

Using Behavior Monitoring to Identify Privacy Concerns in Smarthome Environments

Atheer Almogbil*, Momo Steele*, Sofia Belikovetsky*, Adil Inam†, Olivia Wu*, Aviel Rubin*, Adam Bates†,

*Johns Hopkins University

†University of Illinois at Urbana-Champaign

Abstract—The rise in the adoption of Internet of Things (IoT) has led to a surge in information generation and collection. Many IoT devices systematically collect sensitive data pertaining to users' personal lives such as user activity, location, and communication. Prior works have focused on uncovering user privacy and profiling concerns in the context of one or two specific devices and threat models. However, user profiling concerns within a complete smart home ecosystem, under various threat models, have not been explored. In this work, we aim to analyze the privacy and user-profiling concerns in smart home environments under varying levels of threat models. We contribute an analysis of various IoT attacks existing in literature that enable an adversary to access data on IoT devices. Based on this analysis, we identify user behavior based on data accessed by such attacks. Our work reveals the extent to which an adversary can monitor user behavior based on information collected from smart households under varying threat models.

I. INTRODUCTION

The Internet of Things is rapidly expanding. Globally, it is estimated that there will be around 50 billion IoT devices by 2030 [74]. This increase can be attributed to the declining costs of IoT devices, making them available to a wider range of consumers. Most devices have been made 'smart' to benefit from remote control and automation features. Appliances [53], [70], [72], [73], security systems [4], [38], [61], [69], wearables [9], [40], inventory trackers [10], [62], [67], and farming [44] and factory equipment [8], [64] have all been developed into smart devices. However, to achieve each IoT device's full potential, the devices are meant to connect to each other to create an IoT ecosystem. This ecosystem allows data to be generated, shared, and stored across multiple devices.

The rise in IoT adoption consequently has led to a surge in generated information. IoT devices have become increasingly intertwined in consumers' everyday life, becoming somewhat of a necessity. This entanglement has provided devices with unrestricted access to consumers' personal lives, systematically tracking and recording user activity, location, and communication with or without the user's knowledge. The data itself might not be sensitive. However, if extracted from multiple devices and correlated, the data can be used

to uncover sensitive information about the user and others within the same household. While IoT devices collect and share sensitive information, the security measures in place often fail to adequately reflect the crucial importance and value of safeguarding such private data. The collection and exchange of private information, as well as the ability to gain access to an entire IoT ecosystem such as a smart home setting, has made IoT users the target of malicious attacks. Nevertheless, IoT security is commonly found to be an afterthought during development [77]. Due to this, IoT devices are susceptible to numerous attacks.

A variety of studies [48], [58], [76], [86], [87] are concerned with end-user security and privacy issues associated with IoT devices. [43] defines user profiling as "the process of collecting information about a user in order to construct their profile." In this paper, we refer to this definition when using the terms 'monitoring' or 'profiling.' Most research has studied privacy concerns in the context of specific devices such as Amazon Alexa [17], [45], [63], [85], Google Assistant [5], [22], [84], and Smart TVs [12], [54]. Another focus has been on specific access levels such as network access [1], [80], [88] and physical access [17], [85]. However, user behavior monitoring concerns and implications in the context of a complete smart home ecosystem for varying proximity levels have not been explored.

A. Contributions

We make the following contributions:

- We highlight various types of IoT attacks that enable an adversary to monitor user behavior in different settings. (Section II)
- We reveal three threat models based on common characteristics identified for each attack surface. (Section II)
- We set up four settings to analyze the data collected by IoT devices, their companion applications, and their web interfaces at various attack, skill, and access levels. (Sections 3 and 4)
- We identify the importance of data correlation and the implications of behavior monitoring. (Section V)

II. THREAT MODELS

In this section, we discuss our approach for identifying the threat models based on an analysis of IoT attacks in existing literature.

A. Identifying Attacks

To better understand the threat models and attack surface for this research, we identify and classify IoT attacks into three categories of varying levels of proximity: physical, nearby, and remote (shown in Table I). This analysis allows us to gain a better understanding of the type of data an adversary can access at varying proximity levels, enabling us to identify what can be learned from behavior monitoring.

In addition to refereed publications, we considered attacks that were mentioned in blog posts, videos, newspaper articles, and news channels as these sources contained descriptions that were not found in the academic literature. After identifying a collection of IoT attacks, we then filtered the attacks based on whether user information can be accessed or not. This was done using qualitative coding, where literature with terms such as 'obtain', 'retrieve', 'access', 'steal', and so on were assigned a label to indicate information access. As our focus is to understand what an adversary can learn about smart home users through their IoT devices, attacks that did not involve information gain or access were excluded.

Categorizations After identifying IoT attacks relevant to our study, we then created a list of categories or characteristics that can best be used to describe the attacks. Other categories may exist, but for the purpose of our paper, the attacks selected and the information available regarding these attacks guided the selection of categories. The general classification of attacks is determined by the attack surface. First, each researcher independently classified attacks based on the proximity to the IoT device: physical, nearby, or remote. Second, each researcher independently read the same literature, filled in a table similar to Table I for each IoT attack, and independently came up with a value for each category. After conducting these steps, each table was presented and discussed until a unanimous decision was made for each IoT attack. The final set of categories that were revealed through this process include: the level of access an adversary has to the device, the adversary's skills and motivation, the target of an attack, and the data that can be retrieved or accessed from an attack. We consider these factors when discussing the information an adversary can learn about user behavior within each threat model.

Type of Access The type of access is categorized into two parts: time frame and proximity. The time frame refers to the period of time required to carry out an attack in terms of whether the attack can be conducted by the adversary through one-time access or continuous access to the device, application, or network. Proximity refers to the location of the adversary when conducting the attack. If an adversary needs physical access to the device, this is labeled as physical access. If the adversary requires close proximity to the target, the attack is then labeled as nearby. Finally, if the adversary has the ability to conduct an attack from outside of the smart home, without requiring physical or nearby access to the device, then the type of access is labeled as remote access.

