Privacy Notice over Voice*. We strictly followed Qualtrics' guidelines to minimize bias in the user study [23]. We focused on data collection because it is the premise of other behaviors such as data sharing, storage, and retention.

- Are participants aware of the data collection in voice-apps (§ III-B1)?
- What are the participants' attitudes toward the traditional privacy policy and the concept of *Privacy Notice over Voice*? (§ III-B2)?
- What are the participants' attitudes toward the accessibility, user control and consent, and readability of *Privacy Notice over Voice* (§ III-B3)?

1) Awareness of data collection: We first asked users if they were aware that skills might collect their data. Among the 52 participants, 42 (80.8%) users knew that skills could collect their personal information.

Following this, participants were exposed to an example skill - Intern Helper, which would ask the user's name, phone number, and start date (when the user would start

the internship). Name and phone number are categorized as personally identifiable information as they can be used to identify an individual's identity, whereas the start date is not classified as personal information. Subsequently, users were asked whether the skill collects personal information and to identify the personal information it collects. This question is designed to assess users' awareness and discernment regarding privacy-related aspects. Only 1 out of 52 participants believed that the example skill does not collect personal information, which proves that the majority of users are aware of personal data collection.

It is noteworthy that among the participants who believed the example skill collects personal information, 84.3% of participants consider the "start date" to be personally identifiable information (PII) and only 5 (9.8%) participants were able to successfully identify the correct pieces of PII collected ("name" and "phone number"). This suggests that most users may not accurately differentiate between ordinary and personal information in data collection.

Finding 1: Most users are aware of data collection activities in skills. However, they may not discern the specific types of data that can be classified as personal information.

2) Attitude toward privacy policy and privacy notice over voice: In this section, we first asked participants about their attitudes toward the privacy policy of skills and if they met any difficulties. Then we introduced the concept of *Privacy Notice over Voice* to users and asked about their opinions.

We first asked users whether they were aware of Amazon's requirement that all skills need to provide a privacy policy when they collect personal information from users and if they had ever read skills' privacy policies. 34 (65.4%) participants were aware that the Amazon platform has such a requirement. However, only 20 participants (38.5%) reported having read the privacy policies of any skill. After that, we provided participants with two examples of privacy policy. We asked participants to estimate the time it would take to read each and whether they had any difficulties reading them. 11 participants (21.2%) believed they could finish reading each privacy policy within 5 minutes, while 24 (46.2%) thought they would need 5-10 minutes. Surprisingly, 7 participants even estimated that they needed more than half an hour to read each privacy policy. When asked which section of privacy policies they care about most, 50% of participants selected the data collection section, 57.7% chose the data sharing section, and 51.9% prioritized the data retention section.

When reading a privacy policy, only 7 participants found no difficulty, while the majority of participants encountered various difficulties: 82.7% of participants found privacy policies too lengthy and time-consuming to read; 48.1% of participants felt that using legal terminology made the content difficult to comprehend; 40.4% of participants found it challenging to locate essential information, such as details about personal data collection. Additionally, 21.15% of participants expressed concerns about the lack of accessibility features for those who prefer or require auditory information.

Given these difficulties in reading a privacy policy, we asked whether users believed presenting it in a **shorter, more easily understandable**, and **accessible** manner would help them better understand privacy policies. 35 participants (67.3%) selected yes and 17 participants (32.7%) selected maybe. Therefore, we presented participants with the demo skill in voice to illustrate the concept of *Privacy Notice over Voice*. We also provided a screenshot of the skill's output to enhance participants' comprehension. After that, we questioned whether participants believed the *Privacy Notice over Voice* represents a better way to notify users about their privacy. 50 (96.2%) participants acknowledged that *Privacy Notice over Voice* is preferable for accessible and effective privacy notice. We presented representative user feedback about traditional privacy policies and our proposed *Privacy Notice over Voice* in Table I, and the complete feedback from participants can be found in our GitHub repository (https://github.com/CUSecLab/SkillPoV).

Finding 2: 86.5% of users struggled with reading skills' privacy policies. Most participants agreed that our proposed *Privacy Notice over Voice* in a shorter, easily understandable and accessible manner provides a better way to notify users about their privacy.

3) Aspects of the privacy notice that participants care about most: In this section, we asked users for more details about the aspects they care about regarding *Privacy Notice over Voice*. We designed 3 scenarios, each of them focusing on one aspect: Accessibility, User Control and Consent, and Readability. In Scenario 1, we presented two videos to demonstrate users' interactions with Alexa skills. In Scenarios 2 and 3, we first defined an aspect and then presented two skills in audio format: one is the experimental group with the *Privacy Notice over Voice*; the other one is the control group for comparison. After participants listened to the skills, we provided screenshots of the textual versions of both skills to help them better understand the skills' outputs. Lastly, we asked participants which of the two skills better presents the definition of each aspect (*i.e.*, accessibility, user control and consent, and readability).

Scenario 1 - Accessibility: Can the *Privacy Notice over Voice* mechanism be easily accessible by users, potentially benefiting individuals with visual impairments?

In Scenario 1, we used two videos to present how users can listen to privacy notices in the real world (Scenario 1A) and how users can access a skill's privacy policy in the Alexa app (Scenario 1B). Subsequently, participants were required to answer three questions. The first question involves selecting the video that better presents the concept of accessibility. 73.1% of participants thought the first video with *Privacy Notice over Voice* was better. The second question asked users how they felt about the first video, which provides a more accessible privacy notice through voice interaction. 78.8% expressed that they liked the design of *Privacy Notice over Voice*.

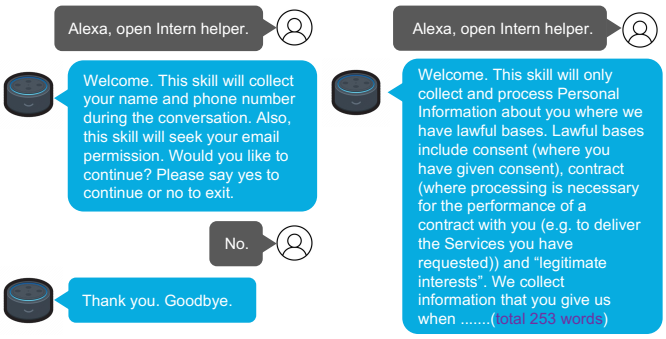

Figure 3: Skill outputs in Scenario 3.

The third question aims to understand participants' views on whether the privacy notice in Scenario 1A can assist people with visual impairments in comprehending the data collection process. 40 (76.9%) users thought the *Privacy Notice over Voice* could be helpful.

Scenario 2 - User Control and Consent: Can *Privacy Notice over Voice* provide users with the option to opt in/out when they become aware of data collection?

In this scenario, the skill with a privacy notice (Scenario 2A) would provide users an option to continue or exit this skill by asking *"Would you like to continue? Please say yes to continue or no to exit."* after playing our generated privacy notice. Meanwhile, the skill in the control group (Scenario 2B) doesn't have such a selection. In the result, most participants (76.9%) thought the skill with *Privacy Notice over Voice* better presents the definition of User Control and Consent. 90.4% users liked the design of *Privacy Notice over Voice* which provides User Control and Consent, demonstrating that such an opt-in or out option is preferable for users.

Scenario 3 - Readability: Can *Privacy Notice over Voice* be short and allow users to easily understand it?

The skills and outputs of Scenario 3 are shown in Figure 3. Our generated privacy notice contains only 20 words to describe data collection (Scenario 3A), while the control group uses the data collection description from a privacy policy with 253 words (Scenario 3B). The majority of participants (67.3%) believed that our generated privacy notice presents better readability. 84.6% users expressed that they liked our design with shorter privacy notice.

Finding 3: The majority of users agree that the option for users, the readability, and the accessibility in *Privacy Notice over Voice* is useful and helpful, potentially benefiting individuals with visual impairments.

Placement of privacy notice: In addition, we asked participants about their preferred placement for the privacy notice, such as at the beginning, in the middle, or at the end of a skill. 67.3% of participants believed that presenting the privacy notice at the beginning is best. Another 30.8% thought that asking about data collection at any moment in the middle of a skill can be useful. The remaining 1.9% users believed the privacy notice at the end is the best time.

In the end, we played the two versions of the Intern Helper, one with and another without the *Privacy Notice over Voice*, and asked for participants' overall opinion on the design of the *Privacy Notice over Voice*. 92.3% of participants expressed that they liked the design of *Privacy Notice over Voice*. More responses from users can be found in Table I.

To evaluate the difference when participants interact with *Privacy Notice over Voice* in media (audio, video, or images) and in actual skills, we conducted an additional user study with 12 university students using actual skills containing *Privacy Notice over Voice* and asked their perceptions. The results are consistent with the findings obtained from participants via Qualtrics.

