



Abuse Vectors: A Framework for Conceptualizing IoT-Enabled Interpersonal Abuse

Sophie Stephenson and Majed Almansoori, *University of Wisconsin–Madison*;
Pardis Emami-Naeini, *Duke University*; Danny Yuxing Huang, *New York University*;
Rahul Chatterjee, *University of Wisconsin–Madison*

<https://www.usenix.org/conference/usenixsecurity23/presentation/stephenson-vectors>

This paper is included in the Proceedings of the
32nd USENIX Security Symposium.

August 9–11, 2023 • Anaheim, CA, USA

978-1-939133-37-3

Open access to the Proceedings of the
32nd USENIX Security Symposium
is sponsored by USENIX.

Abuse Vectors: A Framework for Conceptualizing IoT-Enabled Interpersonal Abuse

Sophie Stephenson[†] Majed Almansoori[†]
Pardis Emami-Naeini[‡] Danny Yuxing Huang* Rahul Chatterjee[†]
[†]University of Wisconsin—Madison [‡]Duke University *New York University

Abstract

Tech-enabled *interpersonal abuse* (IPA) is a pervasive problem. Abusers, often intimate partners, use tools such as spyware to surveil and harass victim-survivors. Unfortunately, anecdotal evidence suggests that smart, Internet-connected devices such as home thermostats, cameras, and Bluetooth item finders may similarly be used against victim-survivors of IPA. To tackle abuse involving smart devices, it is vital that we understand the ecosystem of smart devices that enable IPA. Thus, in this work, we conduct a large-scale qualitative analysis of the smart devices used in IPA. We systematically crawl Google Search results to uncover web pages discussing how abusers use smart devices to enact IPA. By analyzing these web pages, we identify 32 devices used for IPA and detail the varied strategies abusers use for spying and harassment via these devices. Then, we design a framework—*abuse vectors*—which conceptualizes IoT-enabled IPA as four overarching patterns: *Covert Spying*, *Unauthorized Access*, *Repurposing*, and *Intended Use*. Using this lens, we pinpoint the necessary solutions required to address each vector of IoT abuse and encourage the security community to take action.

1 Introduction

Smart, Internet-connected devices (i.e., IoT devices) such as smart speakers, smart plugs, switches, locks, cameras, and thermostats are gaining popularity in the US and around the world. As of 2019, 69% of US households had at least one smart home device, and 12% had multiple smart devices in their household [56]. Smart devices provide convenience (e.g., remotely changing the temperature), efficiency (e.g., automating lighting schedules), and peace of mind (e.g., viewing home security feeds from afar). As a result, the adoption of smart devices in homes, cars, and elsewhere is growing rapidly, and the number of smart devices is expected to double by 2025 [90].

Although smart devices can be beneficial for many households, recent news stories indicate that they can also be used for malicious purposes [11, 93]. For instance, people have used AirTags and Apple Watches to track individuals without

their consent [14, 22, 50, 55, 59]; in another case [23], a man remotely accessed his ex-boyfriend’s Ring doorbell to spy on him and ring the doorbell in the middle of the night. These troubling anecdotes are instances of smart device-enabled *interpersonal abuse* (IPA), where abusive intimate partners, family members, friends, roommates, and coworkers (among others) leverage smart devices to spy on and harass their targets. We refer to this phenomenon as *IoT abuse*.

Despite the prevalence of IoT abuse, academic work in this area is still emerging. Some scholars [49, 51, 53, 62, 80, 81] have begun to study IoT abuse in the context of *intimate partner violence* (IPV) (a particularly prevalent form of IPA [75, 76]); for example, Tanczer et al. reported on real cases of IoT abuse gleaned from interviews with advocates and victim-survivors.¹ However, no work has empirically measured the role of different smart devices in IPA, nor has any work attempted to systematize our understanding of IoT abuse. Without a more thorough understanding of the problem, it is challenging to detect and prevent IoT abuse.

Thus, in this work, we fill this gap by empirically studying evidence of IoT abuse online, then systematizing IoT abuse into four broad patterns (*abuse vectors*). We ask three main research questions:

1. What types of smart devices have been (or can be) involved in interpersonal abuse?
2. Which properties of smart devices enable abuse?
3. How can we conceptualize the problem of IoT abuse to better guide mitigations?

Empirically studying IoT abuse is challenging because such cases are often not reported to or recorded by authorities [11]. Therefore, there is no formal archive of IoT abuse incidents which can be used to systematically understand *how* smart devices are being abused. Following prior work [18, 88], we hypothesize that IoT abuse incidents are posted online, making the web an informal archive of abuse incidents reported by

¹We refer to people who have experienced IPA as “victim-survivors.” This phrasing is inclusive of the varied ways people refer to their experiences of abuse [12], and has been used in prior work [73].

victim-survivors, abusers, and third parties. Accordingly, we learn about IoT abuse by gathering and analyzing online content. We first collect 70,399 unique web pages from 14,491 domains using methodically-generated search queries such as “spy on wife using camera” on Google. We filter irrelevant pages using a simple classifier, then sample 320 web pages for manual investigation.

Our investigation provides a comprehensive archive of IoT abuse. We identify 32 types of smart devices (including the names of several specific devices and vendors) used to surveil or harass individuals. Seven of these are covert spying devices—hidden listening devices, spy cameras, hidden GPS trackers, etc.—that are designed for spying on individuals. We found several descriptions of incidents where abusers hide audio/video recording devices or location tracking devices in a victim-survivor’s home, in a vehicle, or even in a child’s toy. To make matters worse, we found links to spy devices for sale on popular online retailers. On the other hand, we also found that many well-intentioned smart devices are *dual-use* [18], meaning abusers can repurpose them to spy on or harass a target. For example, abusers spy on victim-survivors using doorbells and security cameras; track their locations using item finders, GPS trackers, and smart vehicles; and harass them using smart appliances, door locks, and speakers.

Based on this analysis, we design the first framework to organize the broad space of IoT abuse. Our *abuse vectors* framework divides IoT abuse into four comprehensive components: *Covert Spying*, *Unauthorized Access*, *Repurposing*, and *Intended Use*. By highlighting the four primary vectors of IoT abuse, we pinpoint the tailored solutions required to address each vector and provide a clear path for future work. For example, mitigating *Covert Spying* requires mechanisms to find hidden smart devices, whereas, for *Unauthorized Access*, devices will require changes to authorization and access control mechanisms. Potentially the most worrisome and challenging abuse vector, *Intended Use*, happens when the intended use of the device provides spying and harassment capabilities. By providing this framework for conceptualizing IoT abuse, we hope our work will begin a deeper conversation on the problem and influence safer design choices for smart devices.

In summary, our contributions include:

- We perform a large-scale, systematic survey of web content for evidence of IoT-enabled interpersonal abuse. Our analysis uncovers a wide variety of devices used for abuse, including many not reported in previous work.
- We design a novel framework, *abuse vectors*, which succinctly captures this broad space of IoT abuse.
- Based on our framework, we detail several solutions needed to tackle each vector of IoT abuse.

2 Background & Related Work

Interpersonal abuse (IPA), also called interpersonal violence, is “the intentional use of physical force or power against other persons by an individual or small group of individuals” [58]. IPA can include physical violence, but also emotional or sexual abuse [58]. According to the CDC, one in three women and one in four men in the US have experienced some form of physical IPA by an intimate partner [75, 76]. Lately, abusers are using technology to conduct IPA [10, 45, 57, 68, 83, 88].

2.1 Technology-Facilitated Abuse

Several prior works [18, 26–28, 34, 57, 88] have reported that abusers regularly use seemingly benign applications and web services to spy on, stalk, harass, intimidate, and control their targets. Many spyware tools and benign apps are available for Android and iOS which allow the abuser to collect private data from the victim-survivor’s device, including recorded calls, SMS, social messages, and current location [2, 18]. These applications (a.k.a. “creepware”) exist not only for spying, but also for harassment [68]. Unfortunately, technology-based intimate partner surveillance using these tools is promoted and discussed in online forums [88].

Tech-facilitated IPA (also referred to as *tech abuse* [73]) can be emotionally and psychologically harmful to victim-survivors yet can easily be conducted without technical expertise [26]. To compound the problem, the context of tech abuse is complex [57], and stakeholders like victim service providers (VSPs) are often ill-equipped to handle cases involving technology [27]. Fortunately, technology researchers have begun to design “clinical computer security” [34] solutions for survivors of tech abuse [25, 34, 89], albeit focusing only on mobile device-based abuse.

2.2 Tensions with Shared Smart Devices

We use the term *smart devices*—or *IoT devices*, used interchangeably in this paper—broadly to refer to any consumer device that can directly or indirectly (via another device) communicate information over the Internet. This includes smart speakers, door locks, smart appliances, thermostats, security cameras, vehicles, item finders, and more. We do not include general-purpose computing devices such as laptops, tablets, and smartphones unless they are used in conjunction with another (special-purpose) smart device. Some smart devices such as Tile [84] and AirTag [37] might not have direct Internet access, but they can communicate via other nearby devices, such as a phone, tablet, or an IoT Hub.

As smart devices grow more commonplace in shared environments like the home, it has become clear that these devices exist in a nuanced socio-technical context. For instance, control over smart home devices is imbalanced among users: the person who installs these devices usually has more power over them [16, 29, 94]. This power imbalance can extend

or exacerbate existing power dynamics in (potentially abusive) household relationships [9]. Users have also expressed privacy concerns about smart home devices, such as Alexa and other smart speakers [96]; since these speakers continuously listen to users, people question the data collected by these devices [48]. Although users are aware of these privacy issues, researchers point out that many of them are not concerned enough [94, 95]. Additionally, work on users' privacy perceptions of IoT devices has not considered the threat of interpersonal abusers, as Geeng and Roesner point out [29]. This is a ripe pretext for abuse.

2.3 Smart Devices & IPA

Only recently, researchers have started looking into how smart devices are being used for IPA. Some articles broadly outline the potential for smart device abuse; for example, Levy and Schneier [51] discuss several technological “intimate threats,” many involving smart devices. However, they did not collect empirical data to support their findings. Other work identifies key aspects of this IoT-enabled abuse—or *IoT abuse*—but does not aim to characterize the problem as a whole. Leitão et al. [49] performed a co-design study with survivors of IPV, guiding participants to envision potential vectors for IoT abuse and propose solutions; Parkin et al. [62] evaluated two common smart speakers with a novel usability assessment, revealing opportunities for IoT abuse; and Tanczer et al. [80, 81] and Lopez-Neria et al. [53] interviewed advocates in order to learn about real-world cases of IoT abuse. Finally, Slupska and Tanczer [74] outlined an IPA threat model for device manufacturers to reference during design.

Prior work has approached the problem of IoT abuse from multiple angles. However, crucially, the research community still lacks a strong conceptual understanding of IoT abuse as a whole. We fill this gap by performing a systematic crawl of online content related to IoT abuse (§ 3). Our results enable us to understand the many forms of IoT abuse (§ 4 – § 6), create a framework to conceptualize this broad problem (§ 7), and propose tailored solutions based on this framework (§ 8).

