

Sneaky Spy Devices and Defective Detectors: The Ecosystem of Intimate Partner Surveillance with Covert Devices*

Rose Ceccio[†], Sophie Stephenson[†], Varun Chadha[‡], Danny Yuxing Huang[◇], Rahul Chatterjee[†]
[†] University of Wisconsin—Madison, [‡] Capital One, [◇] New York University

Abstract

Recent anecdotal evidence suggests that abusers have begun to use covert spy devices such as nanny cameras, item trackers, and audio recorders to spy on and stalk their partners. Currently, it is difficult to combat this type of *intimate partner surveillance* (IPS) because we lack an understanding of the prevalence and characteristics of commercial spy devices. Additionally, it is unclear whether existing devices, apps, and tools designed to *detect* covert devices are effective. We observe that many spy devices and detectors can be found on mainstream retailers. Thus, in this work, we perform a systematic survey of spy devices and detection tools sold through popular US retailers. We gather 2,228 spy devices, 1,313 detection devices, and 51 detection apps, then study a representative sample through qualitative analysis as well as in-lab evaluations.

Our results show a bleak picture of the IPS ecosystem. Not only *can* commercial spy devices easily be used for IPS, but many of them are *advertised* for use in IPS and other covert surveillance. On the other hand, commercial detection devices and apps are all but defective, and while recent academic detection systems show promise, they require much refinement before they can be useful to survivors. We urge the security community to take action by designing practical, usable detection tools to detect hidden spy devices.

1 Introduction

Millions of people per year in the United States experience *intimate partner violence* (IPV) [57]. Increasingly, technology has been involved in IPV [19, 40, 41]. Abusers install spyware apps, send harassing messages, and take control of online accounts to spy on, intimidate, and isolate survivors [15, 23, 65]. In a particularly worrying trend, abusers are using hidden surveillance devices such as cameras [2, 13], audio recorders [29], and location trackers [52] to spy on survivors. This *intimate partner surveillance* (IPS) often involves



Figure 1: A photograph from a listing for a GPS tracker. The tracker explicitly claims to “catch cheating spouse.”

devices designed for spying, but it can also involve “dual-use” devices which are *not* intended to be used for spying, yet inadvertently provide such functionalities. For instance, abusers hide AirTags in purses and vehicles to stalk intimate partners as well as strangers [14, 34, 39]. In addition to obvious violations of privacy, this unwanted tracking can escalate to violence, even resulting in homicide [21].

Unfortunately, we find that devices intended for spying, as well as *hideable* dual-use devices, are available for purchase on large, online retailers in the US. Several online articles promoting the use of hidden surveillance devices for IPS and include links to spy devices for sale online, many for less than \$25. Not only does this provide abusers with a quick, affordable way to purchase surveillance tools, but it also implicitly (and sometimes explicitly) condones IPS using these tools. In fact, many of these articles and spy devices justify IPS as a way to catch a cheating spouse (an example of which is shown in Fig. 1). Despite this worrying evidence, prior studies focused on researching IPS via IoT devices (e.g., [32, 44, 45, 63]) and have not looked into the ecosystem of commercial spy devices. As a result, we do not know what spy devices are out there for perpetrators of IPS.

With hidden surveillance devices readily available to abusers, survivors urgently need reliable tools to detect hidden devices. Unfortunately, it is unclear whether the available options work. Commercial detection apps [61, 66] and hardware detectors [5] claim to detect a wide range of hidden

*This paper was published in [USENIX Security 2023](#).

[†]Work done while at University of Wisconsin—Madison

devices including cameras, microphones, and listening bugs. However, there have been no studies assessing their efficacy, and customer reviews for these devices indicate they may not work as advertised [5]. On the other hand, some researchers (e.g., [36, 53, 54]) have built tools to identify hidden devices, but the effectiveness of these devices in real-world scenarios has not been evaluated. This leaves survivors with few reliable tools to combat IPS.

Thus, in this work, we study the spy devices and detection tools that are commercially available. We begin by building a taxonomy of available spy devices to clarify the threat we are facing. We search five large online retailers in the US—Amazon, Walmart, eBay, Best Buy, and Home Depot—for devices intended for spying and “dual-use” [15] devices that *could* be used for spying on other users. Out of 2,228 commercially-available spy devices we gather, we analyze a random sample of 163 devices and find that 29% are advertised for some form of covert surveillance and several are explicitly marketed for spying on an intimate partner. A screenshot of one such listing is shown in Fig. 1. Thus, an abuser looking for ways to spy on an intimate partner can find tailor-made tools on mainstream online retailers. Three of the retailers we searched allow third parties to list products for sale on their websites and none of these marketplaces forbid the sale of covert spy devices in their terms of service [4, 6, 7].

The spy devices in our sample cover a range communication mechanisms. For example, 50 devices (31%) use WiFi communications and could potentially be identified using existing academic tools [53, 54]. However, another 55 devices (34%) use Bluetooth or 4G, and 58 devices (35%) do not use any remote communication and instead rely on local storage. In order to best protect survivors, detection devices must be able to locate spy devices that use any of these types of communication.

Having established a taxonomy of spy devices, we turn our attention to exploring commercial and academic detection tools. We first search on the same set of retailers for devices advertised for detecting spy devices. Then, we crawl Google Play and the Apple App Store looking for detection apps. These apps and tools claim to detect hidden devices using RF detection, magnetometers, and IR lens detectors. To verify their assertions, we purchase a subset of these devices and apps and evaluate them in a laboratory experiment. We find that these commercial detection tools are all but useless, even in a controlled lab environment. This fact is more dire considering that these detectors could add to a survivor’s paranoia by raising false alarms, or worse, provide a false sense of security by failing to detect a hidden device. We additionally test the performance of two academic detection tools, SnoopDog [54] and Lumos [53]. Despite being grounded in promising theory, these tools require more refinement before they can help survivors of IPS.

In summary, although commercial spy devices appear to be effective tools for IPS—and are often explicitly advertised

for an IPS use case—commercial detection devices provide little recourse for survivors. Moving forward, tackling covert spying will require significant effort. Our study makes the following contributions to combat covert IPS using spy devices:

- (1) We are the first to highlight the problem that spy devices used for IPS are available for purchase on popular online retailers in the US.
- (2) We taxonomize the spy devices and dual-use, hideable devices available for sale on online platforms. To direct future work, we also identify the challenges inherent in detecting these devices.
- (3) We examine the effectiveness and usability of existing apps, hardware tools, and academic systems which claim to detect hidden devices. Unfortunately, we show that none of the detection mechanisms work as promised.

We hope that this work will provide guidance and motivation for future work in detecting covert spy devices.

2 Background & Threat Model

Our work centers around intimate partner violence (IPV), which involves “physical, sexual, or psychological harm by a current or former intimate partner or spouse” [3]. IPV is a pervasive issue that affects millions of people per year in the United States [57] and worldwide [46]. IPV is historically a gendered issue. A patriarchal society with gender inequality results in women making up the majority of survivors [46, 57]. However, IPV can happen to anyone regardless of gender, sexuality, race, marital status, or socioeconomic background. Moreover, our work focuses on the technology used within IPV, which is largely unconnected to the gender of the abuser or the survivor. As such, we do not focus on the gendered aspect of IPV in this work. Survivors can experience IPV both during a relationship and after it is over.

2.1 Technology and IPV

Recently, intimate abusers have begun to use technology to assert “digital coercive control” [22] over their partners [9, 11, 15, 18–20, 22, 33, 41, 58, 65]. Abusers install spyware and *dual-use* apps—apps built for a legitimate purpose which can nevertheless be repurposed for non-consensual surveillance [15]—to covertly collect data from a survivor’s phone, including location, messages, calls, and photos [9, 15, 48]. Unfortunately, these tools are discussed and often condoned on online platforms [11, 33, 65]. Beyond spyware and dual-use apps, abusers also leverage access to accounts to see private information [58], impersonate the survivor [20], or post intimate photos without consent [18].

A growing branch of tech-enabled IPV is *IoT-enabled IPV*, where abusers leverage IoT devices such as Bluetooth item finders, smart thermostats, and cameras to spy on or harass survivors (e.g., [12, 14, 39]). Scholars have responded to this new

threat by looking at how smart home devices can be leveraged in abusive ways [31, 32, 35, 44, 45, 56, 59, 60, 62, 63]. For example, Parkin et al. [45] and Leitão et al. [32] evaluated potential threats of smart home devices through usability analysis and workshops with survivors, while Slupska and Tanczer [56] wrote an IPV threat model for IoT device manufacturers.

2.2 Intimate Partner Surveillance Using Hidden Devices

Our work focuses on a subset of IoT-enabled abuse, where abusers *hide* IoT devices to enact *intimate partner surveillance* (IPS). Much like spyware [15], anecdotal evidence indicates that abusers use two types of devices for covert spying: (1) devices that are designed to be used for spying [13], and (2) dual-use devices that are not originally intended to be used for spying [14, 39].

Devices designed for spying. To our knowledge, only one recent paper explores the devices designed for spying (though not in an IPV context). Sathyamoorthy et al. [51] provide an overview of *wireless spy devices* (WSDs) including hidden cameras, audio bugs, and devices disguised as everyday objects (e.g., a pen or a coat hook). They classify WSDs into two types: *active* devices that send data as they capture it and *passive* devices which store data locally.

Dual-use devices. The other type of device used for covert spying is *dual-use* devices, which are not designed for spying but can be used maliciously. Of the many types of dual-use devices, Bluetooth item finders such as AirTag [10] and Tile [64] have received the most attention in prior work. Prior work has tried to bolster protections against unwanted stalking by Bluetooth item finders. Apple, Tile, and Samsung debuted proprietary stalking prevention measures including chirping sounds and in-app detection tools which scan for suspicious devices. However, these tools are not a catch-all solution; for example, Mayberry et al. identified three ways to break Apple’s anti-stalking protections [42], and several “silent” AirTags have been listed on online retailers with their microphones muted [30, 38]. Academic work has attempted to close these gaps. For instance, Heinrich et al. developed the AirGuard app for Android to periodically scan for suspicious AirTags in the background [25].