Type of Adversary An adversary is categorized based on

two factors: motivation and skill level. The motivation of an adversary is considered to be nonmalicious if the adversary unintentionally violates a user's privacy, caused an IoT device or application to leak private information, or intentionally accessed private data out of curiosity without any intention to do harm. A malicious adversary is one that intentionally accesses a user's private information with the motivation to misuse it or do harm (i.e., escalate privilege). An adversary's skill level is categorized into three levels: basic, intermediate, and advanced. An adversary with basic skills might have little to no technical background and is capable of interacting with an IoT device or its companion application. An adversary with intermediate skills has experience with exploiting vulnerabilities of IoT devices and applications using commonly known techniques such as network sniffing, cross-site scripting, social engineering, port scanning, and other methods to gain unauthorized access to an IoT device, web interface, or companion application. An advanced adversary is someone that can reverse engineer or forensically analyze IoT devices, applications, etc.

Target The target of an attack is classified into three categories: device, companion application, and web interface. An IoT device commonly provides users with an application (i.e., mobile and/or desktop) or web interface to interact with the device, configure settings and access control policies, create user accounts, and monitor device activity. The IoT device alone can be vulnerable and insecure; however, the addition of user-friendly UI to an IoT device through applications and web interfaces introduces an entire set of different vulnerabilities. An adversary is not limited to attacks on the device itself but is now able to gain access to other devices on the same network, such as a mobile phone, through exploiting a vulnerability available in one or more of the targets.

Data Obtained Data is categorized based on what can be learned or inferred about the user. Personal data involves a user's schedule, email, credentials, pictures, address, or any other personally identifiable data excluding any financial information. Financial data is classified as personally identifiable financial information such as a user's billing address, purchase history, and banking information. Device data refers to any data that gives insight on the type of device being used, such as the model number, device ID, operating system, protocols, settings, location, and configurations. Activity data consists of user or device activity logs, device behavior, triggers, automation rules, and interactions with other devices. Finally, recordings from IoT devices are a category on their own as user voice recordings can include anything from public information (i.e., weather) to moderately private information (i.e., background noise, events, music taste) to then extremely private information (i.e. banking information and passwords).

B. Direct Physical Access Attacks

Attacks in this category include those that require the adversary to have physical access to the device. This includes being able to download the companion application to pair and interact with the device or reverse engineer the device to obtain

TABLE I
IOT ATTACKS BASED ON PROXIMITY LEVEL.

Proximity	Reference	Access Timeframe		Type of Adversary				Target			Data			
		One time	Cont.	Motivation		Skill Level		Device	App	Web	Personal	Device	Activity	Financial
				Mal.	Non-mal.	Basic	Inter							
Physical	Nardi20 [59]	✓		✓		✓		✓				✓		
	WEWSTV19 [83]	✓		✓			✓	✓	✓	✓	✓		✓	
	Enev11 [21]		✓	✓			✓	✓			✓	✓		
	Akinbi20 [5]	✓			✓			✓	✓		✓	✓		✓
	Boztas15 [12]	✓			✓			✓	✓		✓	✓		✓
	Li19 [49]	✓			✓			✓	✓	✓	✓	✓		✓
	Exploitee.rs16a [24]	✓		✓				✓	✓					
	exploitee.rs17a [31]	✓		✓				✓	✓					
	exploitee.rs15 [23]	✓		✓				✓	✓					
	Exploitee.rs16b [25]	✓		✓			✓	✓						
	exploitee.rs17b [32]	✓		✓			✓	✓						
	Exploitee.rs16c [26]	✓		✓			✓	✓						
	Tierney19 [78]	✓		✓			✓	✓						
Nearby	Tiley16 [79]	✓			✓	✓		✓			✓	✓		
	Protego19 [47]	✓		✓			✓	✓	✓		✓			✓
	Servida19a [68]	✓		✓			✓	✓	✓	✓	✓			
	Servida19b [68]		✓	✓			✓	✓						
	Felch20 [39]	✓		✓				✓	✓	✓	✓	✓		
	Venda18 [82]	✓	✓			✓		✓		✓	✓	✓	✓	✓
	Munro21 [57]	✓	✓		✓			✓			✓			
	Lodge18 [50]	✓		✓			✓			✓	✓	✓	✓	
	Exploitee.rs16d [27]	✓		✓			✓		✓		✓	✓		
	exploitee.rs17c [33]	✓		✓			✓			✓	✓	✓		
Remote	Lomas20 [51]	✓		✓				✓	✓			✓	✓	
	Paul18 [66]	✓		✓			✓			✓	✓	✓		✓
	Bowles18 [11]	✓		✓			✓	✓	✓	✓	✓	✓		
	Barda20 [20]	✓		✓			✓	✓		✓	✓	✓	✓	✓
	Campbell19 [16]		✓	✓			✓	✓						
	Palmer20 [65]	✓		✓			✓	✓			✓			
	Puttaraju16 [52]	✓		✓			✓	✓			✓	✓		
	Albrecht15 [6]	✓		✓			✓	✓			✓			
	Neagle15 [60]	✓		✓			✓	✓			✓	✓		✓
	Gelinias19 [41]	✓		✓			✓	✓	✓	✓	✓	✓		
	Adams20 [2]	✓		✓				✓			✓	✓		
	Brewster16 [13]	✓		✓			✓	✓	✓	✓	✓	✓		
	Munro22 [56]	✓	✓		✓			✓	✓					✓
	Stykas20 [75]	✓	✓	✓				✓	✓	✓	✓	✓	✓	
	Monie18 [55]	✓	✓	✓			✓	✓	✓	✓	✓	✓		
	exploitee.rs17d [34]	✓		✓			✓	✓			✓			
	Exploitee.rs16e [28]	✓		✓				✓			✓	✓		
	exploitee.rs17e [35]	✓		✓				✓	✓	✓	✓	✓		
	exploitee.rs17f [36]	✓		✓			✓			✓	✓	✓		
	Exploitee.rs16f [29]		✓	✓				✓	✓			✓	✓	
	Exploitee.rs16g [30]	✓						✓	✓	✓	✓	✓		
	exploitee.rs17g [37]	✓		✓				✓	✓	✓	✓	✓		✓

information. 92% of IoT attacks that require physical access to the device can take place through one-time access and do not require continuous access to the device. [21] requires continuous access as the attack measures the stability of power supplies' electromagnetic interference (EMI) signatures over time. Furthermore, an adversary that has physical access to an IoT device is usually always malicious. This is true except for the case in which the adversary, such as a legal authority, is obliged by law to obtain personal information from the devices. In terms of skills, an adversary at this level ranges from having basic skills such as popping open a doorbell [59] or drilling a smart lock [78], intermediate skills such as reading EMI signatures [21] and using password dictionaries [83],

to advanced skills including reverse engineering and forensic analysis [5], [12], [49], as well as kernel hijacking [24]. The target of the attack is the device itself 77% of the time, whereas 8% of attacks target both the device and mobile application. Only one out of 13 of the attacks targets the device, mobile application, and web interface. Often, the IoT device is susceptible to both software and hardware vulnerabilities that can easily be identified and exploited, making it a greater target to adversaries that have physical access to the device.

