

Echoes of Privacy: Uncovering the Profiling Practices of Voice Assistants

Tina Khezresmaeilzadeh
University of Southern California
Los Angeles, California, USA
khezresm@usc.edu

Elaine Zhu
Northeastern University
Boston, Massachusetts, USA
zhu.lae@northeastern.edu

Kiersten Grieco
Northeastern University
Boston, Massachusetts, USA
grieco.k@northeastern.edu

Daniel J. Dubois
Northeastern University
Boston, Massachusetts, USA
d.dubois@northeastern.edu

Konstantinos Psounis
University of Southern California
Los Angeles, California, USA
kpsounis@usc.edu

David Choffnes
Northeastern University
Boston, Massachusetts, USA
choffnes@ccs.neu.edu

Abstract

Many companies, including Google, Amazon, and Apple, offer voice assistants as a convenient solution for answering general voice queries and accessing their services. These voice assistants have gained popularity and can be easily accessed through various smart devices such as smartphones, smart speakers, smartwatches, and an increasing array of other devices. However, this convenience comes with potential privacy risks. For instance, while companies vaguely mention in their privacy policies that they may use voice interactions for user profiling, it remains unclear to what extent this profiling occurs and whether voice interactions pose greater privacy risks compared to other interaction modalities.

In this paper, we conduct 1,171 experiments involving 24,530 queries with different personas and interaction modalities during 20 months to characterize how the three most popular voice assistants profile their users. We analyze factors such as labels assigned to users, their accuracy, the time taken to assign these labels, differences between voice and web interactions, and the effectiveness of profiling remediation tools offered by each voice assistant. Our findings reveal that profiling can happen without interaction, can be incorrect and inconsistent at times, may take several days or weeks to change, and is affected by the interaction modality.

Keywords

privacy, profiling, persona, voice assistant

1 Introduction

Voice assistants, including Amazon Alexa [5], Google Assistant [34], and Apple Siri [11], have rapidly gained popularity as versatile tools for accessing information and services, reducing the need for traditional physical interfaces, with more than 142 million users in the US alone [58]. Despite their convenience, voice assistants still raise concerns regarding user privacy, mirroring issues found in other non-vocal interaction modalities such as web browsing. Previous research [1, 14–16, 20, 26, 29, 43] has demonstrated that online services accessed through web browsers or mobile applications often

leverage user data for profiling, aiming to tailor personalized and sponsored content. In the absence of clear disclosures and explicit consumer consent, such profiling constitutes privacy invasion, as it relies on surveillance of consumers' online activities to produce revenue (e.g., via ads). Additionally, recent findings have revealed that certain interactions with third-party skills of the Alexa voice assistant are utilized for profiling purposes [39]. However, as far as we know, it remains uncertain whether and how voice assistants themselves, without utilizing third-party skills, engage in user profiling. Further, the extent to which their profiling methods resemble or differ from other interaction modalities, along with the associated privacy implications, remains unclear.

This highlights the necessity for further privacy research to explore this topic and offer transparency regarding how voice assistants profile their users. Profiling information becomes integrated into the user's data and can have various ramifications, including impacting the user experience by potentially leading to the display of irrelevant advertisements. Additionally, it can influence sponsoring contracts, such as determining inclusion or exclusion from relevant promotional campaigns. Moreover, there is a risk of unwanted or incorrect information being shared with third parties, highlighting the need to understand and address the potential side effects of user profiling in this new context.

The goal of this research is to shed light on the privacy implications due to interactions with voice assistants. This includes investigating the scope, accuracy, and timing of profiling, such as what aspects are being profiled and whether this profiling is pertinent. Additionally, the research will explore differences in profiling between voice interactions and other modalities, such as web browsing, to understand if and how the methods diverge. Finally, we will analyze the effectiveness of the profiling mitigation tools offered by each voice assistant such as opt-out mechanisms and the possibility to delete or rectify incorrect labels.

Unfortunately, attaining such a goal is challenging because the most popular voice assistants are not software that runs on users' premises, but instead they operate as cloud services, making it very difficult to audit their internal operations. While the privacy policies of these platforms disclose the potential for user profiling based on voice interactions [3, 12, 35], they typically lack details on the methodology involved. Fortunately, privacy regulations in certain jurisdictions (e.g., GDPR [28] and CCPA [56]) compel these platforms to allow users to make data disclosure requests [4, 13, 36].

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.
Proceedings on Privacy Enhancing Technologies 2025(2), 71–87
© 2025 Copyright held by the owner/author(s).
<https://doi.org/10.56553/popets-2025-0050>



Through these requests, platforms are required to provide users with all the information they have collected or generated, including user profiles, often described in terms of demographic (e.g., income group) and interest labels (e.g., an interest in cars).

To address the challenge above and achieve our goal, we developed a rigorous experimental methodology and applied it to the three most prevalent voice assistant platforms: Google Assistant [34], Alexa [5], and Siri [11]. Our approach involves the creation of new user accounts, each meticulously trained with a curated set of voice queries designed to simulate various user personas, each one with an expected set of target profiling labels. Then, we leverage the data disclosure features inherent to each voice assistant platform to monitor the appearance of actual profiling labels. For instance, if an account is trained to mimic a “luxury persona” with questions tailored around luxury topics, we anticipate the emergence of labels aligned with this persona, such as “high income.”

After applying this methodology to the three platforms above and completing 1,171 experiments with a total of 24,530 voice queries over the course of 20 months, we discovered that Apple Siri does not produce profiling labels, while in the remaining platforms profiling does occur and can be erroneous and inconsistent at times. For example, we have seen situations in which profiling happens without any user interactions, resulting in arbitrary profiling labels (Google Assistant). Moreover, the process of profiling varied significantly in accuracy and duration, with unexpected labels appearing (Google Assistant), different labels following the same questions at different times (Google Assistant), and profiling labels taking days to appear (Google Assistant and Amazon Alexa). Importantly, we observed that the interaction modality can influence the accuracy and timing of profiling, with different profiling outcomes depending on whether the interaction occurs via voice commands or through the web (Google Assistant and Amazon Alexa). Finally, we tested the mitigation tools each platform offered and whether any label rectification or opt-out request was correctly fulfilled.

Summarizing, the contributions of this paper are the following:

- (1) a new persona-based methodology for producing profiling labels in different voice assistant platforms (Section 4),
- (2) a characterization of voice assistant first-party profiling across several dimensions, such as accuracy and timing (Sections 5.1 and 5.2), along with an understanding of how platforms pre-populate profiling labels on fresh accounts.
- (3) a comparison of profiling across different interaction modalities (Section 5.3),
- (4) an assessment of the efficacy of mitigation mechanism such as profiling opt-out and label rectification (5.4).

To ensure reproducibility and enable new research, we made our research artifacts, including code and datasets, publicly available at: <https://github.com/NEU-SNS/voice-assistants-pets25.git>.

2 Related Work

Voice assistant privacy. Voice assistants and the devices typically used to access them (e.g., smart speakers) have raised privacy and security concerns, as discussed in several recent surveys [24, 44, 63]. A large amount of existing privacy research focuses on risks that are likely due to technological limitations such as smart speakers

misactivations [27, 32], unintentional software bugs that sometimes led to unauthorized sharing of voice recordings [19, 31, 54], and other instances of data exposure due to third-party apps or skills [23, 39, 64]. Another line of work has shown that it is possible to infer user labels from the characteristics of the recording instead of just from its transcription [55], thus creating more profiling opportunities for voice assistant vendors. While these works focus on important risks, we are complementary with respect to them since we focus on the measurement and mitigation of intentional data collection in terms of profiling labels to the first party (i.e., the voice assistant vendor). Specifically, we focus on targeting the labels of each specific platform using explicitly biased utterances while maintaining voice characteristics as neutral as possible.

Persona-based inference of labels. In this work, we use a persona-based methodology to produce labeling in the context of voice assistants and then use data disclosure requests to obtain user labels. This methodology was inspired by the ones used in a large body of prior work (e.g., [1, 14–16, 20, 26, 29, 43]), where different synthetic profiles (personas) are used to study how data is propagated to third parties in contexts such as the web and mobile apps. These works typically use personas to determine what data is shared with third parties [29], server-side data exchanges [16, 26], and other online tracking and advertising behavior [1, 20, 43]. Although our work also uses a persona-based approach, a limitation of most existing work is that it focuses on limited modalities (e.g., the web) that do not include voice assistant. One recent persona-based work [39] focuses on voice assistants, but it has a different scope with respect to ours since it is limited to third-party skills, it only focuses on the Alexa voice assistant, it does not target the voice assistant vendor itself (i.e., the first party), and obtains labeling data indirectly and in a probabilistic way through the analysis of the ads ecosystem rather than using deterministic data disclosure requests as we do.

Other labeling determination approaches. Prior work in the broader context of Internet of Things (IoT), including smart speakers, has traditionally addressed the labeling problem by understanding what data is shared to their cloud platforms (both belonging to the device’s vendor and to third parties). This has been typically done by relying on device modification or logging capabilities of specific platforms [22, 33, 40, 62]. Other work used black-box experiments to identify how the viewing habits gathered by smart TVs [30, 45] and smart speakers interactions [39] were shared with advertisers. Finally, other studies that focus on web [18] and mobile apps [49] have relied on instrumentation [37, 50], automation, and network traffic interception [51, 52] not available in the context of voice assistants. In this work, we specifically address voice assistants, which are cloud services accessed through smart speakers or apps, and not devices themselves. Moreover, profile inference typically happens in the cloud and not the device itself. For this reason, we cannot use device modification or network traffic analysis as done in prior related work. Similarly to some of the work above, we use a black-box approach, where we only access entry points (i.e., our voice queries) and exit points (i.e., user labels) from data disclosure requests mandated by law. The main difference is that we do not rely on third parties such as advertisers, since we target first-party (i.e., vendor) labels and not how such labels are shared or used.

3 Context, Definitions, and Goals

We now describe the context, definitions, and goals of this work.

3.1 Context

The integration of voice assistants into daily life has revolutionized how we interact with technology, offering convenience and efficiency in performing a wide array of tasks. However, this interaction brings significant privacy implications, particularly concerning the profiling capabilities of these systems. Profiling by voice assistants involves analyzing voice commands, searches, and interactions to construct detailed profiles of users [3, 12, 35]. These profiles can reveal private personal information, including demographic details, that users might prefer to keep confidential [4, 13, 36].

The implications of such unintentional data sharing are multifaceted. Firstly, there is the risk of unauthorized access to sensitive information, where personal data could be exposed to third parties or inadvertently leaked through security breaches [19, 31, 54]. Secondly, the demographic and personal information obtained from interactions can be used to tailor marketing efforts in a highly targeted, and potentially intrusive, manner [39].

Even though users, thanks to privacy laws [28, 56], can now access their profile and even opt out [4, 13, 36], the opaque nature of data processing and profiling practices by voice assistant providers complicates users' ability to control what information is collected about them and how it is used [41]. Many users are unaware of the extent to which their data contributes to their digital profile or the ways in which this profile influences their interactions with technology and content online [42, 46, 60].

Understanding the privacy implications of interacting with voice assistants is important for recognizing the potential risks to personal privacy. It highlights the importance of transparent data practices, robust privacy settings, and the need for users to be informed and cautious about the information they share through these increasingly ubiquitous services.

