

Location-Enhanced Information Flow for Home Automations

McKenna McCall*[†]
Colorado State University
mckenna.mccall@colostate.edu

Ben Weinschel*
Carnegie Mellon University
bweinschel@cmu.edu

Kunlin Cai
University of California, Los Angeles
kunlin96@ucla.edu

Ying Li
University of California, Los Angeles
ying.li@ucla.edu

Eric Zeng[†]
Georgetown University
eric.zeng@georgetown.edu

Devika Manohar
Carnegie Mellon University
dmanohar@andrew.cmu.edu

Lujo Bauer
Carnegie Mellon University
lbauer@cmu.edu

Limin Jia
Carnegie Mellon University
liminjia@andrew.cmu.edu

Yuan Tian
University of California, Los Angeles
yuant@ucla.edu

Abstract

Smart-home automations enable users to customize smart devices to react automatically to people, the environment, and more. For example, an automation might adjust the lights when people are at home or enable a garage door to open by voice command. While automations offer convenience and accessibility, they can also inadvertently expose users to security and privacy risks, such as leaking sensitive data or allowing untrusted parties to control users' devices. Prior work has shown that information flow analysis is a promising technique for identifying these kinds of risks, hypothesizing that the analysis would be yet more effective if it could differentiate between devices located in different places in the home.

We tested this hypothesis by developing a tool that extends prior information flow analysis approaches to account for device location. We conducted an interview study with 22 participants to build a dataset of home automations to establish a ground truth to evaluate the tool. We found that incorporating device location leads to an improved analysis that identifies more of the vulnerabilities users care about (F1 score 0.74) compared to prior work (F1 score 0.29). Our results demonstrate the feasibility of incorporating device location into an information flow analysis and, perhaps more importantly, suggest additional ways to prevent security and privacy risks beyond controlling potentially unsafe information flows.

Keywords

information flow analysis, smart home, home automation, trigger-action programs

1 Introduction

Smart-home devices are increasingly used to monitor and control the home, including lighting, appliances, heating and cooling, and security systems. In addition to viewing and controlling devices from a smartphone app, smart homes also allow users to automate

functionality using *home automations*: programs or integrations that allow user interactions, state changes in devices, the physical environment, and online services to trigger an action on another device [30, 38]. For example, users with limited mobility can control devices through commands to a voice assistant (e.g., Amazon Echo), or a user can configure their smart speaker to alert them when products they frequently buy are on sale.

While home automations offer convenience and accessibility, they can also lead to security and privacy risks. In the automation mentioned above, audio alerts tracking sales can leak a user's personal shopping habits to guests in the home. Similarly, an unauthorized user could interact with a voice assistant to control a user's device, like a smart oven, without their permission [31].

When an automation is run, information is sent from the triggering device to the action device, allowing people who observe the action to infer information about the trigger. Due to this relationship between trigger and action, one approach to detect vulnerabilities created by home automations is information flow analysis [11]. Prior work has used information flow analysis to identify potential risks in IFTTT automations [10, 34] but their datasets lacked details to determine whether real users would actually perceive the violations as threats. In particular, the data they examined did not include information like the location of devices within the home and the purpose of automations (e.g., which might signal that users intended to expose the information). Prior work hypothesized that incorporating additional real-world context, such as device location, might lead to more accurate analyses.

In this work, we investigate if an information flow analysis incorporating smart device location is better aligned with what users consider to be security and privacy threats. We pose the following research questions:

RQ1: How do real users automate their smart homes?

RQ2: Which of the potentially unsafe information flows that we identify in home automations are users concerned about and why?

RQ3: Does incorporating device location lead to a more accurate information flow analysis?

To answer these questions, we developed an information flow analysis tool that uses information about where smart devices are located in a user's home to more accurately identify potentially risky automations. We interviewed 22 smart-home users recruited from

*Both authors contributed equally to this research.

[†]This work was completed while at Carnegie Mellon University.

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Proceedings on Privacy Enhancing Technologies YYYY(X), 1–28

© YYYY Copyright held by the owner/author(s).