In terms of the type of data that can be obtained at this proximity, a majority of the attacks identified are able to access device data, as it is normally stored on the device itself. [12] investigates the Samsung smart television through flash

memory and the Samsung Development Kit application. User information such as pictures, connected devices, and visited websites can be easily recovered. These studies demonstrate that digital traces are available across various types of devices and platforms. However, an analysis of information learned from the data retrieved is not mentioned in such studies. Depending on the type of device, other types of data can be accessed. For instance, all but one attack, [59], access activity data. Four out of 13 attacks access personal data, and none access financial data. Moreover, a subset of voice recordings can be found on the flash memory of devices such as the Amazon Echo Dot and Google Home [5], [49], as well as smart TVs [12].

A common theme across attacks at this level of proximity is the use of default root credentials [24] or using a UART adapter to access root privileges [25], [26], [31], [32]. This facilitates the ability to access the root user by performing a quick Internet search for the brand and model of the IoT device. However, it is important to note that most forensic analysis of IoT devices published in literature focus on smart speakers and TVs, with little to no investigative studies of other types of IoT devices.

C. Nearby Attacks

This set of attacks can be conducted without the need for the device to be physically present and in direct reach of the adversary. The attacks may use WiFi or Bluetooth connections to connect to the device from outside of the house. In other words, this category includes any attack that can reach over Bluetooth or WiFi. Nine out of 11 attacks conducted within this range require one-time access only, with a majority of these attacks having a malicious goal. 27% of attacks require continuous access, such as eavesdropping by connecting to a children's toy phone over Bluetooth or monitoring a user's Google calendar through their smart fridge [82]. 64% of the nearby attacks require intermediate skills to carry out a SQL injection [47], Man-in-the-Middle (MITM) attacks [82], monitoring traffic [50], and creating a reverse SSH tunnel [27]. Advanced skills such as reverse engineering devices to find default keys and exploiting vulnerabilities to access root when on the same LAN are required to carry out two nearby attacks [39], [51]. [79] requires minimal skill, i.e., the knowledge of operating a smart speaker by using a voice command, to unlock the front door from outside of the household or request emails to be read out loud. The target of most nearby attacks are the device itself, in which Bluetooth and WiFi connections are used to interact with the device. Based on the number of attacks that exploit these interfaces, it is clear that default WiFi credentials and the placement of IoT devices near windows and doors play a major role in the facilitation of such attacks [27], [79]. Additionally, the attacks in this category target the vulnerabilities in companion applications. Often, user-friendliness outweighs security and privacy when designing these web and mobile applications. This leads to vulnerabilities such as unauthenticated voice commands [79], lack of input validation, and sending keys in plain-text [50].

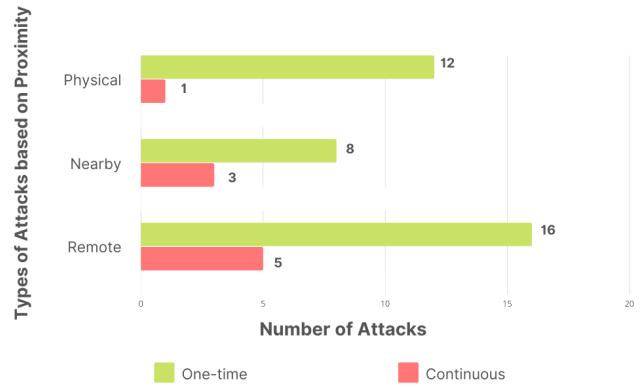


Fig. 1. IoT attacks based on proximity and timeframe.

Unlike physical attacks, a majority of nearby attacks are able to retrieve personal information. Financial data [47], as well as activity and device data [50], [51], [68], can also be obtained. However, unlike physical attacks, no recordings were able to be accessed.

D. Remote Attacks

The third and final attack category can be conducted remotely over the Internet from any distance. The key difference between nearby and remote attacks is that nearby attacks use directly wireless access, whereas remote attacks come through the router over the Internet. A majority of the attacks fall in the remote proximity category making up 38% of the attacks identified. As shown in Figure 1, 15 out of 21 remote attacks are conducted with one-time access, indicating that only one successful attempt can lead to major security and privacy implications. Only one attack is considered to be non-malicious, as Amazon employees are obligated to access voice recordings for training and speech recognition improvement purposes [16]. Aside from this exception, all remote attacks are deemed to be of malicious nature. Furthermore, most remote attacks require intermediate skills such as Cross-Origin Resource Sharing (CORS) misconfiguration and Cross Site Scripting (XSS) [20], traffic monitoring [65], brute force credentials [2], [6], [41], MITM [60], and the use of hacking tools such as AirCrack and WiFite [52]. Other remote attacks are commonly carried out by abusive partners and require little to no effort or technical skill [11], [66]. Unfortunately, it is evident that most IoT companion applications are misused to abuse or terrorize victims in an abusive relationship, even after the relationship has ended. Further exploration of this matter is presented in greater detail in Chapter Four. 86% of remote attacks target the device itself, whereas only 14% target mobile applications, and 43% target web applications. Similar to nearby attacks, remote attacks can feasibly access various types of data as more than half of the attacks retrieve personal, device, and activity data. This underscores the significant costs

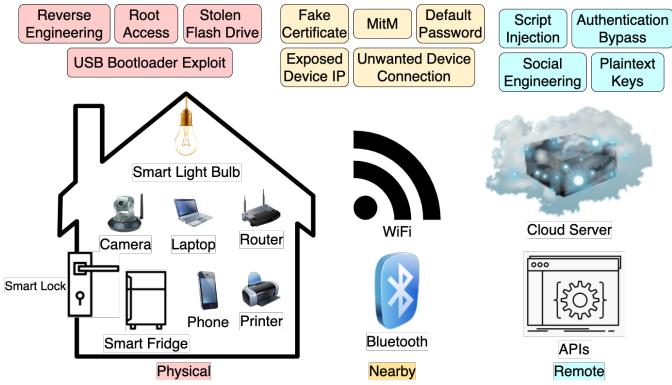


Fig. 2. IoT attack vectors at varying levels of proximity.

associated with such attacks. With online videos and websites providing step-by-step instructions on executing these attacks [81], IoT users face an elevated level of risk.