3.2 Definitions

In the context of our study on the privacy implications and user profiling capabilities of voice assistant platforms, we refine our definitions to ensure clarity and precision in our analysis.

Voice assistants. We define voice assistants as software components designed to interpret and execute voice commands, enabling them to perform various tasks. Voice assistants can receive voice inputs from users and provide audible feedback or take actions based on the analysis of these inputs. The voice assistants we consider in this study, selected based on their popularity in the US [57], are the following: Google Assistant [34], Alexa [5], and Siri [11]. All these voice assistants are cloud-based, meaning that the user-controlled devices used to access them (e.g., smart speakers) simply forward the recordings to them, with the processing (and profiling) happening in the cloud. This process is typically detailed in the privacy policies pages of each voice assistant.

Profiling labels. We define profiling labels as the key components within a user profile, derived from user data and interactions with the voice assistants. These labels demonstrate specific characteristics that distinguish users from one another. We categorized them into two types: demographic and interest labels. Demographic

labels capture the user's demographic attributes, such as relationship status, educational background, etc. These labels can either be based on explicit user data or be derived from user's interactions. In contrast, interest labels offer insights into the user's preferences and interests in products or ad categories. These labels usually are inferred from the user's interactions. Both types of labels enable personalized experiences and targeted content delivery.

Persona. A key aspect of our experimental methodology is to train voice assistant user accounts as a certain *persona*. Taking inspiration from previous work [1, 14–16, 20, 26, 29, 43], we conceptualize a persona as a composite profile consisting of one or more specific profiling labels that collectively define the perceived characteristics, preferences, or demographic details of a user as interpreted by a voice assistant. To construct a specific persona with targeted labels, we curate a series of tailored queries and ask them to the voice assistant. For instance, one might create a persona in Google's ecosystem that is identified as "Single" or "Renter," or in Amazon's ecosystem as having an interest in "Fashion." The definition of a persona is contingent upon the profiling mechanisms of each voice assistant, which use interaction data to assign labels according to their user characterization methods.

Training queries. We define training queries as a curated set of queries intentionally designed to shape or guide the development of a targeted persona. The purpose of these queries is to direct the profiling algorithms of the voice assistant toward converging on a specific label that accurately reflects the intended persona.

Modalities of interaction. We define modalities of interaction as the interaction methods through which users can access information from each voice assistant company. These interactions can occur through voice inputs via voice assistant interfaces or text-based queries on web interfaces. Note that we do not distinguish modalities between vocal voice assistant interactions and textual non-web voice assistant interactions using accessible interfaces designed for people who cannot speak, since during our preliminary experiments, we observed that such interactions are processed in the same way and produce the same results (Appendix B.3).

3.3 Goals

The goal of this study is to explore if and how voice assistants profile their users with labels—both demographic (e.g., "Married") and interest-based (e.g., "Fashion"). Specifically, we will address the following research questions (RQ): *RQ1. Before experimentation, does the voice assistant have any prepopulated labels?*

This research question aims to determine if the system begins with any predefined labels on newly created accounts before any interaction with the system. Having wrong initial labels may affect user experience (e.g., irrelevant recommendations) and mislead advertisers, with the potential of wrong labels being propagated to third parties. We address this question by creating 200 accounts, and then waiting for labels to appear.

RQ2. After experimentation, does profiling actually happen and to what extent?

We aim to quantify the extent of user profiling by voice assistants in terms of accuracy and consistency, assessing the depth and breadth

of data collection. The purpose is to improve transparency over profiling practices. To answer this question, we perform 460 persona-based voice interaction large-scale experiments using 19 different personas for Google, 7 for Amazon, and 5 for Apple (see Section 4 for a detailed description of our experiments and justification for the selected number of personas per platform).

RQ3. Does the interaction modality affect profiling?

This question investigates whether different modalities of interaction (e.g., web vs. voice) have distinct impacts on the scope and accuracy of user profiling. This is important to understand whether users are profiled (or misprofiled) differently when accessing information through a voice assistant with respect to another modality such as the web. To achieve this goal, we perform web persona-based experiments performing the same queries, but with a different modality, and then compare the results with our voice results.

RQ4. Are the opt-out options provided by these platforms effective?

This research question focuses on identifying whether the current strategies to mitigate the privacy risks associated with user profiling by voice assistants are sufficient enough to protect users from being misprofiled or from being profiled against their will. To answer it we investigate whether we can effectively remove or change profiling labels from our persona accounts and the level of granularity (e.g., partial vs. complete opt-out). In the next sections, we detail the methodology employed to investigate these RQs (Section 4), followed by an analysis of the findings (Section 5), and a discussion on the implications of our results (Section 6).

3.4 Ethics

This research project involves only the authors of the paper and their direct collaborators listed in the acknowledgment section; no external research subjects were recruited. The experiments were meticulously performed one at a time to minimize any disruption or disturbance to the voice assistants being tested. While Google, Amazon, and Apple prohibit misrepresentation (such as using fake data), these rules primarily target activities that harm or affect other users, such as fraud, impersonation, or identity misrepresentation. Our experiments did not impact other users or disrupt services and did not involve fraud. Our approach aligns with established practices in similar studies, see, for example, [2, 17, 21, 25].

We did not need to disclose our findings with the companies hosting the voice assistants since their privacy policies already acknowledge the profiling activities associated with their services. Instead, our focus has been to provide greater transparency regarding these profiling activities and their implications for privacy.

4 Methodology

In this section we describe the methodology we used for answering our research questions. Figure 1 shows a summary of our methodology, divided into steps. In the first step we prepare the voice assistant for our experimentation, i.e., we create fresh accounts and reset the smart devices we use to access them. We then perform Steps 2–3 to understand prepopulated labels (RQ1); in particular, in Step 2 we perform a data disclosure request to obtain such prepopulated labels, while in Step 3 we analyze the presence of such labels. In Steps 4–6 we turn the fresh accounts into profiled accounts using

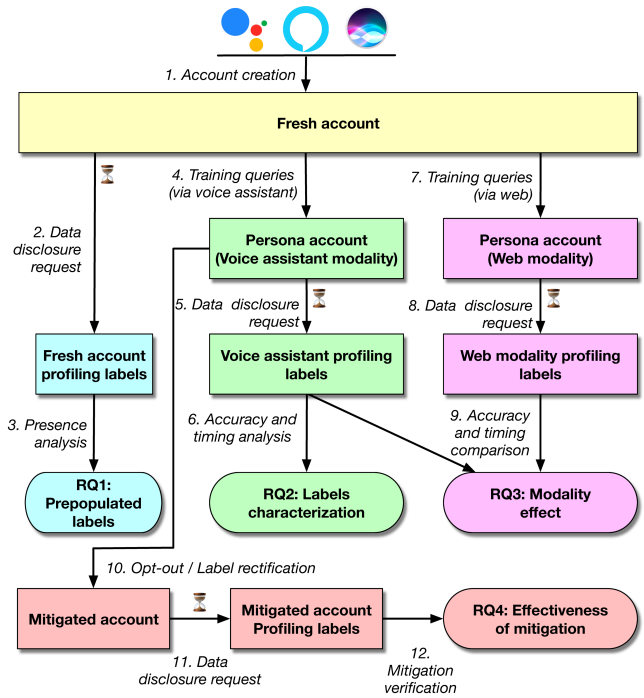


Figure 1: Methodology overview. Step 1 is the creation of fresh accounts. We perform Steps 2–3 to determine prepopulated labels before experimentation (RQ1). We perform Steps 4–6 to characterize profiling after experimentation (RQ2). We perform Steps 7–9 to analyze the effect of the modality (RQ3). Finally, we perform Steps 10–12 to evaluate profiling mitigation (RQ4). Steps marked with “x” include a one-week waiting period to allow sufficient time for profiling to occur.

persona-based voice interactions and then analyze their profiling labels (RQ2). Specifically, in Step 4 we perform training voice queries for different personas; in Step 5 we obtain profiling labels following a data disclosure request; and in Step 6 we analyze such labels, checking their accuracy and time needed to profile. Steps 7–9 are similar to 4–6, but using the web modality instead of the voice modality, so that we can see whether the modality has an effect on profiling (RQ3). Finally, we perform Steps 10–12 to understand the effectiveness of the profiling mitigation tools offered by each voice assistant (RQ4). In Step 10 we use opt-out and label rectification mechanisms to remove or amend the labels; in Step 11 we perform a data disclosure request for obtaining the new mitigated labels (after re-enabling profiling and waiting 10 days); and in Step 12 we verify that profiling mitigation actually worked (i.e., ensuring that labels are still absent or amended). Note that before proceeding with the data disclosure request steps, we allow a one-week wait time for profiling to occur, as profiling is not instantaneous.

In the next paragraphs, we introduce our testing environment (Section 4.1), our platform-independent approach for preparing for the experiments and ensuring minimal account contamination (Section 4.2), our platform-agnostic experimental approach (Section 4.3), and the platform-specific experimental approach for Google assistant (Section 4.4), Amazon Alexa (Section 4.5), and Apple Siri

(Section 4.6). We then conclude this section with a brief discussion of early experiments (Section 4.7) and with a summary and timeline of all our experiments per RQ (Section 4.8). Note that during the 20 months of experimentation we conducted three major batches of experiments, the first in Spring 2023, the second in Fall 2024, and the third in Summer 2024, details of which are shown in Table 1. For the rest of the paper we collectively refer to the first two early batches as preliminary experiments, and to the third most recent one as the large-scale experiment.

4.1 Testing Environment

Our testing environment, located in the United States, for evaluating voice assistant devices comprises three main components: (i) the voice assistant platforms being tested; (ii) a PC equipped with a Python script to automate the playback of audio recordings; and (iii) a pair of non-smart speakers connected to the PC, tasked with playing the recordings to the voice assistant.

To access Google, Alexa, and Siri voice assistants, we used both smart speakers and smart phones/tablets. (In preliminary testing, we found that profiling labels remained consistent regardless of device type, and thus likely have no impact on profiling.) Specifically, we use the following voice-assistant enabled devices:

- **Google Assistant Access:** (i) 3 iPhone 7 devices running iOS 15.8.3, (ii) 1 iPhone 8 running iOS 16.7.8, and (iii) 1 iPhone 13 running iOS 16.0. The Google Home app version on all iPhones was 1.9.115.901.
- **Alexa Access:** (i) 1 Amazon Echo Dot (1st generation), (ii) 3 Amazon Echo Dot (3rd gen.), and (iii) 1 Amazon Echo Dot (5th gen.). The Alexa Home app was installed on an iPhone 13 Pro running iOS 17.5.1, with Alexa app version 2024.15.
- **Siri Access:** 1 iPhone 13 Pro Max running iOS 16.0.

For each voice assistant, we created distinct personas, each associated with a persona-specific set of English training queries. To guarantee the consistency of our experiments across multiple accounts of a persona, we pre-recorded the queries we intended to ask the voice assistants using the voice of one of the authors. Then, we developed an automated Python script to systematically play back these recordings. The process involves the script first triggering the voice assistant device’s wake word (e.g., “Hey Siri”), followed by a single query. It then pauses for a one-minute interval, which during preliminary testing proved to be a sufficient window for all our voice assistants to process and respond to the query. Upon completion of this interval, the script proceeds to play the next query. During the whole experimentation, we manually supervised the automatic execution of the experiments to ensure that the voice assistant under test always activates (i.e., it recognizes the wake word) and always understands the voice query correctly (i.e., its answers are relevant to the questions). This methodical approach ensures our experiments are repeatable under identical conditions.