Detecting hidden devices. To combat covert IPS, survivors need to be able to detect hidden devices in their vicinity. Several commercial apps (e.g., [61, 66]) and hardware detectors (e.g., [5]) claim to be able to detect hidden devices, but to our knowledge, no prior work has evaluated the effectiveness of these tools. Sathyamoorthy et al. [51] are the only academic work to discuss commercial detection devices, listing RF detectors, spectrum analyzers, and nonlinear junction detectors (NLJD) as potentially useful tools.

On the other hand, many researchers have recently developed *academic* detection systems. Several papers [28, 43, 49,

55] focused on classifying IoT devices in a home network; others [16, 24, 36, 37, 50, 53, 54] have built tools to identify unknown devices, either by analyzing network traffic or by using hardware tools. For instance, SnoopDog and Motion-Compass correlate a user’s motion with network traffic to identify spying sensors [24, 54], while EEye uses millimeter wave (mmWave) sensing to identify even hidden electronic devices [36]. While these efforts are exciting, it is unclear whether they are effective and usable in the real world.

2.3 Threat Model

Our threat model extends the *UI-bound adversary* introduced by Freed et al. [19] in the context of mobile surveillance apps. A UI-bound adversary is “authenticated but adversarial” in that they may use a system’s intended functionality in adversarial ways, but may not modify the system or use advanced privileges [19]. We apply this concept to physical spy devices. In our threat model, the adversary is limited to using commercially available spy devices. We assume an average adversary will not modify purchased spy devices, nor can they manufacture new spy devices. When interacting with the data collected by these devices, they may only use the UI of the software provided by the device manufacturer. They may exercise creativity when hiding the device and may use devices not intended for spying to conduct IPS.

Additionally, we assume that the adversary has at least one-time physical access to either the home or personal possessions of the survivor, and the physical access must be long enough for the adversary to set up (e.g., by providing WiFi access, if applicable), hide, and activate a spy device. This is a reasonable assumption within an IPV context. Abusers and survivors often live together, share personal spaces, or share children, pets, vehicles, or possessions which require periodic physical contact to exchange custody. This means abusers will have ample opportunities to gain the one-time access they require to install these devices. Revoking an abuser’s access to the home or shared physical possessions typically requires costly, onerous legal interventions such as a restraining order or a divorce settlement.

Because our threat model considers physical spy devices rather than mobile spyware tools, there are key differences from the UI-bound adversary. First, while spyware apps in a mobile device are governed by the operating system (Android or iOS), there is no such entity controlling the operations of the spy devices. There is no equivalent notion of “app permissions” for spy devices, and thus it is harder to detect their presence in a physical environment. Moreover, unlike official app platforms (Google Play or Apple App Store) that oversee app stores and the distribution of mobile applications, there is no single authority that directly controls the sale of IoT devices and spy devices. This makes existing mitigations such as blocking IPS search terms on app stores [15] infeasible for spy devices. Further, individuals can prevent access to their

phone by changing passwords/PINs and setting permissions. In contrast, as discussed above, revoking an abuser’s access to the home and physical possessions is difficult and cannot be relied upon to prevent abuse. Finally, in many IPV contexts, it might not be possible to revoke access at all (e.g., if the survivor and their abusive partner still live together), as is explored by Stephenson et al. [59].

3 Analysis of Commercial Spy Devices & Detection Tools

Prior works [15, 65] have shown that abusers in IPV often use easily-available tools (such as mobile applications) to spy on or harass their victims. In this work, we investigate the spy devices that abusers can easily find and purchase through popular online retailers in the US, as well as the detection tools survivors can find through these retailers. To this end, we search on those retailers using queries that a potential abuser might use when searching for devices to conduct IPS or that a survivor might use when searching for detection tools to combat IPS. Fig. 2 summarizes the number of spy devices, detection devices, and detection apps in our dataset at each stage of our analysis.

3.1 Gathering Devices & Detectors

We selected five online retailers in the US: Amazon, Walmart, eBay, Home Depot, and Best Buy. Amazon, Walmart, and eBay are the top three general-purpose retailers in the US, while Home Depot and Best Buy are the top two US retailers which sell general electronic devices [17]. These retailers together have more than 50% of the online market share in the US [17]. We do not search dedicated spy device retailers. We believe that abusers and survivors would prefer to buy devices from large retailers than from lesser-known websites due to large retailers’ convenience and strong customer support.

Creating search queries. We first created a set of queries to search for spy devices and detectors on these retailers’ websites. To create the set of search queries, we used an established technique known as *snowball searching* (e.g., [15]). With snowball searching, a set of seed queries is manually chosen and used as inputs to a search suggestion API. The suggestions are then retrieved from the API and added to the set of queries to be used. The process is repeated with the new queries used as inputs to the suggestion API until there are no new queries returned by the API or a predefined number of queries has been gathered. We added an additional step where new queries are ignored if they are unrelated to spy devices. We did this by constructing a block list composed of strings found in irrelevant queries, such as “spyware” to exclude software and “novel” to exclude books. If any substring of a query was in the block list, the query was ignored.

Stage of analysis	# Products		
	Spy devices	Detectors	Detection apps
Initial crawl (§ 3.1)	6,403	1,313	51
Filtering (§ 3.2)	2,228	700	43
Qualitative analysis (§ 3.3)	163	148	43
In-lab testing (§ 3.4, § 3.5)	11	1	11

Figure 2: Spy devices, hardware detectors, and detection apps in our dataset through the different stages of our analysis: the initial crawl, filtering, qualitative analysis, and in-lab testing.

To search for spy devices, we began with a manually curated list of 167 seed queries. These seed queries were generated by placing *targets* (such as “wife” or “boyfriend”) into *templates* (such as “spy on my {target}” or “best hidden camera for catching cheating {target}”). This resulted in queries like “spy on my wife” or “best hidden camera for catching cheating boyfriend.” Initially, we planned to conduct a single snowball search and utilize the results across all retailers. However, we discovered that queries from a different retailer’s suggestion API were unlikely to yield any relevant products. As a result, we conducted separate snowball searches for each retailer, utilizing their respective (hidden) suggestion APIs. This procedure returned 525 queries for Amazon, 289 queries for Walmart, 179 queries for Best Buy, 240 queries for eBay, and 176 queries for Home Depot.¹

We used a similar procedure to search for detection devices. To generate queries for this search, we performed another set of snowball searches with seed queries aimed at detecting devices, such as “hidden camera detector” and “bug finder” This time, we generated 101 queries for Amazon, 61 queries for Walmart, 31 queries for Best Buy, 42 queries for eBay, and 30 queries for Home Depot.

Gathering search results from online retailers. Using the queries gathered in the snowball search, we searched each retailer for seven days in July 2022. For each retailer, we collected the URLs of up to 10 products returned by each query. We then gathered metadata about each product using the collected URLs. Specifically, we collected the product name, sale price, description, user reviews, and user-provided star rating. In total, we gathered 6,403 spy devices and 1,313 detection devices (Fig. 2).

Identifying iOS and Android detection apps. In addition to hardware detection devices, survivors may also turn to detection *apps* available on the Google Play store and Apple App Store, that claim to be able to detect hidden devices and are often free to use. Thus, we used a similar procedure as before to collect a set of iOS and Android apps which claim to detect hidden devices. For our queries, we used the set of seed queries as we did in the hardware detector crawl. In total, we collected 51 apps, 15 from the Apple App Store and 36 from the Google Play Store (Fig. 2).

¹ All the seed query templates, final query lists, and other research data is contained here: <https://github.com/ceccio247/IPV-Spy-Device-Study>.

3.2 Filtering and Sampling

Not all of the products we gathered are, in fact, spy devices or detection devices. Some are not even electronic devices. Thus, we devised a filtering procedure to identify relevant devices based on the product listings.

We considered a device relevant for this research if (a) it is an electronic device, (b) gathers sensitive information such as audio, video, or location, and (c) can be hidden (either by design or due to their form factor). We considered a device able to be hidden if its approximate volume is less than 64 cubic centimeters (4 centimeters to a side). All three criteria must be satisfied for a product to be considered a spy device. To identify relevant spy devices, we first used a heuristic filtering algorithm followed by a logistic regression classification. Appendix A.1 provides details about these classifiers. In all, these classifiers helped us filter our data down to 2,228 spy devices out of 6,403 products scraped.

Similarly to the spy devices, we performed heuristic filtering on our dataset of detection devices (Appendix A.2). After developing and applying our filter, there were 700 detection devices in our dataset. We also filtered the set of detection apps through manual examination. We removed two apps that are not detection apps, as well as six apps that had been deleted shortly after our crawl. We were left with 43 apps meant for detecting hidden devices (14 from the Apple App store, 29 from the Google Play store).

Sampling. After filtering our dataset, we took a stratified random sample of the products for further analysis (Fig. 2). Of the spy devices, we sampled 200 (60 each from Amazon and eBay, 50 from Walmart, 20 from Best Buy, and 10 from Home Depot). Of the detection devices, we sampled 150 (50 each from Amazon, eBay, and Walmart—Best Buy and Home Depot yielded no detectors after filtering). Because our sample of detection apps was small, we analyzed all 43 detection apps.