E. Threat Models based on Proximity

After thorough investigation of the types of attacks that are possible and their shared characteristics, we were able to identify three possible threat models based on attack proximity. Other threat models can be identified, however we limit our scope to those presented shown in Figure 2

The first adversary investigated is under the physical proximity category, P. The adversary in threat model P has direct, physical access to the device. As the adversary would have to physically be within the household, the IoT devices can easily be interacted with. Some of the devices an adversary would have access to are logged in to, or the adversary has been granted access on the device from previous interactions. Potential examples of this type of adversary would be a neighbor, housekeeper, babysitter, or a malicious house guest. All of these example adversaries would have been given access to the devices at one point previously and could now snoop around the device accounts and companion applications to learn personal information about the owner of the device. Additionally, a person who is not authorized to access the household (i.e., burglar or unexpected guest), authorities (i.e., forensic investigators), and second-hand owners of IoT devices are considered in this category. Based on our analysis of existing literature, an adversary in this model is restricted to one-time, physical access to the IoT device(s) and has intermediate to advanced skills. This adversary targets the device itself and aims to retrieve personal, device, and activity data.

The second adversary investigated falls under the nearby proximity category, N. An adversary in threat model N does not have physical access to the device, but is close enough to connect to the device through the LAN or the Bluetooth interface. Similar to threat model P, this adversary is malicious and is restricted to one time access to the device(s). However, in this case, the adversary does not have physical access, but is close enough to connect to the device. Similarly, the target of the attack is the device itself and the adversary is able to

retrieve the same types of data as the previous threat model. However, the adversary is limited to intermediate skills only.

The third adversary is one that remotely attacks IoT devices within a smart home, R. This type of adversary is the most common, as it is apparent from our analysis. This adversary in threat model R is characterized as being malicious and limited to one-time access. This is a common theme across all three threat models. In this threat model, the adversary has intermediate skills and targets personal, device, activity, and recording data on both the device and web interface.

III. METHODOLOGY

After systematically analyzing existing IoT attacks, we interacted with IoT devices to observe the extent of data that can be accessed and the information inferred. Our approach can be largely divided into three steps. We go over each in detail.

A. Data Sources

To evaluate how much an adversary can learn about members of a smarthome, we set up a pilot study to investigate the data collection and inference process of IoT devices in realistic environments. Each environment differed in the number of members, residence type (i.e., single, townhouse, apartment), location, demographic, and number and type of IoT devices. The environment settings ranged from only using one type of IoT device to using multiple, varied IoT devices. This variation was implemented to compare the information that can be obtained when using different types and numbers of devices. In a study of 2000 U.S. households, it was found that 56% own one device, 18% own two devices, and 26% own three or more devices. Moreover, more than 50% of multiple device owners were more inclined to buy a hub to interact and connect with the devices [3]. The most popular categories of IoT devices owned were found to be security and safety (i.e., security cameras and doorbells) and utility management (i.e. thermostats and light bulbs). The type and number of IoT devices in each setting were selected to be representative of such findings and are as follows:

- Setting 1: Two Amazon Echo Dots
- Setting 2: Google Home, Philips Hue Smart Light Bulbs, Nest Thermostat
- Setting 3: Ring Doorbell, Four Nest Thermostats, Google Home, Three Arneti cameras
- Setting 4: Amazon Echo, ADT home security system, Fios router

We familiarized ourselves with the functionality and features offered by each IoT device to understand the devices' capabilities and to learn what information is likely to be collected. Then, we used the devices for 30 days and incorporate them into our daily routines. The selection of a 30-day time frame is based on the average duration during which a device retains data before it is automatically erased.

B. Data Collection Setup

To understand how to collect data from different IoT devices, we initially experimented with a range of IoT devices. The purpose of this exploration was to gather information about device features, as well as data population, storage, and collection. The lab used for experimentation included a diverse range of IoT devices, including speakers, hubs, switches, light bulbs, IP cameras, streaming devices, plugs, thermostats, and doorbells. The information was collected through an Android mobile emulation, Android Studio, allowing access to the companion applications of each IoT device. The web interface of an IoT device, if applicable, was also accessed to export data logs. This initial experimentation revealed that most IoT devices retain activity logs, as well as collect personal and device data. Results are shown in Table II. This information guided the collection and analysis of data obtained from IoT devices in the four settings mentioned previously.

C. Data Analysis Methods

After the 30 day time frame ended, the data collected by the devices, their companion applications, and their web interfaces was then analyzed at various attack, skill, and proximity levels based on the three threat models revealed in Section II-E. The analysis' primary goal was to find all direct inferences that can be drawn from the data. The secondary aim was to find any high level inferences that can be indirectly drawn from the data collected by IoT devices. To reach a consensus on both direct and high-level inferences, each author independently analyzed the findings from all four settings and noted direct and indirect inferences. Then, the findings were presented and discussed until a unanimous decision was made.

IV. FINDINGS

In this section we discuss the results of experimentation. The findings are categorized into two groups: direct inferences and high level inferences. The former refers to knowledge gained by the adversary without making any assumptions. The latter requires assumptions, which may or may not be true. Both categories increase the knowledge an adversary has on user behavior within each settings under different threat models.

A. Direct Inferences Drawn from Data

After analyzing the data from the IoT devices, their companion applications, and web interfaces, the following direct inferences were drawn. It is noteworthy that a significant portion of the collected information is classified as personal data.

Users. Users that interact with a device were identified through a list of connected devices, account usernames, and/or emails associated with the devices. Voice-activated assistants retained logged voice recordings that can be used by an adversary to identify or enumerate users.