4.2 Experiment Preparation

To ensure that our results accurately reflected the profiling behavior of voice assistants without the influence of external variables, we adopted a rigorous method to maintain the independence and isolation of each device and account. While we cannot guarantee

perfect isolation, our method constitutes a best effort to control as many relevant factors as possible.

For each voice assistant platform, we created several fresh accounts (one for each persona instantiation) using the Chrome web browser (due to its popularity [59], and likely use by many consumers) in private mode to minimize the chances of being tracked.

Since the account creation process required a phone number for SMS verification, we obtained a new non-VoIP phone number for each account to ensure independence. Each phone number was used only once per platform, so no two accounts shared the same phone number. By using a browser in private mode and by not reusing phone numbers across accounts within the same platform, we minimize any potential cross-contamination of data or profiling influences that could compromise the integrity of the experiment.

Additionally, we reset the device used to factory settings to access the voice assistant before logging in with a fresh account. This ensured that no residual data from the previous account could influence the voice assistant’s behavior.

Finally, to prevent our IP address from affecting the profiling process, we connected the smart devices used to access voice assistants through a VPN using the same IP address for all experiments, thus ensuring that any observed differences in profiling were attributable to the interactions specific to each account and not to external network factors. Each persona in our study had unique behaviors, and our results showed distinct profiling outcomes across personas, with no evidence of cross-contamination. Thus we have evidence that profiling differences were driven by each account’s unique interactions, rather than any influence from shared IP usage.

4.3 Common Experimental Approach

We employed a common experimental approach across all the voice assistants we tested, as explained below. The following subsections highlight the specific variations in our methodology tailored to each voice assistant.

4.3.1 Profiling and Data Access. Our study began with determining which profiling labels are observable and how they can be accessed on each platform. Since each platform employs different approaches to categorize user information, we systematically searched for accessible profiling labels on each, exploring their data access tools and documentation to prepare for evaluating their profiling behaviors.

4.3.2 Persona Choices. After identifying the accessible profiling labels of each system, we designed personas, each targeting a specific profiling label. This strategy enabled us to develop tailored queries that guide the system towards recognizing each profiling label, and allowed us to study the system’s behavior for each label. For each platform, we first identified the possible profiling labels and then created a distinct persona for each label. For each persona, we created 10 accounts per persona to ensure a meaningful sample size while efficiently managing computational resources and time constraints. Previous studies in similar contexts have shown that 10 samples per persona are sufficient to detect consistent patterns and trends (shown in Appendix B.2 and B.3).

4.3.3 Query Choices. To develop the queries for our experiments, we employed prompt engineering with the goal of generating relevant questions that align with typical user interactions and engage

with each voice assistant’s profiling system. We first used ChatGPT-4o [47] to assist in the process of query generation, ensuring that the queries were not influenced by the biases of the authors. We then carefully reviewed and manually refined the queries based on consensus among the authors. The ChatGPT-4 prompts were crafted with trial and error, using direct, concise instructions with relevant keywords to clearly convey the profiling labels and ensure consistency across different personas. The tailored prompts used for this purpose are included in Appendix A.

4.3.4 Waiting Periods. Based on our preliminary experiments (Appendix B.2), we implemented an initial “prepopulation” one-week waiting period to allow any prepopulated labels to appear, as we observed that these labels typically emerge within this timeframe. This approach also enabled us to identify any pre-existing labels, ensuring that we excluded accounts where the target label was already present from the corresponding experiments.

After this initial waiting period, we deployed a set of training queries to trigger account profiling. Profiling by the voice assistants is not immediate either, and each platform requires a “profiling” waiting period before profiling labels start becoming observable. We determined this waiting period experimentally, and chose a standardized one-week waiting period for all platforms. This decision was informed by our preliminary experiments conducted in the Spring and Fall of 2023 (see Section 4.7), which indicated that profiling labels typically emerge within this timeframe (by 5.9 days for Google, and by 5.5 days for Alexa).

In our most recent, large-scale experiments, we observed that profiling happened slower than in our preliminary experiments (see Section 5 and more specifically Figure 2). Motivated by this, if some target labels were not achieved within the initial one-week period, we generated a new set of queries using the prompts described in Appendix A and waited for a new one-week period before re-observing the profiles. We continued this monitoring process for a total of four weeks, at which point we observed that profiles were stabilized. (This approach additionally helped maintain consistent interaction levels, reducing the risk of accounts being flagged due to insufficient activity.)

4.3.5 Modality Experiments. Profiling via web and voice modality could be similar, given that transcripts of voice interactions can be used similarly to web input. Investigating the nature of this similarity is important, both to understand relevant privacy implications and because consumers may have different expectations from different modalities [48]. To assess the impact of different interaction modes on user profiling, we conducted experiments using two querying methods: web searches and voice commands. For each service, we replicated persona experiments using a web browser, logged into new accounts, and compared the results to those obtained via voice interactions. This reveals how each modality influences the creation and updating of user profiles.

4.3.6 Mitigation Experiments. We conducted mitigation experiments beginning with accounts associated with personas that had successfully reached their target profiling labels. Once these accounts have been profiled, we opted out the accounts to stop data collection and profiling. Later, we opted the accounts back in and monitored their profiles over time to see if opting out had a lasting impact on the system’s ability to profile based on past interactions.

In the rest of the paper, for experiment counting purposes, we count account actions, such as checking for prepopulated labels in an account (RQ1), asking a batch of queries to an account (RQ2, RQ3), and opting out of profiling for an account (RQ4), as one experiment each.

4.4 Google Assistant Experiments

4.4.1 Profiling and Data Access. Google’s ad personalization strategy utilizes comprehensive user interaction data from its services and it enables users to modify their advertisement preferences in their Google Account settings [35]. Google categorizes this information into multiple demographic categories, outlined as follows:

- *Relationship*: whether the user is single, married, etc.
- *Household Income*: estimated income bracket.
- *Education*: highest level of education achieved.
- *Industry*: sector of employment.
- *Employer Size*: rough size of the user’s place of employment.
- *Homeownership*: whether the user owns or rents their home.
- *Parenting*: whether the user has children.

Each of the above demographic categories has several possible options referred to as profiling labels. Google allows users to view their profiling labels in real-time through its My Ad Center interface [36], encompassing services like Google Assistant.

In addition to demographic categories, Google also profiles users based on interest labels (as defined in Section 3.2). We monitored the impact of interactions on the generation of these interest labels by systematically reviewing the My Ad Center interface [36] over a 4-week period. However, we did not observe any evidence of interest labels being assigned through interactions with Google Assistant, suggesting that Google likely does not derive interest labels from voice interactions. Consequently, we focused our analysis of Google profiling on demographic labels, as they were the only ones demonstrably affected by voice interactions.

4.4.2 Persona Choices. Given Google’s demographic-based profiling system, our selection process for personas included all possible demographic labels to ensure comprehensive targeting¹. We considered each label as a potential persona. This allowed us to create a set of 19 distinct personas for Google profiling experiments.

4.4.3 Query Choices. For Google-specific experiments, we tailored prompts specifically for Google Assistant to align with Google’s profiling system. The prompt used for this purpose is included in Appendix A.1. For each persona (profiling label), ten distinct Google account samples were created. We designed 20 queries for each profiling label. This number was chosen as it was sufficient to trigger the target profiling labels based on insights from preliminary experiments (Appendix B.2). This allowed us to systematically evaluate the impact of query type on Google’s profiling accuracy.

4.4.4 Mitigation Experiments. In addition to opting-out, we also investigated the ability to selectively amend and delete labels offered by Google. In this case, our methodology does not include a waiting period since label deletion and editing are processed immediately (i.e., the change is immediately visible in Google My Ad Center [36]).

¹ *Industry* and *Employer Size* demographic categories were excluded due to their generalized nature, which complicates precise targeting.

4.5 Amazon Alexa Experiments

4.5.1 Profiling and Data Access. On Amazon’s website [8], there is no explicit way for users to visualize their personal data categorization, except for explicitly self-assigned labels available through the Amazon Profile Hub [6]. (The Amazon Profile Hub allows users to manually enter their demographic and interest labels. We observe that it does not automatically set or infer these labels based on user behavior.) To access inferred profiling data, Amazon facilitates Data Subject Access Requests (DSAR) [10], allowing users to obtain specific data collected by Amazon.

From preliminary testing, we found out that information about user profiling is mainly found in the “Advertising” category, highlighting where users can find information related to how they are profiled for advertising purposes. After downloading information from the Advertising category, users are given the following files: (i) *Advertising.OptOut.csv* which tracks user opt-outs of profiling by country, (ii) *Advertising.AmazonAudiences.csv* which lists categories of products that interest the user according to Amazon’s analysis, (iii) *Advertising.3PAudience.csv* which identifies user interests based on third-party data, (iv) *Advertising.AdvertiserAudiences.csv* which details which advertisers consider the user part of their target audience, and (v) *Advertising.AdvertiserClicks.csv* which records instances where the user clicked on ads.

Upon interaction with an Amazon profile, the “AmazonAudiences” file reveals specific interest labels. These labels provide the categories of content and products that align with the user’s interests. We did not find evidence of demographic labels except for gender, only when explicitly declared by the user. For this reason, when analyzing Amazon, we focus on interest-based labels that are directly tied to user behavior and product engagement.

4.5.2 Persona Choices. Given the fundamental difference between the labeling systems of Google and Amazon—demographic vs. interest-based—we needed to select personas that align with Amazon’s focus on specific user interests rather than demographic categories. Our selection was based on the set of all relevant interest labels we identified from available data, providing a solid foundation for a robust experimental analysis. This approach ensures that the personas are tailored to the nuances of Amazon’s profiling system.

Amazon does not disclose all the interest labels it assigns to user profiles. To find such labels to target in our experiments, we performed Data Subject Access Requests (DSARs) for all the authors of this study. Upon receiving the data, we analyzed the “AmazonAudiences” file in each author’s profile, where we identified a range of interest labels. We selected all In-Market first-level interest labels that appeared across the profiles, ensuring that our chosen personas reflected a broad spectrum of interests within Amazon’s system.

The seven personas derived from this process are as follows:

- **Video Entertainment Persona:** Targeted at users interested in movies and television shows.
- **Fashion Persona:** For users passionate about fashion.
- **Beauty and Personal Care Persona:** Designed for users interested in beauty and personal care products.
- **Electronics Persona:** For users interested in electronics.
- **Toys and Games Persona:** For users interested in toys/games.
- **Pet Supplies Persona:** Designed for users who frequently inquire about pet products.

- **Books and Magazines Persona:** Focused on users interested in books and magazines

4.5.3 Query Choices. For Amazon, we customized prompts for its voice assistant to align with their profiling system. Prompt details are in Appendix A.2. Since Amazon Alexa supports command queries related to e-commerce, unlike Google Assistant, we structured our queries into two distinct types to analyze their effects on profiling: (i) *General Queries* from which users seek factual information that Alexa answers without any additional action, and (ii) *Command Queries* which require Alexa to perform specific tasks, such as adding items to a shopping list or making a purchase. This categorization helps us examine how each type of query affects profiling on Amazon, highlighting its e-commerce strengths. For Google Assistant, which does not provide command queries related to demographics like Amazon, we did not split the queries, focusing instead on general demographic profiling.