<https://doi.org/XXXXXXXX.XXXXXXX>



Prolific to develop a dataset of their automations, including details about where devices are deployed and their reactions to scenarios with potentially unsafe information flows. We used the interview data to develop a dataset of 373 home automations. To the best of our knowledge, ours is the largest cross-platform home-automation dataset collected from real users. By conducting interviews instead of a larger-scale survey, we gained rich insights into how and why participants automate their homes. The interviews also gave us flexibility to adjust our data analysis to include built-in automations, which we did not anticipate being relevant before the interviews. For the potentially unsafe information flows discussed during the interviews, we assigned ground-truth labels reflecting if the user actually perceived them as concerning, which we used to evaluate our location-enhanced analysis against prior work [10].

Our evaluation using this dataset demonstrates that incorporating device location into an information flow analysis improves the analysis—allowing it to achieve an F1 score of 0.74, compared to prior work’s 0.29—and identifies more of the vulnerabilities that users care about. While still imperfect, our analysis also encouraged people to consider their automations’ safety, showing it is a promising approach for communicating with users about the potential security and privacy risks in their smart homes.

In summary, our paper makes the following contributions:

- The largest, up-to-date dataset of smart-home automations collected from real users via interviews. Unlike datasets collected by prior work, our dataset is multi-platform and includes both end-user and built-in automations, locations of relevant devices, and ground-truth data about which security and privacy violations users actually perceive as risky.
- A novel strategy for incorporating smart device location into and information flow analysis, which detects many violations missed by prior work.
- New insights about how people automate their homes and the role of information flow analysis in identifying security and privacy risks.

The rest of the paper is organized as follows: Section 2 describes background and related work. We describe our location-enhanced information flow analysis in Section 3 and our interview protocol in Section 4. Section 5 describes the resulting dataset of automations (RQ1) and participants’ reactions to our information flow scenarios (RQ2). Our Prolog implementation and its evaluation (RQ3) are described in Section 6. In Section 7, we describe other findings from our interviews about how visitors to the home interact with smart-home devices, and in Section 8, we present additional discussion.

2 Background and Motivation

In this section, we define terminology (Section 2.1), summarize the prior work most closely related to ours (Section 2.2), and describe the information flow analysis used by prior work, which we illustrate through a series of examples (Section 2.3).

2.1 Home Automation Terminology

Home automations, where devices in the home take action in response to events, such as a smart speaker turning on lights in response to a user request, or heat being turned on as a result of a preconfigured schedule, are core features of popular smart-home

platforms. User-configured automations often follow the *trigger-action programming* paradigm, where users select a trigger for the automation (e.g., a person leaving home), and then choose actions that run in response (e.g., turn off the lights). Automations written as trigger-action programs are typically formatted as *IF trigger THEN action*, which we also use for all of the automations in this paper. Trigger-action programming is a feature of many smart-home platforms, including Amazon Alexa (routines), Google Home (routines), and Apple Home (automations). Platforms like IFTTT also offer trigger-action programming [2, 4, 14, 16, 19].

Smart-home platforms offer other automation features, such as scenes, which allow users to group devices to control their state. We refer to automations written through trigger-action programming or custom scenes as *end-user programming*. Additionally, voice assistants allow users to control devices on demand. Even though users do not specifically configure these interactions, they follow the same automation paradigm, where devices take action in response to commands. We refer to these as *built-in automations*.

One of our main contributions is a location-aware information flow analysis. For brevity, we shorten locations in the home to one or two letters, such as “B” for Bedroom and “LR” for Living Room.

2.2 Related Work

Smart-home security and privacy. Smart-home devices present many privacy and security risks to users. Always-on cameras, microphones, and other sensors capture privacy-sensitive data [1, 3, 21, 29], which users may be unaware of [35, 44]. Security vulnerabilities in smart-home devices expose users to remote attackers [20]. Moreover, smart homes are multi-user systems that affect the privacy of not only the primary user but also of household members, guests, and people outside the home [9, 13, 31, 37, 41].

Prior work has explored potential security and privacy solutions, such as network connection monitoring and blocking [17], privacy disclosures and nutrition labels [6], and sensor controls [12]. Researchers have acknowledged that commercially available smart-home platforms have limited access control capabilities and proposed mechanisms to improve privacy and security, including controls that vary based on user, device type, and device location [15, 42] and “optimistically” via notifications to other users [23].

Trigger-action programs. While trigger-action programming offers flexibility, reasoning about its behavior can be difficult [18, 40]. Researchers have explored how to identify issues in trigger action programs, such as comparing program rules with a user-defined policy [5, 22, 43], identifying “anti-patterns” [39], or using automation rules to generate realistic evaluation scenarios [24, 25].