Name. Names of users that interacted with the devices were easily found in logs. It is noted that the primary user's name is repetitively found in logs, in comparison to other users'. For settings with a smart assistant, the name of the device

owner can also be revealed through a simple query "What's my name?" or "Who Am I?". Moreover, since the users had their email accounts linked with the smart assistants, requesting these devices to read out the emails can also reveal the owner's name. This is due to the method in which emails are read, usually starting with greetings followed by the name of the receiver.

Gender. Based on voice recordings, settings 1 and 2 show the presence of higher-pitched and lower-pitched users indicating differing genders. In some instances, we were able to link voice recordings to a user. In setting 1, some voice recordings that belonged to the female user were logged along with the user's name. Moreover, music selection (i.e., music played on smart speakers) can also serve as an indicator for identifying the gender of the users [14], [18]. In setting 3, the primary user's profile picture indicated a male user. Whereas in setting 4, videos from the security cameras indicated the presence of a male and female adult.

Profession and Education. For settings 1 and 2, email content indicated the primary users are affiliated to a private and public university within the U.S., respectively. Data from settings 3 and 4 did not identify the profession or education of users.

Accommodation. In setting 1, the device configuration settings indicated that the accommodation has a living room and a dining room. Based on recordings logged, when a user spoke to one device, the other was also activated but logged the recording as "not intended for this device." This indicates that the two devices are close in proximity, which implies that the living room and dining room are in close proximity. Similarly, for setting 2, the devices were present in the living room, dining room, kitchen, bedroom A and bedroom B. A quick Internet search of the address logged on the devices can be used to identify the type of residence.

Music Taste. Users of settings 1 and 2 listened to Hip-Hop, Pop, and "Top Hits" playlists. This finding aids our analysis of indirect inferences discussed in the next section.

B. High Level Inferences Drawn from Data

After analyzing the data from the IoT devices, their companion applications, and web interfaces, the following high level inferences were drawn. The data retrieved belong to the personal and activity categories. The inferences drawn from the data analysis aim to comprehend how adversaries can potentially tailor their attacks to more effectively target users. This involves understanding the collected information and narrowing down possibilities. Whether accurate or not, these assumptions increase an attacker's likelihood of success by curating targeted attacks. If the assumptions prove false, no harm is done, and the attacker can refine their approach in future attacks by eliminating these assumptions.

Location. Time zone and addresses logged were used to identify the device's location. Distance and temperature units for settings 1, 2, and 3 identified regions in which the devices were located, i.e., setting 1 was outside of the United States,

whereas settings 2 and 3 were within the United States. Furthermore, the temperature settings of smart thermostats used in certain settings can be used to indicate which region devices are located in. For instance, setting 3 showed high heating temperatures during the months of December to February and cooler temperatures from June to August. Moreover, interests in relocating can be inferred from setting 1 as recordings captured queries of apartment rentals and prices in specific U.S. cities.

Relationship. Voice recordings captured queries that can infer the type of relationships that exist between users in each setting. For instance, personal data from setting 1 indicates that there are two users. Queries logged searches for one-bedroom apartments, which may indicate that the users are couple. This can be inferred using published statistics regarding living situations within the region. For instance, 3.4 million U.S. apartment households made up of two occupants are found to be in a romantic relationship, compared to the mere 5% that are roommates [19]. Similar data from setting 2 indicates two bedrooms and two users. This can indicate that the members are most likely not romantically linked. Device data from setting 3 shows a list of user devices logged under first and last names. Users had the same last name, indicating a family dynamic, which makes up 26.5 million U.S. households [19].

Age. The user's age can be estimated based on correlating data from different categories and devices (i.e. emails from social media platforms linked to an account). Based on the 2021 demographic statistics of platforms such as Snapchat and Instagram, an adversary can estimate the primary user's age to be around 13 - 34 years old for primary users in settings 1 and 2 [71]. However, based on the user's education, an adversary can narrow the age range to 18-34 years old. The male user's age in setting 1 can be estimated to be around 20 - 34 based on preferring pop and hip-hop playlists [15] and reminders of COVID-19 vaccination appointments. Similarly, based on the user's education and daily activities in setting 2, the age of the user can be estimated to be around 20-30 years old.

Health. Playlists and reminders set by each user in settings 1 and 2 indicated the type and pace of exercises performed, as well as vaccination and health status. Shopping lists obtained from setting 2's smart assistant offered insight into the diet and health of the users. No such findings were found for settings 3 and 4.

Lifestyle. Logs from settings 1 and 2, such as daily and weekly reminders, give insight on the preference of following a routine. Timestamps of voice recordings as well as other logged activities indicate an estimate of daily sleeping patterns and activity. Activity logs from devices can be correlated to reveal when users were present or away. The male user in setting 2 was away from the device every weekend for 5-6 hours based on Google Takeout data. Access and motion logs taken from the smart doorbell and security cameras from setting 3 indicate high activity during the mornings and afternoons. Daily routines of setting 4 can be learned from the locking and unlocking status logs from the security system. Figure 3 contains a snapshot analysis of the activity routine

in setting 4. The logs were extracted from a security system and reveal information through various triggers of the system. Figure A shows the "Arm Stay" log times. This mode is usually triggered when people are staying indoors and wish to get alerts on doors or windows being opened. The data shows consistency regarding when this mode is triggered in setting 4. It seems that it is used mainly at night, mostly around 10 pm with rare occasions that it is triggered around midnight. This sheds light both on the daily routines and on the rare occasions that might indicate some abnormal activity. Figure A shows motion alerts in a specific locations. The data shows a pattern that indicates a presence at a certain location for several hours throughout the day. Figure A shows logs of the locking and unlocking of a door. This data indicates vast activity throughout the day, suggesting that someone is present throughout the day. This data correlates with the presence activity mentioned previously and can suggest that the user may work remotely.

Political Views. For setting 1, the routines set on the device were requested to play flash briefings and news from People Magazine and CNN. Similarly, for setting 2, the primary user played flash briefings on a daily basis. The type of news channels played could be used to indicate political affiliations. No indications for settings 3 and 4 were found.