For each persona, 10 distinct Amazon account samples were created. For each query type we designed 30 queries. This number was chosen based on preliminary experiments, where 30 queries were sufficient to trigger the target profiling labels on Amazon (Appendix B.2). In contrast, we used 20 queries for the Google experiments, as this number was found to trigger profiling labels in that context. The differences in query numbers were tailored to the specific characteristics of each platform.

4.5.4 Modality Experiments. Since users can interact with Amazon services through voice and web interfaces, we compare these two modalities to verify how differently they affect profiling. However, we limit the comparison to command queries because the Amazon web interface primarily searches for similar products and allows actions such as adding items to the cart or shopping list, but does not support general queries as the voice interface does. Note that the command queries were conducted on the Amazon webpage [7] by searching for products and adding them to the cart.

4.5.5 Mitigation Experiments. Unlike Google, Amazon does not provide a direct method to modify or selectively delete its interest labels. Instead, it offers an indirect approach through the “Improve My Recommendation” page [9]. However, since this page is intended for enhancing recommendations and does not allow for direct visualization, removal, or editing of interest labels, we do not take it into account in this study.

4.6 Apple Siri Experiments

4.6.1 Profiling and Data Access. Our research methodology included submitting Data Subject Access Requests (DSARs) through Apple’s designated platform [13]. Despite our efforts, the data retrieved from these requests did not reveal explicit profiling information related to Siri interactions neither during preliminary experiments nor by inspecting the personal Apple accounts of the authors. For this reason, our experiment methodology could not include persona experiments targeting Apple Siri labels. Instead, since Siri uses third-party services such as Google Search or Bing Search to answer generic voice queries, we perform indirect persona experiments targeting such third-party services.

To address RQ2, which involves characterizing profiling activity, we focused on a third-party entity responsible for generating search results for Siri. We hypothesized that this third party may also

conduct profiling activities on behalf of Apple. Since we already had an established methodology for Google profiling, we considered the integration with Google search and linked newly created Google Accounts to the device used to access Siri.

4.6.2 Persona and Query Choices. For the indirect experiments targeting Siri, we selected the five personas from the Google experiments that demonstrated the highest accuracy. These personas were characterized by the labels: “Single,” “Married,” “Renters,” “Advanced Degree,” and “In a Relationship.” We chose these five personas based on results from Google experiments, described in Section 5.2.1, which indicated that the remaining personas had a low chance of being accurately profiled. We used the exact same queries designed for those personas in the Google experiments.

4.6.3 Modality Experiments and Mitigation Experiments. Since Apple does not offer a web modality to access Siri information, we could not perform modality experiments for Apple Siri. Regarding mitigation experiments, Apple offers an opt-out mechanism similar to the other two platforms; however, we could not test it since Siri did not produce profiling labels.

4.7 Preliminary Experiments

The previous discussion on our platform-specific experimental approach is based on our most-recent, large-scale experiments. Here, we briefly discuss some variations between these and our preliminary experiments. For a detailed presentation of these preliminary experiments, see Appendix B. Most notably, the preliminary experiments were smaller in scale, for example, the Amazon experiments only studied two of the seven interest labels that we studied in our current experiments. As another example, the Google and Apple experiments involved a much smaller number of accounts, queries, and labels. Last, note that in our preliminary experiments we also studied both single-label and multi-label personas. These preliminary experiments laid the groundwork for our final methodology, allowing us to make informed decisions on persona construction, query design, interaction modalities, and wait times.

4.8 Experiments Summary and Timeline

To summarize, we performed a total of 1,171 experiments and 24,530 training queries over a period of 20 months. The definition of experiment is mentioned in Section 4.3. An overview of all the experiments we performed and their dates and relationship with each research question is reported in Table 1, showing the number of labels, experiments per label, and queries per label, along with the timeline of execution. Note that for each experiment used to answer RQ1, RQ2, and RQ3 we used a fresh account per experiment, while for RQ4, we used previously profiled accounts.

5 Results

In this section we discuss the results of our experiments and the answers to our research questions. Unless otherwise stated, we are referring to the results of the most-recent, large-scale experiments.

5.1 Prepopulated Labels in Profiles (RQ1)

5.1.1 Google. We created 200 fresh Google accounts in summer 2024 and did not use them for one week to observe evidence of

VA	RQ	Time	# Labels	# Exp.	# Queries
Google	RQ1	Spr 23	-	16	-
		Fall 23	-	30	-
		Sum 24	-	200	-
	RQ2	Spr 23	4	32	20
		Fall 23	8	16	20
		Win 24	8	16	20
		Sum 24	19	190	60
	RQ3	Spr 23	8	16	20
		Sum 24	19	95	60
	RQ4	Win 24	8	16	-
		Sum 24	19	95	-
	Amazon	RQ1	Sum 23	-	16
Sum 24			-	70	-
RQ2		Sum 23	2	16	30
		Win 24	2	12	30
RQ3		Sum 24	7	70	30
		Win 24	2	12	30
		Sum 24	7	35	30
RQ4		Win 24	2	12	-
		Sum 24	7	70	-
Apple		RQ1	Sum 23	-	12
	Sum 24		-	50	-
	RQ2	Sum 23	4	24	20
		Sum 24	5	50	20
Total=				1171	24530

Table 1: Overview of our experimental efforts, showing the number of labels, experiments, and queries per label for each research question (RQ) and voice assistant (VA). Experiments refer to number of account actions, see Section 4.3. Total number of queries (with repetitions) and total number of experiments are also included.

profiling labels prior to any user interaction, as mentioned in Section 4.3. Our observations showed several prepopulated labels upon creation. As shown in Table 2, the tags “Homeowners” and “Not Parents” were the most frequently assigned prepopulated labels, appearing in 195 out of 200 accounts. Other categories such as “Relationship,” “Household Income,” “Education,” “Industry,” and “Employer Size” also exhibited varying degrees of prepopulated information. For instance, 92 accounts were labeled “In a Relationship,” and 194 accounts were tagged with “Moderately High Income.”

This phenomenon suggests an initial profiling stage where accounts are assigned characteristics that could influence subsequent personalization strategies.

5.1.2 Alexa. During all our experiments with fresh accounts, we have not seen prepopulated interest profiling labels. The only change we have seen is the appearance of an empty Advertising.AmazonAudiences file after a few days, confirming the absence of any interest label.

5.1.3 Siri. We found that Apple Siri does not assign any profiling labels to users before any interaction, which shows there is no early stage of profiling activity in Apple upon account creation and prior to any user interaction.

Takeaway: Google may assign prepopulated profiling demographic labels even before any user interaction occurs, while that does not happen for Alexa and Siri.

Category	Label	Count/200
Relationship	Not Enough Info	104
	Married	4
	Single	0
	In a Relationship	92
Household Income	Not Enough Info	5
	High Income	1
	Moderately High Income	194
	Average or Lower Income	0
Education	Not Enough Info	196
	Highschool Diploma	4
	Attending College	0
	Bachelor's Degree	0
Industry	Advanced Degree	0
	Not Enough Info	200
	Other	0
Employer Size	Not Enough Info	5
	Small Employer	0
	Large Employer	195
	Very Large Employer	0
Homeownership	Not Enough Info	5
	Homeowners	195
	Renters	0
Parenting	Not Enough Info	5
	Not Parents	195
	Other	0

Table 2: Distribution of demographic tags assigned to new Google accounts during the pre-population phase. The table categorizes 200 accounts with different labels, with the frequency of each label presented.

5.2 Characterization of Profiling (RQ2)

For each voice assistant, we analyze the following profiling aspects: (i) *Accuracy*, defined as the fraction of persona accounts that are correctly labeled after a given time period following a persona experiment, and (ii) *Time needed for profiling*, defined as the time required for the accounts to transition from an initial label to a different label. (In case of no change, we define this metric as “N/A.”)

5.2.1 Google. In summer 2024, we conducted our experiments with Google to assess its profiling capabilities, as detailed in Section 4.4. We used fresh Google accounts and began by waiting one week to allow any prepopulated labels to appear, as described in Section 4.3. After this initial waiting period, we selected accounts that did not have the target profiling labels prepopulated and proceeded with the experiments by asking the queries.

To ensure a fair comparison across different labels, we excluded experiments for the “Not parents,” “Moderately High Income,” and “Homeowners” personas because these labels were consistently present across all accounts and we did not want to count those as successful profiling outcomes. On average, it took 18.0 ± 4.1 days for all profiling labels to appear in our Summer 2024 experiments, based on Table 5.

Accuracy results. The accuracy of the accounts after a four-week monitoring interval time is shown in Table 3. The experimental findings for each profiling category are as follows: (i) *Relationships Category*: Initially, no prepopulated labels existed for the relationship status of the accounts. Post-experiment, all labels achieved positive accuracy levels. The label “Married” stood out with the

highest accuracy at 70%, making it the most reliably identified status within this group. The other labels, such as “Single” and “In a relationship,” also displayed positive accuracy but were less accurate in comparison to the “Married” label, achieving 30% and 10% respectively. (ii) *Education Category*: In the Education category, all accounts were initially marked with “Not enough info.” As the study progressed, the label “Advanced Degree” showed an increase in accuracy, achieving a 50% accuracy rate, marking it as the most accurately identified educational status in the experiment. In contrast, the other educational labels—“High school diploma,” “Attending college,” and “Bachelor’s degree”—remained at “Not enough info” throughout the monitoring period. (iii) *Homeownership Category*: All accounts were initially labeled as “Homeowners.” Despite this starting point, the “Renters” label was adjusted to an 80% accuracy rate by the end of the experiment, demonstrating a noticeable correction from the initial label. (iv) *Household Income and Parenting Categories*: These categories maintained strong label adherence. The initial labels “Moderately high income” and “Not parents” persisted without any changes throughout the duration of the study, indicating stability in these profiling categories.

The profiling results demonstrate distinct outcomes across various categories. In the “Household Income” and “Parenting” categories, the same prepopulated labels—“Moderately high income” and “Not parents”—were observed across all personas and remained unchanged throughout the experiments. Conversely, in the “Homeownership” category, adjustments were made to prepopulated labels, indicating changes based on profiling activities. Furthermore, all labels in the “Relationships” category achieved positive accuracy, showcasing effective profiling, while in the “Education” category, three personas maintained an accuracy of 0%, indicating no effective profiling for these attributes.

These results confirm that profiling is actively occurring. The variability in profiling outcomes illustrates that, despite using a uniform methodology across all categories, the accuracy of the profiling results varies over different personas (profiling labels).

Takeaway: Google profiles users based on their interactions, with profiling accuracy varying across different labels.

Comparison with Preliminary Experiments Results. In Spring 2023, we conducted preliminary experiments focusing on the profiling accuracy of four labels: “Married,” “Bachelor’s degree,” “Single,” and “High school diploma.” During that period, voice experiments resulted in an average accuracy of 80% across these labels, with an average time to label change of 4.8 ± 1.1 days, as detailed in Appendix B. In our final Summer 2024 experiments, there have been some changes. The average accuracy for voice experiments for these four labels has decreased to 20%, while the average time to label change has increased to 15.9 ± 4.3 days for the same four labels, as shown in Table 3. Figure 2 further highlights this shift, showing a noticeable delay in the timing of label convergence in the Summer 2024 experiments compared to Spring 2023.