To assess the prevalence of risky trigger-action programs, prior work has collected data about users’ automations through various sources (e.g., scraping publicly available IFTTT and Zapier rules) and used that data to evaluate tools [7, 8, 34, 39]. Other work has sourced automations through crowdsourcing IFTTT data exports [10] and surveying technical experts about possible automations they may create [24, 25]. Trigger-action program research has generally focused on IFTTT, which has a large public dataset of automations crawled by prior work [7, 8, 30, 34, 38, 39]. While these datasets include a large number of programs, they are not necessarily reflective of users’ full smart-home usage, limiting the




Figure 1: Simple lattices for secrecy and integrity. Information is always allowed to flow upwards; if any information flows downwards (e.g., a private input sharing data with a public output), that constitutes a violation.

generalizability of analyses of the data. These datasets are based on publicly available, crowdsourced, and/or synthetic data, generally include data from a single platform, and do not include detailed data about the home layout or automations users chose not to publish. Moreover, these datasets do not include the location of devices and they tend to overlook built-in functionality. To address this gap, our dataset includes both end-user and built-in automations from several smart-home platforms, as well as data about device locations and ground-truth data about which potential information flow violations are concerning to users.

Prior work has also examined whether security tools that analyze risks in trigger-action programs can produce usable feedback for end users [27]. They found that users were capable of using tools that assisted them in identifying programming anti-patterns, specifying invariants for smart-home state, and repairing programs that violated those invariants.

2.3 Information Flow Analysis

Information flow analysis identifies security and privacy risks by examining how the inputs to a program influence its outputs. In particular, information flow analysis can identify *secrecy violations*, where a secret input influences a public output, and *integrity violations*, where an untrusted input influences a trusted output. In the context of home automations, an automation that sends a notification to family members when someone arrives home causes a secrecy violation, as the user’s private location is broadcast to a more-public output. An automation that allows someone to ask an Amazon Echo to add items to a private shopping list causes an integrity violation, as the personal shopping list is more trusted than the people who can operate the Echo, which is accessible to anyone physically present in the home.

To perform an information flow analysis, the triggers and actions of every automation should be assigned *labels* to indicate their sensitivity (i.e., how secret the trigger or action is or who is authorized to learn about the trigger or action) and integrity (i.e., how trusted the trigger or action is or who is authorized to perform the trigger or action). The labels are arranged in a *lattice* to indicate which trigger or action is more secret or trusted. Information is only allowed to flow up the lattice, so more secret labels will appear at the top of the secrecy lattice and more untrusted labels will appear at the top of the integrity lattice. For example, a simple secrecy lattice involving only two labels `public` and `private` with `private` as the top-most label would not permit information flows from a trigger labeled `private` to an action labeled `public` (Figure 1). Importantly, an information flow lattice specifies which information flows are




Figure 2: Lattice of information flow labels used by Cobb et al. [10], which we call a *coarse-grained lattice*.

permitted. Any information flow from a trigger to an action that does not have the same label or cannot be reached by traversing the path up the lattice is *disallowed*.

Prior work has explored using information flow analysis for trigger-action programs on IFTTT. Surbatovich et al. proposed specialized lattices for IFTTT applets that use four secrecy labels and six integrity labels [34]. Cobb et al. expanded the lattice to include two additional labels and unions of labels (Figure 2) [10]. In the remainder of this section, we demonstrate how these lattices can be used to determine if an automation has a secrecy or integrity risk and discuss their limitations.

Benign information flow. The automation “IF the time is 7AM THEN turn on lights” is a benign automation that does not cause a secrecy or integrity violation. Using the labels in the coarse-grained lattice, the trigger (time is 7AM) would have the secrecy label `public` because anyone can know the time and the integrity label `trusted other` because it is determined by a trusted source other than the user. Meanwhile, the action would be labeled `restricted physical` for both secrecy and integrity. Thus, the information flows up the lattice from a more public/trusted trigger to a less public/trusted action and the automation is considered safe.