Before diving into our analysis, we manually investigated the spy devices, detection devices, and detection apps to identify false positives. Of the spy devices, 37 were false positives (19%). These false positives were mostly cameras that could potentially be used for IPS, but are too large to be reasonably hidden. Other false positives included SIM cards for GPS trackers and smart watches that claim to be trackers but, upon inspection, do not actually track location. This left us with 163 true positive spy devices. Of the 150 detection devices, one was a false positive (a spy camera with the phrase “Hidden Camera Detectors” in its product title); thus, we analyze 149 hardware detectors.

3.3 Qualitative Analysis of Product Listings

First, we characterized the 163 spy devices, 148 hardware detectors, and 43 detection apps in our sample (Fig. 2) using the information on the product listings. By looking at the

product name, description, specifications, and accompanying photos—and occasionally the product reviews, which helped to clarify the actual product capabilities—we gathered information about the devices. For spy devices, we gathered (a) the type of information the device claims to collect, (b) communication medium used by the device, (c) advertised use case(s) of the device, (d) covertness, and (e) other metadata, such as price and device manufacturer. We also collected the user reviews of the spy devices to search for anecdotal evidence that these devices are being used to conduct IPS. In total, we collected 15,139 user reviews. We then searched for IPS keywords within these reviews and identified 43 reviews relevant to IPS. For detection tools, we gathered the types of spy devices the detector claims to detect, the technology used by the detection tool, and metadata.

To collect this information, we used Collaborative Coding to reduce the overhead of coordinating among researchers and speed the process of reaching a consensus for our analysis [47]. Two authors divided the devices and apps evenly and inspected them one by one. For spy devices, the authors occasionally disagreed on how to classify whether a product is able to be hidden; this was resolved by refining our definitions to rely on concrete physical dimensions. All authors met together multiple times to resolve conflicts and confusion. Though most devices are fairly clear about their capabilities and advertised usage, we used our best judgement for some devices. For instance, several cameras in our dataset of spy devices do not explicitly say whether they recorded audio as well as video; for these, we assumed the camera did not record audio unless proven otherwise.

As we will discuss in Section 4.2, several of the spy devices are marketed toward general covert surveillance and cite a long list of potential targets (either using text or photos). For these devices, we did not list every possible target; rather, we considered the device to be advertised only for general covert surveillance. We separately noted whether intimate partners are listed among the possible targets.

3.4 In-Lab Testing of Spy Devices

Next, to evaluate whether these devices indeed enable IPS, we purchased a representative sample of spy devices for in-lab testing. The details of the methodology for these experiments can be found in Appendix A.3.

Sampling. We sampled spy devices with two goals in mind. First, we chose devices that appear to be *functional*, based on positive customer reviews. When possible, we selected devices that had reviews that implied past usage of IPS; these reviews are discussed in more detail in Section 4.2. Second, we chose a *representative* set of devices that covers all types of data collected (e.g., audio, video, location) and communication technologies (e.g., Wifi, LTE, etc.) we observed in our dataset (Section 5.1). In total, we collected a sample consisting of 11 devices: 2 cameras using local storage, 3 cam-



Figure 3: The canonical hardware detection device in our dataset. 40 hardware devices in our sample look very similar to this. We test this particular device in Section 5.2.

eras using WiFi, 2 microphones using local storage, a battery powered GPS tracker using cellular networks, an OBD2 port powered GPS tracker using cellular networks, and 2 trackers using Bluetooth mesh (an AirTag and a Tile).

Evaluating recording devices. To evaluate the collected audio and video recording devices, we inspected the physical products to gauge their effectiveness. We observed whether the devices would readily power on and collect the audio/video information they claim to collect, as well as whether the products could feasibly be hidden.

Evaluating location trackers. To evaluate the location tracking devices, we devised a field experiment. While work by Heinrich et al. [27] has performed experiments to evaluate the effectiveness of Apple AirTags, little is known about Tile devices or cellular network GPS trackers; moreover, none of these devices have been tested in the context of IPS. With this in mind, we designed an experiment to simulate the walking pattern of a survivor living in a moderately sized city. The 2.1-mile experiment included both walking and driving paths through a public park, residential zones, and commercial zones. Three authors followed this path while carrying an AirTag, a Tile, and a battery-powered cellular network GPS tracker. (Additionally, an OBD2 cellular network GPS tracker recorded the driving portion of the experiment.) We recorded the location reported by the tracking devices then compared the data recorded by these devices to the data recorded by a smartphone’s internal GPS.

Usability evaluation. For all eleven spy devices, we also performed a brief usability evaluation. For each device, we examined the setup instructions and wrote down the number of steps required to set up the device. Using these steps, as well as our own experiences setting up the spy devices, we compared the setup process of the spy devices to the setup process for benign IoT devices. We also made notes on our experiences using the spy devices and compared them to the experience of using benign IoT devices.



Figure 4: Our experimental setup. We created nine total equal-height place markers at 1 foot, 5 feet, and 10 feet away from a spy device and recorded each detector’s reading at each place marker. The 10 foot place markers are not shown.

3.5 In-Lab Testing of Detection Tools

We evaluated the effectiveness of 12 commercially available device detection tools: a sample of 11 detection apps, and one physical detection device.

Sampling. When sampling detection apps, we focused on those that claimed non-network-based ‘bug detection’, since we observed that the network-based detection apps simply list all devices connected to the network. We started by installing all apps with over 1,000 reviews (9 apps). These apps claimed to detect hidden cameras, general bugs, and general RF signals; we added 2 additional apps to ensure we had a representative sample of apps meant for detecting all of these devices and designed for both iOS and Android.

For our sample of the physical detection devices, we chose one representative device (shown in Fig. 3) which carries the three technologies we observe in our sample of crawled detectors: an RF scanner, a magnetometer, and an infrared lens detector. We selected this single detector due to its popularity: 40 devices in our sample (26%) appear to be the same detector sold by different brands.

Experiment design. We designed an experiment to test the functionality of commercial detection devices and apps. Our goal was to evaluate the detectors in a best-case environment. We selected a room in our building with as little existing technology as possible to reduce the chance of false positives. Before the experiment, we marked where we would place the spy devices as well as positions that were 1, 5, and 10 feet away from the spy device. At each distance, we picked three points that are at a 0, −30, and 30 degree angle from the spy device (nine positions in total). To reduce inconsistency in our measurements, we ensured that the detector was placed at the same vertical height in the room for each measurement. Fig. 4 shows a photo of our experimental setup.

For each detector, we evaluated its effectiveness at detecting each of the technologies in our representative sample of spy devices. To do this, we tested the detector in the presence of six devices: a WiFi-connected spy camera, a spy camera

with local storage, a GPS tracker that uses cellular networks, a GPS tracker that uses Bluetooth mesh networks, a microphone that uses local storage, and a WiFi-connected smart home camera (for validation). For each detector, we recorded its measurements at each position when there is no spy device in the testing area. Then, for each spy device, we powered on the device, put it in active transmission mode (if applicable), and placed it at the designated location for taking measurements from the detectors.

3.6 Ethical Considerations

Our methodology was constructed to emulate an abuser performing searches on online retailers, similar to prior works [15, 51]. Although all information we collected is easily found online, aggregating such information can be harmful. We recognize that the descriptions of the spy devices and their titles may point an abuser toward them. We therefore do not publish the URLs of any of the spy devices we gathered in our analysis. We do not collect any PII; we remove the names of people leaving reviews for products we are interested in before storing in our database. We consulted our institution’s IRB regarding the ethical collection and storage of this data. While this work does not involve human subjects and as such does not fall under their purview, they found the measures we had taken to be more than sufficient.

3.7 Limitations

We were only able to examine, taxonomize, and test a sample of device detectors and spy devices. This sample did not include device detectors that cost more than a hundred dollars. Specifically, we found some Non-linear Junction Detectors that claim to detect arbitrary circuit boards, but cost thousands of dollars; we did not test such a device as they are beyond the budget of an average survivor.

We could only work with prototypes of academic works. If these works were given more development time and budget, it is possible that they would work better than they do.

The experiments to evaluate the efficacy of detectors took place in a computer science building, which may have more signals and interference than a typical single-family home. We did not perform experiments with real users and thus can only approximate the usability of the spy devices and detectors.

4 Taxonomy of Commercial Spy Devices

We collected a set of spy devices available on popular online retailers and analyzed these devices from multiple angles. We characterize the types of commercial spy devices abusers can find online. These devices are often advertised for IPS, and reviews indicate that customers are indeed using them for IPS (Section 4.2). We confirmed through in-lab testing that these varied spy devices are effective tools for IPS (Section 4.3).

Category	Value	# Devices (%)
Retailer	Amazon	51 (31%)
	Ebay	47 (29%)
	Walmart	45 (28%)
	Best Buy	11 (7%)
	Home Depot	9 (6%)
Information collected	Video	74 (45%)
	Audio	64 (39%)
	Location	59 (36%)
Communications used	Local storage	58 (36%)
	4G	50 (31%)
	Wifi	47 (29%)
	Bluetooth	5 (3%)
	Unclear	5 (3%)
Covertness	Intended to be hidden	82 (50%)
	Can be hidden	46 (28%)
	Disguised	35 (21%)

Figure 5: Characteristics of our sample of 163 potential spy devices. A device could collect multiple types of information and use multiple types of communication. A device’s communication was labeled “unclear” when it could have used two or more methods but the description did not indicate what it used. For example, a small camera that could have used WiFi or purely local storage was marked “unclear”.

4.1 Characteristics of Commercial Spy Devices

The spy devices we found encompass a range of capabilities, communication technologies, and types of covertness. Fig. 5 presents the characteristics of the spy devices in our sample. By identifying the different varieties of commercial spy devices, we can inform the design of future detectors—ideally, detection tools should be able to detect all of these varieties of spy devices.