Takeaway: The accuracy and speed of profiling by Google voice assistant have decreased over time, with noticeable differences across the four labels tested, despite using the same methodology.

5.2.2 Alexa. Accuracy results. In our analysis of Alexa’s profiling accuracy for different interest labels, we observed changes only in the “AmazonAudiences” file. The impact of the training questions

Profiling Category	Profiling Label	Accuracy (%)
Relationships	Married	70%
	Single	30%
	In a relationship	10%
Household Income	High Income	0% (inaccurate)
	Moderately high	-
	Average or low	0% (inaccurate)
Education	High school diploma	0% (no tags)
	Attending college	0% (no tags)
	Bachelor's degree	0% (no tags)
	Advanced degree	50%
Homeownership	Homeowners	-
	Renters	80%
Parenting	Not Parents	-
	Infants	0% (inaccurate)
	toddlers	0% (inaccurate)
	preschoolers	0% (inaccurate)
	grade schoolers	0% (inaccurate)
	teenagers	0% (inaccurate)

Table 3: Accuracy of label assignment for Google experiments across different profiling labels (in percentage): “no tags” indicates all labels were categorized as “Not Enough Info,” while “inaccurate” denotes incorrect label assignment.




Figure 2: Comparison of target label achievement between Spring 2023 and Summer 2024 experiments, only including accounts that successfully reached their target labels. The Spring 2023 experiments show faster profiling progression, while the Summer 2024 experiments indicate a slower rate of achieving target labels.

on profiling accuracy is detailed in Table 4. On average, it took Alexa 7.7 ± 1.3 days to generate profiling labels after our persona experiments. No evidence of profiling from general questions was found in the experiments, while command-based queries resulted in 100% profiling accuracy. These results support our initial hypothesis that for Amazon to assign relevant personas to accounts, the queries must be directly related to its product offerings rather than consisting of general inquiries.

Takeaway: Users are profiled only based on command queries, such as adding items to a shopping cart or making a purchase, and such profiling has been 100% accurate in our experiments.

Query	Persona	Acc (%)	Avg Time (Days)
General	Fashion	0	N/A
	Video	0	N/A
	Beauty	0	N/A
	Electronics	0	N/A
	Toys	0	N/A
	Pet Supplies	0	N/A
	Books	0	N/A
Command	Fashion	100%	8.2 ± 1.0
	Video	100%	7.6 ± 1.4
	Beauty	100%	6.2 ± 1.5
	Electronics	100%	8.4 ± 1.4
	Toys	100%	7.6 ± 0.8
	Pet Supplies	100%	8.4 ± 0.5
	Books	100%	7.6 ± 1.0

Table 4: Comparison of profiling accuracy and the mean time required for accurate label convergence for Alexa’s general and command queries across different personas.

Comparison with Preliminary Experiments Results. In both our preliminary and final experiments, command-based queries achieved 100% profiling accuracy. However, the time required for accurate label convergence increased slightly, from 5.0 ± 0.82 days in the preliminary experiments (as shown in Table 10 in Appendix B) to 7.7 ± 1.3 days in the final experiments (Table 4). This suggests a slight variation in timing, though accuracy remains consistent.

Takeaway: For the experimented personas, the accuracy of Amazon’s profiling through command-based queries has remained consistent over time, while the time required for profiling has slightly increased, indicating a slower response in the profiling process.

5.2.3 Siri. Direct profiling. As discussed in Section 6.2, Apple’s privacy policies indicate that they may infer user interests for targeted advertising [12]. However, our experimental submissions of Data Subject Access Requests (DSARs) for various persona-based Apple accounts yielded no profiling information. This outcome suggests that Apple does not profile users based on their Siri interactions.

Takeaway: Despite Apple’s policy statements allowing profiling, there is no evidence that Apple engages in profiling activities based on voice interactions through Siri.

Indirect profiling. We next investigated whether there is any profiling via Apple Siri’s integration with third-party search engines for answering user queries. Given our established methodology for Google profiling, we focused on the integration with Google search and linked newly created Google Accounts to the device used to access Siri. We used the queries for the top five selected personas of Google voice experiments (see Section 5.2) to uncover any indirect profiling by Google, as evidenced by changes in the Google privacy dashboards linked to the Siri interactions. Despite allowing a waiting period of four weeks after the completion of each of these experiments, we observed no new profiling activities in the Google accounts associated with the device used to perform the Siri queries. This lack of observed profiling activity suggests the absence of indirect profiling by Google for Siri voice searches.

Takeaway: There is no evidence of indirect profiling by Google based on Siri voice queries when Siri uses Google searches.

Modality	Persona	Acc (%)	Avg Time (Days)
Voice	Married	70%	16.1 ± 4.9
	Single	30%	15.3 ± 2.6
	In a relationship	10%	25
	High school	0%	N/A
	Attending college	0%	N/A
	Advanced degree	50%	16.5 ± 1.5
Web	Renters	80%	19.6 ± 2.5
	Married	80%	1.8 ± 1.3
	Single	60%	1.0 ± 0.0
	In a relationship	60%	1.0 ± 0.0
	High school	20%	1.0 ± 0.0
	Attending college	60%	5.7 ± 0.5
	Advanced degree	60%	4.0 ± 2.2
	Renters	100%	1.0 ± 0.0

Table 5: Comparison of profiling accuracy and the mean time required for accurate label convergence across Google voice and web modalities. The labels that did not reach any level of accuracy (0%) in both modalities were excluded in this table.

5.3 Impact of Interaction Modality (RQ3)

To evaluate the impact of interaction modalities on user profiling practices, we systematically conducted experiments across different modalities for Google and Alexa, as detailed in Sections 4.4–4.5. This approach enables direct comparisons of how interaction modalities contribute to the profiling capabilities of the voice assistants.

Note that all the voice assistants we examined provide an option to input queries via typing, serving as an alternative for individuals who may be unable or unwilling to use voice queries. However, we do not categorize these textual interactions as distinct modalities. Unlike the web modality discussed in this section, these inputs are directly processed by the voice assistant.

5.3.1 Google. In our investigation, we compared the profiling efficacy of voice and web modalities on Google’s platform using the same persona queries detailed in Section 4.4.

The results, summarized in Table 5, illustrate the differences in mean accuracy and the average time required for changes in Google’s label profiling to be observed following the experiments with each modality. The web modality yielded a higher average accuracy of 62.86%, compared to the voice modality, which exhibited a lower accuracy of 48% across all labels. Furthermore, the average label convergence time for web modality was 2.2 ± 2.0 days, which was lower than the one for voice modality, which averaged 18.0 ± 4.1 days. These findings highlight the superior performance of web modality in both accuracy and efficiency.

Takeaway: Interactions with Google via web modality are more accurate and faster than those conducted through voice.

5.3.2 Alexa. We explored the impact of different interaction modalities on user profiling by Amazon, focusing on voice and web queries as outlined in Section 4.5. It is important to note, as mentioned in Section 4.5, that we excluded web-based queries for the Video Persona due to the lack of a direct web interface equivalent for the voice command “Add to Watchlist” used for the Video Persona.

Table 6 shows the web modality interaction results, where we used the same command queries using the web interface. (We could not use the web interface for general questions requiring information, since Amazon website does not provide such a service.) The

Modality	Persona	Acc (%)	Avg Time (Days)	
Voice	Fashion	100%	8.2 ± 1.0	
	Video	100%	7.6 ± 1.4	
	Beauty	100%	6.2 ± 1.5	
	Electronics	100%	8.4 ± 1.4	
	Toys	100%	7.6 ± 0.8	
	Pet Supplies	100%	8.4 ± 0.5	
	Books	100%	7.6 ± 1.0	
	Web	Fashion	100%	7.4 ± 0.5
		Video	100%	8.6 ± 1.4
		Beauty	100%	8.2 ± 1.2
Electronics		100%	6.4 ± 1.5	
Toys		100%	6.4 ± 0.8	
Pet Supplies		100%	5.8 ± 0.7	
Books		100%	7.0 ± 0.9	

Table 6: Comparison of profiling accuracy and the mean time required for accurate label convergence across Amazon voice and web modalities.

web results show the same 100% accuracy and an average time of 7.1 ± 1.4 days for labels to change in web experiments, which is similar to the results of voice interactions (7.7 ± 1.3 days).

Takeaway: For Alexa, the web modality profiles similarly to the voice modality and with the same 100% accuracy in our experiments.

5.3.3 Siri. Since Siri does not offer the option to search through a web interface, the way Google and Amazon do, we were unable to investigate how different modalities affected Siri profiling.

5.4 Profiling Mitigation (RQ4)

5.4.1 Google. Google provides users with two approaches to manage profiling via voice assistant interactions, tailored for those who prefer not to share their information for ad personalization and those who prefer to have fine-grained control over it.

Direct modification of demographic labels. This feature allows users to personally adjust and confirm their demographic labels within the Google Privacy Dashboard [36]. By navigating to this site, individuals have the flexibility to review and modify their demographic information to align with their actual identifiers. This functionality empowers users to directly influence the profiling mechanism to better reflect their preferences.

After testing this feature, we confirmed that users can either change their existing demographic labels or confirm an existing tag of their choosing, ensuring that the ad content is closely aligned with their specified characteristics.

Selective opt-out from profiling. Google not only offers a general opt-out from profiling but also allows users to specifically disable profiling for certain categories [36]. To evaluate the impact of opting out on user profiles, we conducted an experiment with all accounts being profiled using the highly accurate web modality (Table 5) for all personas. Following this, we opted out of profiling for these accounts. After a one-week break, we reactivated ad personalization and monitored the accounts daily for more than 10 days.

Observations revealed that immediately after re-enabling ad personalization, all accounts were assigned the “Not Enough Info” label across all demographic fields. This condition remained unchanged

after four weeks, evidencing that previous interactions prior to opting out did not influence subsequent profiling post-opting in.

Takeaway: Both mitigation mechanisms of Google (changing the demographic labels directly from the Google Privacy Dashboard and opting out from profiling) work as advertised and therefore are useful to mitigate mislabeling and unwanted profiling.

5.4.2 Alexa. To investigate whether Amazon has safeguards in place to enhance user privacy, we selected “Do not show me interest-based ads provided by Amazon” and deleted personal information from Amazon’s ad systems, as found on Amazon’s ad preferences page [4]. After monitoring the accounts for four weeks, we found that the recommendations page showed generic ads, rather than personalized ads. A limitation of Amazon’s approach is that there is no selective way to remove or update the labels once they appear.

Takeaway: Also for the Amazon Alexa platform, choosing to opt out works as expected and is an effective way to limit profiling. However, Amazon does not provide tools to limit mislabeling by allowing users to selectively remove or rectify their labels.

5.4.3 Siri. According to Apple’s privacy policy [12], users have the ability to deactivate Personalized Ads by adjusting their device settings, which prevents the delivery of ads tailored to their interests. However, we could not test this since, as discussed in Section 5.2.3, we did not observe any direct or indirect user profiling by Apple with respect to Siri interaction.