Information flow causing a secrecy violation. The automation “IF new sale THEN play an audio notification on the kitchen smart speaker” violates the secrecy of the user’s shopping habits. Using the labels in the secrecy lattice, the trigger is `private` because the user’s shopping habits should only be known by them, and the action is `restricted physical` since anyone in the physical vicinity could overhear the audio notification. The analysis would flag this as a secrecy violation because information flows down the lattice from a more private trigger to a more public action.

Information flow causing an integrity violation. The automation “IF someone says *Alexa add to my shopping list* THEN add to my shopping list” violates the integrity of the user’s shopping list. The integrity label of the trigger is `restricted physical` because anyone nearby can speak to Alexa, and the action is `restricted`

online to reflect the user’s personal shopping list. In this case, the information does not flow down the lattice, but is still an integrity violation because the trigger is not at least as trusted as the action.

Violations not captured by prior work. While the coarse-grained information flow analysis from prior work is capable of identifying these unsafe information flows, there are some secrecy and integrity violations that this technique cannot identify. For example, in the automation “IF someone says *Alexa turn the upstairs thermostat to [temperature] degrees* THEN adjust temperature,” the secrecy and integrity labels of both the trigger and action would be restricted physical. However, this automation could still be risky. In a home with private bedrooms on the upper floor and a voice assistant in the living room on the lower floor, a person adjusting the upstairs thermostat from the living room is fundamentally different from a person adjusting it from the physical thermostat on the upper floor. The voice assistant in the living room is more likely to be accessible to guests in the home than the thermostat upstairs. Therefore, this automation is a possible integrity violation because someone (a guest) could interact with a device to change its settings in a way that was not possible without automation. These cross-location risks cannot be captured by the existing coarse-grained information flow analysis.

Violations that are not concerning. The technique from prior work is also likely to flag information flows that pose little risk to the user as unsafe. For instance, its labeling scheme only labels things coming from the user’s personal devices as trusted. This means that the seemingly benign automation “IF the time is 7AM THEN send me a good morning notification” would be an integrity violation because the trigger (trusted other) would be considered less trusted than the action (trusted). In this work, we develop a new lattice that incorporates device locations and how people might interact with them. We will refer to the information flow analysis involving this proposed lattice as the *location-enhanced analysis* (using the *location-enhanced lattice*). Our evaluation compares the run-time performance and analysis accuracy of our proposed approach against the one used by Cobb et al. [10], which we will refer to as the *coarse-grained analysis* (using the *coarse-grained lattice*).

3 Location-Enhanced Analysis

In this section, we describe our location-enhanced information flow analysis. At its core, an information flow analysis compares the relative confidentiality and integrity of the source (i.e., automation trigger) and sink (i.e., automation action) of an information flow. A lattice of information flow labels determines which information flows are permissible and which are potentially risky. The challenge in designing a location-enhanced information flow analysis, then, is determining how to incorporate location into this label comparison.

Recall the potentially risky automation from Section 2.3 “IF someone says *Alexa turn the upstairs thermostat to [temperature] degrees* THEN adjust temperature.” In a home with private bedrooms on the upper floor, this automation should be identified as an integrity violation since someone in the living room could control the temperature in a more private space. The coarse-grained analysis would not identify this as an unsafe information flow because the information flow lattice is not expressive enough to distinguish between devices in different rooms.

To address this challenge, we develop a location-aware information flow analysis. Naïvely, we could take a traditional information flow approach and assign different locations in the home labels based on their relative sensitivity and trustworthiness. However, this would require each user to answer unintuitive questions, such as whether their basement is more sensitive than their office. Instead, we propose a novel location-enhanced lattice (shown in Figure 3), which is based on where people can interact a device, rather than the sensitivity and trustworthiness of the device location.

As the layout of every home differs, acceptable information flows between rooms naturally differ, too. In a secrecy analysis, an information flow from trigger to action is permitted if the triggering event can be learned from the same or more places in the home as the action. Similarly, in an integrity analysis, an information flow is permitted if the triggering event can be activated from the same or fewer places as the action. As such, while there are labels that will appear in every location-enhanced lattice (e.g., the top-most label of the integrity lattice is always Least Trusted and the bottom-most label is Most Trusted), the location-enhanced lattice also has a location sub-lattice specific to a user’s home. The labels in the sub-lattice are a set of rooms, which reflect where devices can be interacted with. Labels are arranged such that information is allowed to flow from l to l' if the rooms in l' are a subset of l .