V. DISCUSSION

A. Primary Users vs Secondary Users

Results indicate that information retrieved from IoT devices within a smart home better reflect a primary user. A primary user is the user that interacts with the devices the most, sets up the configurations, and has linked their email to the devices [46]. Some indications of other users can be found, but not comparable to that of the primary user. In other words, the collection of data on different users is not balanced. Phone numbers, emails, and usernames were used to identify primary users. We noticed that IoT devices were linked to only one user account for each setting. This can be due to a lack of access control capabilities offered by IoT applications, which is an entire field of research in itself. An abundant amount of data can be found that allows an adversary to easily profile the primary user. It was much more difficult to profile secondary users or users that did not interact with the IoT devices as much. This gives us insight as to why a greater amount of information about the primary user has been identified in comparison to the secondary user.

B. Emails and Recordings

An unlimited amount of data and insight can be obtained by gaining access to users' emails and voice recordings. A vulnerability exploited in any of the IoT devices, their companion applications, or web interface can ultimately allow an adversary to escalate privileges and gain access to a user's email account. It was evident that an adversary even at threat model P and N is able to read users' emails and reply to them using the voice-activated assistant available within the smart home. This alone can tell a lot about users such as social

media platforms used, newsletters subscribed to, academic or work-related progress and updates, sensitive information from health providers, insurance companies, financial institutions, or even romantic relationships. Moreover, under threat model R, attackers can leverage email access to compromise any device or account through the password recovery process. Similarly, in the context of recordings, smart speakers learn users' voices and indicate who the person is by replying with their name. Call logs (i.e., contact and time) and message logs (i.e., contact, content, and time) are found within recordings, only if a user used a voice command to initiate the call or message. Other information included emails (i.e., content and sender), weather, background noise, music playlists played, news briefings, and data from other categories, such as personal, activity, device, and financial. Although Amazon does not explicitly state that it has access to the data in other categories, the company does have access to users' voice recordings to train employees and their natural language processing systems [16]. Access to voice recordings alone can provide companies with great insight on users, as described in Section ??.

C. Financial Data

Financial information was not explicitly found on the devices, applications, or web interfaces, as we chose not to share that type of information with the devices. However, user credentials obtained from device and personal data can be used to access a user's Amazon or Google account and view their purchase history and banking information [7], [42]. Moreover, addresses can be used to infer the standard of living based on property type and cost in that area.

D. Companion Applications

An adversary at threat level R can misuse authorized access to the IoT devices through the use of companion applications, such as Arlo and Nest mobile applications. If a user's access control is not restricted, he or she has full access and control to the devices and the data stored on the applications. For instance, guest users such as a housekeeper or neighbor that were meant to access an IoT devices during a single visit will have continuous, unrestricted access to a live stream of any IP camera set up around the house. This allows an adversary to observe and learn information that violates user privacy, such as when users or guests are active, whether or not a user is present, who is present, and other private videos or pictures of users recorded with or without their consent. An unlimited amount of information can be learned about users through continuous access of companion applications.

E. Correlation of Data

In our analysis, we observed that in the context of user monitoring, the sum is greater than the parts. In other words, correlating data from different sources and devices leads to better inferences in certain cases. The analysis of data from several devices can be utilized to support evidence obtained from individual devices. For instance, by correlating the location data from different sources, the user's location and

whereabouts can be more concretely identified at a bigger scale. Moreover, by identifying important patterns, information can be linked to improving user behavior inferences. For example, an adversary can better estimate when the user enters or leaves a house by correlating the temperature history data from a Nest Thermostat with the recording-based interaction data from an Amazon Echo. Similarly, we observed that the user in setting 2 turned off the Philips smart bulbs and at the same time changed the Nest Thermostat temperature to a night-based setting. This can be used to infer when the user goes to sleep and when they wake up. In setting 4, it was possible to correlate several logs from the security system to learn about the routines of the users. Additionally, it is possible to extract router logs with DHCP logs and correlate them to door activity. Thus, understanding not only when someone is present, but also who is present.

F. Targeted Attacks and Other Implications

Not only does monitoring the behavior of IoT users lead to serious privacy breaches, but also adversaries can further use this data to carry out customized and targeted attacks on device users. For a user interested in politics, adversaries can craft spear-phishing attacks with emails regarding politics or a newsletter the user subscribed to. Similarly, avid smarthome users can be targeted to receive fake promotions about new IoT devices that are tampered with a backdoor or other vulnerabilities that the adversary intends to exploit to leak information. The identification of primary users and the types of devices within a smart homes can allow attackers to further customize spear phishing emails (i.e. software upgrade notifications) to targeted victims regarding their own devices. These details make the emails appear to be more realistic and trusting to the user. There are several other implications such as selling sensitive user information to third parties (e.g, marketing/promotion companies, hacker forums, etc.), which can lead to serious consequences in itself.

VI. LIMITATIONS AND FUTURE DIRECTIONS

While the settings and devices used in our pilot study are representative of common U.S households, they are still limited in number. We aim to extend this study by increasing the number of devices and realistic settings. Additionally, there is a scarcity of literature when it comes to attacks under threat model P. Consequently, our analysis is limited in that regard.

VII. CONCLUSION

In our study, we analyzed the privacy concerns in several smart home environmental settings under a set of threat models. We categorized various IoT attacks in existing literature based on varying levels of proximity that can be conducted to gain access to data on IoT devices and their applications. We analyzed the data collected by IoT devices and their applications to understand how much an adversary can infer about users through their IoT devices.

ACKNOWLEDGMENT

This research results from the SPLICE research program, supported by a collaborative award from the National Science Foundation (NSF) SaTC Frontiers program under award numbers 1955172 and 1955228. The views and conclusions are those of the authors. Any mention of specific companies or products does not imply any endorsement.