3 Collecting and Analyzing Web Content

IoT abuse is difficult to survey because many such cases are not formally reported to authorities, nor does a database of these incidents exist. We hypothesize that victim-survivors (and abusers) might share their experiences with smart devices online, making the web an informal archive of information about IoT abuse. Indeed, other work has collected similar data [18, 88]. Accordingly, we study IoT abuse by collecting and analyzing relevant web content.

3.1 Search Query Generation

First, we did a preliminary Google search to understand what types of web pages we could find online. Using 19 queries

like “spy on spouse smart home,” we found 55 relevant web pages, including several forum posts, news articles, and blog posts discussing specific IoT abuse incidents and potential ways smart IoT devices can be used for IPA. We also found anecdotes from abusers claiming abuse and blog posts on how to weaponize smart devices to catch a cheating spouse. These examples formed the basis of our search queries.

Next, we used these preliminary search results to build a set of seed queries. From the 55 web pages, we gathered 97 illustrative quotes such as “spurned husband uses Wi-Fi thermostat to take lingering, ghostly revenge.” Then, from these quotes, we extracted a list of relevant *devices*, *actors*, and *verbs* (e.g., “Wi-Fi thermostat,” “husband,” and “take revenge”). Finally, we merged these components into 50 query templates which combine different actors, devices, and verbs. For example, one template reads “spy on [actor] using [device].” To cover a wide range of relevant queries, we bolstered our list of devices with 55 popular devices reported by Ren et al. [65] and devices like “smart thermostat” which had appeared in news reports but were not in our preliminary search. In all, we combined 50 query templates, 78 devices, and 8 actors for a total of 21,897 seed queries.²

Prior work has often used snowballing techniques to expand lists of search queries using automation [2, 18, 35]. Thus, we further expanded our set of queries with suggestions from Google. We submitted our 21,897 seed queries to the Google query suggestions service, analyzed the suggested queries, and filtered obviously irrelevant ones (e.g., queries about health tracking rather than tracking a target). We added suggested queries until we reached 25,000 search queries.

3.2 Crawling Google Search Results

In October 2021, we searched Google using our 25,000 queries. We used Selenium [78] to query and Beautiful Soup [20] to parse the HTML pages. On average, we made two queries per second, which we believe had a negligible effect on Google’s regular operation. For each query, we recorded the title, URL, and snippet—a fragment of the page content, usually containing part of the search query—of all results shown on the first page of the Google search results (at most ten results per query).

In total, we collected 181,991 search results. Different search queries often yielded the same URLs; after removing duplicate search results from the dataset, we were left with 70,399 unique web pages hosted on 14,491 domains. For each unique web page, we downloaded the raw HTML and noted how *popular* the page was by computing the number of search queries which had returned that web page in the first page of results. As we will discuss, the most popular web pages were more likely to be relevant to our research

²Not all queries contain both an agent and a device, resulting in less than $50 \times 78 \times 8 = 31,200$ queries. We share the actors, devices, and templates we used, along with our codebook, at <https://go.wisc.edu/k7ai9p>.

Full dataset		After filtering	
Domain	# URLs (%)	Domain	# URLs (%)
books.google	14,817 (21.1%)	amazon	1,083 (4.1%)
amazon	3,303 (4.7%)	reddit	945 (3.6%)
reddit	1,956 (2.8%)	bestbuy	679 (2.6%)
quora	1,373 (2.0%)	books.google	496 (1.9%)
bestbuy	726 (1.0%)	youtube	270 (1.0%)
Total	70,399	Total	26,286

Figure 1: Top-five domains in our dataset or URLs before and after applying our automated relevance classifier (§ 3.3).

questions. For example, the most popular web page (returned by 923 search queries) is a page recommending which spy gadgets to use to catch a cheating spouse.

3.3 Filtering Irrelevant Web Pages

As with any web search, our search queries did not always return relevant web pages. For example, our search results contain several mobile spyware websites and general IoT product pages, which are not relevant to our precise research questions. Thus, the first step in our analysis was to filter out some of the irrelevant pages in our dataset.

We began by analyzing two distinct samples of web pages from our dataset: the 25 most popular web pages and a random sample of 25 web pages. We chose these pages to solidify our definition of relevance, but also to determine whether focusing on popular web pages might yield more relevant results than a random sample. Three authors read the web pages, marked which ones were relevant, and discussed the requirements for a web page to be relevant. We defined relevance as follows: a web page is relevant if it (1) mentions smart devices and (2) mentions intentional, tech-enabled IPA. Under this definition, only 3 web pages in the random sample were relevant (12%), compared to 12 web pages out of the 25 most popular (48%).

Based on our definition of relevance, we trained a Random Forest classifier to identify relevant pages. Our training and testing data were pulled from a hand-labeled sample of 300 web pages from our dataset, of which 23% were relevant. Appendix A contains more details of how we trained the classifier. Our final model had 80% accuracy with 93% recall, averaged across a five-fold cross-validation. The precision of the classifier was low (68%), but it helped us filter 44,113 pages from our dataset. After filtering using this classifier, we had 26,286 web pages for further analysis. Fig. 1 depicts the top domains in our dataset before and after filtering.

We designed our classifier to be conservative, with high recall and low precision, in order to avoid filtering out any relevant web pages. Thus, several irrelevant pages remained in the filtered dataset. To address this, before analyzing a web page, we first manually identified if the page was relevant by looking for references to smart devices and IPA in the page title, snippet, and content. If the page was irrelevant, we

Structural code	Example sub-codes
Domain Type	Forum, News
Domain Topic	Tech, Divorce law
Page Type	Forum post, News article
Anecdotes	First-person, Survivor POV
IPA Targets	Intimate partner, Family member
IPA Strategies	Spy through camera
Anti-IPA Strategies	Reset devices
General Tech	Car, Smart speaker
Specific Tech	Tesla, Amazon Echo
Motivations	Infidelity

Figure 2: Our structural codes, along with example sub-codes.

ignored it and moved onto the next result. Of the 320 total web pages we investigated during analysis (§ 3.4), 162 web pages (51%) were relevant (Fig. 3).

3.4 Qualitative Descriptive Analysis

Using the filtered web pages we collected, we aimed to understand and characterize the different forms of IoT abuse. Thus, we employed *qualitative descriptive analysis*, a technique “oriented toward summarizing the informational contents of data” [70]. To code individual pages, we used structural coding [69, Ch. 3]. We defined ten structural codes based not only on the data, but also on our a priori knowledge and research questions [67]; then, within each structural code, we generated subcodes organically from the data. By extracting information from the relevant web pages—e.g., the target of IPA, the technology discussed, method of IPA described—we produced a taxonomy of IoT abuse that explains our dataset.

Codebook generation. We developed the structure of our codebook using the 25 most popular pages, plus a random sample of 25 pages (same as § 3.3). Three researchers made notes on the content of the web pages, then discussed notes to solidify the structural codes. Fig. 2 depicts the structural codes we defined, as well as some examples of sub-codes we generated later in the process. We gathered information about the type and topic of the page’s domain, the type of web page (e.g., forum post), and the IoT abuse mentioned in the page—the target of the IPA, the presence of real-world anecdotes, the strategies used to enact (or combat) IPA, the types of technology mentioned, and the motivations given.

Next, three researchers independently coded the 100 most popular web pages within these structural codes. We coded information about the domain by looking at the domain’s “About” section (or equivalent), and coded for page content by reading the page carefully. For most web pages, we focused on the main content of the page, excluding user comments; for forum posts, we coded both the original post as well as the first page of user replies. After this process was complete, we discussed as a group to resolve any disagreements and converged on a codebook.

Analyzing web pages. With this solidified codebook, three

	Web Pages	Reddit	Quora	Comments	Domains
Notation	W###	R##	Q##	C##	—
Total	70,399	1,956	1,373	N/A	14,491
Filtered	26,286	945	195	N/A	8,402
Investigated	220	50	50	32	152
Relevant	113	28	21	14	85

Figure 3: Summary of the Google Search results through our initial crawl, filtering, sample, and final set of relevant results. The “Investigated” row refers to all pages we looked at during analysis, including those we manually labeled irrelevant.

researchers equally divided and analyzed the most popular web pages in order until we had coded 100 relevant web pages. (We needed to look through the 207 most popular web pages before we found 100 relevant pages.) To evaluate whether our codebook was saturated, we then looked at the rest of the 500 most popular web pages and coded any pages that mentioned a device we had not seen or mentioned a new way of using a device for IPA. In this process, we coded only 13 more pages, indicating that the top 100 relevant pages captured much of the content we gathered. In total, we investigated 220 web pages, out of which 113 (51%) were relevant (see Fig. 3).

Forums & comments. We noticed that web pages coming from forums often contained valuable first-hand anecdotes of IoT abuse. To dig deeper, we performed a separate analysis of web pages we had collected from Reddit and Quora. We selected these two forums because they were the top two general-purpose forums which appeared in our dataset (the filtered dataset contains 945 pages from Reddit and 195 pages from Quora). We thus investigated the top 50 most popular search results from both Reddit and Quora, using the same method and codebook as before. In total, 28 of the Reddit posts and 21 of the Quora posts were relevant (Fig. 3), bringing our total number of coded web pages to 162.

We also noticed that the web pages often contain relevant comments. Thus, we coded the comment sections of several pages. We scanned the top 100 most popular pages in our dataset and identified 32 which contained comment sections. Then, we coded the first page of the comment section, treating the collection of comments as another page in our dataset. In total, we coded 14 relevant comments sections (Fig. 3).

3.5 Assessing Cross-Platform Saturation

Though Google has dominated the search engine market for years [43], it can be biased in coverage and ranking of web pages [30, 91]. To assess the generality of our results, we compared our taxonomy with a new set of web pages obtained via the DuckDuckGo search engine [21]. Using the same 25,000 queries as before, we crawled DuckDuckGo and extracted 87,666 unique search results. Out of these URLs, only 12% were identified in the prior Google search. This is not surprising given the differences in the two search engines and

the fact that the crawls were done 8 months apart. Next, we used our machine learning classifier and manual pruning to identify the 100 most popular relevant search results from the DuckDuckGo crawl. Of these 100 URLs, we had already analyzed 44 in the initial study; of the rest, 26 appeared in the Google crawl but were not a part of our sample. Though the total overlap was small, many of the most popular and relevant results remain constant across both crawls.

We analyzed the 56 *new* relevant results from the DuckDuckGo search using the method detailed in § 3.4. Nearly every result repeated the same IoT-abuse strategies we had already seen in our Google crawl and captured in our previous analysis. We found only one new IoT-abuse strategy, which uses smart speakers; we append this new strategy to our discussion of spying in § 5. Nevertheless, we already identified several other ways smart speakers can be used to spy, as shown in Fig. 5 and § 5. Our analysis of these DuckDuckGo crawl results emphasizes that though specific search results may differ across platforms, the overall takeaways from those results (i.e., the types of IoT-enabled abuse one can find information on) remain constant. Thus, we believe our Google crawl methodology provides a good representation of the IoT-enabled abuse strategies being discussed on the Internet.

3.6 Ethics & Limitations

Because all web content used in this study is publicly available, this study does not qualify for an IRB review. Even so, we are mindful that our work involves a vulnerable population: victim-survivors of interpersonal abuse. In particular, some of the forum posts might contain accounts of ongoing abuse. Thus, we do not collect any personally identifiable information such as names and email addresses, nor we compile messages or posts from the same author in our data. To avoid providing potential abusers with an archive of IPA-related content, we do not publish the list of web pages we coded (though they may be shared with researchers on request) and use only IDs when referring to IPA-relevant web pages. We are aware that the paper’s content may itself provide ideas to potential abusers. However, a potential abuser may find the same information via Google Search. Similarly to prior influential papers (e.g., [18, 68]), the goal of this research is to raise awareness among stakeholders to prevent IoT abuse.