Information collected. The spy devices in our sample collect three basic types of information: video (74 devices), audio (64), and location (59). These devices are typically cameras with video (and usually audio) capabilities, devices designed for audio recording only, and devices meant for tracking location only. Only three devices combined both audio/video recording *and* location tracking: a GPS tracker with audio recording capabilities, a body camera with GPS, and a children’s smart watch that can track location in addition to audio and video recordings.

Communication technologies. The spy devices in our sample utilize four methods of sending information. Several devices—including every location tracker in our sample—share the information they collect using cellular networks (50 devices), WiFi (47), or Bluetooth (5). In contrast, some cameras and recording devices rely on local storage only (such as an SD card) for storing information (58).

A device’s communication method and storage method impact that device’s utility to an abuser. Spy devices that share data remotely allow an abuser to place the device once,



Figure 6: A photograph from a listing for a spy camera. The camera is disguised as an air freshener.

then continue to view data from that device without being physically near the survivor. Since abusers routinely share children, homes, or assets with survivors, it is not unlikely that an abuser would have one-time physical access to the survivor’s surroundings. On the other hand, devices that use only local storage require an abuser to have routine physical access to the device to look at the data it collects; typically, this would not be possible unless an abuser has periodic access to the survivor’s personal spaces. Though we should prevent spying with either type of device, devices that share data remotely might be of concern to more survivors.

The communication and storage methods a device uses also impact detection. Many academic detectors can only identify devices that communicate over WiFi, and some commercial tools that rely on RF sensing can only find devices that send data externally. Ideally, survivors should be able to easily detect all types of devices with a single detector.

Covertness. We identified 117 devices built to be hidden and 46 devices not meant to be hidden, but still hideable: they are small and unobtrusive enough to be hidden (we use the heuristic of less than 4cm to a side) while not specifically advertised as being a hidden device. These include devices such as small cameras that are advertised to be used like an action camera and excludes devices such as large smart home cameras and surveillance cameras that would be difficult to hide. Of the devices built to be hidden, 35 devices are disguised as everyday objects (as shown in Fig. 6), while the rest were simply small enough to be discreet. We note that devices that are intended to be hidden are not necessarily intended to be spy devices; for example, many GPS trackers are advertised for protecting a vehicle and must be hidden to be effective against theft.

Knowing the covertness of these devices impacts how we address them. To spy with dual-use devices and spy devices intended to be hidden, abusers must hide them so that they are difficult to find, but can still capture their intended data. For cameras, this means the lens must be visible; for audio recorders, there cannot be too much covering the microphone;




Figure 7: CDF of device prices from our sample. The median price for spy devices in our sample was \$30.99, while the median price of detection devices was \$34.94 (not including apps). The most expensive spy device costs \$399, while the most expensive detector costs \$1,010.

and for location trackers, they must simply be able to capture the signals needed for location tracking. As such, many of the product listings we looked at advertise specific places to hide these devices in homes or cars. On the other hand, some devices are disguised as pens, photo frames, alarm clocks, and even bottles of Mountain Dew; these devices could be hidden in plain sight. Thus, when survivors try to find these devices, they need detection tools that can detect devices hidden in many different places, from a photo frame in the survivor’s living room to deep under the hood of a car.

Price. In addition to the physical properties of these devices, we look at their price. The distribution of device prices is shown in Fig. 7. The median price for spy devices in our sample is \$34.94 (with standard deviation (SD) of \$85.25). While the most expensive spy device in our sample—a hidden camera disguised as a wall outlet—costs \$399, half of the devices are available for under \$30, and the cheapest spy device costs \$8.98. Thus, an abuser looking to purchase a spy device will find many devices at an affordable price.

Advertised use cases. The devices we found are advertised for various use cases ranging from monitoring home or office, vehicle, small items, pets, children, to covertly monitoring employees or even intimate partners. The distribution of advertised use cases of different types of devices—Audio recorders, Cameras, and GPS trackers—is shown in Fig. 8.

4.2 Advertising for IPS Use Case

In the previous section, we describe the different types of spy devices an abuser could find on online retailers. However, these characteristics alone do not tell the whole story. Not only are many types of potential spy devices available to abusers, but many are also advertised for covert spying and often explicitly for IPS (Fig. 8). Further, product reviews show

Advertised purpose	# Devices (%)		
	Recorders	Cameras	Trackers
<i>Malicious use cases</i>			
General covert surveillance	10 (16%)	25 (34%)	8 (14%)
Catching an intimate partner	0	0	8 (14%)
Catching an employee	0	3 (4%)	0
Tracking an employee	0	0	1 (2%)
<i>Non-malicious use cases</i>			
Home or office security	2 (3%)	54 (73%)	0
Protecting a vehicle	0	1 (1%)	26 (44%)
Recording for business	10 (16%)	6 (8%)	0
Protecting a child	0	1 (1%)	15 (25%)
General recording	5 (8%)	3 (4%)	0
Recording oneself	5 (8%)	1 (1%)	0
General protection	0	0	4 (7%)
Fitness tracking	0	0	3 (5%)
Protecting pets	0	0	3 (5%)
Protecting family members	0	0	2 (3%)
Protecting small items	0	0	2 (3%)

Figure 8: The advertised use cases of spy devices. Devices are often advertised for more than one use case. Though many devices are not advertised for malicious purposes, they are often used maliciously (as we describe in Section 4.2).

that customers are, in fact, using many of these devices for IPS. Thus, an abuser looking online can not only find devices that *could* be used for IPS, but they can find devices that are *encouraged* to be used in IPS.

Dual-use devices. Many of the devices in our sample are dual-use devices, advertised for uses like home security (55), vehicle protection (27), recording business events (16), and protecting children (17) or pets (3). Some of these devices are intended to be hidden; for example, GPS trackers that are meant to prevent vehicle theft are intended to be hidden inside the car. Though these devices are not promoted for an IPS use case, we note that they (a) appeared in searches that used explicit IPS-related search terms and (b) are capable of being used for covert IPS.

Devices designed for spying. In contrast, many devices are specifically advertised for covert surveillance. We see 35 devices advertising themselves for general covert surveillance—e.g., they are advertised as a hidden spy device and list several possible targets the device could be used to surveil. Many of these devices include intimate partners among the listed targets, either in text (5), pictures (6), or both text and pictures (2). For example, one GPS tracker on eBay is described as “Black shell, easy to hide, perfect for tracking vehicles, teens, spouses, elderly persons or assets.” However, the tracker’s description also states that it “doesn’t encourage the product use for illegal purposes.”

In addition to general covert spy devices, we see seven devices advertised specifically for spying on an intimate partner. These devices, all of which are GPS trackers, include such blatant titles as “GPS Tracking Device for Cheating Spouse

Boyfriend” and cite suspicions of cheating as the reason for surveillance. Such devices are tailor-made for an abuser interested in IPS and are available for \$121.41 on average.

Evidence of use for IPS. Interestingly, none of the devices advertised specifically for IPS have received any customer reviews or ratings. However, several other devices have customer ratings that imply that the review writer used the device for IPS, even if the device is not advertised explicitly for it. These reviews are largely positive, and almost all of them discuss how the device helped them catch an intimate partner engaging in infidelity. A reviewer for a camera wrote “*I caught my wife cheating in the act. It worked great but the setup did not accept the special characters like !@#%& for the WiFi password.*”, and a review for a microphone was titled “*To catch a cheater*”. Moreover, during our investigation of spy devices, we find many more examples of reviews for these devices that imply IPS. A reviewer for a GPS tracker wrote:

“My girl has been giving me “trust issue vibes.” I luckily found this item, and purchased it, and I must say it works EXCEPTIONALLY well! On the first day, I caught her up lying about which Popeyes chicken location she went to LOL, and of course when I confronted her, she says she lied to me to prevent a negative confrontation with me, so yah it’s always the males fault no matter what I guess lol!”

Another person reviewing a different product wrote:

“Been tracking my husband and now I’m tracking his lies of what he’s doing and where he’s going... SMH. I snuck it in the back pocket of his drivers side seat and it works perfectly.”

While we do not have precise statistics regarding the usage of these devices in IPS, this anecdotal evidence shows that these devices are being used for IPS.

4.3 Efficacy of Spy Devices

We purchased a subset of 11 devices (7 recording devices, 2 traditional GPS trackers, and 2 Bluetooth mesh GPS trackers) to evaluate their efficacy. We find that these devices are indeed powerful tools for IPS.

Evaluation of recording devices. Among 7 recording devices we tested, five can gather sensitive data when used in our lab. The picture and audio recorded by the recording devices are clear, and both the cameras and the microphones can record hours of sensitive data when supplied with large-capacity Micro SD cards. Finally, all of the recording devices are either disguised as common household objects or are small and unobtrusive enough to be easily hidden by an abuser. The devices that do not work are WiFi cameras that cannot be paired with a smartphone app to view the data it records. Despite this failure, the remaining devices work and are indeed effective are enabling IPS.

Evaluation of location trackers. All 4 of the GPS trackers we test are able to effectively communicate their location when used in our lab. We find that while the traditional GPS trackers offered superior accuracy and data density compared

to the Bluetooth mesh trackers, the data recorded by the Bluetooth mesh trackers is more than sufficient to allow an abuser to reconstruct the path taken by the theoretical survivor. We note that we collect the Bluetooth mesh tracker data from Apple and Tile’s hidden APIs as opposed to the official Apple and Tile smartphone apps. When comparing the data we collected from these APIs to the data presented in the apps, we notice small discrepancies between the coordinates reported by the APIs and the ones shown on the phone app’s interface; we have no explanation for this discrepancy.