Feedback	Users/Developers' Responses
Traditional privacy policy in text form	It must be short and straightforward, and authentic for the customers, no hidden policies must be there.
	It should be readable and easy to understand all the data that is being collected.
	It should be brief and easy to understand, avoid lengthy explanations and legal terms that can be confusing. Describe exactly what is being collected and ask for permissions!
	Just to keep it simple so that it is understandable and use a language that most would understand.
	We want to be transparent, easy to understand, and in short words.
Proposed privacy notice over voice	Privacy Notice through Voice can be useful for impaired persons.
	The privacy notice through voice is better choice for every individual, it should be in first place and ensure the user understand it properly before proceeding further.
	Privacy notice through voice looks more secure and alerting the user before the conversation is very helpful.
	It would be great if we can change the welcome message in the skill. *
	If it is too long, people will not want to listen, also will not want to hear it every time permission is needed.
	There needs to be a happy medium between merely telling you what is being collected and being transparent about why and how that information is used. Simply saying "we are collecting your data" is also not helpful.
	It's great, looks really good, pretty impressive. *

Table I: Feedback from users and developers. * represents developers' responses.

C. Skill Developers' Perspectives

In addition to users, we designed a survey to understand developers' views on the *Privacy Notice over Voice* and if they are interested in a tool that can help generate it. To recruit real-world Alexa skill developers, we extracted 1880 email addresses from skill descriptions in the current Amazon Alexa platform and dispatched invitations to these email addresses, which included a link to the survey we designed. As a result, we received 21 valid responses.

Compared to regular users, of whom 34 participants (65.4%) were aware that the Amazon platform requires skills to have a privacy policy when collecting personal information, all developers believed that a privacy policy should always be provided when user information is collected. When presenting *Privacy Notice over Voice* to developers, 16 (76.2%) developers found it helpful for their skills. Regarding visually impaired users, when asked if *Privacy Notice over Voice* could facilitate communication with visually impaired users about data collection in skills, 19 (90.5%) developers expressed a need for

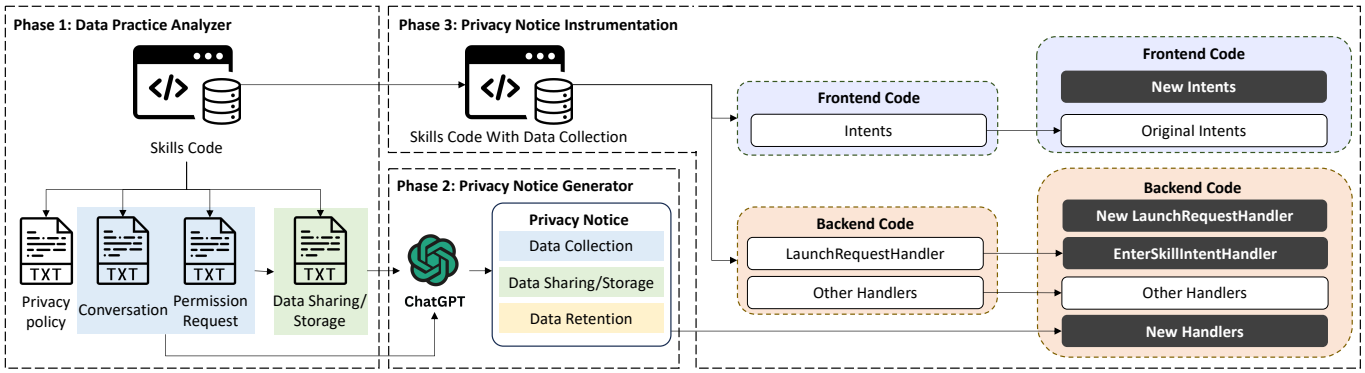

Figure 4: System overview of SKILLPOV.

this and only 2 developers were uninterested. On the question, "Would you like to use a tool that assists in generating privacy notices for users in your skill development?" 71.4% of developers indicated they would consider using it. The remaining developers expressed concerns about security or privacy issues associated with *Privacy Notice over Voice* or indicated a preference for visual-based privacy notices (e.g., traditional privacy policies) over other formats. More responses from developers can be found in Table I.

Finding 4: Compared to regular users, developers have a better understanding of the privacy policy requirements on the Alexa platform. Most developers are interested in a tool that facilitates generating a privacy notice.

IV. SYSTEM OVERVIEW

Design Guidelines. Informed by our user study results in Section III, it is desirable to enhance existing skill code with the new *Privacy Notice over Voice* mechanism with user control and better readability for Alexa skill users. To provide user control and consent, we incorporate options to learn about which data practice at the beginning of skills and provide choices for users to continue or exit with each privacy notice. For better readability and to help users easily understand, we restrict the generated privacy notice to few sentences, avoiding using complex sentence structures and technical terms.

For developers, we design and develop an automated tool named SKILLPOV (Skill's Privacy Notice over Voice) to generate privacy notices and automatically integrate privacy notices into skill code without impacting original skill functionalities. To automatically generate privacy notices, we utilize the LLMs and design prompts to ensure the reliability of generated privacy notices. To ensure the accuracy and completeness of privacy notices, we employ static analysis of the skills' code to detect actual data collection behaviors. For ensuring that the original functionality of the skill remains unchanged, we preserve all previous code structures and only add necessary intents and intent handlers. Note that SKILLPOV only generates a reference implementation of *Privacy Notice over Voice* and developers can further customize the privacy notice content based on their needs (such as users' rights, regulatory compliance, and contact details).

Overview: Figure 4 presents the design overview of the SKILLPOV, which includes three components: Data Practice Analyzer (§ V-A), Privacy Notice Generator (§ V-B), and Privacy Notice Instrumentation (§ V-C). In the Data Practice Analyzer, SKILLPOV first scans the skill code provided by developers and detects data collection behaviors within the skill code. For data collection in conversations, SKILLPOV extracts all the possible conversations from the back-end code and applies an NLP-based method to detect data collection behaviors. For data collection permissions, SKILLPOV captures them from the manifest file. In addition, we employ large language models to detect data sharing/storage in skill code. Next, we generate short, easy-to-understand, and accurate privacy notices based on data practices using the LLMs. We design prompts to deal with different cases of data practices and limit the length of generated sentences. At last, we instrument the privacy notices into the original skill code. To smoothly present the privacy notice to users, we need to redesign the skill interaction flow, insert a privacy notice into the skill code, and provide users an option to continue or quit the skill after listening to the privacy notice. This needs the additional intents and intent handlers (black boxes in Figure 4) to process different cases of user preferences while keeping the original skill functionalities unaffected. Finally, we evaluate the effectiveness of SKILLPOV by measuring the quality of generated privacy notices, such as functionality and reliability, and comparing generated privacy notices with traditional privacy policies regarding readability, accuracy, and completeness.

V. SYSTEM DESIGN

A. Data Practice Analyzer

Data Practice Analyzer is designed to analyze skill codes and identify data collection, data usage (data sharing/storage) and data retention. We first define the sensitive data types considered in our work, which are presented in Table II. For data collection in conversations, we consider 16 types of personally identifiable information obtained from a NIST document [15] and Amazon's definition [2]. For data collection permissions, we mainly consider seven types of permissions listed in Table II. For data usage, we focus on the external APIs, URLs, and databases where data is shared/stored. The

output of this module is the data practices and the type of sensitive data being collected/shared/stored, such as "conversation data collection: [data noun]", "permission data collection: [data noun]", and "data usage: [data noun], entity: [entity]". These outputs will be used as the input for Privacy Notice Generator (§ V-B).

Personal sensitive data [2] [15]	Address, Name, Email, Birthday, Account, Location, Phone number, Passport number, Driver license number, Bank account number, Debit card number, Credit card number, Credit card verification code, Taxpayer identification number, Social Security number (SSN), Vehicle identification number (VIN)
Data collection permissions	Name, Given_name, Email, Mobile_number, Address, CountryAndPostalCode, Geolocation

Table II: Keywords related to personal data collection.

1) Detecting Data Collection Behaviors in Conversation: Skills can directly ask for user data in a conversation. For example, skills can provide an output such as "Tell me your name" and obtain user information through user responses. To find such data collection behaviors from skill code, we first search for all strings in skills' back-end code after removing comments. This is done by employing regular expressions to capture strings enclosed within single (' '), double (" "), and backticks (` `) quotes, which are indicative of textual outputs in skill codes. After obtaining all the strings, which are potential skill outputs in conversations, we detect data collection behaviors in these sentences using an NLP-based method, similar to previous works [38], [61], [47]. We check whether any type of sensitive data in Table II is used as a noun semantically following the word "your" using the Spacy [22] library. If a sentence contains "your + sensitive data", we consider it as a data collection behavior. Previous work [48] demonstrates that this method achieved an accuracy of 98% for identifying data collection behaviors from Alexa skill outputs. To enhance the accuracy and remove potential false positives, we also include a list of common sentences of personal data collection, such as "How old are you" and "What can I call you". In addition, we remove the sentences with a negative word, such as "I will not collect your data". For example, in Listing 2, we identify two data collection behaviors in skill outputs: "What is your name?" in Line 2 and "What is your phone number?" in Line 7.

```
1 const LaunchRequestHandler = {
2     const speakOutput = 'Hello! What is your name?';
3     return handlerInput.responseBuilder.speak(
      speakOutput).reprompt(speakOutput).getResponse()
      ;};
4
5 const CaptureUserNameHandler = {
6     const userName = handlerInput.requestEnvelope.
      request.intent.slots.name.value;
7     speakOutput = "Thanks ${userName}, what is your
      phone number?";
8     return handlerInput.responseBuilder.speak(
      speakOutput).reprompt(speakOutput).getResponse()
      ;};
```

Listing 2: Data collection behaviors in conversations.