Study limitations. First, we rely on web pages for reports of IoT abuse and assume each reported incident on the web is accurate and truthful, which might not always be the case [47]. For example, some device manufacturers overstate their capabilities—a common attempt by spyware developers [18] to allure abusers—and several articles include results that mention *potential* ways that smart devices could be used to enact IPA. As long as they are plausible, we consider these devices and methods in our study. We do this to ensure that we learn as much as possible about how smart devices can be used for IPA, even if an abuse strategy is not widely re-

ported. We do not verify whether all of the spying/harassment capabilities are, in fact, true; this would require lab testing, beyond the scope of this paper. Finally, though our results are stable across search engines and over time (§ 3.5), Google search might not be a complete representation of the web. For instance, our search results were limited to pages in English. Therefore, our results reflect a lower bound of the smart devices being abused for IPA.

4 Overview: IoT Abuse in Our Dataset

We analyzed 320 Google Search results and identified 162 relevant web pages, including 49 forum posts and 14 comment sections (Fig. 3). These pages revealed rich information about the smart devices used for interpersonal abuse and the specific ways abusers (ab)use these devices.

Notation. In the following sections, we refer to web pages using four types of IDs (Fig. 3). *W###* refers to a webpage from the top 500 most popular web pages; *R##* refers to a Reddit post; *Q##* refers to a Quora post; and *C##* refers to a comment section. Within each of these groups, IDs are in order of frequency (e.g., *R01* is the most popular Reddit post).

Result characteristics. The IPA mentioned in our search results is directed towards several targets. The most frequent by far is intimate partners, mentioned in 118 web pages (76%) we coded; we also found that others, including family members, roommates, neighbors, and hired workers (e.g., nannies), are regularly targets of IPA (Fig. 4). Our dataset contains informational pages aimed at raising awareness, pages advertising the use of covert spy devices, anecdotes from worried victim-survivors or potential abusers, news stories of real-life instances of IoT abuse, and more.

Types of smart devices used in IPA. We identified 32 types of devices that can be used for interpersonal surveillance and harassment. These devices can be separated by context: (1) devices that are typically shared and (2) devices owned and used by a single person. Of the shared devices, we found home control devices (meant for controlling smart home devices), smart appliances, security systems, and vehicles. For personal devices, we identified tracking devices, devices meant for entertainment, and covert spying technologies (designed specifically for illicit use). Fig. 5 outlines the devices we found and marks the spying (§ 5) and harassment (§ 6) strategies associated with each device.

Some devices enable spying or harassment only in a limited way. For instance, some devices allow audio/video surveillance, but with caveats that the abuser is within Bluetooth range of the device or must be able to return to the physical device in order to access the footage. Similarly, motion-sensing devices can indicate when a target is in proximity to that device, providing a clue to the target’s location, but cannot move with the target to track their location continuously. We note these limited capabilities in Fig. 5 with a blue icon.

Target	Count (%)	Target	Count (%)
Intimate partner	118 (75.6%)	Guest	6 (3.8%)
Family member	20 (12.8%)	Employee	3 (1.9%)
Roommate	16 (10.3%)	Coworker	2 (1.3%)
Neighbor	10 (6.4%)	Friend	2 (1.3%)
Hired worker	8 (5.1%)	Client	2 (1.3%)
Vague	13 (8.3%)		

Figure 4: Distribution of targets. Hired workers include nannies, maids, and other domestic workers. Results often mentioned multiple targets of IPA.

5 IoT-Enabled Interpersonal Surveillance

We identified four ways that abusers surveil their targets with IoT devices: audio surveillance, video surveillance, location tracking, and accessing private data.

5.1 Audio Surveillance

Abusers regularly use smart devices such as smart speakers, AirPods, and smart TVs to eavesdrop on victim-survivors. This type of spying can include monitoring previously-recorded audio and even eavesdropping on a live conversation.

Using home control systems. Home control systems help users control different smart home devices from one place via voice or touch interfaces. Examples of these devices include smart speakers, such as Google Home, Nest Home, Amazon Echo, Echo Dot, and Echo Show, as well as smart home control tablets. These devices often record surrounding audio when they hear a wake-word such as “Alexa” [61]. In our search, we found that abusers can exploit this feature to spy on victim-survivors. For example, in one instance, a woman searched through the command history of her Amazon Echo to look for evidence that her partner was unfaithful (W053). Similarly, smart TVs (such as ones with inbuilt Alexa or Google Home apps) are equipped with microphones so users can interact with them using voice commands. Interactions with these smart TVs are recorded just like smart speakers, and abusers can use this capability for spying (W219).

On Alexa-enabled devices, a second feature that enables eavesdropping is Drop In [4]. Fourteen different pages that we coded note the dangerous potential of Drop In. The Drop In feature is meant to facilitate intercom-type communication; a user can “drop in” to any device they own (or on a device owned by a friend who pre-authorizes access) and immediately project their voice through the device and hear any response. However, because the receiving device *automatically* accepts the call, Drop In allows potential abusers to seamlessly eavesdrop using Alexa-enabled devices. The devices make a chime sound and flash a green light to announce to the surrounding users that Drop In is being used. However, these announcements may be easily missed if the device is quiet or muted (potentially by an abuser) or is not in the

Context	Category	Device	Discussed Strategies		Abuse Vectors
			Spy	Harass	
Shared-use devices	Home control	Smart speaker	🔊📹 - 🔒	🏠 - 📧	- U R -
		Control tablet	🔊📹 - 🔒	- - -	- U R -
	Smart appliances	TV	🔊📹 - -	🏠 - -	- U R I
		Thermostat	- - 📍 -	🏠 ⚠️ -	- U R -
		Lights	- - 📍 -	🏠 - -	- U - I
		Router	- - - 🔒	🏠 - -	- - R I
		Plug	- - 📍 -	🏠 - -	- - R I
		Kettle	- - - -	🏠 - -	- - - I
		Smoke alarm	- - - -	🏠 - -	- - R -
		Fridge	- - 📍 -	- - - -	- - R -
	Mattress	- - - 🔒	- - - -	- - R -	
	Security systems	Doorbell	🔊📹 📍 🔒	🏠 - -	- U - I
		Security camera	- 📹 - 🔒	- - - -	- U - I
		General camera	🔊📹 - -	- - - -	C U - I
		Baby monitor	🔊📹 - -	- - - -	C - - I
		Lock	- - 📍 -	🏠 ⚠️ -	- U - I
		Motion sensor	- - 📍 -	- - - -	C - - I
		Presence sensor	- - 📍 -	- - - -	C - - I
		Garage door opener	- - - -	- ⚠️ -	- - R I
	Vehicles	Car	- - 📍 -	🏠 ⚠️ -	- - - I
Car accessory		- - 📍 -	- - 📧	C - R -	
Personal-use devices	Tracking devices	Watch	🔊📹 📍 🔒	- - - -	- - R -
		Item tracker	- - 📍 -	- - - -	C - - -
	Entertainment	Bluetooth headphones	🔊 - 📍 -	- - - -	C U R -
		Smart toy	🔊 - - -	- - - -	- - R -
	Covert spying technologies	Hidden camera	🔊📹 - -	- - - -	C - - -
		Spy drone	- 📹 📍 -	- - - -	C - - -
		Thermal camera	- 📹 - -	- - - -	C - - -
		Listening device	🔊 - - -	- - - -	C - - -
		Landline recorder	🔊 - - -	- - - -	C - - -
		GPS tracker	- - 📍 -	- - - -	C - - -
USB keylogger	- - - 🔒	- - - -	C - - -		

Full spying: Remote audio (🔊) and video (📹) surveillance, precise location tracking (📍), and accessing private data (🔒).
Limited spying: Distance-limited audio/video spying (🔊, 📹); location tracking with stationary device (📍).
Harassment: Disrupting the home environment (🏠), threatening physical safety (⚠️), and manipulating private data (📧).
Abuse vectors (discussed in § 7): Covert Spying (C), Unauthorized Access (U), Repurposing (R), Intended Use (I)

Figure 5: The smart devices found in our web crawl, along with their associated abuse strategies and *abuse vectors* (§ 7).

victim-survivor’s immediate line of sight.

Unfortunately, Drop In is not the only way an abuser can use smart home control systems for live eavesdropping. One webpage discusses how an abuser can ask a smart speaker to call their phone, answer the call, and leave the line open to provide a live channel for eavesdropping (DuckDuckGo result D14; see § 3.5). In another instance, an abuser was found guilty of eavesdropping on a victim-survivor through a wall-mounted iPad they used to control their smart home. He was able to do so by connecting to the microphone on the iPad through an app on his phone (W005, W110).

Using AirPods and Apple Watches. AirPods, Apple’s wireless Bluetooth headphones, have a “Live Listen” feature [41] which allows the user to listen to sound coming through the

microphone of their paired iPhone. This feature is meant to facilitate conversation in noisy areas or across a large room—unfortunately, it also allows an abuser to place their phone near the victim-survivor, walk away, and listen to the sounds picked up by the phone (R23). This approach is limited: the abuser must be physically present to place the device, then must remain in Bluetooth range (around 40 meters [15]). The phone’s lock screen also indicates that Live Listen is on. Nevertheless, these physical limitations are trivial for abusers who live with their target, and the abuser can hide the Live Listen notification by simply placing it face-down.

An Apple Watch, too, can be converted into a subtle eavesdropping device. Watch-compatible apps like Just Press Record [77] allow a user to press a button and record audio using the Watch (W216). Further, before Apple updated it in

2019, Apple Watch’s Walkie Talkie feature [42] allowed users to connect to a contact’s iPhone and listen to sound picked up by its microphone (W215). Fortunately, the updated Walkie Talkie feature requires each user to press and hold a button on the Watch in order to send audio during a call [42].

Using other smart home devices. While cameras are primarily used for video surveillance (§ 5.2), some also have audio recording capabilities. For instance, smart video-enabled doorbells such as Ring can be used to listen to conversations remotely—in one web page, the author described how he used a Ring to listen to his partner outside the door (W047). We also found that smart toys can have audio recording capabilities, like Hello Barbie (now discontinued), which used natural language processing to have basic conversations with children. With this type of toy, abusers could remotely listen to victim-survivors without their knowledge (W57, W67).

Using covert audio recording technologies. Lastly, abusers use listening devices designed for spying (a.k.a. “bugs”) to surveil victim-survivors. In one Quora answer, the author describes how a family member placed a recording device in her husband’s truck to prove he was cheating (Q16). These devices are sometimes disguised as everyday objects such as pens, key fobs, thumb drives, and tissue boxes (W001, W64, W184), making them more difficult for victim-survivors to detect. Abusers may also hide a listening device in a child’s toy (W086); if the victim-survivor and abuser share custody of a child, the abuser can surveil the victim-survivor this way without having physical access to the victim-survivor’s home.