Takeaway: Apple provides tools to limit profiling in terms of opt-outs, but we could not test them because we never experienced profiling with Siri. For the same reason, we could not investigate the presence of tools for rectifying and selectively removing labels.

5.5 Results from Preliminary Experiments

So far, we have focused on reporting results from our most recent larger-scale experiments. Here we summarize additional results from our preliminary experiments; see Appendix B for more details.

RQ1: Our experiments revealed that Google accounts may have prepopulated labels, whereas Amazon and Siri did not show any signs of prepopulated labels.

RQ2: Our experiments revealed that profiling does indeed happen. Interestingly, we found that the fewer the number of labels per persona, the faster and more accurate the labeling is. Specifically, while one- and two-label personas demonstrated good accuracy, four-label ones required repeated queries a few weeks after the first round of queries to show good accuracy. In our recent experiments, informed by these insights, we targeted one label across all personas. This approach allowed us to cover all labels with more number of accounts per label and conduct large-scale experiments.

RQ3: Our experiments revealed that Google’s web modality outperformed voice, while Amazon showed similar behavior between web and voice, both consistent with our recent findings. Note that since Siri does not offer the option to search through a web interface, we did not test the web modality for that as well.

RQ4: Our experiments revealed that mitigation tools are effective, consistently with our recent experiments. Specifically, for Google, after opting out of profiling categories and reactivating ad personalization, we found that personas did not regenerate their previous labels. For Amazon, after disabling interest-based ads and deleting

personal information, we observed a shift to generic recommendation ads. We could not test Siri’s profiling mitigation tools due to the lack of profiling by Siri.

6 Discussion

This section describes privacy implications of our results, an overview of the systems’ privacy policies, and limitations of our study.

6.1 Implications of Our Findings

Prepopulated labels. In our experiments with Google Assistant, we observed that some newly created accounts had prepopulated labels, a phenomenon not observed with other platforms we studied. The exact mechanisms by which these labels are generated are unclear. It is possible that the platform is using indirect methods, e.g., IP address geolocation, to infer these labels [61]. Because these labels could be inaccurate, they could lead to incorrect ad targeting or personalized content. It is worth noting that these prepopulated labels tend to be replaced as users begin interacting with the platform, as shown in Section 5.2.1, indicating that Google may not have high confidence in the initial labels. If so, we recommend better transparency by sharing confidence levels of labels with users.

Profiling accuracy. We observe that Google Assistant and Alexa conduct profiling differently. Alexa’s approach is deterministic, predicting user interests based on explicit information, resulting in 100% accuracy and 100% consistency (Table 4). However, Google’s profiling appears probabilistic, possibly depending on factors like demographic labels, timing of experiments, and interaction modality, as shown in Tables 3 and 5. This introduces the risk of mislabeling, leading to similar concerns with prepopulated labels as discussed above. The lack of information regarding how these systems operate and the undisclosed rationale behind label assignments make it hard to know the causes of this probabilistic behavior.

Profiling timing. On Amazon and Google’s platforms, profiling typically began after user activity started, with the profiling process taking approximately 7.7 days on Amazon and 18.0 days on Google, as stated in Section 5.2. During this time, a user’s profile is being formed but not yet finalized. This presents a privacy concern: if a user requests their data too early (i.e., before the profiling period is complete), they may receive no information about profiling, leading them to believe they have not been profiled or have been profiled incorrectly. To promote transparency, we suggest that platforms disclose the profiling timeline and explore options to reduce and disclose the time required to complete a user’s profile.

Modality. We found that Google exhibits varying profiling outcomes depending on modality, impacting both timing and accuracy, and suggesting different profiling methods for the web and voice modality. Specifically, users of voice assistants face a higher risk of mislabeling and longer delays in profiling (Tables 5 and 6). This finding aligns with prior research indicating that modality can influence the accuracy and effectiveness of user profiling [38].

Mitigation. All platforms provided options to opt out from profiling, which helps address privacy concerns related to profiling. However, there are times when users might want to be profiled, such as to receive more relevant content during voice searches. In these cases, the ability to selectively edit or remove labels is important for reducing risks associated with profiling (e.g., due

to mislabeling). In our study, we found that only Google offered the option to fully opt out of profiling while also allowing users to manage their labels. Note that Amazon’s Profile Hub [6] provides a method for users to declare their own labels. However, these user-declared labels are added to the inferred ones, and the inferred labels cannot be removed unless the user opts out entirely. These tools align with prior work that emphasizes the importance of user control over their data to enhance privacy and personalization [53].

6.2 Privacy Policies vs. Our Results

Google Privacy Policy. An in-depth review of Google Assistant’s privacy policy [35] indicates that personalized advertisements are fundamentally rooted in user-specific data, such as the data we provide to Google when creating accounts, our activity using Google services, and our location information. The policy clearly excludes sensitive details like race, religion, sexual orientation, and health from ad targeting. Importantly, all profiling activities we observed are allowed by the policy.

Amazon Privacy Policy. Amazon’s privacy policy [3] contains many references to personalized advertising, with ads being tailored based on a mixture of recommended products/services and personal user information. The policy covers a wide range of personal data collected, including basic identification and contact details, interaction logs, and personal identifiers. The profiling we observed is a small subset of the potential scope mentioned in the policy.

Apple Privacy Policy. Apple’s privacy policy [12] states that user interactions are the basis for targeted advertising. While we did not observe profiling, the policy acknowledges the profiling potential, which can happen anytime in the future. Therefore, even if current experiments indicate lower privacy risks from profiling compared to competitors, the underlying potential for profiling remains comparable.

6.3 Limitations

Our study faced multiple limitations that narrowed the depth and scope of our analysis, each presenting unique challenges:

Scale of the experiments. The requirement to create numerous new accounts for controlled experiments with voice assistants was the foremost challenge in conducting the experiments. Although some account creations did not need any phone numbers, some accounts needed phone numbers, and since we wanted to have completely independent accounts, we had to use a unique phone number for each account. This process was time-consuming, and significantly slowed down our research progress.

Profile contamination. While we used different phone numbers for each account, identifiers such as IP addresses and device IDs remained unchanged across different experiments due to the reuse of the same devices for multiple experiments (and, to avoid IP-address-dependent profiling). This can be a source of contamination.

Opaque box assumption and dynamism. We are not testing static software, but cloud services that are constantly being updated, and therefore every experiment may be a test of a different system. This variability could provide an alternative explanation for observing different results or behaviors at different times.

Query generation. Our voice assistant queries are simple in nature to make it trivial to profile our personas; however, these do not

necessarily represent how real people issue voice commands. Thus a limitation of our study is that a company’s algorithms may react differently to our personas than to real people. That said, it would be infeasible to run controlled experiments on real people who have never used the Internet before our study.

Voice characteristics. We assume that voice characteristics are not used for profiling purposes, but previous research [55] has shown that there is potential for this kind of profiling. Thus, our results could be influenced by the use of voice characteristics profiling. To mitigate this limitation, we used the same voice for all experiments. For reasons mentioned above, it is also infeasible to run controlled experiments on real people.

Undisclosed labels. Our reliance on the visibility of labels for compliance with privacy legislation does not account for the potential of undisclosed profiling by voice assistants, which are essentially opaque boxes. Although we aim to provide a lower bound for profiling based on disclosed information, we acknowledge the limitation of not being able to detect undisclosed profiling activities.

Language specificity. Our study used only the English language for conducting experiments. The profiling behaviors and results may vary with other languages, so our findings do not generalize to non-English contexts.

Despite these limitations, our study provides valuable insights into the profiling of voice interactions by Google, Alexa, and Siri. This lays the groundwork for future research to explore these areas further and motivates greater transparency for voice assistants.

7 Conclusion

In this study, we provided an overview of profiling practices of voice assistants. Through an extensive persona-based experimentation spanning 1,171 experiments and 24,530 voice queries over a period of 20 months, we confirmed that Google Assistant and Amazon Alexa engage in user profiling based on voice interactions, while evidence of such profiling was not found for Apple Siri, despite Apple’s disclosure in its privacy policy.

We observed distinct profiling methods between Google and Alexa, with Google employing a probabilistic approach that continuously refines demographic profiling labels based on user interactions, while Alexa adopts deterministic interest profiling labels primarily from explicit user commands within the Amazon shopping ecosystem. Our investigation also revealed different profiling behavior between voice and web interactions for Google, with lower profiling accuracy and longer profiling periods for the voice modality. Last, while all services offer opt-out mechanisms, only Google allows users to selectively delete or amend labels.

Through this study, we hope to motivate voice assistant companies to improve transparency and refine tools for mitigating misprofiling and empowering users to edit their profiles.

Acknowledgments

This research was partially funded by the NSF under the ProperData Frontier award (SaTC-1955227, SaTC-1956435) and by Consumer Reports Innovation Lab. We thank the reviewers for their valuable feedback and Ganesh Danke, Shriya Dhaundiya, Akshata Lolayekar, Tilak Patel, Sairam Ramasubramanian, Shantanu Sachdeva, Landyn Sparacino, and Kaveh Waddell for their contributions.