For example, in the location-enhanced integrity lattice for the automation “IF someone says *Alexa turn the upstairs thermostat to [temperature] degrees* THEN adjust temperature” (Figure 3b), if the thermostat in the upstairs hallway (UH) can also be triggered from the voice assistant in the living room (LR), the trigger would be labeled {UH, LR}. Meanwhile, the action would typically only be able to be triggered by people interacting with the physical thermostat, so it would have the label {UH}. Since different people would be able to activate the trigger than the action, this automation would be flagged as an integrity violation.


We also revisit how some triggers and actions are labeled. Prior work labeled common triggers like time trusted other [10], which led to false positives when automations triggered trusted behaviors at the same time every day. In the last example from Section 2.3, the automation “IF the time is 7AM THEN send me a good morning notification” would be flagged as a potential violation since a less-trusted trigger (labeled trusted other) triggers a more-trusted action (labeled trusted). In our location-enhanced analysis, we label time trusted, so this benign automation would no longer be flagged as a potential violation.

4 Interview Methods

To evaluate our new information flow lattice with a more detailed and generalizable dataset of automations than previously available, we built a cross-platform dataset of real home automations from 22 smart-home users. In this section, we describe our recruitment strategies, interview structure, and data analysis methods.

4.1 Recruitment

We used a series of screening surveys to recruit participants from Prolific (Figure 4). We screened participants and invited them for interviews in batches. Screening began in March 2024 and interviews were conducted between March and July 2024.



(a) Example secrecy lattice. Labels toward the bottom are for triggers and actions that can be viewed or controlled by more people or from more places (i.e., more public); labels toward the top are for triggers and actions that can be viewed or controlled by fewer people and in fewer places (i.e., more private).

(b) Example integrity lattice. Labels toward the bottom are for triggers and actions that can be viewed or controlled by fewer people and in fewer places (i.e., more trusted), while labels toward the top are for triggers and actions that can be viewed or controlled by more people and in more places (i.e., less trusted).

Figure 3: Location-enhanced lattices. The dark gray boxes include device locations specific to each user. The letters in the location lattice represent rooms in the home. Here, B represents a bedroom, UH the upstairs hallway, and LR the living room.




Figure 4: Our recruitment process involved two screening surveys to find participants who use home automation features and were willing to participate in an interview.

Initial screening survey. Our first screening survey was designed to identify smart-home users interested in doing an interview study. We recruited participants in several batches. The first batch recruited were adult users located in the US who indicated on their Prolific profile that they speak English fluently and own at least one smart device. To ensure we had a diverse dataset, we used Prolific’s built-in prescreeners in subsequent recruitment batches to recruit people with different smart-home devices than previous batches. We recruited a gender-balanced sample to attempt to recruit a demographically diverse sample.

In the survey, we explained what home automations are, asked whether the participant currently uses home automations, and if they are interested in participating in an interview.¹ If they indicated that they used automations and were interested, we asked how many automations they have installed. The median time for the first screening survey across all batches was between 1 and 2.5 minutes, and participants were paid \$0.75, regardless of whether

¹After the first 18 interviews, we recruited another batch of participants to find more people with at least 10 automations. For this batch, only people with at least 10 automations who were willing to share a screenshot of their automations were invited to the next screening survey.

they qualified for the next screening survey. We screened a total of 1,729 participants.

Secondary screening survey. Participants who completed the first screening survey and reported having at least five home automations were eligible for the second round of screening. The purpose of the second screening survey was to gather more data about their smart-home environment. Again, participants were invited to the second screening survey in batches, starting with people who had the most automations.

In this survey, participants were asked to share more details about their home setup and use of automations, including which platforms they used, the types of devices and automations, their use cases for automations, and details about their home layout (e.g., single family home or apartment). We used their responses to prepare our interview script. For example, if participants told us during the screening survey that they lived in a shared apartment building and owned a voice assistant, we included questions about whether their voice assistant responded to people in the hallway. Of the 156 participants we invited, 83 completed the second screening survey. The median time was 4–7 minutes across all recruitment batches and participants were paid \$2.50, regardless of whether we invited them to schedule an interview.

Participant selection and scheduling. We then reviewed responses to invite people to our interviews. Participants with many (at least 10) automations, less common devices (i.e., something other than a voice assistant or smart light bulb), or unique automation use cases (e.g., accessibility) were invited to interviews. Additional invitations were given to participants who wrote thoughtful responses to our question about what they use home automation for.