2) Detecting Data Collection Permissions: Developers can also request permissions for certain types of data from the user's account in the Amazon Alexa platform. In this approach, when users enable a skill for the first time, they need to grant permission to the skill to access the information associated with their Alexa accounts. Table II shows seven permissions in Alexa skills related to personal data that we consider. When developers request permissions in skills, such information is stored in the skill manifest file [13], which is a JSON file that includes diverse types of skill metadata such as skill name, description, category, permission, and privacy policy link. Listing 3 shows the code snippet of a skill's manifest file. This skill requests for permissions of user name (`alexa::profile:name:read`) and email (`alexa::profile:email:read`).

```
1 "manifestVersion": "1.0",
2 "permissions": [
3     {"name": "alexa::profile:name:read"},
4     {"name": "alexa::profile:email:read"}],
5 "privacyAndCompliance": {
6     "allowsPurchases": false,
7     "locales": {
8         "en-US": {"privacyPolicyUrl": "https://www.
      alexa.com/help/privacy"}}}
```

Listing 3: A skill requests for data collection permissions in the skill's manifest file.

3) Detecting Data Usage Related to Data Collection: Skills can share and store collected data with 1st or 3rd party services. To detect how user data is used, we employ Large Language Models (LLMs) to analyze the data flow as recent works have demonstrated the great potential of LLMs in software analysis [36]. We first pick out the files of skills' code which contain data collection behavior. We then instruct LLMs to identify data sharing and storage activities in the code. We use the prompt: "Your mission is to identify data sharing activities about {data_collection_content} in the code. If there are data sharing activities about {data_collection_content}, output {data_collection_content} is being shared to which entity. The format should be 'data usage: [], entity: [], code: []'." `data_collection_content` refers to the data collection results obtained from Section V-A1 and V-A2. Using this prompt, ChatGPT can identify the data flow of a specific type of data and the entity it is shared, such as "data usage: [name], entity: [api.razorpay.com]".

In addition, users may want to learn about other information, such as the data retention period, which is infrequently mentioned in the skill code. Therefore, we provide developers with an interface to enter and provide such customized information to users.

B. Privacy Notice Generator

After detecting data practices in skill code, we aim to generate reliable, short, easy-to-understand, and accurate privacy notice content for users. However, it is hard to apply existing privacy policy generation tools for our task because of the unacceptable length of the generated privacy policies. To tackle this issue, we leverage LLMs because of their excellent

performance in handling natural language processing (NLP) tasks.

Prompt Design: To ensure accuracy and specificity in our privacy notices, we employ ChatGPT to generate privacy notices that are precise and customized to reflect the particular data types collected during user interactions. Hence, we create specific prompts based on the identified data types from both conversations (as detailed in § V-A1) and permissions (as specified in § V-A2). Below, we describe three specific conditions of data collection, each supported by prompt formulations to generate concise privacy notices:

- **Data Collection in Conversations:** Our privacy notice aims to inform the users briefly about the data collected during interactions with the skill. For example, if SKILLPOV identifies that a skill collects a user's name *[data noun1]* and age *[data noun2]* during the conversation, the prompt would be formulated as follows: *If my input is 'conversation data collection: name, age', your output should be "This skill will collect your name and age during the conversation."*
- **Data Collection with Permissions:** In cases where the skill requires user permission to access specific user data, the privacy notice would explicitly clarify this requirement to make it transparent. For instance, if SKILLPOV determines that a skill requires user permission to access a user's email *[data noun1]*, the prompt would be formulated as follows: *If my input is 'permission for data collection: email', your output should be "This skill needs your permission to access your email."*.
- **Data Collection in Both Approaches:** We also aim to inform the users when the same data type is collected through both conversation and permissions. For example, if it is identified that a skill collects a user's name *[data noun1]* during conversation and also seeks permission for the same data (name) *[data noun1]*, the prompt would be formulated as follows: *If my input is 'conversation data collection: name; permission data collection: name', your output should be "This skill will collect your name during the conversation. Also, this skill will need your permission to access your name."* In cases where the same data type is collected through both conversational interaction and permissions, we direct ChatGPT to generate two separate sentences linked by *"Also"* to avoid ambiguous privacy statements that might confuse users.

Figure 5 presents different prompts and corresponding answers from GPT-4. We initially employed a zero-shot prompting strategy to generate accurate privacy notices. However, the outputs often lacked the necessary precision, as they sometimes included irrelevant or generalized information. To address this issue, we transitioned to "Structured-based Prompting" to guide the language model to generate privacy notices aligned with our standards for specific data collection scenarios. This approach assures that each type of data collection was explicitly mentioned in a concise manner. Nonetheless, when dealing with overlapping data collection scenarios where the same data types, such as the name, were collected both through conversation and permissions, it fails to mention that the skill will also seek permission for the name, leading to incomplete privacy notices. Therefore, we optimized the prompt further ("Optimized Structured-Based Prompting" in Figure 5) to allow the explicit inclusion of each data type collected through both methods without omitting any critical details. In addition, our method allows developers to customize and easily adjust the content of privacy notices by modifying the prompts based on specific requirements.

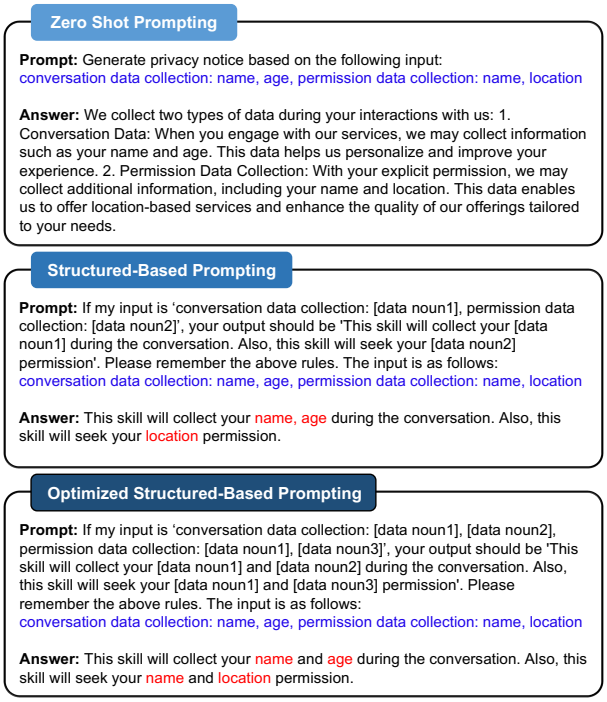

Figure 5: Prompt strategies for generating privacy notices.

C. Privacy Notice Instrumentation

After generating privacy notices with LLMs, we instrument privacy notices into skills' code. To keep the original skill functionalities and provide better services for users, we need to design new interaction flows to process users' different preferences.

Original interaction flow: We take the skill "Intern Helper" as an example to present how we instrument the skill interaction flow. The left side of Figure 6 shows the original interaction flow. The blue text represents the user's input, while the red text indicates Alexa's response. When users invoke this skill with the utterance "Alexa, open Intern Helper", the back-end code will be triggered first and the `LaunchRequestHandler` will generate a user greeting: "Hello! Welcome to Intern Helper. What is your name?" When a user replies, such as "Bob", the front-end code will obtain the user's response and find the most suitable intent, *i.e.*, `CaptureUserNameIntent`, to extract the user name "Bob". After that, the corresponding intent handler in the back-end code, named `CaptureUsernameHandler` starts

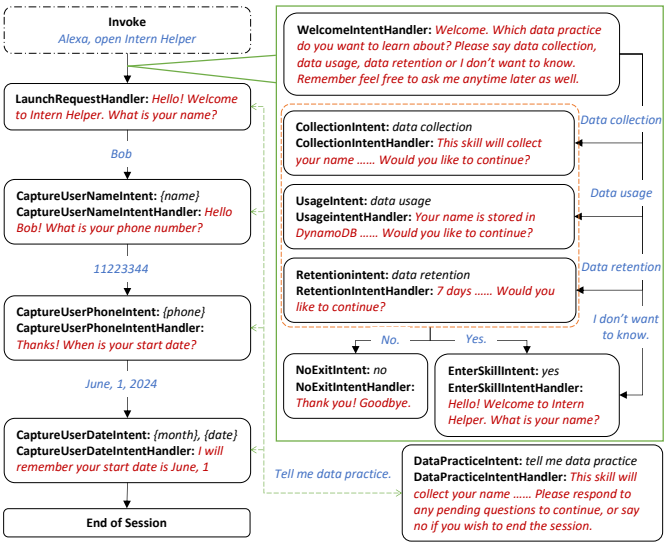

Figure 6: Interaction flow instrumentation. The user's input is shown in blue text, while Alexa's response is in red text. The newly added intents and corresponding intent handlers are shown in the right side.

to work and process data. After performing certain skill functionality, it will generate a response to users, "Hello Bob! What is your phone number?" This cycle repeats with user requests triggering intent activation, intent handler processing data, and generation of new responses to the user.