References

- [1] Pushkal Agarwal, Sagar Joglekar, Panagiotis Papadopoulos, Nishanth Sastry, and Nicolas Kourtellis. 2020. Stop tracking me Bro! Differential Tracking of User Demographics on Hyper-Partisan Websites. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) (WWW '20). Association for Computing Machinery, New York, NY, USA, 1479–1490. <https://doi.org/10.1145/3366423.3380221>
- [2] Pushkal Agarwal, Sagar Joglekar, Panagiotis Papadopoulos, Nishanth Sastry, and Nicolas Kourtellis. 2020. Stop tracking me Bro! Differential Tracking of User Demographics on Hyper-Partisan Websites. In *Proceedings of The Web Conference 2020* (Taipei, Taiwan) (WWW '20). Association for Computing Machinery, New York, NY, USA, 1479–1490. <https://doi.org/10.1145/3366423.3380221>
- [3] Amazon. 2023. Amazon.com Privacy Notice. <https://www.amazon.com/gp/help/customer/display.html?nodeId=Gx7NJQ4ZB8MHFRNJ>. (Accessed on 02/29/2024).
- [4] Amazon. 2024. Advertising Privacy and Preferences. <https://www.amazon.com/adprefs>. (Accessed on 02/29/2024).
- [5] Amazon. 2024. Alexa Voice Assistant. <https://developer.amazon.com/en-US/alexa>. (Accessed on 02/29/2024).
- [6] Amazon. 2024. Amazon Profile Hub. <https://www.amazon.com/slc/hub> [Accessed: 2024-08-27].
- [7] Amazon. 2024. Amazon.com. Spend less. Smile more. <https://www.amazon.com/> [Accessed: 2024-11-12].
- [8] Amazon. 2024. Data and Privacy. <https://www.amazon.com/>. (Accessed on 02/29/2024).
- [9] Amazon. 2024. Improve Your Recommendations. <https://www.amazon.com/gp/yourstore/iyr>. (Accessed on 02/29/2024).
- [10] Amazon. 2024. Request Your Data. <https://www.amazon.com/hz/privacy-central/data-requests/preview.html>. (Accessed on 02/29/2024).
- [11] Apple. 2024. Siri Voice Assistant. <https://www.apple.com/siri/>. (Accessed on 02/29/2024).
- [12] Apple Inc. 2023. Apple Advertising & Privacy. <https://www.apple.com/legal/privacy/data/en/apple-advertising/>. (Accessed on 02/29/2024).
- [13] Apple Inc. 2024. Data and Privacy. <https://privacy.apple.com/>. (Accessed on 02/29/2024).
- [14] Paul Barford, Igor Canadi, Darja Krushevska, Qiang Ma, and S. Muthukrishnan. 2014. Adscape: harvesting and analyzing online display ads. In *Proceedings of the 23rd International Conference on World Wide Web* (Seoul, Korea) (WWW '14). Association for Computing Machinery, New York, NY, USA, 597–608. <https://doi.org/10.1145/2566486.2567992>
- [15] Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. 2016. Tracing Information Flows Between Ad Exchanges Using Retargeted Ads. In *Proceedings of Usenix Security*. Austin, TX.
- [16] Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. 2016. Tracing information flows between ad exchanges using retargeted ads. In *25th USENIX Security Symposium* (USENIX Security 16), 481–496.
- [17] Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, and Christo Wilson. 2016. Tracing Information Flows Between Ad Exchanges Using Retargeted Ads. In *25th USENIX Security Symposium* (USENIX Security 16). USENIX Association, Austin, TX, 481–496. <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/bashir>
- [18] Muhammad Ahmad Bashir and Christo Wilson. 2018. Diffusion of User Tracking Data in the Online Advertising Ecosystem. *Proceedings on Privacy Enhancing Technologies* (PoPETS) 2018, 4 (July 2018).
- [19] S. Burke. 2017. Google admits its new smart speaker was eavesdropping on users. <http://money.cnn.com/2017/10/11/technology/google-home-mini-security-flaw/>. (Accessed on 02/29/2024).
- [20] Juan Miguel Carrascosa, Jakub Mikians, Ruben Cuevas, Vijay Erramilli, and Nikolaos Laoutaris. 2015. I Always Feel like Somebody's Watching Me: Measuring Online Behavioural Advertising. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies* (CoNEXT '15).
- [21] Juan Miguel Carrascosa, Jakub Mikians, Ruben Cuevas, Vijay Erramilli, and Nikolaos Laoutaris. 2015. I always feel like somebody's watching me: measuring online behavioural advertising. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*. 1–13.
- [22] Z Berkay Celik, Leonardo Babun, Amit Kumar Sikder, Hidayet Aksu, Gang Tan, Patrick McDaniel, and A Selcuk Ulugac. 2018. Sensitive information tracking in commodity IoT. In *27th {USENIX} Security Symposium* ({USENIX} Security 18), 1687–1704.
- [23] Long Cheng, Christin Wilson, Song Liao, Jeffrey Young, Daniel Dong, and Hongxin Hu. 2020. Dangerous Skills Got Certified: Measuring the Trustworthiness of Skill Certification in Voice Personal Assistant Platforms. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security* (CCS '20). Association for Computing Machinery, Virtual, 1699–1716.
- [24] Peng Cheng and Utz Roedig. 2022. Personal Voice Assistant Security and Privacy—A Survey. *Proc. IEEE* 110, 4 (2022), 476–507. <https://doi.org/10.1109/JPROC.2022.3153167>
- [25] John Cook, Rishab Nithyanand, and Zubair Shafiq. 2019. Inferring tracker-advertiser relationships in the online advertising ecosystem using header bidding. *arXiv preprint arXiv:1907.07275* (2019).
- [26] John Cook, Rishab Nithyanand, and Zubair Shafiq. 2020. Inferring Tracker-Advertiser Relationships in the Online Advertising Ecosystem using Header Bidding. *Proceedings on Privacy Enhancing Technologies* (PETS '20) 1 (2020), 65–82.
- [27] Daniel J Dubois, Roman Kolcun, Anna Maria Mandalari, Muhammad Talha Paracha, David Choffnes, and Hamed Haddadi. 2020. When Speakers Are All Ears: Characterizing Misactivations of IoT Smart Speakers. In *Proceedings on Privacy Enhancing Technologies*, Vol. 2020 (4). Sciencio, Virtual, 255–276.
- [28] European Commission. 2018. General Data Protection Regulation (GDPR). <https://gdpr-info.eu/>. (Accessed on 02/29/2024).
- [29] Shehroze Farooqi, Maaz Musa, Zubair Shafiq, and Fareed Zaffar. 2020. Canary-Trap: Detecting Data Misuse by Third-Party Apps on Online Social Networks. *Proceedings on Privacy Enhancing Technologies* (PETS '20) 4 (2020), 336–354.
- [30] Financial Times. 2019. Smart TVs sending private data to Netflix and Facebook. <https://www.ft.com/content/23ab2f68-d957-11e9-8f9b-77216eb1f17>. (Accessed on 09/18/2023).
- [31] Geoffrey A. Fowler. 2019. Alexa has been eavesdropping on you this whole time. <https://www.washingtonpost.com/technology/2019/05/06/alexa-has-been-eavesdropping-you-this-whole-time/>.
- [32] Gary Horcher. 2018. Woman says her Amazon device recorded private conversation, sent it out to random contact. <https://www.kiro7.com/news/local/woman-says-her-amazon-device-recorded-private-conversation-sent-it-out-to-random-contact/755507974/>.
- [33] Dennis Giese and Guevara Noubir. 2021. Amazon echo dot or the reverberating secrets of IoT devices. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*. 13–24.
- [34] Google. 2024. Google Assistant. <https://assistant.google.com/>. (Accessed on 02/29/2024).
- [35] Google. 2024. Google Privacy Policy. <https://policies.google.com/privacy>. (Accessed on 02/29/2024).
- [36] Google. 2024. My Ad Center. <https://myadcenter.google.com/controls>. (Accessed on 02/29/2024).
- [37] Catherine Han, Irwin Reyes, Álvaro Feal, Joel Reardon, Primal Wijesekera, Narseo Vallina-Rodriguez, Amit Elazar, Kenneth A Bamberger, and Serge Egelman. 2020. The price is (not) right: Comparing privacy in free and paid apps. *Proceedings on Privacy Enhancing Technologies* (PoPETS) 2020, 3 (2020).
- [38] Abdelhamid Hachouat, Slimane Bellaouar, Attia Nehar, and Hadda Cherroun. 2023. Modality Influence in Multimodal Machine Learning. *arXiv preprint arXiv:2306.06476* (2023).
- [39] Umar Iqbal, Pouneh N Bahrami, Rahmadi Trimananda, Hao Cui, Alexander Gamero-Garrido, Daniel J. Dubois, David Choffnes, Athina Markopoulou, Franziska Roesner, and Zubair Shafiq. 2023. Tracking, Profiling, and Ad Targeting in the Alexa Echo Smart Speaker Ecosystem. In *Proc. of the Internet Measurement Conference* (IMC'23) – to appear.
- [40] Yunhan Jack Jia, Qi Alfred Chen, Shiqi Wang, Amir Rahmati, Earlene Fernandes, Z Morley Mao, Atul Prakash, and Shanghai JiaoTong University. 2017. ContextIoT: Towards Providing Contextual Integrity to Appified IoT Platforms. In *Proceedings of the 21st Network and Distributed System Security Symposium* (NDSS'17).
- [41] Monica Kowalczyk, Johanna Gunawan, David Choffnes, Daniel Dubois, Woodrow Hartzog, and Christo Wilson. 2023. Understanding Dark Patterns in Home IoT Devices. In *Proceedings of ACM Human Factors in Computing Systems* (CHI 2023). Hamburg, Germany.
- [42] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, are you listening? Privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. *Proceedings of the ACM on human-computer interaction* 2, CSCW (2018), 1–31.
- [43] Mathias Lecuyer, Riley Spahn, Yannis Spiliopolous, Augustin Chaintreau, Roxana Geambasu, and Daniel Hsu. 2015. Sunlight: Fine-grained targeting detection at scale with statistical confidence. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security* (CCS '15). 554–566.
- [44] Jingjin Li, Chao Chen, Mostafa Rahimi Azghadi, Hossein Ghodosi, Lei Pan, and Jun Zhang. 2023. Security and privacy problems in voice assistant applications: A survey. *Computers & Security* 134 (2023), 103448. <https://doi.org/10.1016/j.cose.2023.103448>
- [45] Sapna Maheshwari. 2018. TVs That Find An Audience For Your Data. *The New York Times* (Jul 2018), A1. <https://www.nytimes.com/2018/07/05/business/media/tv-viewer-tracking.html>
- [46] Nathan Malkin, Joe Deatrack, Allen Tong, Primal Wijesekera, Serge Egelman, and David Wagner. 2019. Privacy attitudes of smart speaker users. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (2019).
- [47] OpenAI. 2024. ChatGPT. <https://chat.openai.com/chat> [Accessed: 20-May-2024].
- [48] Kentrell Owens, Johanna Gunawan, David Choffnes, Pardis Emami-Naeini, Tadayoshi Kohno, and Franziska Roesner. 2022. Exploring Deceptive Design Patterns in Voice Interfaces. In *Proceedings of the 2022 European Symposium on*

- Usable Security* (Karlsruhe, Germany) (*EuroUSEC '22*). Association for Computing Machinery, New York, NY, USA, 64–78. <https://doi.org/10.1145/3549015.3554213>
- [49] Eileen Pan, Jingjing Ren, Martina Lindorfer, Christo Wilson, and David Choffnes. 2018. Panoptispy: Characterizing Audio and Video Exfiltration from Android Applications. In *Proc. of PETS*.
- [50] Joel Reardon, Álvaro Feal, Primal Wijesekera, Amit Elazari Bar On, Narseo Vallina-Rodríguez, and Serge Egelman. 2019. 50 ways to leak your data: An exploration of apps' circumvention of the android permissions system. In *28th USENIX security symposium (USENIX security 19)*. 603–620.
- [51] Jingjing Ren, Martina Lindorfer, Daniel J. Dubois, Ashwin Rao, David R. Choffnes, and Narseo Vallina-Rodríguez. 2018. Bug Fixes, Improvements, ... and Privacy Leaks – A Longitudinal Study of PII Leaks Across Android App Versions. In *Proc. of NDSS*.
- [52] Jingjing Ren, Ashwin Rao, Martina Lindorfer, Arnaud Legout, and David Choffnes. 2016. ReCon: Revealing and Controlling Privacy Leaks in Mobile Network Traffic. In *Proc. of MobiSys*.
- [53] Jorge Eduardo Rivadeneira, Jorge Sá Silva, Ricardo Colomo-Palacios, André Rodrigues, and Fernando Boavida. 2023. User-centric privacy preserving models for a new era of the Internet of Things. *Journal of Network and Computer Applications* 217 (2023), 103695. <https://doi.org/10.1016/j.jnca.2023.103695>
- [54] Hamza Shaban. 2018. Amazon Alexa user receives 1,700 audio recordings of a stranger through 'human error'. <https://www.washingtonpost.com/technology/2018/12/20/amazon-alexa-user-receives-audio-recordings-stranger-through-human-error/>.
- [55] Rita Singh. 2019. *Profiling humans from their voice*. Springer.
- [56] State of California. 2018. California Consumer Privacy Act (CCPA). <https://oag.ca.gov/privacy/ccpa>. (Accessed on 02/29/2024).
- [57] Statista. 2022. Most used voice assistants in the United States in 2022, by device. <https://www.statista.com/statistics/1368148/voice-assistants-use-by-device-united-states>. (Accessed on 02/29/2024).
- [58] Statista. 2022. Voice assistant users in the United States from 2022 to 2026. <https://www.statista.com/statistics/1384575/voice-assistant-users-united-states/>. (Accessed on 02/29/2024).
- [59] StatCounter Global Stats. 2024. Desktop Browser Market Share Worldwide. <https://gs.statcounter.com/browser-market-share/desktop/worldwide> [Accessed: 2024-11-12].
- [60] Madiha Tabassum and Heather Lipford. 2023. Exploring privacy implications of awareness and control mechanisms in smart home devices. *Proceedings on Privacy Enhancing Technologies* 1 (2023), 571–588.
- [61] Imdad Ullah, Roksana Boreli, and Salil S Kanhere. 2023. Privacy in targeted advertising on mobile devices: a survey. *International Journal of Information Security* 22, 3 (2023), 647–678.
- [62] Qi Wang, Wajih Ul Hassan, Adam Bates, and Carl Gunter. 2018. Fear and Logging in the Internet of Things. In *Proceedings of The Network and Distributed System Security Symposium (NDSS 2017)*.
- [63] Chen Yan, Xiaoyu Ji, Kai Wang, Qinrong Jiang, Zizhi Jin, and Wenyuan Xu. 2022. A survey on voice assistant security: Attacks and countermeasures. *Comput. Surveys* 55, 4 (2022), 1–36.
- [64] Jeffrey Young, Song Liao, Long Cheng, Hongxin Hu, and Huixing Deng. 2022. SkillDetective: Automated Policy-Violation Detection of Voice Assistant Applications in the Wild. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, 1113–1130.