Apple has taken steps to limit the use of AirTags for stalking. An AirTag that has been separated from the paired iPhone will emit a sound when moved, and an iPhone that detects the extended presence of a “lost” AirTag will alert its user. The iPhone can also help locate AirTags, and can trigger the AirTag’s sound [8]. However, a survivor has to own a modern iPhone to be alerted by the phone, and they have to recognize the sound made by the AirTag to realize they are being stalked. Moreover, Heinrich et al. found that a person with some technical skill can create a third-party bluetooth tracker that connects to the Find My mesh network [26]. An abuser dedicated to getting around the AirTag’s noise-based safety features could pay someone to create a third-party Find My mesh device that does not produce any sound. As such, AirTags and the Find My mesh as a whole enable IPS despite Apple’s efforts to mitigate the potential for tech abuse.

Usability. In our usability analysis, we found that the spy devices are no more difficult to set up and use than typical IoT devices. Two spy devices require an associated app to view the collected data; for these devices, the user needs to download that app, register an account, and bind the device to that account. (For AirTags, this process is easier because the user’s iPhone already contains the app and iCloud account required to use the device.) Other than that, the only actions required by the user are plugging in the device or charging it, and usually flipping a switch on the device to start recording.

4.4 Summary

We identified 163 commercial devices that abusers could use for covert surveillance. These cameras, recorders, and location trackers come in many forms and generally work as advertised—and often, they are advertised explicitly for IPS. Thus, an abuser looking for IPS tools can easily find functional, effective spy devices for an affordable price. In the next section, we contrast this with the space of commercial detection devices.

5 Evaluation of Available Detection Tools

In response to the rising availability of commercial spy devices, many spy device detectors are now on the market as well. These commercially-available detectors take the form of

Category	Value	# Devices (%)	# Apps (%)
Retailer	Amazon	49 (33%)	—
	Ebay	49 (33%)	—
	Walmart	50 (34%)	—
	Google Play	—	29 (67%)
	Apple App Store	—	14 (33%)
Claimed detection	All bugs	117 (79%)	18 (42%)
	Cameras only	31 (21%)	23 (53%)
	RF signals	0	1 (2%)
	Networked devices	0	1 (2%)
Technology used	IR lens detector	133 (90%)	17 (40%)
	RF detector	117 (79%)	1 (2%)
	Magnetometer	73 (49%)	28 (65%)
	Network scan	0	11 (26%)
	AI image recognition	0	5 (12%)

Figure 9: Characteristics of our sample of 148 detection devices and 43 detection apps. Many detectors offered multiple detection technologies.

both physical detectors and smartphone apps and make claims ranging from highlighting hidden lenses to detecting any hidden electronics in a room. In addition to these commercially available spy device detectors, researchers have built systems that detect and locate hidden spy devices. Systems such as Lumos [53] and SnoopDog [54], for example, sniff WiFi packets to identify and locate hidden devices. All these detection solutions claim to help users detect spy devices, but they may not deliver on these claims. We pose two research questions regarding these detectors: (1) Are these spy device detectors technologically capable of detecting the spy devices used in IPS? (2) What challenges do we anticipate survivors might face while using these detectors in real-world scenarios?

To answer these questions, we repeat a similar crawl to Section 3—this time looking for detection devices—and further, we crawl the Google Play Store and App Store to find detection apps. Using a sample of these detection tools, we identify characteristics of commercial detection tools available to survivors. We then design and perform a laboratory experiment to test the efficacy of a representative sample of detection tools. Unfortunately, we find that off-the-shelf detection tools are, for the most part, entirely ineffective in finding any kind of spy device. Additionally, in close collaboration with the original authors, we evaluate two prominent academic works on device detection: SnoopDog [54] and Lumos [53]. We find that both systems have promising theoretical grounding but require more polish and broadening of scope before they can be released as general purpose detection systems.

5.1 Taxonomy of Commercial Detectors

The majority of detectors clearly state what they claim to detect and what techniques they use for detection. In a handful of cases, the devices claim to detect magnetic fields while implying they were only equipped with an RF scanner. In these cases, we add magnetometer to the list of claimed technolo-

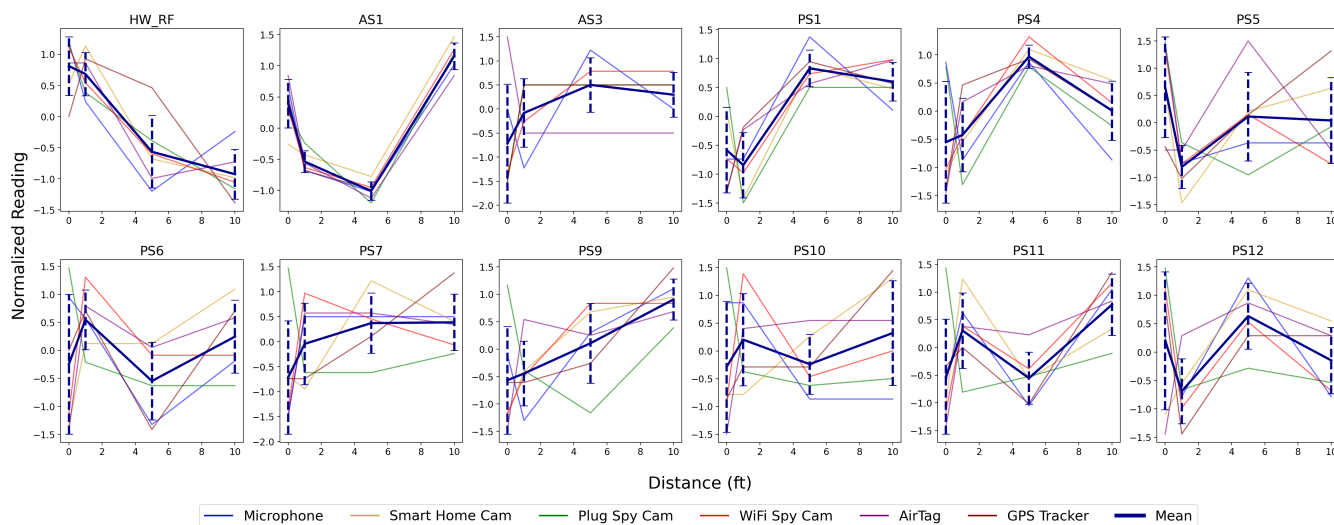


Figure 10: The performance of different detector devices for various spy devices we tested in the lab (shown in different color lines). The X-axis shows the distance of the detector from the spy device in feet (ft), and the Y-axis shows the normalized difference of the device reading compared to the setting when there was no spy device present in the room. Thick blue line shows the average of the measurements across all spy devices.

gies, as they *claim* to detect something only a magnetometer could detect. Fig. 9 summarizes our findings.

Detection devices. Of the 148 hardware detectors in our sample, most (117 devices) are advertised with a blanket claim of detecting all bugs. These devices generally claim to use RF detection in combination with a magnetometer and/or an IR lens detector. These detectors are quite homogenous; 40 devices appear nearly identical to the photo in Fig. 3. The remaining 31 physical detectors claim to detect only cameras and include only an IR lens detector.

Detection apps. The detection apps have a similar breakdown. Of the 43 detection apps we collected, 23 of them claim to detect hidden cameras using the phone’s camera, and 18 claim to detect all types of hidden bugs using the phone’s built-in magnetometer. The other two apps claim to detect general RF signals and all network attached devices. The apps employ IR lens detectors (17), RF detectors (1), and magnetometers (28), similar to the physical detection devices; additionally, some detection apps claim to find devices using network scanning (11) and AI image recognition (5).

Price. Similarly to spy devices, the median price of the hardware detectors is \$34.94 (Fig. 7). They cost as little as \$7.79, but three devices cost over \$800, with the most expensive detector listed at \$1,010. Survivors looking for detection tools may not be able to afford hardware detection tools at this price. In contrast, most detection apps are free. Only two apps cost money—\$1.99 and \$4.99, respectively—and two additional apps have an upgrade option for \$3.99 to unlock more detection capabilities.

5.2 Poor Efficacy of Commercial Detectors

Commercial detection tools claim to be able to find several types of hidden devices using several different techniques. Unfortunately, through a lab experiment, we find that nearly all of the detectors are ineffective and unusable. We summarize the results of our experiment in Fig. 10. We graph the normalized difference between the measurements recorded by the detectors when a device is present and the measurements recorded by that detector in the no device setup, represented by the low-opacity lines. We then take the mean of those measurements and plot them with error bars representing standard deviation. The graph labeled “HW_RF” refers to the readings from the hardware RF scanner. The graphs labeled AS and PS represent the readings from the Apple App Store apps and the Google Play Store Apps, respectively.

Detection apps. The values reported by the apps when scanning for arbitrary devices were very difficult to read. They are frequently in the range of 150-160 microteslas with unpredictable fluctuations. The apps give messages based on their readings, but they almost always report some variation of “high radiations detected.” We hypothesize that the messages are based entirely on the measurement falling within predefined ranges. The only time we can positively confirm that the apps were reacting to the presence of a device is when the phone is held within two centimeters of the device; when this is done, the reading spikes dramatically. We hypothesize that this happens because the magnetometer is acting as a crude metal detector. However, this behavior is unhelpful when attempting to find a device, as it would require the user to perform a brute force search to find hidden devices.

During our experiment, we observed that most of the apps do not detect our spy devices. As can be seen in Fig. 10, none of the apps (excluding AS1 and PS1) display a coherent pattern when a device is present, making them useless for detection. AS1 and PS1 are the only apps that appear to be consistently detecting something. However, we stress that all the apps, including AS1 and PS1, are extremely sensitive to their position within the room, with minor positioning changes resulting in fluctuating values. We also note that these patterns are only visible because we graph the difference between the measured reading and the no devices reading. This difference is in the range of -2 to 2, which would be difficult for an average user to notice when the detector’s reading is in the range 150-160. These apps are barely usable in a lab setting, and would likely be useless for a survivor of IPS.