5.2 Video Surveillance

Video surveillance, the most prevalent form of spying we observed, appears in 94 pages we analyzed (58%). Abusers use many types of cameras—security cameras, baby monitors, video doorbells, etc.—and even devices like Apple Watches to spy on victim-survivors with and without their knowledge.

Using security and home control systems. Abusers regularly use video-enabled devices like doorbells, security cameras, and Amazon Echo Show displays to surveil victim-survivors. Doorbells and security cameras can not only allow remote viewing, but they can also record (and notify) if a certain activity is detected, such as the presence of a human. Additionally, on Echo Show devices, the aforementioned Drop In feature can also provide video surveillance. Echo Show devices have an inbuilt camera and video display, so abusers can use Drop In to not only eavesdrop on audio but also see through the device’s camera. Alexa-enabled smart TVs also allow video surveillance through this Drop In feature.

These devices are typically present in the home and known to the victim-survivor, but it is hard to detect when they are being used to spy. As mentioned earlier, Amazon devices give little notification when Drop In is used, and victim-survivors who notice the notification may not know what it means. Ring Doorbells and security cameras often do not provide

any notification on the device when the live feed is viewed remotely, leaving victim-survivors unaware of surveillance. For example, a Reddit post reads, “*I think my husband is spying on me. How can I tell when he is watching on the indoor camera?*” (R35). Detecting spying activities on these devices is a difficult problem (§ 8).

Using covert cameras. In addition to overt cameras, 42 web pages discuss surveillance via hidden cameras, including cameras designed for spying, discreetly-placed security cameras, and even stealthy baby monitors. Fourteen pages go so far as to recommend that abusers use hidden cameras to catch a cheating spouse (e.g., W001, W007). One such camera is a “spy drone,” or a drone with a built-in video camera: for example, one unsuspecting woman at a backyard barbecue looked up and saw a drone piloted by her ex-husband (W219). Additionally, abusers have been known to purchase security cameras and baby monitors, then place them out of view to use as a hidden camera. One such abuser asks on Quora: “*What should I do? My girlfriend found my security camera in her apartment*” (Q19). In contrast to the previous section, once the victim-survivor uncovers one of these devices, its presence is immediately concerning; thus, detection is crucial.

Using Apple Watches. Though less common, pages also documented the use of Apple Watches for video surveillance. While in Bluetooth range of its paired iPhone, an Apple Watch can use the Camera Remote feature [40] to open the phone’s camera app and see the current video feed (W218, R27). If the victim-survivor discovers the phone, they only see that the camera app is open; they have no way to tell that the abuser is watching the camera feed through their Apple Watch.

5.3 Location Tracking

Beyond audio and video surveillance, abusers use IoT devices to track victim-survivors’ movements. Our results surfaced two types of location tracking: precise location tracking via GPS or Bluetooth mesh networks (e.g., [38]) and tracking the victim-survivor’s proximity to a device, like a doorbell.

Using dual-use tracking devices. In our dataset, we found many instances of abusers using benign tracking devices to spy. Item trackers like Tile and AirTag provide abusers with an easy, affordable, and discreet way to track someone’s location without their knowledge. Like AirTags, AirPods and Apple Watches are also connected to the Find My [38] network to help users find lost devices; as a result, abusers also repurpose these devices to track victim-survivors. For example, R30 reads, “*My Ex connected to my AirPod [P]ros the other day and is using them to track me. How do I stop this?*”

While AirTags (and other Find My devices) can only communicate through nearby Apple devices, other tracking devices (e.g., GPS-enabled pet trackers (W204)) often have a GPS sensor to identify their location and a Mobile LTE connection to share location data. Abusers with access to this type of device can track the victim-survivor’s whereabouts while

they go out, say, for a dog walk. Although this seems like a short period of tracking, it can be dangerous; an abuser could use this information to confront the victim-survivor while they are outside, without home locks or security systems to protect them.

Using covert tracking devices. Among the web pages we coded, 38 mention the use of *covert* GPS trackers to monitor victim-survivors. Unlike covert audio recorders and cameras, these devices are not usually camouflaged as everyday objects; instead, they are meant to be hidden from view entirely. In our results, we observed that abusers often attach tracking devices to victim-survivors' vehicles or hide them inside teddy bears and other toys (W139, W161). As long as the victim-survivor is ignorant of the tracker's presence, this provides the abuser with fine-grained location tracking. Spy drones can also provide discreet location tracking if an abuser uses the drone to follow the victim-survivor from a distance (W154).

Using smart vehicles and accessories. While covert tracking devices are being used to track cars, many smart vehicles provide inbuilt GPS tracking capabilities, often via mobile applications which allow remote tracking. Abusers regularly exploit these technologies to track victim-survivors. In many cases, victim-survivors and abusers share ownership of a car, or the abuser is the sole owner (e.g., if the abuser is a parent). Therefore, the abuser can have easy access to information about the car's location on the app, even if the abuser is no longer in the victim-survivor's life. For example, one woman's husband bought her a Tesla, then used the app to track her location (W029).

Not just cars, but also their accessories can be used to track a victim-survivor. EZ-Pass devices [33], for example, are designed for electronic toll collection, so the driver does not have to stop to pay tolls. Information from a car's EZ-Pass can help an abuser identify different tolls and highways the victim-survivor has traveled on, providing clues about their location history (W050). Abusers also leverage GPS-enabled vehicle tracking devices which can be attached to a car's on-board diagnostic (OBD) ports to track its location (W219).

Using security systems and smart appliances. More locally, an abuser can use smart home clues to tell whether a victim-survivor is in the house and identify their habits. With access to devices like smart locks, motion sensors, security tabs, and video doorbells, an abuser can tell when the victim-survivor enters or leaves the house by monitoring when they trigger motion sensors, open smart locks, or appear in the doorbell's video feed. This is dangerous information that allows the abuser to plan a time to harass the victim-survivor or to enter the house while the victim-survivor is away (R18). Even data from appliances and furniture can reveal a victim-survivor's movement: if the abuser has access to data from devices like a thermostat, light, fridge, or even a smart mattress, they can tell when the victim-survivor is arriving home for the day, heading to bed, or having a meal (e.g., W005, W036).

5.4 Accessing Other Private Data

Smart home devices collect a wealth of private information about their users, including a person's health, activities, habits, and more. Unfortunately, they may also provide abusers with this data. With access to the doorbell, an abuser can monitor visitors to the house, which can be dangerous for both the victim-survivor and the visitors. This is especially true when the abuser is motivated by suspicions of infidelity. Health tracking devices also hold very private information: smart watches capture heart rate and other fitness/movement stats, while smart mattresses capture data on sleep and other activity in the bed. Unfortunately, abusers can use this data maliciously. For example, they might track the activity on a watch or a mattress to monitor for signs of cheating: one woman "*caught an ex-boyfriend cheating when their synced exercise devices informed her that he was, erm, 'active' late at night* (W065). One smart mattress, the Smarttress, is even explicitly marketed for uncovering infidelity (W045).

Home control devices like smart speakers and routers hold even more information. Routers can enable abusers to monitor a victim-survivor's browsing activity (R11). An abuser might use this capability to find out if the victim-survivor is trying to seek help, for example. Smart speakers are additionally connected to Amazon, Google, or other accounts and therefore can have access to a victim-survivor's to-do lists, reminders, shopping lists, music, etc.—information that is typically private, but can be easily accessed by another person in the home through the speaker. For example, the author of web page W051 placed an Echo speaker in the living room he shared with roommates. Later, he learned they could access his calendar, reminders, etc., and could also manipulate this data (as we describe in § 6.3).

6 IoT-Enabled Interpersonal Harassment

Smart devices are used not just for spying but also for interpersonal harassment. As detailed in this section, harassment via smart devices can be psychologically taxing and even physically dangerous for victim-survivors.

6.1 Disrupting the Home Environment

The opportunities for harassment via smart home devices unfortunately seem limitless. We identified eleven devices, including nearly every smart appliance we identified, which can be used to harass victim-survivors by disrupting the smart home environment (Fig. 5). Our dataset contained two types of disruptions: those aimed at disturbing the victim-survivor's surroundings and those aimed at creating barriers for the victim-survivor.

Uncontrollable surroundings. We found several examples of abusers misusing smart home devices in order to disrupt the victim-survivor's environment. Abusers change the temperature with smart thermostats (R01), turn lights on and off

(W036), boil water with a smart kettle (W057), blast music over a speaker (W081), turn the TV on and off (W018), trigger alarms with a smoke alarm or lock (W108), and ring the doorbell from afar (W057), to name a few examples. In one illustrative anecdote, which we saw multiple times in our results (e.g., W018, W031), a woman's former partner used her smart home devices to blare music, flicker lights, and turn the TV on and off during the night to harass her. The partner was able to do this because he had set up the smart home system while he lived in the house.

Our results also document the use of smart speakers and AirTags to harass victim-survivors. Abusers can use smart speakers to broadcast their own voices to the home; in one incident, a woman remotely stalked her ex-partner through security cameras and yelled through the Amazon speaker in her ex-boyfriend's house to tell his new girlfriend to leave (W112). Additionally, abusers use AirTags as a harassment tool as well as a tool for surveillance. One web page (identified in earlier stages of our search) details how a plasterer hid an AirTag in a wall he was repairing, then later pinged the AirTag to annoy the client. These are examples of ways abusers take advantage of devices' secondary functionality to enact IPA. We discuss this and other *abuse vectors* in § 7.

Creating barriers. In addition to disrupting the victim-survivor's surroundings, smart home devices can also create barriers for the victim-survivor. Smart plug users can turn the plug on and off by using the associated app. With access to smart plugs in the victim-survivor's home, an abuser can prevent the victim-survivor from accessing power (W067). We also found abusers with access to a smart lock can easily change the locks to fully restrict the movement of the victim-survivor in or out of the house (W005, W011) or to exert control. Further, they could use the router to (remotely) block a victim-survivor from accessing the Internet or certain websites (R41). Without WiFi, some victim-survivors could not work, communicate with support systems, or seek help.

Along these lines, abusers who live with the victim-survivor often prevent victim-survivors from controlling smart devices, making them feel powerless in their own homes (W067). The design of many smart devices does not help. They often require a single administrator account, which allows one person (usually the abuser) to take control (W005). Combined with the environmental disruptions discussed above, this lack of control could be detrimental to an individual's well-being.

6.2 Threatening Physical Safety

Though the aforementioned strategies can take a deep emotional toll on a victim-survivor, we uncovered other harassment strategies with urgent safety implications. An abuser can make the home environment unsafe to live in by reducing the temperature to too low or too high, as shown in R01. An abuser with access to a smart lock can also remotely lock the victim-survivor in or out, which can be dangerous to the phys-

ical safety of the victim-survivor. The same abuser would also be able to bypass home security and enter the home of the victim-survivor at any time. So can an abuser with access to a smart garage door opener such as myQ [60], which allows the abuser to open the garage door remotely using a smartphone app (W047, R21).

Smart vehicles are also very dangerous under an abuser's control. With access to the victim-survivor's car or associated app (e.g., the Tesla app [82]), an abuser can control the car's temperature, horn, lights, and locking mechanisms. This access could allow the abuser to disrupt the car's internal environment while the victim-survivor is driving, potentially creating a dangerous level of distraction. Further, web pages mention that an abuser can exploit a car's computer system (à la [32]) or use external devices to gain access to critical systems like the brakes and the fuel supply (W108, W219). When we investigated whether the latter devices do exist, we found numerous remote kill switches for car batteries [3, 6] and even a device marketed for spying which can reportedly disable the car's ignition remotely, in addition to tracking its GPS location [79]. This type of access could give an abuser the power to prevent a victim-survivor from fleeing or, in the worst case, to cause an accident.