Figure 5: Our interviews included four sections to collect a dataset of automations and ground truth for our information flow analysis.

Participants we selected were invited to a final Prolific task, where they completed the consent form for the interview and scheduled a time for the interview using Calendly. After completing the interview, which took a median time of 55 minutes, participants were paid \$20.00. We interviewed participants until we reached saturation with respect to the use cases and themes discussed and then interviewed an additional four participants with many (at least 10) automations to ensure we had a large dataset for our evaluation. We invited a total of 54 participants to schedule an interview, scheduled 30 interviews, and completed 22.

The complete screening survey questions and more detailed information recruitment may be found in Appendix A.

4.2 Interview

We piloted our interview with three volunteers from the lead author’s institution. With permission, we recorded our pilot interviews to test microphones and transcription software. The data collected were used to refine the interview questions and the analysis procedures. The results of the sample analysis are not included in this paper.

Two authors conducted the semi-structured interviews. One acted as the primary interviewer, while the other took notes and asked occasional clarifying questions. The interviews were conducted remotely and recorded using Zoom. We used Otter.ai to generate transcripts, and two researchers manually verified the output. We also recorded the call video, so we could review the parts of the interview where participants shared their screens with us. The interview structure is outlined in Figure 5 and summarized below; a full script may be found in Appendix B. While most of the interview was scripted, interviews sometimes deviated from the script to ask follow-up questions and when we discussed information flow scenarios.

Background. The interview began with introductions, reviewing the consent form, and explaining the interview would be recorded. We asked participants background questions about their home automation experiences: how they create automations, problems they encounter with automations, and how guests interact with devices and automations. This helped us build rapport with participants and learn how they might benefit from a security analysis tool.

Home automations. Next, we asked participants to show us their home automations. Participants (who were able to) shared their screen and described their automations, scenes, routines, and schedules in the apps for various smart-home platforms they used.

We used this information to create the structured dataset of home automations described in Section 5.

Home layout. In the third part of the interview, we asked participants for more detail about the layout of their home, such as where they could talk to their smart speakers from. These data were used to assign labels to automations using our location-enhanced lattice. The labeling process is described in Section 4.3.

Information flow scenarios. Finally, we discussed scenarios with potential privacy violations that could be detected by an information flow analysis tool. Because the interviewees did not know which automations participants used until the interviews, they had to review participants’ automations in real time to identify possible privacy violations. The interviewees described the potential violations and asked participants how they would feel in that scenario. For example, if a participant had a smart speaker placed near a window that could control their air conditioner, we asked how they would feel if someone standing outside their window used the smart speaker to change their thermostat. The data from this section of the interview served as the ground truth for comparing our location-enhanced analysis to the coarse-grained analysis. We describe the results of the ground-truth labeling in Section 5.2 and the evaluation in Section 6.

4.3 Qualitative Coding and Data Extraction

From our interview transcripts and screen recordings, we extracted three types of data: (i) a list of participants’ home automations, (ii) secrecy and integrity labels for the automations in our dataset, and (iii) a list of information flow violation scenarios posed to participants and qualitative codes for participants’ reactions. Based on guidance from prior work, measuring inter-rater reliability was not necessary in our case [28]. Beyond developing a structured dataset for evaluating our tool, we also identified common challenges, goals, and other experiences with home automation that participants told us about during interviews.

4.3.1 Constructing a dataset of home automations. To build a structured dataset of home automations formatted for analysis, we manually extracted participants’ automations from screen recordings of their home automation apps and their verbal descriptions of their automations. We extracted each automation’s trigger(s), action(s), and associated device(s) and platform(s). Sometimes we failed to collect device information, specifically device brand and location. In these cases, we either left it unspecified (for device brand) or made an educated guess (for device location, e.g., “living room” was typically used for thermostat).

4.3.2 Information flow labeling. As mentioned in Section 2.3, in order to do an information flow analysis, the triggers and actions of every automation should be assigned a label. The goal of our study was to develop a new information flow analysis tool that is more accurate than what was used by prior work, so we needed to label everything twice: once using a coarse-grained lattice [10] and once using labels from our location-enhanced lattice. Prior to labeling the data from the interviews, one of the authors trained two other authors to do information flow labeling, which they practiced by labeling data collected during pilot interviews. Below, we describe the procedures for labeling the full interview data.