Hardware detector. Our hardware detector has three functions: an RF scanner, a magnetometer, and an IR-based lens detector. Much like AS1 and PS1, the RF scanner has rather consistent readings signifying that it is detecting something, and unlike the apps, the range of its readings are only 0-10, meaning differences are noticeable. However, the RF scanner is still very sensitive to its placement in the room, as our lab (much like most commercial homes) has a variety of RF signals contributing noise to the reading. We also stress that the RF detector can only detect signals when the spy device in question is transmitting data. Devices that infrequently send data will be difficult to detect using RF, making the RF scanner unsuitable for real-world use.

We separately evaluate the lens detector and the magnetometer in our lab, and find both to be disappointing. The lens detector works by shining IR light that reflects off convex camera lenses, thus appearing as a shining spot when viewed through the viewport on the device. Unfortunately, we find that many reflective, convex things in our lab resulted in an overwhelming amount of false positives. The magnetometer only registers the presence of a magnet when placed within an inch of the magnet, meaning a user using the magnetometer would have to manually search their entire living space or car with it for the magnetometer to be effective. Neither of these tools live up to their claims of effective device detection.

Repeated Trials. To ensure that our findings were consistent across multiple trials, we performed a second test with the detector app PS5 and all of the spy devices used in the initial test. PS5 was selected as it is extremely similar in functionality and layout to the other apps under investigation. In this second test, the app did not perform any better than it did in the initial test. Thus, we are confident that these detectors are *technologically incapable* of detecting spy devices.

5.3 Efficacy of Academic Detection Tools

As the prevalence of WiFi connected devices has increased, multiple academic works have proposed systems that detect

and localize hidden devices [16, 24, 36, 37, 50, 53, 54]. To evaluate whether these academic tools can help detect and localize the spy devices used in IPS, we replicated two of these systems, SnoopDog [54] and Lumos [53]. We select these systems because they present the two most refined procedures in the literature; much of the other detection work is a version of these two systems. In both cases we reach out to the authors of the papers and request access to any working prototypes of their systems.

SnoopDog. After contacting Singh et al. [54] we receive a prototype of the detection portion of SnoopDog. To evaluate the system, we replicate the detection procedure described by the authors. The user installs an app that records accelerometer movements on their phone and activates packet capture software on a computer with a network chip in monitor mode. The user then alternates standing still with performing jumping jacks. The resulting packet capture and IMU data are processed by SnoopDog and the user is informed whether there is a device sensing them.

We were able to use SnoopDog to detect a camera in our lab, but there are issues with the system that prevent it from being widely usable. SnoopDog relies on two metrics to determine whether the user is being spied upon: a mathematical correlation metric and an overlap metric. The correlation metric works consistently, but the overlap metric is calculated using parameters that must be fine-tuned to the location the detection is taking place in. This would be very difficult for a survivor to perform, as a survivor who wants to use SnoopDog in a room does not know whether there is a device present. In addition, if the camera recording a room is not actively streaming data (i.e., it is storing data on an SD card to be accessed later), SnoopDog is incapable of detecting it.

Lumos. We similarly contact Sharma et al. [53] and attempt to use their prototype software for Lumos. As instructed by the authors, we install the prototype software on a Raspberry Pi and an iPhone XR. The Raspberry Pi captures encrypted 802.11 frames for device classification and the iPhone captures IMU data for localization as well as providing an interface for the user to interact with the Pi. After detecting the presence of device by analyzing captured frames, the iPhone app instructs the user to walk around the edge of the room under investigation. The user can then activate an AR mode that visualizes the location of hidden devices.

Using Lumos, we can detect the presence of the test WiFi camera we connected to our lab network. Due to the incomplete nature of the prototype, however, there were numerous false positives and the localization system was incomplete. Moreover, the system still has fundamental problems. While Lumos is compelling in theory, in addition to the existing prototype being incomplete, it shares one of SnoopDog’s limitations: Lumos requires cameras to be actively streaming data to detect them. As such, a more comprehensive solution to device detection is still necessary.

6 Discussion

By analyzing product listings and physical devices, we determined that spy devices are both **available** to abusers and are **effective** at conducting IPS. From our lab experiments, we determined that both commercial and academic detectors are **ineffective** at mitigating these devices. We discuss these findings and propose solutions to the threat of spy devices.

Spy devices are available and effective. We show that spy devices are *easy to find* and *affordable*. Using simple searches on popular retailers in the US, we found over 2,228 listings of spy devices for sale (Section 4.1). Retailers often guide searching for such products through query completion, for example “catch a cheater” query is completed on Amazon with “catch a cheater spy devices car.” As we demonstrate in Fig. 7, many of these devices can be purchased for under \$50.

In Section 4.3, we find that nearly all of the devices we tested are capable of IPS. Generally, an abuser and a survivor of IPS are either *living together* or *living apart* (this is a simplification of the three-phase framework introduced by [41]). When the abuser and survivor are living together, the abuser has access to the survivor’s physical surroundings and belongings. Therefore, the abuser can even use devices that require regular retrieval, such as devices that rely on rechargeable batteries or devices that record exclusively to an SD card. These devices are much harder to detect compared to devices relying on wireless communication (Section 5.2). When the abuser and survivor are living apart, the abuser can still continue conducting IPS using devices that use WiFi or 4G.

Detection tools are ineffective. In Section 5.2 we find that the majority of the commercially available detectors we examined are unusable. The apps we tested primarily relied on the magnetometer built into smartphones, which allowed them to act as crude metal detectors, but we observe that this functionality was largely useless at distances more than a few inches from the spy device. The physical detector we purchased is equally ineffective. The infrared lens detectors yielded large amounts of false positives and the magnetometer only detected magnets within a few inches of its sensor, requiring a brute-force search for spy devices. The RF detector can only pick up RF signals while the device is transmitting data, requiring the abuser to be actively streaming video to be useful. Finally, both the RF detector and the phone magnetometers are extremely sensitive to the environment they are being used in, making their results difficult to interpret.

Academic works show more promise, but they are very limited (Section 5). Both SnoopDog [54] and Lumos [53] can only detect WiFi-enabled devices, and even then they can only detect the devices if they are actively streaming.

As such, current detection techniques are not up to the task of detecting the spy devices available to abusers. Devices that don’t use WiFi, such as LTE GPS trackers or cameras that only use local storage, cannot be detected and localized with any

current technique. Finally, WiFi devices that are not streaming are also difficult to localize with current technologies.

Directions for Future Detection Tools. The spy devices we found can be categorized into four groups based on the communication technology they rely on for data transmission: WiFi, BLE, LTE, and no transmission / local storage. Detection solutions are therefore required for all of these categories, as all four types of devices have the potential to be used for IPS. Here, we outline the challenges in detecting devices in each category, as well as some potential approaches to overcome those challenges.

WiFi. Some prior works [53, 54] have studied how to detect and localize WiFi devices, but none has tried to make a *usable* suite of detection and localization tools that an IPV survivor can use. Current tools such as Lumos and SnoopDog are still prototypes and can only detect cameras in specific situations. Future research must move beyond technical questions regarding WiFi device localization and work towards making an app or physical tool that is usable and workable in real world IPS situations. Preferably, researchers will work with survivors of IPS and their advocates in the design process to ensure that any resulting tools are easy and safe to use.

BLE. In regards to BLE localization, iPhones claim to localize AirTags but the resolution is low. As a result, it is very difficult to find a hidden AirTag in a car. Moreover, there is no practical localization tool for Android users, and no localization tool for Tiles. Future work must make a universal tool with sufficient resolution in spaces such as cars to enable a survivor to find these devices.

LTE. LTE devices are difficult to detect due to the large range of wavelengths used by different LTE carriers, cell towers, and countries. Also, our experiments have found that LTE based GPS trackers polled very slowly — once every 30 seconds or slower. Any detection and localization system must be built to cope with a wide range of potential wavelengths and with the fact that small amounts of data is sent by the trackers. This is an unexplored area of research, so future work must be exploratory in nature.

Local Storage. Most detection techniques we have examined rely on detecting wireless transmissions. As such, devices that do not transmit information pose a unique challenge. Devices such as Non-linear Junction Detectors (NLJDs) claim to detect arbitrary circuit boards, but are prohibitively expensive. Moreover, they require a brute-force search of an area to ensure that it is clean. Future work should focus on making a reliable solution that is cheaper than a NLJD, as well as attempting to avoid the necessity of brute-force techniques.

7 Conclusion

In this paper we demonstrated that spy devices pose a threat and there are few mitigations against this threat. Spy de-

vices are widely available to the average consumer for small amounts of money. These devices effectively enable IPS without the abuser having any specialized technical skill. Commercially available device detection tools are unusable and often fail to detect anything. Academic detection tools, despite being grounded in promising theory, are far from ready for wide-scale deployment. We encourage the research community to expand upon existing detection techniques while ensuring these techniques are usable for average people, and we encourage online retailers and lawmakers to put rules in place that curtail the sale of devices intended to enable IPS. As a community of security researchers, we have a responsibility to ensure the security of vulnerable populations. This work provides us with an opportunity to fulfill that responsibility.

Acknowledgements

We thank the anonymous reviewers and the shepherd for their insightful feedback. This research was partly funded by University of Wisconsin–Madison Office of the Vice Chancellor for Research and Graduate Education with funding from the Wisconsin Alumni Research Foundation.