6.3 Manipulating Private Data

An abuser may use the access provided by a smart speaker to create unwanted reminders, appointments, or alarms. As W051 reports, "*I linked the Echo to my own Amazon account [...] As my roommates learned—once they discovered the new device after the living room TV accidentally triggered Alexa—they could speak any number of commands to the Echo to f*** with me.*" Abusers may also add things to a victim-survivor's shopping list, or worse, make unauthorized purchases that the victim-survivor must cover (W051, W099). Manipulating private data in these ways could only add to a victim-survivor's distress, and unauthorized purchases could exacerbate existing financial abuse.

6.4 Psychological Control

In tandem with the above methods of IoT-enabled harassment, abusers employ several psychological tactics to further manipulate victim-survivors. First, smart devices enable abusers to easily *gaslight* victim-survivors by "denial, misdirection, contradiction and lying to destabilize a victim-survivor" (W207). An abuser may gaslight the victim-survivor by repeatedly disrupting the home environment while pretending it is the victim-survivor's fault. For example, the abuser could change the code for the front door lock, persuade the victim-survivor that they simply forgot the right code, and change the code again the next day (W207). Another example of gaslighting is Q34, where the author writes, "*While organizing the house, I discovered that my boyfriend had hidden a webcam. This was after weeks of me having a feeling of being watched, feeling*

unsafe in my own home and him telling me I was just being paranoid.” This type of psychological abuse can make the victim-survivor question what they are experiencing and can take a substantial psychological toll.

An abuser may also use their abuse as a way to assert power over the victim-survivor. For instance, web page W005 reads, “If one person is the account administrator of everyday items like the heating system and washing machine, they can each be used as tools of coercion and control.” Further, abusers may use smart devices to collect evidence of a survivor’s drug use or extramarital relationships, compounding the abuse with potential legal implications (W139). These psychological tactics only exacerbate the abuse enabled by smart devices.

7 The Abuse Vectors Framework

Our analysis surfaced many troubling strategies for IoT-enabled spying and harassment. Here, we synthesize recurring themes in the ways abusers use smart devices to abuse others. Using three simple axes, we divide abuse strategies into a set of *abuse vectors*, which we visualize as a decision tree (Fig. 6). Our novel framework simplifies the complex ecosystem of IoT abuse and reveals a set of necessary solutions.

7.1 Axes of Consideration

Based on our findings, we divide abuse strategies along three axes: *covertness*, *ownership*, and *functionality*.

Covertness. Covertness refers to whether or not the victim is aware of the physical presence of the device used for abuse. We consider a device *covert* if the victim is unaware of the device’s presence or if the victim is aware the device exists but does not know its location. Anything else—the victim’s personal devices, any device the victim shares with others in the house, someone else’s personal device that the victim knows about (e.g., the abuser’s Apple Watch)—is *overt*.

Ownership. The abuser needs access to the IoT devices to carry out spying or harassment. However, abusers often access devices they do not own and, therefore, should not have control over them. We say that an abuser is an *owner* (or co-owner) of a device if it is one of the abuser’s personal devices or is a device shared with others in the abuser’s home.

Functionality. Although smart devices are typically designed with a *primary functionality*, such as capturing video feed (for cameras) or changing the temperature remotely (for thermostats), many have *secondary functionalities* that an abuser could exploit for abuse. To tell if an abuser is using a device’s primary functionality, we ask: would the device still be usable if we remove that functionality? For instance, consider an abuser who uses a security camera to capture video of someone. If we were to remove the video capture functionality of the security camera, it would no longer be usable. Thus, this is an example of the abuser (ab)using the device’s pri-




Figure 6: Decision tree of the four abuse vectors we observed.

mary functionality. An example of an abuser using secondary functionality is an abuser who uses Apple AirPods to track their victim’s location; if Apple removed the location tracking functionality, AirPods would still be able to fulfill their primary function as headphones. When abusers use a secondary functionality, we call it *repurposing*.

7.2 Abuse Vectors

Based on combinations of these three axes, we characterize the methods used for IPA in four *abuse vectors*. Similarly to an attack vector, an abuse vector denotes the mechanism or path used by the abuser to conduct the IPA. Based on the abuse methods we observed in our dataset, we define four abuse vectors: *Covert Spying* (C), *Unauthorized Access* (U), *Repurposing* (R), and *Intended Use* (I), named to indicate the dominant property of each vector.

Fig. 6 presents the four abuse vectors as a decision tree based on the three axes, while Fig. 5 shows the different vectors associated with each type of smart device. Importantly, depending on the situation, smart devices can be involved in several abuse vectors.

Covert Spying. A repeating theme in our results has been the difference between abuse via overt devices and abuse via covert devices. If abuse occurs using a covert device, we call this the *Covert Spying* vector. We found several instances of covert spying using devices designed for being stealthy such as hidden cameras, listening devices, and GPS trackers (as noted in Fig. 5), as well as devices not designed for spying but which can nevertheless be hidden (e.g., security cameras, baby monitors, and AirTags). One example of covert harassment came up in the early stages of our search—the abuser hid an AirTag in the victim’s wall, then pinged it repeatedly to annoy them [92]—but this was not a recurring theme in our results.

Unauthorized Access. When an abuser uses a device that is *not* covert, we have three possible vectors. The *Unauthorized Access* vector describes when the abuser utilizes a device they do not own and should not have access to at the time of the abuse. One example is R01, where the abuser moved out of their ex-partner’s house but was still able to manipulate her thermostat. The abuser might have been an owner of the device in the past—perhaps when the abuser lived in the house—but after the physical access to the house is “revoked,”

the digital access to the smart home should also be revoked.

Repurposing. The *Repurposing* vector describes when an abuser is an owner of an overt device and leverages a secondary functionality of the device (i.e., repurposes the device) for conducting IPA. This is a less common vector because, as we will discuss, many smart devices allow abusers to enact IPA using their primary functionalities. Examples of the *Repurposing* vector in action include using an Apple Watch as an audio and video surveillance device or using Echo Drop In to eavesdrop on private conversations.

Intended Use. The final abuse vector—and the trickiest to overcome—is *Intended Use*. This vector describes when an abuser uses an overt device that they own and uses the device’s primary functionality for IPA. We call this the *Intended Use* vector because in this vector, abusers are able to commit IoT abuse by using the device essentially as it was intended by the manufacturer (e.g., by using a thermostat to change the temperature). We observed copious examples of this vector: abusers using security cameras to surveil victims, using thermostats to turn the temperature too low, using an AirTag to track a victim, and more. This vector often occurs when the abuser has sole control over shared devices or brings a device into the home without the victim’s consent (e.g., R02), but it can also involve devices both parties agreed to purchase.

8 Discussion

Through a systematic crawl of online content, we identified a large set of IoT devices involved in abusive situations (Fig. 5) and the strategies abusers use to enact IoT abuse using these devices (§ 5 and § 6). In § 7, we defined four *abuse vectors* that characterize this vast problem. Moving forward, these vectors—*Covert Spying*, *Unauthorized Access*, *Repurposing*, and *Intended Use*—are a powerful lens to conceptualize IoT abuse and the interventions required to combat it.

One reason the abuse vectors are helpful is that they reveal which features of smart devices enable unique forms of IPA. First, many smart devices are uniquely integrated in the survivor’s physical space. This allows for an unprecedented level of harassment; abusers can control not only a person’s digital world, but the physical world they inhabit. Second, shared smart devices assume by default that all parties who share the device are non-adversarial, making it difficult to restrict access to abusers who live with the survivor or to revoke access when needed (hence, *Unauthorized Access*). Finally, smart devices are often designed to monitor (e.g., cameras, smart speakers, GPS trackers) or control surroundings (e.g., locks, thermostats, lights) for some legitimate purpose such as convenience. However, these exact capabilities can also be misused by an abuser in the *Intended Use* vector.

While our abuse vectors framework is useful on its own, it can also extend or be used in tandem with prior work. For instance, Slupska and Tanczer [74] demonstrate how IoT de-

vice manufacturers can identify opportunities for IoT abuse during the design process. Our framework clarifies which features may be potentially problematic; to evaluate whether their device enables the *Covert Spying* vector, for example, manufacturers should identify whether their device can be easily hidden. As another example, the abuse vectors can be used in tandem with Matthews et al.’s framework [57] and reveal further patterns: in their Life Apart phase, for example, an abuser is less likely to own devices within the survivor’s home, meaning the *Covert Spying* and *Unauthorized Access* vectors are most concerning.

Paths to mitigation. To tackle IoT abuse, we need a concrete plan on how to protect IPA survivors from all of these vectors. Here, we consider the fundamental challenges to mitigating each vector and lay a path for future work.

Though it was not a focus of our study, we noticed that several web pages included suggestions for mitigating IoT abuse. Appendix B summarizes the types of suggestions we found. More research should look into the accuracy and usefulness of online suggestions for combating IoT abuse.

8.1 Mitigating *Covert Spying*

To prevent covert spying, three directions are needed: (a) creating effective tools to detect hidden devices, (b) preventing the sale of spy devices on mainstream retailers, and (c) preventing dual-use devices like AirTags from being easily used for covert spying.

Design effective detection tools. The key feature of the *Covert Spying* vector is that the devices are hidden from the victim-survivor. As a result, the necessary solution for *Covert Spying* is an accurate and usable detection tool. Unfortunately, survivors currently have limited options for detection tools. Prior work has attempted to detect hidden cameras [17, 19, 72] and other hidden devices [64, 72]. However, these works often cover only one type of device (such as those which communicate over WiFi), and no studies have evaluated whether academic detection tools are usable for concerned survivors. Additionally, several of our search results discussed commercial tools for find hidden devices (e.g., using mobile apps or off-the-shelf devices; see Appendix B), but we are not aware of any research investigating their efficacy. Future studies should validate the usability and effectiveness of existing detection tools in the context of IPA.

Prevent the sale of spy devices. In the *Covert Spying* vector, abusers often utilize devices that were *designed* for spying. These devices are available for sale on websites dedicated to spying activities (e.g., [79]); however, several results (e.g., W001, W044) also link to dedicated spy devices being distributed on mainstream marketplaces such as Amazon. Many are marketed as a stealthy way to “Keep An Eye On Your Unfaithful Partner” (W044). Retailers should remove such products from their platforms as soon as possible and ensure that sellers do not advertise their devices for an IPA use

case in the future. The security research community could encourage this action by studying and raising alarms about the availability of spy devices on large online retailers.

Make dual-use devices discoverable. On the other hand, abusers leverage dual-use devices such as webcams and AirTags for covert surveillance. These dual-use devices often fail to “announce” their presence, meaning they can be hidden without much effort. Thus, vendors who do not intend their devices to be used for covert spying should provide ways to easily detect the presence of these devices in the vicinity and specify steps to disable them. For example, Apple—after public outcry—added a feature to AirTags to not only alert a user if they are being stalked but also a way to find the device by providing directional search [13, 39]. Manufacturers could follow Apple’s lead by implementing announcements for their devices which are noticeable to users (e.g., lights or sound) and/or invisible to users (e.g., Bluetooth packets).