New interaction flow: To integrate the generated privacy notice into the skill code without impacting original skill functionalities and offering a better user experience, *e.g.*, providing users an option to exit the skill after listening to the privacy notice, we designed a list of intents to capture different types of user responses and the corresponding intent handlers to provide outputs to users to continue the conversation. The new interaction flow, integrated intents, and intent handlers are shown in the right side of Figure 6. Figure 7 in Appendix A presents the interaction of the instrumented skill.

According to our user study results (as shown in Section III-B3), most users prefer to have the privacy notice at the beginning of skills. Therefore, we integrate the privacy notice into the skill code at the beginning. However, since not every user wishes for a notice covering all data practices (data collection, sharing/storage, and retention), we introduce a new output that prompts users at the start of the skill interaction to decide which data practices they want to know about or to enter the skill directly without listening to the privacy notice. To prevent users from repeatedly listening to the privacy notice, the question and privacy notice will be hidden if this skill has already been invoked. We design a new intent handler:

- `WelcomeIntentHandler`: This handler will become the first triggered handler when a skill is invoked. For doing that, we need to rename it as the "LaunchRequestHandler". When a skill starts, this handler will ask users to choose data practices, such as "Welcome. Which data practice do you want to learn about? Please say data collection, data usage, data retention or I don't want to know. Remember feel free to ask me anytime later as well."

After that, users can select which data practice they want to learn about. New intents and intent handlers for processing user responses are:

- `CollectionIntent/UsageIntent/RetentionIntent`: These intents will capture corresponding users' responses to the question in the `WelcomeIntentHandler`, such as "data collection", "data usage" or "data retention".
- `CollectionIntentHandler/UsageIntentHandler/RetentionIntentHandler`: These intent handlers will provide users with the corresponding privacy notice content. After that, they will ask users, "Would you like to continue?"

Following the privacy notice, users have the option to continue interacting with the skill or exit the skill based on their preference regarding data collection. For users who decide to continue the skill after listening to the privacy notice or who don't want to listen to the privacy notice, they will enter the original skill functions.

- `EnterSkillIntent`: This intent will capture the user's negative responses to the question in the `WelcomeIntentHandler`, such as "i don't want to know", and the user's positive responses to the question in the `CollectionIntent/UsageIntent/RetentionIntentHandler`, such as "yes" or "yes continue".
- `EnterSkillIntentHandler`: This intent handler is the original `LaunchRequestHandler` of this skill, which is the first function of the original skill. After this intent handler, the skill functions will return to the original flow.

For users who decide not to enter the skill after listening to the privacy notice, they can exit the skill.

- `NoExitIntent`: This intent will capture the user's negative responses to the question in the `CollectionIntent/UsageIntent/RetentionIntenthandler`, such as "no" or "exit".
- `NoExitIntentHandler`: This intent handler will provide conversation closing words, such as "Thank you! Goodbye." and end this skill.

In addition, users may want to get data practices information at any time during the conversation. In this case, we design the following intent and corresponding handler so that users can listen to the privacy notice at any point during the interaction.

- `DataPracticeIntent`: This intent will capture the user's request to the privacy notice at any time during the conversation, such as "tell me data practice".
- `DataPracticeIntentHandler`: This intent handler will continue the previous conversation so that the skill function and user experience will not be impacted. It will output "Please respond to any pending questions to

continue, or say no if you wish to end the session." When users respond to previous questions, the corresponding intent will be triggered and the skill will continue the original interaction flow.

D. Implementation

We implemented SKILLPOV in Python and made it an automatic tool for developers to scan a skill's code, generate a privacy notice, and instrument it into the skill's code. SKILLPOV is available at https://github.com/CUSecLab/SkillPoV. For the Data Practice Analyzer, SKILLPOV will scan the skill codes first, including back-end code and manifest files, to detect data collection behaviors. Then, SKILLPOV employs ChatGPT based on GPT-4, using JavaScript files and data collection results as inputs. For the Privacy Notice Generator, we utilize GPT-4 and design prompts for generating privacy notices. For the Privacy Notice Instrumentation, SKILLPOV will automatically integrate new intents and the corresponding intent handlers into the front-end and back-end code of a skill. Normally, the skill's front-end code is written in JSON format and it facilitates the integration of new intents. After scanning and locating the original intents, we can effortlessly add new intents to the original JSON file. The intent name, slot, and samples of each newly added intent are displayed on the right side of Figure 8 in the Appendix A. Since we only need to capture different conditions of users' replies instead of personal data, these intents don't include any slots. The sample utterances in the new intents are designed based on all possible replies, and we ensure they don't overlap to avoid confusion.

Since most skills' back-end code is developed in JavaScript (82.5% in our skill code dataset, which will be introduced in Section VI), we mainly focus on integrating intent handlers in back-end code to JavaScript files. We first locate the original handlers in the code, *e.g.*, `LaunchRequestHandler`. Then, we insert our proposed intent handlers, such as `WelcomeIntentHandler`, `Collection/Usage/RetentionIntentHandler`, `EnterSkillIntentHandler`, and `NoExitIntentHandler`, with all necessary code into the existing skill's back-end code. To redefine the entry point of the skill interaction, we need to rename the `WelcomeIntentHandler` to `LaunchRequestHandler` and rename the original `LaunchRequestHandler` to `EnterSkillIntentHandler`. In addition, we need to check if the user invokes the skill for the first time. This is implemented by calling Amazon's `attributesManager` function to get the user session attribute and record users' visits. Detailed instrumented code is provided in Figure 9 in Appendix A.

VI. EVALUATION

Since SKILLPOV is designed for developers and will run locally to integrate privacy notices into their skill code, we assume the source code of skills is provided by developers as the input of SKILLPOV. To evaluate the performance of SKILLPOV, we first searched and obtained 1,459 open-source skill codes from GitHub. Then, we checked whether they had data collection behaviors and used the JavaScript language for the back-end code. As a result, we identified 143 skills with data collection, and we used them to build a benchmark dataset. For each skill, SKILLPOV automatically generates a privacy notice and integrates it into the skill code. Then, we evaluate the performance of SKILLPOV from the following perspectives:

- **Functionality:** Whether SKILLPOV can successfully generate privacy notices for skills and integrate them into skill code (§ VI-A).
- **Reliability:** Whether the privacy notice content generated by SKILLPOV using our designed prompts always follows the expected sentence pattern (§ VI-B).
- When comparing the privacy notice content generated by SKILLPOV against the traditional privacy policies, how about the performance of these two approaches in terms of the **Readability**, **Accuracy** and **Completeness** (§ VI-C)?
- When comparing the privacy notice content generated by SKILLPOV against the privacy policies generated by existing online generators, how about their performance in terms of the **Readability**, **Accuracy** and **Completeness** (§ VI-D)?

A. Functionality

We first tested the functionality of SKILLPOV by ourselves. To this end, we manually checked the source code of 143 skills and found that SKILLPOV successfully generated privacy notices for all skills and instrumented privacy notices into the skill code. After that, we checked whether skills worked well after integrating privacy notices using SKILLPOV. Since many skills with code in the wild are used for learning purposes only or as toy skills and cannot be executed, we randomly selected 20 skills from our dataset and manually uploaded them to the Alexa Developer Console to verify them. As a result, we got 10 executable skills, which we used for testing. After analyzing data collection behaviors and generating privacy notices using SKILLPOV, we successfully integrated privacy notices into skill codes. These modified skills were then uploaded to the Alexa Developer Console again and tested to see if they worked well. As a result, all these skills work well when informing users about data collection and providing users with the option to opt out without affecting the functionality of the skills.

After that, we invited 3 researchers (who are not authors of this work) in our research team to perform as developers and test the tool. All of them have experience in developing Alexa skills and have published more than one skill in the skill marketplace. As a result, all of them successfully completed the whole process within 20 minutes in Linux/MacOS environments. Finally, we published SKILLPOV and asked Alexa skill developers to provide feedback after testing it, which is part of the developer user study we mentioned in Section III-C. Of the 21 developers who responded, 17 successfully generated privacy notices using SKILLPOV, and half of them finished the testing within 10 minutes. The other 4 developers failed to run SKILLPOV. For the developers who provided their encountered issues, one developer used the Windows platform that we didn't support (we will discuss it in the limitation section); another developer mentioned the generated files were empty,

which is possible due to the skill that the developer used does not contain any data collection. After testing SKILLPOV, developers were asked to upload their skills to the Amazon Developer Console. Over 80% of developers completed the process within half an hour. We also received a commendation from one developer, who stated, *"It's great, looks really good, pretty impressive."*

B. Reliability

As we discussed in Section V-B, sometimes ChatGPT may not strictly follow the designed pattern and will generate unexpected outputs. To assess the effectiveness of our designed prompt and the reliability of generated sentences, we manually checked the privacy notice sentences generated by ChatGPT for all 143 skills in our dataset and compared whether these sentences were the same as our designed sentence pattern. Since ChatGPT may provide different responses for the same input, we undertook three rounds of testing. As a result, we found that most of the generated privacy notices were the same as what we designed and the accuracy rate is 97.67%. The results demonstrate the effectiveness of our designed prompt for ChatGPT in generating privacy notices. The error cases are because of the incorrect sentence format, such as missing or additional commas and characters.