A Query Generation Prompts

All queries we used are available at <https://github.com/NEU-SNS/voice-assistants-pets25.git>.

A.1 Google Query Generation Prompt

A.1.1 Relationships Sample Query Generator Prompt. Design 20 queries to ask a Google Voice Assistant that clearly identify me as a person that is single (Not married and not in a relationship). Ensure the questions reflect my specific relationship status accurately and distinctly. the queries should be simple and short, with the use of related key words.

A.1.2 Household Income Sample Query Generator Prompt. Design 20 queries to ask a Google Voice Assistant that clearly identify me as a person with a high income (not moderately high income and not average or lower income). Ensure the questions reflect my specific household income status accurately and distinctly. the

queries should be simple and short, with the use of related key words.

A.1.3 Education Sample Query Generator Prompt. Design 20 queries to ask a Google Voice Assistant that clearly identify me as a person with a high school diploma. These queries should not suggest any level of education higher or lower than a high school diploma. Ensure the questions reflect my specific educational background accurately and distinctly. the queries should be simple and short, with the use of related key words.

A.1.4 Homeownership Sample Query Generator Prompt. Design 20 queries to ask a Google Voice Assistant that clearly identify me as a person that is renter (not homeowner). Ensure the questions reflect my specific homeownership status accurately and distinctly. the queries should be simple and short, with the use of related key words.

A.1.5 Parenting Sample Query Generator Prompt. Design 20 queries to ask a Google Voice Assistant that clearly identify me as a person that is parent of infants (not parent of toddlers, preschoolers, grade schoolers, or teenagers). Ensure the questions reflect my specific parenting status accurately and distinctly. the queries should be simple and short, with the use of related key words.

A.2 Amazon Query Generation Prompt

A.2.1 Command Query Generator Prompt. Imagine you are a person interested in [specific product category] products. Generate 30 simple and specific command queries for an Amazon Voice Assistant to show your interest. all queries must be either adding a product to cart or adding a product to shopping list. do not ask general queries. do not have any subscribe in your queries. we do not have any specific cart we have only “add to my cart” but we have add to my x list (it means for lists you can have names but for cart we have only add to cart command.

A.2.2 General Query Generator Prompt. Imagine you are a person interested in [specific product category] products. Generate 30 simple and specific queries for an Amazon Voice Assistant that express your interest in this category. Ensure the queries are general in nature, avoiding any command-like phrasing such as 'Add to cart' or 'Add to list.' Focus on inquiries that reflect curiosity, desire for information, or exploration of related products.

B Details of Preliminary Experiments

The details of preliminary experiments for the four research questions have been mentioned hereunder:

B.1 Prepopulated Labels in Profiles

In Spring 2023, we found no pre-populated labels in 16 new Google accounts, with all labels showing “Not enough info.” However, in Fall 2023, many new Google accounts were created with pre-populated labels, notably “Homeowners” and “Not Parents.” The results are shown in Table 7.

During all our experiments with fresh accounts in Alexa and Siri, we have not seen prepopulated profiling labels.

Category	Label	Count/30
Relationship	Not Enough Info	16
	Single	8
	Married	4
	In a Relationship	2
Education	Not enough info	24
	Bachelor’s degree	3
	High school diploma	3
Homeownership	Not enough info	7
	Homeowners	23
Parenting	Not enough info	4
	Not Parents	26

Table 7: Distribution of demographic tags assigned to new Google accounts during an initial profiling phase in Fall 2023.

B.2 Characterization of Profiling Activities

B.2.1 Summary of Persona Experiments. The study conducted a series of experiments to evaluate Google’s profiling accuracy using both simple and complex personas. The personas were designed as follows:

- **Persona 1:** “Married,” “Bachelor’s Degree.” Targeted the relationship and education demographic labels.
- **Persona 2:** “Single,” “High School Diploma.” Also targeted relationship and education labels.
- **Persona 3:** “Married,” “Homeowner,” “Parent,” “Bachelor’s Degree.” This complex persona targeted four demographic labels: relationship, homeownership, parenting, and education.
- **Persona 4:** “Single,” “Renter,” “Not a Parent,” “High School Diploma.” This persona similarly targeted the same four demographic labels.

Simple Persona Experiments: In both Spring and Fall 2023, we tested the accuracy of Google’s profiling using Personas 1 and 2. The results, as shown in Table 8, were consistent across both periods for the education label, but there were differences in profiling accuracy for the relationship label. The experiments in both periods show the evidence of profiling based on user interaction with Google assistants.

Complex Persona Experiments: In Fall 2023, we introduced complex personas (Personas 3 and 4) to explore how Google handles more detailed profiling. The results showed generally lower accuracy for relationship and education labels compared to simple personas. For the homeownership and parenting labels, accuracy varied significantly—one persona had nearly perfect accuracy, while the other showed none.

Accuracy Experiments with Repeated Training: To understand if the low accuracy in the complex persona experiments was due to insufficient training, we conducted additional experiments with repeated queries:

- **Short-Term Repetition:** This involved repeating all training questions for each persona five times daily over one week. The results were mixed—Persona 4 showed an increase in relationship accuracy (from 30% to 75%), but Persona 3’s relationship accuracy dropped to 0%. Education accuracy for Persona 3 improved slightly (from 0% to 25%), while Persona 4’s education

Time	Persona	Relationship	Education
Spring 2023	Persona 1	50%	62.5%
	Persona 2	62.5%	100%
Fall 2023	Persona 1	0% (no tags)	66.67%
	Persona 2	100%	100%

Table 8: Comparison of accuracy metrics for Google’s Simple Persona experiments in Spring and Fall 2023. The accounts were freshly created for each experiment. “no tags” indicates the absence of tags, with all profile labels concluding as “Not Enough Info.”

accuracy decreased to 0%.

- **Long-Term Repetition:** Conducted in January 2024, this experiment used the same Google accounts from Fall 2023. The accounts were re-exposed to the complex persona training questions after several months. This long-term repetition led to a significant improvement, with nearly all labels achieving 100% accuracy.

The results, as shown in Table 9, indicate that while Google’s profiling accuracy may initially be low for complex personas, repeated interactions, especially over a longer period, can substantially improve the accuracy of demographic labels.

To analyze Alexa profiling accuracy we created two personas targeting the Video Entertainment and Fashion interest labels. We conducted our initial Amazon experiments during Summer 2024, utilizing a total of 12 new Amazon accounts, allocating six to each persona. The findings of these experiments are presented in Table 10, showing similar results as depicted in Section 5.3 showing that general questions do not result in profiling, while command-based queries result in 100% profiling accuracy.

Experiments with simple personas linked to Google accounts also showed no indirect profiling by Google through Siri voice searches, suggesting Apple does not share user identifiers that lead to profiling on Google’s end.

B.3 Impact of Interaction Modality

To evaluate the impact of different interaction modalities on user profiling, experiments were conducted across various modalities for Google and Alexa. These experiments compared the profiling accuracy and speed of voice versus web interactions using simple personas.

For Google, results summarized in Table 11, showed that web interactions were more effective for profiling, achieving 100% accuracy for both two-label personas, while voice interactions had lower accuracy (56.25% for Persona 1 and 81.25% for Persona 2). Additionally, web interactions resulted in relatively faster tag changes (around 3.5 days) compared to voice interactions (around 4.9 days).

We studied the impact of voice and web interaction modalities on Amazon’s user profiling, focusing on “Fashion” and “Video” Personas. For “Fashion,” both modalities were tested, while for “Video,” only voice commands were used due to the lack of a web equivalent. Results, summarized in Table 11, showed that web interactions achieved 100% accuracy and relatively similar label changes compared to voice interactions. We were unable to investigate how different modalities affected Siri profiling due to not offering the search through a web interface.

Complex Persona Experiments	Persona	Relationship	Education	Homeownership	Parenting
Fall 2023 Initial New Accounts	Persona 3	20%	0% (no tags)	90%	0% (inaccurate)
	Persona 4	30%	30%	0% (inaccurate)	100%
Repeating 5x Daily New Accounts	Persona 3	0% (no tags and inaccurate)	25%	25%	0% (inaccurate)
	Persona 4	75%	0% (no tags)	0% (inaccurate)	100%
January 2024 Fall Initial Accounts	Persona 3	100%	100%	100%	0% (inaccurate)
	Persona 4	100%	100%	100%	100%

Table 9: Accuracy of label assignment for Google Complex Persona experiments across three scenarios (in percentage): initial assessment on fresh accounts (Fall 2023), repeated query exposure (5x daily for one week), and a follow-up evaluation (January 2024) on initially queried accounts. “no tags” indicates all labels were categorized as “Not Enough Info”, while “inaccurate” denotes incorrect label assignment.

Query	Persona	Avg Acc.	Mean Time (Days)
General	Fashion	0	N/A
	Video	0	N/A
Command	Fashion	100% ± 0%	4.67 ± 0.58
	Video	100% ± 0%	5.33 ± 0.58

Table 10: Comparison of accuracy metrics for Alexa’s general and command queries across Fashion and Video Personas in Summer 2023 and January 2024.

VA	Mode	Persona	Accuracy	Time (Days)
Google	Voice	Persona 1	56.25% ± 8.84%	4.915 ± 0.83
		Persona 2	81.25% ± 17.68%	4.895 ± 0.39
	Web	Persona 1	100% ± 0%	3.5 ± 0.71
		Persona 2	100% ± 0%	3.435 ± 0.27
Alexa	Voice	Fashion	100 ± 0	4.67 ± 0.58
	Web	Fashion	100% ± 0%	3.5 ± 0.71

Table 11: Comparison of mean profiling accuracy and the mean time required for accurate label convergence across various modality types for Google and Alexa voice assistants.