REFERENCES

- [1] Abbas Acar, Hossein Fereidooni, Tigist Abera, Amit Kumar Sikder, Markus Miettinen, Hidayet Aksu, Mauro Conti, Ahmad-Reza Sadeghi, and Selcuk Uluagac. Peek-a-boo: I see your smart home activities, even encrypted! In *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pages 207–218, 2020.
- [2] Chase Adams. How I hacked my airbnb neighbor’s smart lock (and you can, too!) [u], Jan 2020.
- [3] Kabir Ahuja and Matt Patel. The connected home market, 2016.
- [4] Alhomesolutions, 2021.
- [5] Alex Akinbi and Thomas Berry. Forensic investigation of google assistant. *SN Computer Science*, 1(5):1–10, 2020.
- [6] Katherine Albrecht and Liz McIntyre. Privacy nightmare: When baby monitors go bad [opinion]. *IEEE Technology and society magazine*, 34(3):14–19, 2015.
- [7] Inc. Amazon. Amazon.com privacy notice, feb 2021. amazon account information.
- [8] Amper, 2021.
- [9] Apple. Apple watch, 2021.
- [10] Axzon, 2021.
- [11] Nellie Bowles. Thermostats, locks and lights: Digital tools of domestic abuse, Jun 2018.
- [12] Abdul Boztas, ARJ Riethoven, and Mark Roeloffs. Smart tv forensics: Digital traces on televisions. *Digital Investigation*, 12:S72–S80, 2015.
- [13] Thomas Brewster. Hackers could have turned vulnerable smart teddy bear into demon toy, 2016.
- [14] Chris Buckman. Music and gender, 2017.
- [15] Published by Statista Research Department and Jan 8. Favorite music genres among consumers by age group in the u.s. 2018, Jan 2021.
- [16] Mikey Campbell. Thousands of amazon workers are listening in on echo audio, report says [u], Apr 2019.
- [17] Hyunji Chung, Junghyun Park, and Sangjin Lee. Digital forensic approaches for amazon alexa ecosystem. *Digital Investigation*, 22:S15–S25, 2017.
- [18] Ann Colley. Young people’s musical taste: Relationship with gender and gender-related traits 1. *Journal of applied social psychology*, 38(8):2039–2055, 2008.
- [19] National MultiFamily Housing Council. Household characteristics, Nov 2020.
- [20] Barda Dikla. Keeping the gate locked on your iot devices: Vulnerabilities found on amazon’s alexa, Aug 2020.
- [21] Miro Enev, Sidhant Gupta, Tadayoshi Kohno, and Shwetak N Patel. Televisions, video privacy, and powerline electromagnetic interference. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 537–550, 2011.
- [22] Steven Engelhardt. Smart speaker forensics, 2019.
- [23] Exploitree.rs15. Hisense android tv, 2015.
- [24] Exploitree.rs16a. Greenwave reality bulbs, 2016.
- [25] Exploitree.rs16b. Epson artisan 700/800, 2016.
- [26] Exploitree.rs16c. Amazon firetv, 2016.
- [27] Exploitree.rs16d. Exploiting nest thermostats, 2016.
- [28] Exploitree.rs16e. Summer baby zoom wifi, 2016.
- [29] Exploitree.rs16f. Netgear ntv200-100nas, 2016.
- [30] Exploitree.rs16g. Wink hub, 2016.
- [31] Exploitree.rs17a. Belkin wemo, 2017.
- [32] Exploitree.rs17b. Google chromecast, 2017.
- [33] Exploitree.rs17c. Samsung allshare cast, 2017.
- [34] Exploitree.rs17d. Zmodo greet, 2017.
- [35] Exploitree.rs17e. Belkin n300, 2017.
- [36] Exploitree.rs17f. Muzo cobblestone, 2017.
- [37] Exploitree.rs17g. Samsung smartcam, 2017.
- [38] Eyelock, 2021.
- [39] Felch. Reverse engineering a smart lock, Aug 2020.
- [40] Fitbit, 2021.
- [41] Gelinas. How to stop your smart home from getting hacked, Nov 2019.
- [42] Inc. Google. Get a summary of data in your google account, Sep 2021. google account information.
- [43] Omar Hasan, Benjamin Habegger, Lionel Brunie, Nadia Bennani, and Ernesto Damiani. A discussion of privacy challenges in user profiling with big data techniques: The eexcess use case. pages 25–30, 06 2013.
- [44] Herddogg, 2021.
- [45] Catherine Jackson and Angela Orebaugh. A study of security and privacy issues associated with the amazon echo. *International Journal of Internet of Things and Cyber-Assurance*, 1(1):91–100, 2018.
- [46] Vinay Koshy, Joon Park, Ti-Chung Cheng, and Karrie Karahalios. “we just use what they give us”: Understanding passenger user perspectives in smart homes. pages 1–14, 05 2021.
- [47] Protego Labs. Alexa sql injection hack into unsecured app, Sep 2019.
- [48] Hosub Lee and Alfred Kobsa. Understanding user privacy in internet of things environments. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, pages 407–412. IEEE, 2016.
- [49] Shancang Li, Kim-Kwang Raymond Choo, Qindong Sun, William J Buchanan, and Jiuxin Cao. Iot forensics: Amazon echo as a use case. *IEEE Internet of Things Journal*, 6(4):6487–6497, 2019.
- [50] David Lodge. Hijacking philips hue, 2018.
- [51] Alex Lomas. Pwning smart garage door openers, 2020.
- [52] Amith Raj Megalapete Puttaraju, Rekha Jayaram, Bindu SM, and Sneha HR. A study on 802.11 wireless routers hacking techniques and security encryption levels, 05 2016.
- [53] Moen. Smart faucet, 2021.
- [54] Hooman Mohajeri Moghaddam, Gunes Acar, Ben Burgess, Arunesh Mathur, Danny Yuxing Huang, Nick Feamster, Edward W Felten, Prateek Mittal, and Arvind Narayanan. Watching you watch: The tracking ecosystem of over-the-top tv streaming devices. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 131–147, 2019.
- [55] Alan Monie. Tracking and snooping on a million kids, 2018.
- [56] Ken Munro. Domestic cctv and audio recording, 2017.
- [57] Ken Munro. Audio bugging with the fisher price chatter bluetooth telephone, 2021.
- [58] Pardis Emami Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. Privacy expectations and preferences in an iot world. In *Thirteenth Symposium on Usable Privacy and Security ({}SOUPS} 2017)*, pages 399–412, 2017.
- [59] Tom Nardi. Pop open your neighbor’s front door with 12 volts, Jul 2020.
- [60] Open Source Community By Colin Neagle and Colin Neagle. Smart refrigerator hack exposes gmail account credentials, Aug 2015.
- [61] Nuki, 2021.
- [62] Odoo, 2021.
- [63] Douglas A Orr and Laura Sanchez. Alexa, did you get that? determining the evidentiary value of data stored by the amazon® echo. *Digital Investigation*, 24(3):72–78, 2018.
- [64] Palletech, 2021.
- [65] Danny Palmer. Smart vacuum flaws could give hackers access to camera feed, say security researchers, Feb 2020.
- [66] Kari Paul. People are using this smart doorbell to spy on each other, May 2018.
- [67] SCNSoft. Iot inventory management, 2021.
- [68] Francesco Servida and Eoghan Casey. Iot forensic challenges and opportunities for digital traces. *Digital Investigation*, 28:S22–S29, 2019.
- [69] SimpliSafe, 2021.
- [70] Smarter. Smart coffee, 2021.
- [71] Sprout Social. Social media demographics to inform your brand’s strategy in 2021, 2021.
- [72] Williams Sonoma. Smart toaster, 2021.
- [73] Statista. Family hub refrigerator, 2021.
- [74] Statista. Worldwide connected devices, 2021.
- [75] Vangelis Stykas. Hacking smart devices to convince dementia sufferers to overdose, 2020.
- [76] Madiha Tabassum, Tomasz Kosinski, and Heather Richter Lipford. ” i don’t own the data”: End user perceptions of smart home device data practices and risks. In *Fifteenth Symposium on Usable Privacy and Security ({}SOUPS} 2019)*, 2019.
- [77] Lo’ai Tawalbeh, Fadi Muheidat, Mais Tawalbeh, Muhamad Quwaider, et al. Iot privacy and security: Challenges and solutions. *Applied Sciences*, 10(12):4102, 2020.