Though they would provide clear benefits to survivors, these announcements must be designed with care. One issue is that the goal of making devices more discoverable may be at odds with the original purpose of the device, such as theft prevention. When this is the case, users may devise creative methods to get around device announcements; for example, in response to Apple’s anti-stalking features, users started selling modified AirTags with their speakers disabled, which muted the audible safety alerts [44, 54]. Further, though, the manufacturers themselves may undermine their own anti-stalking protections. In 2022, Tile introduced their Scan and Secure feature [85], which allows users to scan for unknown Tiles nearby; then, in 2023, Tile added Anti-Theft mode [86], which users can enable to *prevent* their Tile device from showing up in Scan and Secure. As this saga illustrates, manufacturers must strike a difficult balance between protecting users from the *Covert Spying* vector and preserving the intended use case of these devices. Tile’s approach, which requires that Anti-Theft users must provide identification and agree to share personal information with law enforcement if stalking is suspected, indicates that the way forward may be to enable anti-stalking features by default and introduce hurdles for disabling these features.

8.2 Mitigating Unauthorized Access

In the *Unauthorized Access* vector, abusers leverage access to a device they should not be able to control. Thus, this vector necessitates changes to access control protocols on IoT devices. The first step is making it easier for users to recognize a need to revoke access and to perform the revocation. In tandem, legislation should be updated to support victim-survivors in this situation.

Make access revocation easier. In the current paradigm of IoT access control, shared devices remain shared unless one user takes action to revoke another’s access. Unfortunately, users do not always realize that they are responsible for revok-

ing access. For example, in R01, when the abuser moved out of their ex-partner’s house, the ex-partner did not revoke the abuser’s access to the thermostat. Better consumer education on proper access control and authentication procedures will help with this. More practically, though, researchers should investigate how devices could reduce the burden on the individual by detecting when access might need to be revoked, perhaps by identifying users who have not been present in the house for a certain amount of time. We foresee studies which collaborate with users to identify (1) how to tell when a user should no longer have access to smart devices and (2) how best to raise this concern with the user.

Once users identify the need to revoke access, it should be easy for them to do so. Though we did not find evidence that current revocation mechanisms are hard to use, vendors should ensure that users can easily manage access to their devices by providing usable, fine-grained device-sharing mechanisms. For example, August Lock provides the option to invite users, set different access levels, and easily revoke access when needed [7]. If these sharing mechanisms are not provided, users may use less controllable methods, such as sharing the credentials for a single account.

Provide legal options. Even when survivors take the correct steps to revoke access, their attempts can be futile. For example, one story depicts a man who changed the password on his Ring Doorbell twice, only to find out that his ex-partner could still access the device (W080). Amazon claimed that access should have been revoked after an hour, but this did not happen. In these situations, legislation can provide another way for survivors to take back control. For example, in New York State, courts can now order an abuser to “refrain from remotely controlling any connected devices affecting the home, vehicle or property of the person protected by the order” [1]. However, this provision is rarely used in practice [46]. Security researchers need to collaborate with legislators to (a) encourage the use of this provision in New York State and (b) advocate for similar legislation in other states.

8.3 Mitigating Repurposing

To prevent abusers from repurposing a device’s secondary functionality, device manufacturers need to identify ways that their device could be misused, then close those gaps. For example, to prevent users from abusing Apple Watch’s Camera Remote function, Apple could show a clear notification on the phone and show the camera feed for only a short time on the Watch. Moving forward, manufacturers should pinpoint risky auxiliary functionalities in the design stage before a device is released to customers. The research community could spur this effort by studying potentially harmful secondary functionalities in existing devices.

8.4 Mitigating *Intended Use*

The *Intended Use* vector is a difficult problem to solve. Manufacturers cannot just identify and remove the functionality used for abuse, as with the *Repurposing* vector, nor can victim-survivors simply revoke an abuser's access. Mitigations should include new access controls that prevent a single person from monopolizing control over devices in the home and designing notifications and logs for smart devices that reveal abusive activity.

Guarantee access to all users in the home. One potential solution is to make it more difficult to operate smart devices without the consent and control of all users in the house. Towards this goal, Zeng and Roesner [95] prototyped a system that enables fine-grained access control, such as the option to restrict users from controlling smart devices unless they are physically near the device. However, as Zeng and Roesner emphasize, access controls in an adversarial environment are inherently dual-use—the same capabilities that could prevent an abuser from remotely manipulating smart devices could also enable an abuser to assert control by restricting the victim-survivor's power over devices [95]. Future iterations of these protocols will need to be designed with this risk in mind.

Improve notifications. Smart devices must provide clear and understandable notifications about their activity, especially when they are collecting data and sharing that data live with a remote party (e.g., an abuser). Currently, there is no industry standard for such notifications. A particularly concerning example is Amazon Echo's Drop In feature. To indicate Drop In has started, the device plays a quick chime sound, and the lights turn green, but it does not say anything about who is "dropping in." Users unaware of the feature would be confused about what is going on and might therefore miss the subtle notification [63] or even brush it off as just another erroneous trigger [71]. Researchers should work to design notifications for smart devices which are clear and useful without overwhelming the user.

Create easily accessible logs. The aforementioned notifications might not be feasible for some IoT devices, such as smart plugs or kettles. To provide transparency, these devices must keep an activity log and allow authenticated users to view the logs easily. Currently, smart devices have minimal logging—for example, Google's Nest Thermostat keeps detailed logs, including the users who initiated changes, but only for the past 10 days [31, 52]. Other devices like Philips Hue light bulbs have no easily-accessible logs, leaving some users to devise their own logging scripts [8, 87]. Victim-survivors of IoT abuse are therefore left with few resources when investigating and understanding the abuse they are experiencing.

To best help survivors of IPA, the security community should work to design effective logs for smart devices. These logs could be a lifeline for a victim-survivor of IPA by providing concrete evidence of abuse to use in legal proceedings or simply by validating the victim-survivor's experience. How-

ever, on the other hand, the transparency provided by activity logging capabilities could enable an abuser with another way to track the victim-survivor's presence and activity in the house. Researchers must work to design activity logs for smart devices which provide transparency to victim-survivors while minimizing the potential for abusers to misuse the logs for further surveillance.

8.5 Additional Mitigations

As we have discussed, many smart devices can easily be repurposed for abuse. Frustratingly, it is likely that these devices' potential for abuse would have been obvious to victim-survivors had they been consulted in the design process. Thus, while designing smart devices, it is integral that device manufacturers consider an IPA threat model (such as the one created by Slupska and Tanczer [74]) and consult IPA victim-survivors and advocates in order to carefully evaluate the risk of their devices being used to harm others. This paper illustrates what can happen when manufacturers do not appropriately consider the threat of IPA during design.

In addition, our results stressed the importance of individual IoT knowledge in detecting and preventing IoT abuse (e.g., W057, W067). Toward creating awareness of the capabilities and risks of IoT devices, policymakers should enforce that device manufacturers release a statement about what harms the device could cause in the context of IPA and what steps they have taken to reduce that harm. These statements could be appended to the existing concept of privacy nutrition labels [24]. This way, individuals can be well-informed about how their devices may be misused.

9 Conclusion

Smart devices provide convenience, efficiency, security, and other benefits to users and thus are gaining fast adoption in the US and worldwide. But, unfortunately, as our findings show, they also enable a new class of tech-enabled interpersonal abuse (IPA). Through a systematic study of web content, we uncovered evidence that 32 types of smart devices are being used to spy and harass IPA victim-survivors, often in surprising ways. Then, we defined four *abuse vectors*—*Covert Spying*, *Unauthorized Access*, *Repurposing*, and *Intended Use*—which characterize the ways IoT devices are being used for interpersonal surveillance and harassment. We outlined the primary challenges and potential mitigations for each vector in our framework. Tackling this new variant of tech-enabled IPA will require a multi-pronged approach from many stakeholders, including policymakers, device manufacturers, and, importantly, the research community. It is our responsibility to work to mitigate IoT-enabled IPA and make smart devices safer for everyone.

Acknowledgements

We thank the anonymous reviewers for their insightful feedback, which helped improve the paper significantly. This work was supported in part by the University of Wisconsin—Madison Office of the Vice Chancellor for Research and Graduate Education with funding from the Wisconsin Alumni Research Foundation, as well as a Consumer Reports Digital Fellowship. Additional thanks to Nick Ceccio, Blaine Hoak, Mazharul Islam, and Max Zinkus for their help.