C. SKILLPOV's Privacy Notices vs. Skills' Original Privacy Policies

Next, we evaluated and compared the readability, accuracy, and completeness of privacy notices generated by SKILLPOV with the skills' original privacy policies. We manually labeled data collection behaviors in 143 skills as ground truth. For each skill, the privacy policy link is stored in the manifest file if it is provided (*e.g.*, Line 8 in Listing 3) and we could easily obtain it. Among the 143 skills in our dataset, 59 skills provide a privacy policy link, of which 9 of them are inaccessible. Therefore, we used the remaining 50 accessible skills' privacy policies for comparison. Since a privacy policy may include other meaningful information, such as data retention, data collection purposes or user rights, *we only compare the generated privacy notice with the sentences related to data collection in a privacy policy*. We present the comparison results of part of the skills in Table III.

Readability: In our user study (Section III-B2), most users complained that "privacy policies are too lengthy" or "the content is difficult to comprehend". Therefore, we considered the length and complexity of privacy notices as important factors in evaluating their readability. We first used an online tool, Word Counter [28], to count the word number in privacy policies' data collection section and privacy notices generated by SKILLPOV. As shown in Table III, the length of generated privacy notices by SKILLPOV is much shorter than the original privacy policies and more stable. While the data collection section in privacy policies tends to be extensive, with lengths ranging from 27 to 1750 words, the privacy notices generated by SKILLPOV exhibited remarkable conciseness, maintaining lengths between 7 to 25 words. For most skills, the word decrease rate for privacy notices is more than 80% and the overall word decrease rate is 93.2% for all skills with a privacy policy link. Such results underscore the readability of our generated privacy notices regarding the length.

We used the Flesch reading ease score [43], [9], [34], a widely used score for evaluating how difficult a passage in English is to understand, to evaluate the complexity of conventional privacy policies and SKILLPOV's privacy notices. A higher score indicates that the text is easier to read. The formula for the Flesch reading-ease score test is:

$$\text{score} = 206.835 - 1.015\left(\frac{\text{total words}}{\text{total sentences}}\right) - 84.6\left(\frac{\text{total syllables}}{\text{total words}}\right)$$

As a result, the average Flesch reading ease score for skills' privacy policies is 46.2, while a score of 30-50 is college-level and means the text is difficult to read. For privacy notices generated by SKILLPOV, the average score is 68.4, and 60-70 means the text is easily understood by 13 to 15-year-old students. Such scores reflect the readability of SKILLPOV's privacy notices and the generated privacy notices are easy to understand, in particular for kid users.

Accuracy: We found that certain skills' privacy policies may omit specific data collection. For example, as shown in Table III, the skill "Feed Me Now" collects multiple types of data, such as name, location, address, and phone number. However, its privacy policy doesn't include the "name" information but claims to collect the "email" information, which is inaccurate. Furthermore, we found that some skills, such as "Mega Food", will redirect to the company's privacy policy links, which tend to describe data collection in a broader context, often leading to imprecise descriptions. In contrast, our generated privacy notices are designed for individual skills and specific data types, ensuring a higher degree of correctness and relevance.

When comparing skills' privacy policies with their actual data collection behaviors, we found that the privacy policies had only 35.3% precision. Meanwhile, the privacy notices generated by SKILLPOV achieved a precision of 91.3%, which proves the accuracy of the SKILLPOV's privacy notices. The false positives are due to our NLP-based algorithm incorrectly identifying certain sentences as data collection, such as "your address is ...".

Completeness: On the other hand, SKILLPOV can also provide a more complete summary of data collection behaviors than the original privacy policies. Notably, several skills, such as Skilltober Hacktober, Date a Voice, and Polar Cast in Table III, claim that no personal information would be collected in their privacy policies. However, these skills collect users' data like name, email, and number in their code, showing that their privacy policies don't fully disclose these data collection behaviors. The completeness of privacy notices generated by SKILLPOV is 96.4%. while skills' original privacy policies only have 50% completeness.

Skill Name	# of words in Privacy Policy	# of words in SKILLPOV	Decrease Rate	Privacy Policy - Data Collection	SKILLPOV - Data Collection	Ground Truth - Data Collection
Polar Cast	-	9	-	-	name	name
date.a.voice	56	7	87.5%	-	name	name
Reddit Deals	67	7	89.6%	name	name	name
Feed Me Now	89	25	71.9%	location, address, phone number, email	name, location, address, phone number	name, location, address, phone number
Psy's Quiz Game	59	7	88.1%	name	name	name
Skilltober Hacktober	95	10	89.5%	-	phone number, email, name	phone number, email, name
How High Am I in Utah	27	17	37.0%	-	address	address
Megafood's Wellness Assistant	1750	9	99.5%	age and 12 other data	age	age

Table III: Comparison between privacy notices generated by SKILLPOV and skills' original privacy policies for example skills.

D. SKILLPOV*'s Privacy Notices vs. Privacy Policies from Online Generators*

To demonstrate the need of our LLM assisted approach to generate the privacy notice content, we compared SKILLPOV's privacy notice with privacy policies from online generators. Developers often lack knowledge about privacy policies and select to employ *Online Automated Privacy Policy Generators* to generate privacy policies. However, the quality of them varies. This section compares our generated privacy notices with the privacy policies generated by these online generators.

We obtained 10 online privacy policy generators from [52] and tested them. After removing the generators that would generate totally the same privacy policies, we kept four generators, as shown in Table IV. We standardized the same input parameters across all generators, including country, state, nature of the platform, entity type, and the necessity for professional legal compliance. Since it is time-consuming to manually input data collection types and other necessary information to the privacy policy generators, we only selected two skills that collect multiple types of data, including name, email, phone number, address, etc., and generated privacy policies for them using all the 4 privacy policy generators.

Generator \ Skill Name	**Feed Me Know**	**Skilltober Hacktober**
Terms Feed [25]	284 (91.2%)	196 (94.9%)
App Privacy Policy Generator [4]	158 (84.2%)	158 (93.7%)
Websitepolicies [27]	86 (70.9%)	86 (88.4%)
Termly [24]	159 (84.3%)	152 (93.4%)
SKILLPOV's Privacy Notice	25	10

Table IV: Comparison between the length of privacy notices generated by SKILLPOV and privacy policies generated by online generators.

Table IV presents the comparison between the four generators and our generated privacy notices for two skills. We only focus on the length of the data collection section in privacy policies. Notably, multiple generators would generate a privacy policy with over 150 words in the data collection section, which is significantly longer than the privacy notices generated by SKILLPOV. The numbers in the brackets show the decreased rate of privacy notice when compared to a privacy policy generator, and the overall decrease rate is 87.6%, showing the readability of our generated privacy notices is better. For complexity evaluation, we still use the Flesch reading ease score, and the average score for the data collection section in the generated privacy policies is 34.0, which means the text is difficult to understand. Since privacy policy generators require developers to input the data they collect, the accuracy of generated privacy policies is 100%. Surprisingly, we discovered that the generated privacy policies may miss some developers' inputted data and the completeness rate is 75%, while the completeness of SKILLPOV's privacy notices is 100%.

E. Accuracy of Data Sharing/Storage Detection in SKILLPOV

In addition to data collection, we also evaluated the accuracy of data sharing/storage. Out of 143 skills with data collection, SKILLPOV detected 44 skills that exhibited data sharing or storage behaviors. We manually validated all these skills and checked their data usage. As a result, we found that SKILLPOV achieved a precision of 82.4% in detecting data sharing/storage. Some false positives occurred due to ChatGPT's incorrect entity recognition. For example, in the skill *starport-75*, ChatGPT mistakenly identified the skill name as an external network port, leading to an incorrect assumption of data sharing behavior. In addition, we compared the accuracy between SKILLPOV and the privacy policies. Among these 44 skills, 10 provided a privacy policy. Surprisingly, all 10 privacy policies stated that they neither share nor store users' information.

VII. DISCUSSION

A. Implication

Our proposed *Privacy Notice over Voice* mechanism and SKILLPOV tool can benefit various stakeholders, including users, developers, and visually impaired individuals. For users, the *Privacy Notice over Voice* provides a new approach that allows users to conveniently request and receive privacy notices directly through voice interaction with skills rather than accessing skills' privacy policies through mobile apps or websites. Although it frees users from the hassle of reading potentially inaccurate but lengthy privacy policies, it is worth mentioning that our proposed *Privacy Notice over Voice* mechanism is compliant with other privacy notification approaches, such as the privacy policy, and a complete privacy policy is still needed to comply with laws and regulations. For developers, they can utilize the SKILLPOV to automatically generate privacy notices and insert them into skill codes within 5 to 10 minutes without impacting existing functionalities. It could

also enhance the transparency of privacy practices, thereby increasing user satisfaction and trust. Moreover, developers can significantly reduce legal risks associated with privacy non-compliance issues.

Privacy Notice over Voice has potential benefits for visually impaired individuals. While there already exist some tools available to assist visually impaired individuals in accessing text-based content, such as screen readers [20] and braille displays [5], it's important to note the challenge posed by lengthy words, especially for privacy policies. For example, privacy policies for 75 prominent mobile applications and websites have around 4,000 words on average [18]. Based on typical reading speeds, it would take approximately 15-30 minutes to read aloud at a normal pace, which highlights the significant time commitment and high reading cost, making it impractical for daily use. Therefore, we introduce the privacy notice in the voice-interactive interface that enables visually impaired individuals to quickly and accurately understand what data is being collected, streamlining their access to vital information without the burden of navigating through lengthy documents.