References

- [1] Strava | Run and Cycling Tracking on the Social Network for Athletes. URL: <https://www.strava.com/>.
- [2] If my husband put a security camera in our house to spy on me, do I have a right to be upset? There is a history of me not being fully honest with him, nothing related to infidelity or anything close to that. Quora, February 2018. URL: <https://www.quora.com/If-my-husband-put-a-security-camera-in-our-house-to-spy-on-me-do-I-have-a-right-to-be-upset-There-is-a-history-of-me-not-being-fully-honest-with-him-nothing-related-to-infidelity-or-anything-close-to-that>.
- [3] Intimate Partner Violence. National Institute of Justice, September 2021. URL: <https://nij.ojp.gov/topics/crimes/violent-crimes/intimate-partner-violence>.
- [4] Prohibited Products Policy and Product Limitations for Marketplace Sellers, October 2021. [Online; accessed 7. Oct. 2022]. URL: https://sellerhelp.walmart.com/s/guide?language=en_US&article=000006005.
- [5] JMDHKK Anti Spy Detector, Bug Detector, Hidden Camera Detectors, GPS Detector, RF Signal Scanner Device Detector for GPS Tracker Listening Device Camera Finder. Amazon, 2022. URL: https://www.amazon.com/Detector-Wireless-Signal-Listening-Scanner/dp/B07B93347H/ref=sr_1_4.
- [6] Prohibited and restricted items, October 2022. [Online; accessed 7. Oct. 2022]. URL: <https://www.ebay.com/help/policies/prohibited-restricted-items/prohibited-restricted-items?id=4207>.
- [7] Restricted products, September 2022. [Online; accessed 7. Oct. 2022]. URL: <https://sellercentral.amazon.com/help/hub/reference/external/200164330>.
- [8] What to do if you get an alert that an AirTag, Find My network accessory, or set of AirPods is with you, Sep 2022. URL: <https://support.apple.com/en-us/HT212227>.
- [9] 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.
- [10] Apple. AirTag. Online, 2022. URL: <https://www.apple.com/airtag/>.
- [11] Rosanna Bellini, Emily Tseng, Nora McDonald, Rachel Greenstadt, Damon McCoy, Thomas Ristenpart, and Nicola Dell. "So-called privacy breeds evil" Narrative justifications for intimate partner surveillance in online forums. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW3):1–27, 2021.
- [12] Nellie Bowles. Thermostats, Locks and Lights: Digital Tools of Domestic Abuse, Jun 2018. URL: <https://www.nytimes.com/2018/06/23/technology/smart-home-devices-domestic-abuse.html>.
- [13] Sinead Butler. Boyfriend says 'I'm not a creep' as he uses a surveillance camera to catch cheating girlfriend. Indy100, June 2021. URL: <https://www.indy100.com/viral/boyfriend-surveillance-camera-cheating-girlfriend-tiktok-b1860463>.
- [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] 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, 2018. doi:10.1109/SP.2018.00061.
- [16] 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, ASIACCS '18*, page 1–13, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3196494.3196509.
- [17] Stephanie Chevalier. Market share of leading retail e-commerce companies in the united states as of june 2022. Statista, August 2022. URL: <https://www.statista.com/statistics/274255/market-share-of-the-leading-retailers-in-us-e-commerce/>.
- [18] Cynthia Fraser, Erica Olsen, Kaofeng Lee, Cindy Southworth, and Sarah Tucker. The new age of stalking: Technological implications for stalking. *Juvenile and family court journal*, 61(4):39–55, 2010.
- [19] 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.
- [20] 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.
- [21] Scott Gleeson. Woman used an AirTag to track boyfriend, then ran over and killed him, police say. USA Today, June 2022. URL: <https://www.usatoday.com/story/news/nation/2022/06/15/woman-airtag-track-boyfriend-death/7632348001/>.
- [22] Bridget A Harris and Delanie Woodlock. Digital coercive control: Insights from two landmark domestic violence studies. *The British Journal of Criminology*, 59(3):530–550, 2019.
- [23] 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.
- [24] Yan He, Qiuye He, Song Fang, and Yao Liu. MotionCompas: Pinpointing wireless camera via motion-activated traffic. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys ’21, page 215–227, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3458864.3467683.
- [25] Alexander Heinrich, Niklas Bittner, and Matthias Hollick. AirGuard - Protecting Android users from stalking attacks by Apple Find My devices. In *Proceedings of the 15th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, WiSec ’22, page 26–38, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3507657.3528546.
- [26] Alexander Heinrich, Milan Stute, and Matthias Hollick. OpenHaystack: a framework for tracking personal bluetooth devices via Apple’s massive find my network. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*, pages 374–376, 2021.
- [27] Alexander Heinrich, Milan Stute, Tim Kornhuber, and Matthias Hollick. Who can Find My devices? Security and privacy of Apple’s crowd-sourced Bluetooth location tracking system. *arXiv preprint arXiv:2103.02282*, 2021.
- [28] Danny Yuxing Huang, Noah Apthorpe, 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.
- [29] Zoe Christen Jones. "American Idol" winner Laine Hardy arrested in Louisiana for allegedly placing recording device in ex-girlfriend’s room. CBS News, April 2022. URL: <https://www.cbsnews.com/news/laine-hardy-arrest-louisiana-recording-device-american-idol/>.
- [30] 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>.
- [31] Roxanne Leitão. Digital technologies and their role in intimate partner violence. In *Extended abstracts of the 2018 CHI conference on human factors in computing systems*, pages 1–6, 2018.
- [32] 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.
- [33] Roxanne Leitão. Technology-facilitated intimate partner abuse: a qualitative analysis of data from online domestic abuse forums. *Human-Computer Interaction*, 36(3):203–242, 2021.
- [34] 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>.
- [35] Karen Levy and Bruce Schneier. Privacy threats in intimate relationships. *Journal of Cybersecurity*, 6(1):tyaa006, 2020.
- [36] Zhengxiong Li, Zhuolin Yang, Chen Song, Changzhi Li, Zhengyu Peng, and Wenyao Xu. E-Eye: Hidden electronics recognition through mmWave nonlinear effects. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*, pages 68–81, 2018.
- [37] Tian Liu, Ziyu Liu, Jun Huang, Rui Tan, and Zhen Tan. Detecting wireless spy cameras via stimulating and probing. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys ’18, page 243–255, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3210240.3210332.
- [38] 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/>.
- [39] Ryan Mac and Kashmir Hill. Are Apple AirTags Being Used to Track People and Steal Cars? . New York Times, December 2021. URL: <https://www.nytimes.com/2021/12/30/technology/apple-airtags-tracking-stalking.html>.
- [40] Tara Matthews, Kerwell Liao, Anna Turner, Marianne Berkovich, Robert Reeder, and Sunny Consolvo. "She'll just grab any device that's closer" A Study of Everyday Device & Account Sharing in Households. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 5921–5932, 2016.
- [41] 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.
- [42] Travis Mayberry, Ellis Fenske, Dane Brown, Jeremy Martin, Christine Fossaceca, Erik C Rye, Sam Teplov, and Lucas Foppe. Who tracks the trackers? Circumventing Apple’s anti-tracking alerts in the Find My network. In *Proceedings of the 20th*

Workshop on Privacy in the Electronic Society, pages 181–186, 2021.

- [43] Markus Miettinen, Samuel Marchal, Ibbad Hafeez, N. Asokan, Ahmad-Reza Sadeghi, and Sasu Tarkoma. IoT SENTINEL: Automated device-type identification for security enforcement in IoT. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, pages 2177–2184, 2017. doi:10.1109/ICDCS.2017.283.
- [44] Moh, Datta, Warford, Bates, Malkin, and Mazurek. Characterizing Everyday Misuse of Smart Home Devices. In *2023 IEEE Symposium on Security and Privacy (SP)*, volume 0, page 1558–1572, Dec 2023. URL: <http://dx.doi.org/10.1109/SP46215.2023.00089>, doi:10.1109/SP46215.2023.00089.
- [45] 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.
- [46] Joint News Release. Devastatingly pervasive: 1 in 3 women globally experience violence. World Health Organization, March 2021. URL: <https://www.who.int/news/item/09-03-2021-devastatingly-pervasive-1-in-3-women-globally-experience-violence>.
- [47] K Andrew R Richards and Michael A Hemphill. A practical guide to collaborative qualitative data analysis. *Journal of Teaching in Physical education*, 37(2):225–231, 2018.
- [48] 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.
- [49] Said Jawad Saidi, Anna Maria Mandalari, Roman Kolcun, Hamed Haddadi, Daniel J Dubois, David Choffnes, Georgios Smaragdakis, and Anja Feldmann. A haystack full of needles: Scalable detection of IoT devices in the wild. In *Proceedings of the ACM Internet Measurement Conference*, pages 87–100, 2020.
- [50] Sriram Sami, Sean Rui Xiang Tan, Bangjie Sun, and Jun Han. LAPD: Hidden spy camera detection using smartphone time-of-flight sensors. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, SenSys ’21, page 288–301, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3485730.3485941.
- [51] Dinesh Sathyamoorthy, Mohd Jalis Md Jelas, and Shalini Shafii. Wireless spy devices: A review of technologies and detection methods. *Defence S&T Technical Bulletin*, 7(2):130–139, 2014.
- [52] Aarti Shahani. I Know Where You’ve Been: Digital Spying And Divorce In The Smartphone Age. NPR, January 2018. URL: <https://www.npr.org/sections/alltechconsidered/2018/01/04/554564010/i-know-where-you-ve-been-digital-spying-and-divorce-in-the-smartphone-age>.
- [53] Rahul Anand Sharma, Elahe Soltanaghaei, Anthony Rowe, and Vyas Sekar. Lumos: Identifying and localizing diverse hidden IoT devices in an unfamiliar environment. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 1095–1112, 2022.
- [54] 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.
- [55] Arunan Sivanathan, Hassan Habibi Gharakheili, Franco Loi, Adam Radford, Chamith Wijenayake, Arun Vishwanath, and Vijay Sivaraman. Classifying IoT Devices in Smart Environments Using Network Traffic Characteristics. *IEEE Transactions on Mobile Computing*, 18(8):1745–1759, 2019. doi:10.1109/TMC.2018.2866249.
- [56] 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.
- [57] 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>.
- [58] Cynthia Southworth, Jerry Finn, Shawndell Dawson, Cynthia Fraser, and Sarah Tucker. Intimate partner violence, technology, and stalking. *Violence against women*, 13(8):842–856, 2007.
- [59] Sophie Stephenson, Majed Almansoori, Pardis Emami-Naeini, and Rahul Chatterjee. “It’s the equivalent of feeling like you’re in jail”: Lessons from firsthand and secondhand accounts of IoT-enabled intimate partner abuse. In *32nd USENIX Security Symposium (USENIX Security 23)*, 2023.
- [60] Sophie Stephenson, Majed Almansoori, Pardis Emami-Naeini, Danny Yuxing Huang, and Rahul Chatterjee. Abuse vectors: A framework for conceptualizing IoT-enabled interpersonal abuse. In *32nd USENIX Security Symposium (USENIX Security 23)*, 2023.
- [61] Royal Tech App Studio. Detect bug - camera microphone. Google Play Store, 2022. URL: https://play.google.com/store/apps/details?id=com.royaltechapps.hiddencameradetector&hl=en_US&gl=US.
- [62] 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.
- [63] 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.