References

- [1] Criminal procedure law (CPL) 530.12: Protection for victims of family offenses. The New York State Senate. URL: <https://www.nysenate.gov/legislation/laws/CPL/530.12>.
- [2] Majed Almansoori, Andrea Gallardo, Julio Poveda, Adil Ahmed, and Rahul Chatterjee. A global survey of android dual-use applications used in intimate partner surveillance. *Proceedings on Privacy Enhancing Technologies*, 4:120–139, 2022.
- [3] Amazon. Car wireless remote control battery switch disconnect latching relay anti-theft, e-kylin dc 12v electromagnetic solenoid valve terminal master kill system, 2022. URL: <https://amzn.to/3yRMBBK>.
- [4] Amazon. How does drop in work?, 2022. URL: <https://www.amazon.com/gp/help/customer/display.html?nodeId=GS3WRTSRKD2U6MCK>.
- [5] Amazon. How to change your wake word, 2022. URL: <https://www.amazon.com/b?ie=UTF8&node=21341305011>.
- [6] Amazon. Upgraded dc12v 230a kill switch for car, anti-theft, remote battery disconnect switch, prevent battery drain, 2022. URL: <https://amzn.to/3ynH04K>.
- [7] August. How to invite guests to your smart lock, 2021. URL: https://support.august.com/how-do-i-invite-guests-r1BR1IJA_M.
- [8] Wesley Baugh. hue-log. Github repository, 2016. URL: <https://github.com/bwbaugh/hue-log>.
- [9] Genevieve Bell and Paul Dourish. Yesterday’s tomorrows: notes on ubiquitous computing’s dominant vision. *Personal and ubiquitous computing*, 11:133–143, 2007.
- [10] Brad Boserup, Mark McKenney, and Adel Elkbuli. Alarming trends in us domestic violence during the covid-19 pandemic. *The American journal of emergency medicine*, 38(12):2753–2755, 2020.
- [11] Nellie Bowles. Thermostats, locks and lights: Digital tools of domestic abuse. The New York Times, Jun 2018. URL: <https://www.nytimes.com/2018/06/23/technology/smart-home-devices-domestic-abuse.html>.
- [12] Kaitlin M Boyle and Kimberly B Rogers. Beyond the rape “victim”–“survivor” binary: How race, gender, and identity processes interact to shape distress. In *Sociological Forum*, volume 35, pages 323–345. Wiley Online Library, 2020.
- [13] Kellen Browning. Apple says it will make airtags easier to find after complaints of stalking. The New York Times, February 2022. URL: <https://www.nytimes.com/2022/02/10/business/apple-airtags-safety.html>.
- [14] Albert Fox Cahn. Apple’s airtags are a gift to stalkers. Wired, May 2021. URL: <https://www.wired.com/story/opinion-apples-air-tags-are-a-gift-to-stalkers/>.
- [15] Steven Carr. How far do airpods reach? Decor Tweaks, 2022. URL: <https://decortweaks.com/how-far-do-airpods-reach/>.
- [16] Marta Cecchinato and Daniel Harrison. Degrees of agency in owners and users of home iot devices. In *CHI’17 workshop: Making Home: Asserting Agency in the Age of IoT*. Association for Computing Machinery (ACM), 2017.
- [17] Anadi Chaman, Jiaming Wang, Jiachen Sun, Haitham Hassanieh, and Romit Roy Choudhury. Ghostbuster: Detecting the presence of hidden eavesdroppers. In *Proceedings of the 24th Annual International Conference on Mobile Computing and Networking*, pages 337–351, 2018.
- [18] Rahul Chatterjee, Periwinkle Doerfler, Hadas Orgad, Sam Havron, Jackeline Palmer, Diana Freed, Karen Levy, Nicola Dell, Damon McCoy, and Thomas Ristenpart. The spyware used in intimate partner violence. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 441–458. IEEE, 2018.
- [19] Yushi Cheng, Xiaoyu Ji, Tianyang Lu, and Wenyuan Xu. Dewicam: Detecting hidden wireless cameras via smartphones. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*, pages 1–13, 2018.
- [20] Crummy. Beautiful soup: We called him tortoise because he taught us., 2022. URL: <https://www.crummy.com/software/BeautifulSoup/>.
- [21] DuckDuckGo. Duckduckgo — privacy, simplified., 2022. URL: <https://duckduckgo.com/>.
- [22] Lisa Eadicicco. Airtags are linked to stalking, and apple can’t solve this problem alone. CNet, February 2022. URL: <https://www.cnet.com/tech/mobile/airtags-are-being-linked-to-stalking-and-its-a-problem-apple-cant-solve-alone/>.
- [23] Melanie Ehrenkranz. Ring’s smart doorbell let a man spy on his ex-boyfriend—even after the password was changed. Gizmodo, May 2018. URL: <https://gizmodo.com/rings-smart-doorbell-let-a-man-spy-on-his-ex-boyfriend-1825963112>.
- [24] Pardis Emami-Naeini. Privacy and security nutrition labels to inform IoT consumers. *USENIX Association*, 2021.
- [25] Diana Freed, Sam Havron, Emily Tseng, Andrea Gallardo, Rahul Chatterjee, Thomas Ristenpart, and Nicola Dell. “Is my phone hacked?” Analyzing clinical computer security interventions with survivors of intimate partner violence. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–24, 2019.
- [26] Diana Freed, Jackeline Palmer, Diana Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. “A stalker’s paradise”: How intimate partner abusers exploit technology. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pages 1–13, 2018.
- [27] Diana Freed, Jackeline Palmer, Diana Elizabeth Minchala, Karen Levy, Thomas Ristenpart, and Nicola Dell. Digital technologies and intimate partner violence: A qualitative analysis with multiple stakeholders. *Proceedings of the ACM on human-computer interaction*, 1(CSCW):1–22, 2017.
- [28] Manuel Gámez-Guadix, Erika Borrajo, and Esther Calvete. Partner abuse, control and violence through internet and smartphones: Characteristics, evaluation and prevention. *Papeles del Psicólogo*, 2018.
- [29] Christine Geeng and Franziska Roesner. Who’s in control? Interactions in multi-user smart homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2019.
- [30] Eric Goldman. Search engine bias and the demise of search engine utopianism. In *Web Search*, pages 121–133. Springer, 2008.
- [31] Google. Thermostat energy history. Google Nest Help, 2022. URL: <https://support.google.com/googlenest/answer/9247300?#zippy=%2Cnest-app>.
- [32] Andy Greenberg. Hackers remotely kill a jeep on the highway—with me in it. Wired, July 2015. URL: <https://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/>.

- [33] E-ZPass Group. E-ZPass group - home, 2022. URL: <https://www.ezpassag.com/>.
- [34] Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart. Clinical computer security for victims of intimate partner violence. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 105–122, 2019.
- [35] Danny Yuxing Huang, Maxwell Matthaios Aliapoulos, Vector Guo Li, Luca Invernizzi, Elie Bursztein, Kylie McRoberts, Jonathan Levin, Kirill Levchenko, Alex C. Snoeren, and Damon McCoy. Tracking ransomware end-to-end. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 618–631, 2018. doi:10.1109/SP.2018.00047.
- [36] Danny Yuxing Huang, Noah Aporthe, Frank Li, Gunes Acar, and Nick Feamster. IoT Inspector: Crowdsourcing labeled network traffic from smart home devices at scale. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(2):1–21, 2020.
- [37] Apple Inc. AirTag, 2022. URL: <https://www.apple.com/airtag/>.
- [38] Apple Inc. Find my network, 2022. URL: <https://developer.apple.com/find-my/>.
- [39] Apple Inc. Locate an airtag or other item in find my on iphone, 2022. URL: <https://support.apple.com/guide/iphone/locate-an-item-iph79f0c10/ios>.
- [40] Apple Inc. Use camera remote and timer on apple watch, 2022. URL: <https://support.apple.com/guide/watch/camera-remote-apda6e61c287/watchos>.
- [41] Apple Inc. Use live listen with airpods or beats, 2022. URL: <https://support.apple.com/en-us/HT209082>.
- [42] Apple Inc. Use walkie-talkie on your apple watch, 2022. URL: <https://support.apple.com/en-us/HT208917>.
- [43] Joseph Johnson. Global market share of search engines 2010–2022. Statista, March 2022. URL: <https://www.statista.com/statistics/216573/worldwide-market-share-of-search-engines/>.
- [44] Michael Kan. ‘Silent AirTags’ with speakers removed pop up on Etsy, eBay. PCMag, 2022. URL: <https://www.pcmag.com/news/silent-airtags-with-speakers-removed-pop-up-on-etsy-ebay>.
- [45] Jeffrey Kluger. Domestic violence is a pandemic within the covid-19 pandemic. Time, February 2021. URL: <https://time.com/5928539/domestic-violence-covid-19/>.
- [46] Anna Kodé. Is your smart home controlling you? The New York Times, February 2023. (Accessed on 02/20/2023). URL: <https://www.nytimes.com/2023/02/17/realestate/smart-home-devices.html>.
- [47] Srijan Kumar and Neil Shah. False information on web and social media: A survey. *arXiv preprint arXiv:1804.08559*, 2018.
- [48] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. Alexa, are you listening? privacy perceptions, concerns and privacy-seeking behaviors with smart speakers. *Proceedings of the ACM on Human-Computer Interaction*, 2(CSCW):1–31, 2018.
- [49] Roxanne Leitão. Anticipating smart home security and privacy threats with survivors of intimate partner abuse. In *Proceedings of the 2019 on Designing Interactive Systems Conference*, pages 527–539, 2019.
- [50] Michael Levitt. Airtags are being used to track people and cars. here’s what is being done about it. NPR, February 2022. URL: <https://www.npr.org/2022/02/18/1080944193/apple-airtags-theft-stalking-privacy-tech>.
- [51] Karen Levy and Bruce Schneier. Privacy threats in intimate relationships. *Journal of Cybersecurity*, 6(1), 2020.
- [52] Craig Lloyd. How to view your nest thermostat’s usage history. How-To Geek, May 2016. URL: <https://www.howtogeek.com/251340/how-to-view-your-nest-thermostat%E2%80%99s-usage-history/>.
- [53] Isabel Lopez-Neira, Trupti Patel, Simon Parkin, George Danezis, and Leonie Tanczer. ‘Internet of things’: How abuse is getting smarter. *Safe – The Domestic Abuse Quarterly*, (63):22–26, 2019.
- [54] Ben Lovejoy. AirTags with deactivated speakers being sold on eBay and Etsy; seller claims not for stalking. 9to5Mac, 2022. URL: <https://9to5mac.com/2022/02/03/airtags-with-deactivated-speakers-being-sold/>.
- [55] Ryan Mac and Kashmir Hill. Are apple airtags being used to track people and steal cars? The New York Times, December 2021. URL: <https://www.nytimes.com/2021/12/30/technology/apple-airtags-tracking-stalking.html>.
- [56] Chuck Martin. Smart home technology hits 69% penetration in U.S. MediaPost, Sept 2019. URL: <https://www.mediapost.com/publications/article/341320/smart-home-technology-hits-69-penetration-in-us.html>.
- [57] Tara Matthews, Kathleen O’Leary, Anna Turner, Manya Sleeper, Jill Palzkill Woelfer, Martin Shelton, Cori Manthorne, Elizabeth F Churchill, and Sunny Consolvo. Stories from survivors: Privacy & security practices when coping with intimate partner abuse. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 2189–2201, 2017.
- [58] James A Mercy, Susan D Hillis, Alexander Butchart, Mark A Bellis, Catherine L Ward, Xiangming Fang, and Mark L Rosenberg. Interpersonal violence: global impact and paths to prevention. *Injury Prevention and Environmental Health*. 3rd edition, 2017.
- [59] Chuck Morris. Man arrested for attaching apple watch to girlfriend’s car. WSMV4, March 2022. URL: <https://www.wsmv.com/2022/03/27/man-arrested-attaching-apple-watch-girlfriends-car/>.
- [60] myQ. Smart garage hub, 2022. URL: <https://www.myq.com/smart-garage-control>.
- [61] Alfred Ng. Amazon alexa transcripts live on, even after you delete voice records. CNet, May 2019. URL: <https://www.cnet.com/home/smart-home/amazon-alexa-transcripts-live-on-even-after-you-delete-voice-records/>.
- [62] Simon Parkin, Trupti Patel, Isabel Lopez-Neira, and Leonie Tanczer. Usability analysis of shared device ecosystem security: Informing support for survivors of iot-facilitated tech-abuse. In *Proceedings of the new security paradigms workshop*, pages 1–15, 2019.
- [63] Rebecca S Portnoff, Linda N Lee, Serge Egelman, Pratyush Mishra, Derek Leung, and David Wagner. Somebody’s watching me? assessing the effectiveness of webcam indicator lights. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 1649–1658, 2015.
- [64] Anthony Rowe Rahul Anand Sharma, Elahe Soltanaghaei and Vyas Sekar. Lumos: Identifying and localizing diverse hidden IoT devices in an unfamiliar environment. In *31st USENIX Security Symposium (USENIX Security 22)*, Boston, MA, August 2022. USENIX Association.
- [65] Jingjing Ren, Daniel J Dubois, David Choffnes, Anna Maria Mandalari, Roman Kolcun, and Hamed Haddadi. Information exposure from consumer iot devices: A multidimensional, network-informed measurement approach. In *Proceedings of the Internet Measurement Conference*, pages 267–279, 2019.
- [66] Ignacio Rodríguez-Rodríguez, José-Víctor Rodríguez, Aránzazu Elizondo-Moreno, Purificación Heras-González, and Michele Gentili. Towards a holistic ict platform for protecting intimate partner violence survivors based on the iot paradigm. *Symmetry*, 12(1):37, 2019.
- [67] Nicholas C Romano Jr, Christina Donovan, Hsinchun Chen, and Jay F Nunamaker Jr. A methodology for analyzing web-based qualitative data. *Journal of Management Information Systems*, 19(4):213–246, 2003.