Ethical Considerations: In the development and deployment of our tool, we placed emphasis on ethical considerations. In our evaluation procedures, we exclusively utilized skill codes that are publicly available on GitHub, ensuring full transparency and adherence to legal and ethical standards concerning copyright. When developers use SKILLPOV to integrate privacy notice-related content into the existing skill codes, we respect the developers' copyrights and do not alter any original functionalities. Furthermore, SKILLPOV is designed to operate locally, negating the necessity to share/store code by developers on cloud servers or other external platforms.

B. Limitations

Although SKILLPOV could help developers and users generate privacy notices and integrate them into skill code, it has several limitations. First, SKILLPOV is primarily focused on the skills written in JavaScript, excluding other programming languages like Python and Java from its scope [17]. This is because our statistical data indicates that a significant majority of developers (82.5%) utilize JavaScript for skill development. Our future work could extend the applicability of our tool to a broader range of programming languages, such as Python and Java. Second, our evaluation exclusively concentrates on modifying skills from GitHub, given the impracticality of accessing real-world skills' code hosted on cloud platforms. However, we ensured the structure of published skills and those from GitHub are similar. The functionality of SKILLPOV is also proved in real-world skills, as corroborated by developer surveys. Third, SKILLPOV only focuses on Amazon Alexa skills, which is the most popular VPA platform. We plan to expand our scope to other VPA platforms such as Google Assistant, Samsung Bixby and Microsoft Cortana, and provide *Privacy Notice over Voice* to more users. Fourth, we didn't conduct a user study with visually impaired individuals due to challenges in recruiting participants with visual impairments. We plan to work on it in the future and obtain more feedback from them to enhance our work and provide more accessible and inclusive privacy notices for them.

VIII. RELATED WORK

Privacy notice in different formats: Privacy policies are often extensive and contain many specialized terminologies, making users challenging to read and understand. Thus, many users seldom read privacy policies [35]. This issue led to the investigation of different formats of privacy notices. Kelley *et al.* [42] proposed the "nutrition label" for privacy which helped users accurately and quickly find information. However, there may exist non-compliance issues among privacy labels [40], [57]. Le *et al.* [45] developed a browser extension that employs smart lights to deliver notifications regarding data collection. The U.S. Department of Health and Human Services (HHS) introduced Notices of Privacy Practices (NPP) that utilized plain language and accessible designs to facilitate patient comprehension and engagement [14]. Researchers proposed the concept of contextual privacy policies (CPPs) to enhance readability and user engagement [37], [56], [50], [51]. These privacy notice formats primarily focus on visual enhancements and they may be limited to individuals with visual impairments. We proposed *Privacy Notice over Voice* for audio enhancement to bridge this gap.

Privacy issues in Alexa skills: Existing research work highlights the prevalence of privacy issues in Alexa skills. Chen *et al.* [31] conducted a user study about Amazon skills and revealed 84.2% participants struggled with understanding technical terms in the skills' privacy policies. Lau *et al.* [44] found Alexa users are unable to fully understand privacy risks and current privacy controls of Alexa are not well-aligned with users' needs. These user studies show users' privacy concerns and difficulty in understanding privacy policies. Meanwhile, the privacy policies provided by developers couldn't fully disclose the data collection behaviors to users [49], [54], [47], [62], [32], [59], [29]. Liao *et al.* [49] measured the effectiveness of privacy policies for 64,720 Alexa skills and 16,002 Google actions. SkillVet [33] collected 199,295 Amazon skills to check data practices and privacy issues. SkillExplorer [38] discovered 1,141 out of 28,904 skills ask for users' personal information without disclosing it in their privacy policies. PICO [58] detects abnormal data collection behaviors in skills. Yan *et al.* [60] analyzed privacy policies of voice apps from timeliness, availability, completeness and readability.

IX. CONCLUSION

In this work, we introduced a new privacy notice mechanism – *Privacy Notice over Voice*, to enable users for privacy-aware interaction with Alexa through the conversational interface. We first conducted a user study to understand users' and developers' attitudes toward *Privacy Notice over Voice* and the majority of participants (95.6%) considered it as a better way than traditional privacy policies. Informed by the user study results, we designed an automated tool named SKILLPOV for developers to analyze data collection behavior, generate

privacy notices and integrate privacy notices into skill code. We demonstrated the effectiveness of SKILLPOV through a comprehensive evaluation.

ACKNOWLEDGMENT

This work is supported by National Science Foundation (NSF) under the Grant No. 2239605, 2228616, and 2145675.

REFERENCES

[1] Alexa Skills Privacy Requirements. https://developer.amazon.com/fr/docs/custom-skills/security-testing-for-an-alexa-skill.html#25-privacy-requirements.

[2] Amazon managed data identifiers. https://docs.aws.amazon.com/macie/latest/user/mdis-reference-quick.html.

[3] Amazon Mechanical Turk. https://www.mturk.com/.

[4] App privacy policy generator. https://app-privacy-policy-generator.firebaseapp.com/.

[5] Braille Displays. https://store.humanware.com/hus/braille-devices/braille-displays.

[6] California Consumer Privacy Act (CCPA). https://oag.ca.gov/privacy/ccpa.

[7] Children's Online Privacy Protection Rule (COPPA). https://www.ftc.gov/legal-library/browse/rules/childrens-online-privacy-protection-rule-coppa.

[8] 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.

[9] Flesch-Kincaid Readability Test. https://secure.ssa.gov/poms.nsf/lnx/0910605105.

[10] General Data Protection Regulation. https://gdpr-info.eu.

[11] 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.

[12] Hey, Alexa! What are you doing with my data? https://www.ftc.gov/business-guidance/blog/2023/06/hey-alexa-what-are-you-doing-my-data.

[13] Manifest. https://developer.amazon.com/en-US/docs/alexa/smapi/skill-manifest.html.

[14] Model Notices of Privacy Practices. https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/model-notices-privacy-practices/index.html.

[15] NIST: Guide to Protecting the Confidentiality of Personally Identifiable Information (PII). https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-122.pdf.

[16] Number of digital voice assistants in use worldwide from 2019 to 2024. https://www.statista.com/statistics/973815/worldwide-digital-voice-assistant-in-use/.

[17] Prerequisites to build a skill. https://developer.amazon.com/en-US/docs/alexa/build/build-your-skill-overview.html.

[18] Privacy policy comparison reveals half have poor readability. https://www.choice.com.au/consumers-and-data/protecting-your-data/data-laws-and-regulation/articles/privacy-policy-comparison.

[19] Qualtrics. https://www.qualtrics.com/.

[20] Screen Readers. https://www.afb.org/blindness-and-low-vision/using-technology/assistive-technology-products/screen-readers.

[21] Smart speakers - statistics facts. https://www.statista.com/topics/4748/smart-speakers/#topicOverview.

[22] spaCy. https://spacy.io/.

[23] Survey bias types that researchers need to know about. https://www.qualtrics.com/experience-management/research/survey-bias/.

[24] Termly. https://termly.io/.

[25] Terms Feed. https://app.termsfeed.com/.

[26] The new Alexa design guide helps developers design skills that keep users coming back for more. https://developer.amazon.com/en-US/blogs/alexa/alexa-skills-kit/2023/03/alexa-design-guide-march-2023.

[27] Websitepolicies. https://app.websitepolicies.com/.

[28] Word Counter. https://countwordsfree.com/.

[29] Mohammed Aldeen, Jeffrey Young, Song Liao, Tsu-Yao Chang, Long Cheng, Haipeng Cai, Xiapu Luo, and Hongxin Hu. End-users know best: Identifying undesired behavior of alexa skills through user review analysis. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 8(3):1–28, 2024.

[30] Carolina Carreira, Cody Berger, Khushi Shah, Samridhi Agarwal, Yashasvi Thakur, McKenna McCall, Nicolas Christin, and Lorrie Faith Cranor. Who's listening? analyzing privacy preferences in multi-user smart personal assistants settings.

[31] Baiqi Chen, Tingmin Wu, Yanjun Zhang, Mohan Baruwal Chhetri, and Guangdong Bai. Investigating users' understanding of privacy policies of virtual personal assistant applications. In *Proceedings of the 2023 ACM Asia Conference on Computer and Communications Security*, pages 65–79, 2023.

[32] Jide Edu, Xavier Ferrer-Aran, Jose Such, and Guillermo Suarez-Tangil. Measuring alexa skill privacy practices across three years. In *Proceedings of the ACM Web Conference (WWW)*, page 670–680, 2022.

[33] Jide S Edu, Xavier Ferrer-Aran, Jose Such, and Guillermo Suarez-Tangil. Skillvet: automated traceability analysis of amazon alexa skills. *IEEE Transactions on Dependable and Secure Computing*, 20(1):161–175, 2021.

[34] Derar Eleyan, Abed Othman, and Amna Eleyan. Enhancing software comments readability using flesch reading ease score. *Information*, 11(9):430, 2020.