- [64] Tile. Find Your Keys, Wallet & Phone with Tile’s App and Bluetooth Tracker Device | Tile, 2022. URL: <https://www.thetileapp.com/>.
- [65] 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.
- [66] Zeehik IT Zon. Hidden Devices Detector. Google Play Store, 2022. URL: https://play.google.com/store/apps/details?id=com.zeehikitzon.hiddendevicesdetector&hl=en_US&gl=US.

A Additional Method Details

A.1 Classifying Spy Devices

It would be infeasible to manually classify all 6,403 products we scraped when searching for spy devices. To deal with this, we developed two classifiers: A heuristic filter that identifies if a product is technology-related or not and a logistic regression classifier that classifies technological products as spy devices or irrelevant products.

Heuristic filtering. To determine whether a given product is *technological* and *gathers information*, we create a heuristic filtering. We assemble a list of phrases that indicate a product is relevant to our study, such as “GPS”, “camera”, and “microphone”. We then assemble a similar list of phrases that indicate a product is irrelevant, such as “game”, “novel”, and “shoes”. Each positive and negative phrase was given a score from 0 to 5. For each product, we found all positive and negative phrases within the title and description of the product and compared the maximum positive score to the maximum negative score. If the negative score was greater than the positive score, or if no positive phrases were found, the product was removed from the sample.

We designed the heuristic filtering to achieve a low false negative rate—so that we do not miss any relevant products. To evaluate its performance, we manually classified 200 products. The final iteration of the heuristic filtering achieved a false negative rate of less than 2%. When used on all the products we gathered, the classifier marked 2,996 products as irrelevant and help remove 47% of irrelevant products.

Logistic regression classification. To ease the task of classifying the products we scraped we made use of a linear regression classification model. To generate training data for the model, we manually classified 869 products, finding 596 of them to be relevant spy devices and 273 of them to be irrelevant. Using these manually classified product listings, we trained our model. We used the product descriptions and the product names as features for the classifier, processing them into buckets of words before training. We randomly selected

half of the manually classified data to be used as training data and the rest to be used as test data. After training, the model achieved 82% precision and 98% recall on relevant spy devices (and 96% precision and 70% recall on irrelevant devices). When used on the 3,407 products marked as technological by the heuristic filtering, the logistic regression classifier marked 2,228 (65%) as relevant spy devices.

A.2 Filtering Detection Devices

Since the number of detection devices was smaller than the number of spy devices in our previous crawl (Section 3), we were able to rely on heuristic filtering methods. We first scanned through the list of detection devices and created an initial accept list and block list. For each detection device, we looked at the device’s URL, name, description, specification, and features as listed on the product page. If the device contained a phrase in the accept list, we marked it as a detector; otherwise, if it contained a phrase in the block list, we marked it as *not* a detector; else, we marked it as undetermined. We then updated the accept list and block list, re-ran our algorithm, and repeated the process until the list of undetermined devices was small enough to manually classify (28 devices).

We tested our heuristic filter on a random sample of 50 detectors and 50 non-detectors as classified by our heuristic. We manually classified these 100 devices and identified 0 false positives and 2 false negatives. Finally, after applying this filter to the list of detection devices and manually classifying the 28 undetermined devices, we had 700 detection devices in our filtered dataset.

All phrases on the accept list and the block list for this heuristic classifier can be found here: <https://github.com/ceccio247/IPV-Spy-Device-Study>.

A.3 Empirical Analysis of Spy Devices

To understand whether these devices are effective for conducting IPS, we perform a series of lab experiments on a representative sample of spy devices. This sample consists of 7 *recording devices* and 4 *tracking devices*.

Recording devices. Our sample of recording device includes 5 cameras, two of which use WiFi to transmit data, and 2 microphones, which only use local storage. To test their effectiveness, we follow the provided instructions to set up the devices and observe what, if anything, they record. We note whether they function at all, what data they record, how much data they can record, and if they have to be manually activated or if they are activated by activity such as motion or speech. These experiments are performed within controlled spaces where the only sensitive data they could record is that of the researchers. All data recorded by these devices was later deleted.

Tracking devices. To determine the effectiveness of our

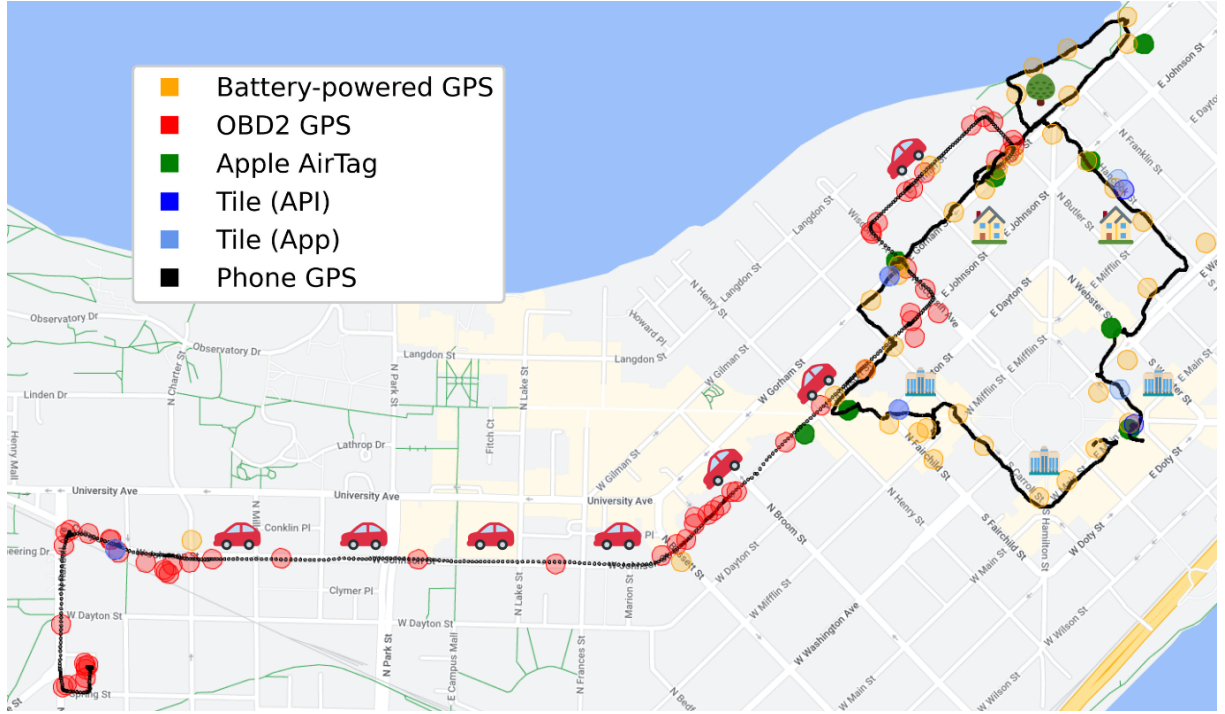


Figure 11: Map depicting our path and the data reported by our recording devices. The black marks represent the ground truth as reported by the Strava app [1] running on an Android phone. The other dots correspond to the reported locations from our tracking devices. Red is the OBD2 GPS tracker, yellow is the battery-powered GPS tracker, green is the Apple AirTag, blue is the Amazon Tile. We include two blues, one dark (representing the data from the Tile API), and one light (representing the data from the Tile app). Note that in some places these points overlap. The emojis indicate the path corresponds to each portion of the experiment: cars for the driving portion, trees for the park, houses for the residential areas, and offices for the commercial areas.

tracker devices, we construct a route that simulates a realistic routine for a survivor of IPV. Our route starts by driving from our lab through areas with multiple apartment buildings to a residential area with single-family homes. The route then makes a walking circuit of a local park followed by a circuit of the downtown of our city. During our time downtown we stop briefly at a store and a restaurant. Fig. 11 depicts our route. The entire route takes about one hour to complete.

During the experiment, we use the two GPS tracking devices and the two Bluetooth Mesh tracking devices and activate them at the beginning of the route. As one of the GPS tracking devices relies on power from a car’s OBD2 port and only claims to track cars, we only use it for the driving portion of the experiment. To approximate ground truth location, we also bring an Android Google Pixel 2 phone that continuously records our GPS coordinates using the Strava application [1].

After two walks, we save and analyze the recorded location information. Both our GPS trackers allow a user to directly download the recorded GPS coordinates, but neither Tile nor AirTag allow their users to download precise location history. To extract the GPS coordinates for these two systems, we used scripts to access Tile’s hidden API and read the unprotected cache files of Apple’s Find My application.