- [78] Andrew Tierney. Drilling open a smart door lock in 4 seconds, 2019.
- [79] Aaron Tilley. Neighbor unlocks front door without permission with the help of apple's siri, Sep 2016.
- [80] Rahmadi Trimananda, Janus Varmarken, Athina Markopoulou, and Brian Demsky. Packet-level signatures for smart home devices. *Signature*, 10(13):54, 2020.
- [81] Hacking Tutorials. Hacking tutorials, 2022.
- [82] Pedro Venda. Hacking defcon 23's iot village samsung fridge, 2018.
- [83] WEWSTV. Ethical hacker shows us how easily smart devices can be hacked and give access to your personal info, Nov 2019.
- [84] Ilkan Yildirim and MS ErkanBostanci. Forensic analysis of amazon alexa and google assistant built-in smart speakers.
- [85] Min-A Youn, Yirang Lim, Kangyoun Seo, Hyunji Chung, and Sangjin Lee. Forensic analysis for ai speaker with display echo show 2nd generation as a case study. *Digital Investigation*, 2021.
- [86] Eric Zeng, Shrirang Mare, and Franziska Roesner. End user security and privacy concerns with smart homes. In *thirteenth symposium on usable privacy and security ({{SOUPS}}) 2017*, pages 65–80, 2017.
- [87] Eric Zeng and Franziska Roesner. Understanding and improving security and privacy in multi-user smart homes: a design exploration and in-home user study. In *28th {{USENIX}} Security Symposium ({{USENIX}} Security 19)*, pages 159–176, 2019.
- [88] Wei Zhang, Yan Meng, Yugeng Liu, Xiaokuan Zhang, Yingqian Zhang, and Haojin Zhu. Homonit: Monitoring smart home apps from encrypted traffic. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1074–1088, 2018.

APPENDIX

APPENDIX A

Figure 3 presents activity logs from Arm Stay.

APPENDIX B

Table 2 summarizes the results observed by the data population and collection experiments conducted initially.

TABLE II
SUMMARY OF DATA COLLECTED BY A RANGE OF IOT DEVICES.

Device	Log Available	Data Category	Data Specifics	Notes
Ring Doorbell	Depends on service purchased	Can't share activity without subscription		
Google Home	Yes	Personal, Device Data, Activity Data	Email, Device Information	Exportable Logs
Amazon Echo	Yes	Personal Data, Recordings	Audio recordings, text transcript, date and time	Available on the app as well as account website
Google Nest Thermostat	Yes	Personal Data, Acitivity data	Heating cycles/routines, timezone, temperature set, email, name	Exportable Logs
Philips Hue	No	No history	Routines, Rooms in the house, home location	
Wemo	No	Personal Data	No history section, but can see scheduled settings for the lights through the app	Can be connected to GoogleHome, Alexa, and IFTTT, so data would be there as well
Arenti Cam	Partial	Personal data, Activity data, Recordings	In app: sharing footage history and motion detected calendar with video footage, username and located region	Easy to delete, but not to share/export
Roku Express	No	Device data, personal data, financial data	Remote control of tv, name, email, address, phone number, birthdate, demographic info, credit card info	"Do not steal my personal information" Option due to the law in California
YI Home Camera	Depends on service purchased	Can't share activity without subscription	Home & Away settings, more security preferences with a higher subscription	
Topgreener Wifi Power Plugs	No	Personal Data	Automation of plugs, rooms in the house	

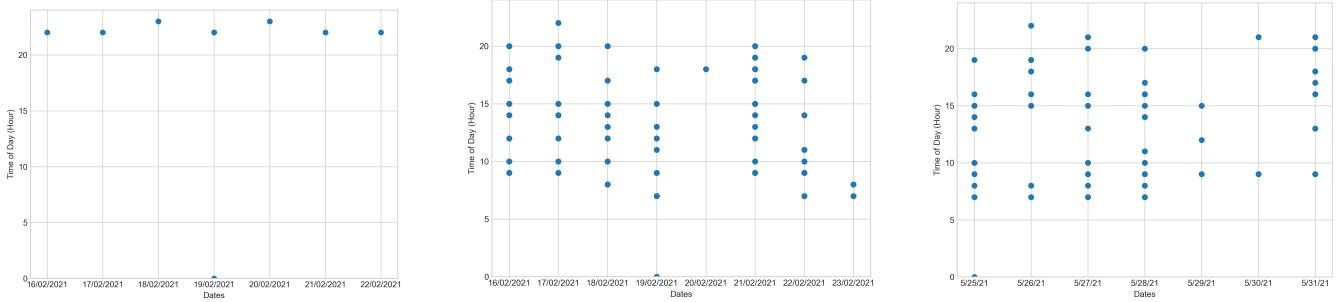


Fig. 3. Snapshot data for (a) Arm stay logs, (b) basement activity, (c) front door activity in setting 4 within a week.