Coarse-grained analysis labeling. All three researchers independently labeled all of the triggers and actions of all of the automations. They followed the guidance from prior work [10] to determine which labels should be applied. After labeling the automations, they compared their labels and discussed any disagreements until they reached 100% agreement.

Location-enhanced analysis labeling. The same three researchers developed a heuristic for the location-enhanced labeling. Rather than double- or triple-labeling the automations, they split the automations and labeled them independently. The researchers met periodically to discuss automations that did not fit into the existing heuristics. The heuristics applied for this round of labeling can be found in Appendix D, and more information about the location-enhanced analysis can be found in Section 3.

4.3.3 Qualitatively coding reactions to information flow violation scenarios. We categorized each information flow violation scenario that we posed to participants during the interviews. One researcher extracted each information flow scenario from the interview in a structured format, listing the triggers, actions, devices, and type of violation (secrecy or integrity). The same researcher labeled participants’ reactions to each scenario. Additionally, they coded the reason users were concerned or not concerned; these codes were derived from the codebook developed in Section 4.3.4. Lastly, they identified which automations (from Section 4.3.1) correspond to the information flow scenarios, generating a mapping between the datasets. One other researcher verified the correctness of the labels and mappings for these reactions.

4.3.4 Qualitative coding for home automation background and experiences. We used an inductive qualitative coding approach to analyze the sections of the interviews about participants’ experiences using home automations. First, each interview transcript was assigned to one of three researchers, who independently generated an initial set of codes using open coding. One researcher merged these codes into a draft codebook.

Next, we iteratively refined the draft codebook. Five researchers independently coded two interviews in MAXQDA and met to seek agreement across our codes. In cases where disagreements were due to ambiguities in the codebook, we removed redundant codes, added missing codes, or clarified the names and definitions of codes. For other disagreements, the coders discussed and agreed on the correct code to use and updated the applied code. One more interview was coded independently by all five researchers, and the coders met to agree on definitions and revise the codebook.

The remaining 19 interviews were randomly divided among coders such that every interview was assigned two coders. The same group of coders coded at most three interviews together. The two coders independently coded the interviews and met to resolve all disagreements, none of which required changes to the codebook.

4.4 Ethical Principles

Our study was approved by the institutional review board (IRB) from the lead author’s institution. We used a multi-tiered screening process to ensure participants were compensated for their time but would not fill out unnecessary surveys or forms if they were not eligible to participate in the interview. Only an IRB-approved subset

of the research team had access to the interview audio/video recordings. We used cloud transcription software to create transcriptions, and one of the authors reviewed the transcripts to correct mistakes and remove any potentially identifying information before they were shared with the rest of the research team.

To benefit future researchers, the code for our tool is available online along with the dataset of automations we built from our interviews [26]. The interview script, screening survey questions, and the codebook we used to analyze our interviews are in the Appendix. Out of an abundance of caution and respect for the privacy of the participants in our study, we will not publish transcripts from our interviews, even after anonymization. The automation dataset was built by hand from our interview data by our authors (including ones who were not present for the interviews), and our tool required many automations to be re-formatted, which minimizes the risk that they could be used to re-identify the participants in our study.

5 Smart-Home Automation Dataset

In this section, we describe the smart home automation dataset we built from our interview data, including the automation platforms and devices participants use, who creates automations, and whether the automations were created through end-user programming or automatically built into the automation platforms. We also describe the results of the manual analysis we performed in real-time during our interviews, which will serve as the ground truth for the evaluation in the next section. The dataset is available online [26].

5.1 Smart Home Automations (RQ1)

In this section, we describe the automation platforms and types of devices each participant uses, how their automations were created, and the number of automations they have. A table summarizing the results described in this section can be found in Appendix C.

Multi-platform configurations are common, but general-purpose platforms are the most used. The most popular platforms used by participants in our study were Amazon Alexa (13 participants), Google Home (6 participants), and Home Assistant (5 participants); all but one of the participants used at least one of these three general-purpose platforms. Of the remaining 18 platforms, 16 were used by only one participant, mainly being used for vendor-specific functionality (e.g., controlling advanced lighting features). About half of the participants in the study used multiple automation platforms, but these participants generally used a single general-purpose platform for the majority of their automations.