- [68] Kevin A Roundy, Paula Barmaimon Mendelberg, Nicola Dell, Damon McCoy, Daniel Nissani, Thomas Ristenpart, and Acar Tamersoy. The many kinds of creepware used for interpersonal attacks. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 626–643. IEEE, 2020.
- [69] Johnny Saldaña. *The coding manual for qualitative researchers*. SAGE publications Ltd, 2021.
- [70] Margarete Sandelowski. Whatever happened to qualitative description? *Research in nursing & health*, 23(4):334–340, 2000.
- [71] Lea Schönherr, Maximilian Golla, Thorsten Eisenhofer, Jan Wiele, Dorothea Kolossa, and Thorsten Holz. Exploring accidental triggers of smart speakers. *Computer Speech & Language*, 73:101328, 2022.
- [72] Akash Deep Singh, Luis Garcia, Joseph Noor, and Mani Srivastava. I always feel like somebody’s sensing me! a framework to detect, identify, and localize clandestine wireless sensors. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1829–1846, 2021.
- [73] Julia Slupska and Angelika Strohmayer. Networks of care: Tech abuse advocates’ digital security practices. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 341–358, 2022.
- [74] Julia Slupska and Leonie Maria Tanczer. Threat modeling intimate partner violence: Tech abuse as a cybersecurity challenge in the internet of things. In *The Emerald International Handbook of Technology Facilitated Violence and Abuse*. Emerald Publishing Limited, 2021.
- [75] Sharon G Smith, Kathleen C Basile, Leah K Gilbert, Melissa T Merrick, Nimesh Patel, Margie Walling, and Anurag Jain. National intimate partner and sexual violence survey (nisvs): 2010–2012 state report. 2017. URL: <https://stacks.cdc.gov/view/cdc/46305>.
- [76] Sharon G. Smith, Xinjian Zhang, Kathleen C. Basile, Melissa T. Merrick, Jing Wang, Marcie jo Kresnow, and Jieru Chen. The national intimate partner and sexual violence survey (nisvs): 2015 data brief – updated release. Atlanta, GA: National Center for Injury Prevention and Control, Centers for Disease Control and Prevention, 2018. URL: <https://www.cdc.gov/violenceprevention/pdf/2015data-brief508.pdf>.
- [77] Open Planet Software. Just press record on the app store. App Store, 2022. URL: <https://apps.apple.com/us/app/just-press-record-voice-audio/id1033342465>.
- [78] Software Freedom Conservancy. Selenium, 2022. URL: <https://www.selenium.dev/>.
- [79] Spy Spot. Spy spot 4g hard wire kill switch gps vehicle tracker - remotely disable the ignition from any location - locator tracking device - black, 2 x 1.8 x 1 inches - us coverage, subscription required, 2022. URL: <https://spy-spot.com/new-4g-lte-hardwired-gps-tracker-with-remote-starter-disable-vehicle-kill-switch/>.
- [80] Leonie Tanczer, Isabel Lopez-Neira, Simon Parkin, Trupti Patel, and George Danezis. Gender and iot research report: The rise of the internet of things and implications for technology-facilitated abuse. STEaPP, 2018.
- [81] Leonie Maria Tanczer, Isabel López-Neira, and Simon Parkin. ‘I feel like we’re really behind the game’: perspectives of the united kingdom’s intimate partner violence support sector on the rise of technology-facilitated abuse. *Journal of gender-based violence*, 5(3):431–450, 2021.
- [82] Tesla. Mobile app, 2022. URL: https://www.tesla.com/ownersmanual/model3/en_jo/GUID-F6E2CD5E-F226-4167-AC48-BD021D1FFDAB.html.
- [83] Kurt Thomas, Devdatta Akhawe, Michael Bailey, Dan Boneh, Elie Bursztein, Sunny Consolvo, Nicola Dell, Zakir Durumeric, Patrick Gage Kelley, Deepak Kumar, et al. Sok: Hate, harassment, and the changing landscape of online abuse. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 247–267. IEEE, 2021.
- [84] Tile. Find your keys, wallet & phone with tile’s app and bluetooth tracker device, 2022. URL: <https://www.thetileapp.com/>.
- [85] Tile. Use tile’s scan & secure feature to stay safe from unwanted tracking, 2022. (Accessed on 02/17/2023). URL: <https://www.tile.com/en-us/blog/tile-introduces-scan-and-secure-feature-unwanted-tracking-safety>.
- [86] Tile. What is anti-theft mode?, 2023. (Accessed on 02/17/2023). URL: <https://www.tile.com/blog/tile-anti-theft-mode>.
- [87] tomh. How to send philips hue activity data to a google sheet. Stack-Overflow, 2018. URL: <https://stackoverflow.com/questions/52611535/how-to-send-philips-hue-activity-data-to-a-google-sheet>.
- [88] Emily Tseng, Rosanna Bellini, Nora McDonald, Matan Danos, Rachel Greenstadt, Damon McCoy, Nicola Dell, and Thomas Ristenpart. The tools and tactics used in intimate partner surveillance: An analysis of online infidelity forums. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 1893–1909, 2020.
- [89] Emily Tseng, Diana Freed, Kristen Engel, Thomas Ristenpart, and Nicola Dell. A digital safety dilemma: Analysis of computer-mediated computer security interventions for intimate partner violence during covid-19. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–17, 2021.
- [90] Lionel Sujay Vailshery. Internet of things (iot) and non-iot active device connections worldwide from 2010 to 2025. Statista, March 2021. URL: <https://www.statista.com/statistics/1101442/iot-number-of-connected-devices-worldwide>.
- [91] Liwen Vaughan and Mike Thelwall. Search engine coverage bias: evidence and possible causes. *Information processing & management*, 40(4):693–707, 2004.
- [92] Albert Wagner. Tradie shares his jaw-dropping revenge on rude client using apple airtag in new reddit post. Verve Times, July 2021. URL: <https://vervetimes.com/tradie-shares-his-jaw-dropping-revenge-on-rude-client-using-apple-airtag-in-new-reddit-post/>.
- [93] Rob Waugh. Spurned husband uses wi-fi thermostat to take lingering, ghostly revenge, Nov 2014. URL: <https://metro.co.uk/2014/11/05/spurned-husband-uses-wi-fi-thermostat-to-take-lingering-ghostly-revenge-4935919/>.
- [94] 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.
- [95] 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.
- [96] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. User perceptions of smart home IoT privacy. *Proceedings of the ACM on human-computer interaction*, 2(CSCW):1–20, 2018.

A Additional Details of Classifier

Training data. Our training data consisted of (a) the 100 most popular results in our sample and (b) a random sample of 200 results, all hand-labeled for relevance. Of these 300 results, 26 had broken URLs or otherwise rejected automatic analysis; we used the other 274 labeled results to train a Random Forest classifier. 23% of the training data was relevant.

Feature selection. When creating feature sets for our results, we first extracted the text from the HTML source code,

then cleaned the text by removing white space, empty lines, numbers, URLs, and punctuation. We then used a normalized bag-of- n -grams model to represent each page. Using feature importance testing with Naive Bayes, we identified the 100 most important n -grams ($1 \leq n \leq 3$) in the training data (e.g., “camera hidden”, “hidden spy”, and “husband cheating”). We combined these 100 features with manually-selected features: 48 smart devices found during preliminary coding, 888 IoT devices from [36], 8 IPV keywords (e.g., “spy”, “eavesdrop”, “snoop”), and flags for the top ten domains. In all, we used 1054 features.

Training the classifier. We trained binary classification models using multiple algorithms, such as Linear Regression, Decision Tree, and Naive Bayes; the Random Forest model performed the best overall. We iteratively tuned the class weights and desired threshold to find a combination that had the desired recall ($> 95\%$) while maintaining good precision; this way, we would be sure to capture most of the true relevant results. Our final classifier used a threshold of 0.12 (as opposed to 0.5) and class weights of 4:1 (relevant:irrelevant). The training and testing accuracy were 89% and 80% respectively, whereas the precision and recall on the test data were 58% and 93% (respectively).

B Suggestions on the Web for Mitigating IoT-Enabled Abuse

In our search, we found several articles presenting ideas for how victims and stakeholders can deal with abuse via smart devices. Here, we highlight mitigation strategies victims might find online and discuss how, in general, they are not very effective. We remind the reader that though these anti-IoT-abuse strategies are a rich finding, we did not explicitly search for such results in our web crawl. Thus, a dedicated study is needed to further explore this angle.

Detecting surveillance devices. Twelve search results center around finding hidden spy devices (e.g., W006, W042). A victim searching for help on these pages would find suggestions like manually searching the environment (R006), using apps designed for detecting spy devices (W006, W052), using RF detectors or high-end hidden camera detectors (W048, W052), or even hiring professional contractors (W123). Detection is quite a challenging problem, and these suggestions are unfortunately not a satisfying solution. Manual searching is unlikely to be effective due to the miniature size and camouflaged appearance of many covert spy devices and may do more harm than good by instilling paranoia. Further, prior work has not confirmed the effectiveness of existing apps or tools designed for detecting hidden devices. If the devices are faulty, they may raise false alarms, or worse, leave victims with a false sense of security.

Resetting devices. To remove an abuser’s remote control

over a device, many results say that victims should simply unplug the device (W008). While this is an effective solution, it also prevents the victim from using the device at all, which is not an option for core devices in the home like locks and thermostats. It will also be easy for an abuser to notice, potentially inspiring escalated abuse (W011). Alternatively, 12 results suggest resetting devices to factory defaults (e.g., W024). Resetting devices alone will not solve a victim’s problems in most cases. If the abuser is connected directly to the device and not an account (e.g., if the abuser has connected to the victim’s AirPods), resetting the device should mitigate the issue. However, oftentimes the abuser’s access is via the account associated with that device; it is possible they know the credentials or that they were added as a user on that account (e.g., by being invited as a guest on an August Lock [7]). In this case, resetting the device will be futile unless the victim *also* revokes this access—for example, by changing the account password—or uses a different account altogether. Fortunately, several results note the importance of these types of account changes (e.g., W047).

Changing the WiFi. Since devices are usually connected via the home WiFi network, changing the WiFi name (SSID) and/or password would remove the devices from the network. This will likely disable their functionality (R18), including any hidden or malicious WiFi-connected devices. However, this also means a victim would have to reconnect all of their smart devices back to the WiFi, which can be frustrating and even dangerous if some safety-critical smart devices—such as a smart lock or Ring Doorbell, which are recommended for a victim’s safety after separation (W126, [66])—are disconnected. It is also possible that the device being used for abuse is not WiFi-connected (e.g., AirTags, AirPods, and most covert spy devices), meaning changing the WiFi is unhelpful.

Changing device settings. More subtle changes involve reducing the device’s functionality by, for example, muting the microphones (W002). These types of changes walk the line between usability and “abusability”; for instance, muting a device’s microphone indeed prevents audio surveillance, but it also removes the user’s ability to use voice commands. For smart speakers, some results suggest victims should delete the command history (W008) or change the device’s wake word (W051). Deleting the command history regularly can be cumbersome and diminish the benefit of having a smart speaker. Changing the wake word—to one of the few possible wake words available to choose from [5]—may only temporarily disrupt an abuser in the house.