[35] Benjamin Fabian, Tatiana Ermakova, and Tino Lentz. Large-scale readability analysis of privacy policies. In *Proceedings of the International Conference on Web Intelligence*, WI '17, page 18–25, New York, NY, USA, 2017. Association for Computing Machinery.

[36] Angela Fan, Beliz Gokkaya, Mark Harman, Mitya Lyubarskiy, Shubho Sengupta, Shin Yoo, and Jie M. Zhang. Large language models for software engineering: Survey and open problems, 2023.

[37] Denis Feth. Transparency through contextual privacy statements. 2017.

[38] Zhixiu Guo, Zijin Lin, Pan Li, and Kai Chen. SkillExplorer: Understanding the behavior of skills in large scale. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 2649–2666. USENIX Association, August 2020.

[39] Umar Iqbal, Pouneh Nikkhah Bahrami, Rahmadi Trimananda, Hao Cui, Alexander Gamero-Garrido, Daniel J Dubois, David Choffnes, Athina Markopoulou, Franziska Roesner, and Zubair Shafiq. Tracking, profiling, and ad targeting in the alexa echo smart speaker ecosystem. In *Proceedings of the 2023 ACM on Internet Measurement Conference*, pages 569–583, 2023.

[40] Akshath Jain, David Rodriguez, Jose M Del Alamo, and Norman Sadeh. Atlas: Automatically detecting discrepancies between privacy policies and privacy labels. In *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 94–107. IEEE, 2023.

[41] Carlos Jensen and Colin Potts. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '04, page 471–478, New York, NY, USA, 2004. Association for Computing Machinery.

[42] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. A" nutrition label" for privacy. In *Proceedings of the 5th Symposium on Usable Privacy and Security*, pages 1–12, 2009.

[43] J Peter Kincaid, Robert P Fishburne Jr, Richard L Rogers, and Brad S Chissom. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. 1975.

[44] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. Alexa, are you listening?: Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. *Proc. ACM Hum.-Comput. Interact.*, 2(CSCW):1–31, 2018.

[45] Tu Le, Zixin Wang, Danny Yuxing Huang, Yaxing Yao, and Yuan Tian. Towards real-time voice interaction data collection monitoring and ambient light privacy notification for voice-controlled services.

[46] Shuai Li, Zhemin Yang, Yuhong Nan, Shutian Yu, Qirui Zhu, and Min Yang. Are we getting well-informed? an in-depth study of runtime privacy notice practice in mobile apps. 2024.

[47] Suwan Li, Lei Bu, Guangdong Bai, Zhixiu Guo, Kai Chen, and Hanlin Wei. Vitas: Guided model-based vui testing of vpa apps. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, pages 1–12, 2022.

[48] Song Liao, Long Cheng, Haipeng Cai, Linke Guo, and Hongxin Hu. Skillscanner: Detecting policy-violating voice applications through static

analysis at the development phase. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 2321–2335, 2023.

[49] Song Liao, Christin Wilson, Long Cheng, Hongxin Hu, and Huixing Deng. Measuring the effectiveness of privacy policies for voice assistant applications. In *Annual Computer Security Applications Conference (ACSAC)*, page 856–869, 2020.

[50] Anna-Marie Ortloff, Maximiliane Windl, Valentin Schwind, and Niels Henze. Implementation and in situ assessment of contextual privacy policies. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*, pages 1765–1778, 2020.

[51] Shidong Pan, Zhen Tao, Thong Hoang, Dawen Zhang, Tianshi Li, Zhenchang Xing, Sherry Xu, Mark Staples, Thierry Rakotoarivelo, and David Lo. {A New Hope}: Contextual privacy policies for mobile applications and an approach toward automated generation. *arXiv preprint arXiv:2402.14544*, 2024.

[52] Shidong Pan, Dawen Zhang, Mark Staples, Zhenchang Xing, Jieshan Chen, Xiwei Xu, and Thong Hoang. Is it a trap? a large-scale empirical study and comprehensive assessment of online automated privacy policy generators for mobile apps.

[53] Zahy Ramadan, Maya F. Farah, and Lea El Essrawi. From amazon.com to amazon.love: How alexa is redefining companionship and interdependence for people with special needs. *Psychology & Marketing*, 38(4):596–609, 2021.

[54] Faysal Hossain Shezan, Hang Hu, Gang Wang, and Yuan Tian. Verhealth: Vetting medical voice applications through policy enforcement. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2020.

[55] Matthew W Vail, Julia B Earp, and Annie I Antón. An empirical study of consumer perceptions and comprehension of web site privacy policies. *IEEE Transactions on Engineering Management*, 55(3):442–454, 2008.

[56] Maximiliane Windl, Niels Henze, Albrecht Schmidt, and Sebastian S Feger. Automating contextual privacy policies: Design and evaluation of a production tool for digital consumer privacy awareness. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2022.

[57] Yue Xiao, Zhengyi Li, Yue Qin, Xiaolong Bai, Jiale Guan, Xiaojing Liao, and Luyi Xing. Lalaine: Measuring and characterizing {Non-Compliance} of apple privacy labels. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 1091–1108, 2023.

[58] Fuman Xie, Chuan Yan, Mark Huasong Meng, Shaoming Teng, Yanjun Zhang, and Guangdong Bai. Are your requests your true needs? checking excessive data collection in vpa app. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering*, pages 1–12, 2024.

[59] Fuman Xie, Yanjun Zhang, Chuan Yan, Suwan Li, Lei Bu, Kai Chen, Zi Huang, and Guangdong Bai. Scrutinizing privacy policy compliance of virtual personal assistant apps. In *Proceedings of the 37th IEEE/ACM international conference on automated software engineering*, pages 1–13, 2022.

[60] Chuan Yan, Fuman Xie, Mark Huasong Meng, Yanjun Zhang, and Guangdong Bai. On the quality of privacy policy documents of virtual personal assistant applications. *Proceedings on Privacy Enhancing Technologies*, 2024.

[61] Jeffrey Young, Song Liao, Long Cheng, Hongxin Hu, and Huixing Deng. {SkillDetective}: Automated {policy-violation} detection of voice assistant applications in the wild. In *31st USENIX Security Symposium (USENIX Security 22)*, 2022.

[62] Nan Zhang, Xianghang Mi, Xuan Feng, XiaoFeng Wang, Yuan Tian, and Feng Qian. Dangerous skills: Understanding and mitigating security risks of voice-controlled third-party functions on virtual personal assistant systems. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1381–1396. IEEE, 2019.

APPENDIX

A. Instrumented Skill, Front-end Code, Back-end Code

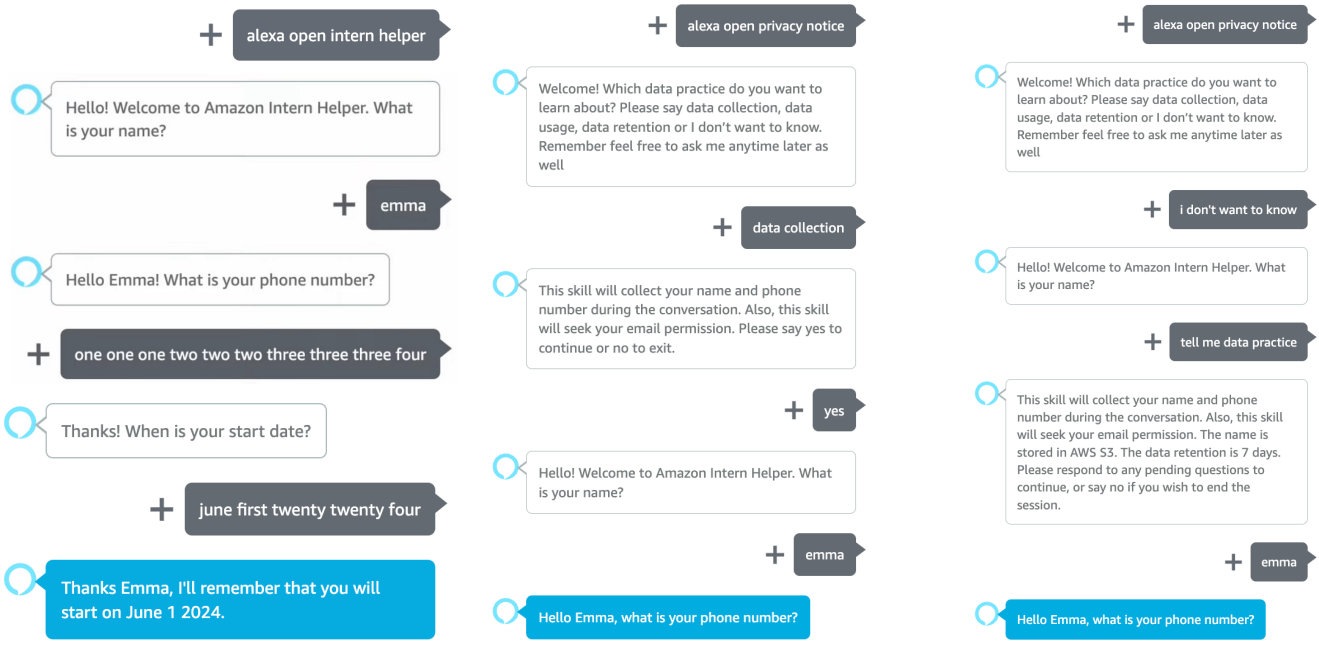

(a) Before modification (b) After modification (c) Asking data collection at any time

Figure 7: Example of the interaction flow modifications. (a) shows the interaction flow before modification. (b) represents interaction flow after modification. When invoking the skill (The new skill is named *Privacy Notice* because reusing the original name *Intern Helper* would result in invoking the existing skill instead.), it initially queries the user whether they wish to be informed about the data practices and which one they want to know. Upon receiving the response, the skill plays the privacy notice and explains the name and email will be collected. It then further inquiries if the user wishes to proceed. With a positive response, the skill formally commences, maintaining consistency with the original skill as illustrated in (a). In (c), users can also inquire about the data practices by the skill at any time. When a user states, "tell me data practice" following any query, the skill will notify users of the data collection, data sharing/storage and data retention (if any). Then, the skill will remind users to answer pending questions and continue using the app. After the user provides their name, the skill returns to its regular operation, proceeding to ask, "What is your phone number?"

```
1.  {
2.   "name": "CaptureUserNameIntent",
3.   "slots": [{
4.     "name": "name",
5.     "type": "AMAZON.FirstName"}],
6.   "samples": [
7.     "{name}", "My name is {name}"]
8.  },
9.  {
10.  "name": "CaptureUserPhoneIntent",
11.  "slots": [{
12.    "name": "phone",
13.    "type": "AMAZON.PhoneNumber"}],
14.  "samples": [
15.    "{phone}", "My phone number is {phone}"]
16. },
17. {
18.  "name": "CaptureStartDateIntent",
19.  "slots": [{
20.    "name": "month",
21.    "type": "AMAZON.Month"},{
22.    "name": "day",
23.    "type": "AMAZON.Ordinal"}],
24.  "samples": [
25.    "My start date is {month} {day}",
26.    "{month} {day}"]
27. }
```

```
1.  {"name": "CaptureUserNameIntent",……},
9.  {"name": "CaptureUserPhoneIntent", ……},        → original intents
17. {"name": "CaptureStartDateIntent", ……},
28. {
29.   "name": "CollectionIntent",                   → newly added intents
30.   "slots": [],
31.   "samples": ["data collection"]                → newly added sample utterances
32. },
33.{
34.   "name": "UsageIntent",
35.   "slots": [],
36.   "samples": ["data usage"]
37. },
38.{
39.   "name": "RetentionIntent",
40.   "slots": [],
41.   "samples": ["data retention"]
42. },
43. {
44.   "name": "EnterSkillIntent",
45.   "slots": [],
46.   "samples": ["yes", "yes, continue", "continue", "I don't want to know"]
39. },
40. {
41.   "name": "NoExitIntent",
42.   "slots": [],
43.   "samples": ["no", "no, exit", "exit"]
44. },
45. {
46.   "name": "DataPracticeIntent",
47.   "slots": [],
48.   "samples": ["tell me data practice"]
49. }
```

Figure 8: Front-end code. The left side displays the front-end code before modification. The right side shows the front-end code after modification. Lines 1-27 on the right side mirror the original front-end code exactly as on the left side. Lines 28-49 include new added intents. Red boxes indicate the names of our newly added intents while green boxes represent newly added sample utterances.

Figure 9: Back-end code. The left side displays the back-end code before modification. The right side shows the back-end code after modification. Lines 1-35 include new intent handlers. Lines 1-11 illustrate the new LauncnRequestHandler which determines whether the user is invoking the skill for the first time and provides different responses accordingly. Lines 32-35 is the original LaunchRequestHandler. Lines 36-61 contain the original intents, which remain unchanged from the code on the left side.

ARTIFACT APPENDIX

In this work, we propose the concept of Privacy Notice over Voice, an accessible and inclusive mechanism to make users aware of the data practices of Alexa skills through the conversational interface: for each skill, we will generate a short and easily understandable privacy notice and play it to users at the beginning of the skill in voice. To realize this concept, we design and develop an automated tool named SKILLPOV (Skill's Privacy Notice over Voice) to generate privacy notices and automatically integrate privacy notices into skill code without impacting original skill functionalities. The submitted artifact consists of three components:

1) Data Practice Analyzer, which scans the skill code to detect data practices and applies an NLP-based method to identify data practices in both conversations and permissions from the manifest file;
2) Privacy Notice Generator, which uses large language models to generate short, easy-to-understand privacy notices based on detected data practices;
3) Privacy Notice Instrumentation, which integrates the generated privacy notices into the skill code and redesigns the skill interaction flow to allow users to continue or quit the skill after hearing the notice.

This artifact aims to reproduce the results described in Sections V-A, V-B and V-C of our paper.

A. Description & Requirements

1) How to access: The artifact is stored in the Zenodo and the DOI link is https://doi.org/10.5281/zenodo.14187988. The artifact can also be accessed by cloning or downloading our publicly available GitHub repository. The repository contains all the necessary data and source code required for reproducing the results. To obtain the artifact, please clone or download the repository using the URL provided in the HotCRP submission.

2) Hardware dependencies: None.

3) Software dependencies: The artifact requires a system running either Linux or macOS, as it is not supported on Windows. Docker (which can be installed from https://www.docker.com/) is necessary to ensure proper environment setup and execution of the provided code. All other dependencies, including Python and relevant libraries, are bundled within the containerized environment. An OpenAI API key is required. An active Alexa Developer Console account is optional, as the execution of the skill takes place within the Alexa Developer Console environment. However, this is not the major claims of this artifact.

4) Benchmarks: We have included part of the examples in the `/dataset/repos` directory, and a complete list of repositories can be found in the `/dataset/all_skills_dataset.csv` file for further reference.

B. Artifact Installation & Configuration

Git clone the artifact access URL: https://github.com/CUSecLab/SkillPoV and open Docker. `cd` into the cloned repository directory. From here:

- Place your OpenAI API key in Line 6 of the following file: `/DockerImage/code/privacy_notice_generator/chatGPT_summary.py`
- To build the system, execute the following commands: `./build.sh`. This step will automatically install all dependencies and libraries.

Please note that each time you upload new repositories or make any changes, rerun `./build.sh`.

C. Experiment Workflow

Run `./run.sh`. Please note that if the `/dataset/results` directory contains output but `/dataset/repos` shows no updates, rerun `./run.sh`. This command will automatically execute all necessary modules, including the Data Practice Analyzer, Privacy Notice Generator, and Privacy Notice Instrumentation, without any further manual steps required.

D. Major Claims

- (C1): SKILLPOV achieved a 100% success rate in generating privacy notices for all skills and instrumenting these privacy notices into the skill code. This is demonstrated by the experiment (E1) whose results are reported in Section VI-A. During the testing process on the Alexa Developer Console, 10 selected skills were able to run successfully and followed the new interaction flow as intended.
- (C2): Most of the generated privacy notices were the same as what we designed and the accuracy rate is 97.67%. This is proven by the experiments (E2) and in Section VI-B of the paper.
- (C3): The privacy notices generated by SKILLPOV achieved a precision of 91.3%. The completeness of privacy notices generated by SKILLPOV is 96.4%. This is proven by the experiment (E3) whose results are reported in Section VI-C.

E. Evaluation

1) Experiment (E1): [Functionality] [5-10 human-minutes]: Whether SKILLPOV can successfully generate privacy notices for skills and integrate them into skill code.

[How to] Manually check the results in the `/dataset/repos` directory.

[Results] A new `index_new.js` file will be generated, which should successfully modify the original `LaunchRequestIntentHandler` and insert new intent handlers. The specific handlers can be found in `/dataset/add_intent.txt`. Additionally, the JSON files will be updated to include the new intents.

2) Experiment (E2): [Reliability] [5-10 human-minutes]: Whether the privacy notice content generated by SKILLPOV using our designed prompts always follows the expected sentence pattern.

[How to] Compare the results in `/dataset/data_collection_results/final` with those in `/dataset/data_collection_results/chatGPT`.

[Results] The folder `/dataset/data_collection_results/final` contains the results of the data collection scan, while the folder `/dataset/data_collection_results/chatGPT` contains the results generated by ChatGPT. The results generated by ChatGPT should ensure consistency with the original scan results, with no omissions or additional information, except for the addition of template content.

3) Experiment (E3): [Accuracy and Completeness] [10-20 human-minutes]: How about SKILLPOV's performance in terms of the Accuracy and Completeness?

[How to] Compare result in `/dataset/data_collection_results/chatGPT` folder with ground truth. The ground truth is obtained through manual scanning of the code and the manifest (skill.json) file.

[Results] To calculate accuracy and completeness, we need true positives, false positives, and false negatives. We have provided all the results in the `/dataset/Skill_Test_Result.csv` file for reference.

F. Customization

In addition to the examples in the `/dataset/repos` directory, we have also provided a complete list of repositories in the `/dataset/all_skills_dataset.csv` file for further reference. Testers can select additional skills from this list for testing. If the skill involves data collection behavior and the original skill functions correctly, testers can further upload it to the Alexa Developer Console for testing to ensure that the skill still runs smoothly and follows the new interaction flow as specified.