Participants have a variety of devices and most own voice assistants and smart lights. Voice assistants and smart lamps/light bulbs were used by the majority of participants (18 and 16 participants, respectively). We selected participants to interview who had multiple types of devices and a diversity of devices, and our data reflect that selection: all participants used at least two types of devices, and 15 used at least four. 13 types of devices were used only by a single participant.

Participants used a variety of techniques to create automations. 17 participants created some or all of their automations themselves. Some used automations created by friends or family (6 participants), while others found automations online or used

Table 1: Summary of participants’ ground-truth concerns about information flow violation scenarios. Each cell shows the [count of participants concerned] / [total number of participants responding] for a given category of scenario and other user. Scenarios marked with a * describe a *malicious* scenario. In cases where participants were asked about multiple scenarios of the same category (e.g., voice commands with different actions), we counted them as concerned if they were concerned about at least one scenario.

Scenario	Other user involved			
	Guest	Household	Stranger	All
Integrity Violations	8 / 13	2 / 2	13 / 15	18 / 22
Non-voice automation	0 / 1	1 / 1	1 / 2	2 / 4
Non-voice automation*	0 / 0	1 / 1	1 / 2	2 / 3
Outside camera	0 / 1	0 / 0	0 / 3	0 / 3
Outside camera *	0 / 0	0 / 0	4 / 5	4 / 5
Voice command	5 / 11	0 / 0	0 / 0	5 / 11
Voice command*	2 / 2	0 / 0	0 / 0	2 / 2
Voice command outside	1 / 2	0 / 0	11 / 12	11 / 12
Secrecy Violations	2 / 4	2 / 7	2 / 5	6 / 14
Observe via automation	1 / 2	1 / 6	2 / 5	4 / 12
Observe on camera	1 / 3	0 / 1	0 / 0	1 / 4
Voice command	0 / 0	1 / 1	0 / 0	1 / 1
All Violations	9 / 15	4 / 8	14 / 15	19 / 22

platform-suggested automations (5 participants each). One participant, P16, used generative AI to create most of their automations.

Not all “automations” come from end-user programming. Participants had between 4 and 51 automations (mean 17, median 13). Automations were from a variety of platforms, including Amazon Alexa, IFTTT, and Home Assistant (mean 14, median 9). Others were built-in automated features, created, for instance, by connecting devices to a voice assistant (mean 3, median 2). The use cases for platforms varied: on Home Assistant, which is a more complex platform to set up and use, all 138 automations were user-programmed. On Amazon Alexa, which emphasizes voice control, 41 of 109 automations used built-in voice controls.

5.2 Ground-Truth Labels (RQ2)

At the end of our interviews, we described hypothetical information flow violations that could occur and asked participants how they would feel if they happened. Their perception of the risks posed by the scenarios (i.e., concerning or not concerning) served as the ground truth for our evaluation. Due to time constraints, we prioritized discussing automations that were likely to be identified as violations by the tool. Across 22 participants, we asked about 89 information flow scenarios covering 114 of the 373 automations in our dataset. These scenarios were personalized to the participant’s devices and home configurations; because everyone had slightly different setups, different participants were asked about slightly different scenarios. We only asked about scenarios that a tool would likely flag as information flow violations, meaning we also had more scenarios to discuss with some than others (4 scenarios on average).

For example, if a participant had a voice assistant that controlled a smart oven and lights, we might discuss scenarios where a person physically outside the home speaks to the voice assistant to control those devices. Table 1 shows a summary of participants’ reactions to these scenarios. We grouped the personalized scenarios into ten categories; unless otherwise noted, we report how many scenarios participants found concerning per category, across all participants.

Scenarios involving built-in automations were more concerning than end-user programmed ones. We discussed scenarios with participants that the interviewers decided in real-time were had potential information flow violations. Participants considered at least one scenario related to 67 automations to be concerning of the 114 automations that we asked about. These scenarios covered the most-used platforms in our study: 38 automations on Amazon Alexa, 34 automations on Google Home, and 18 automations on Home Assistant. However, the rate of concern across platforms was uneven: 34 of the 38 Alexa automations were concerning, but only 8 of the 27 Home Assistant ones were. Similarly, participants were concerned about built-in automations (29 of 35 automations) at a higher rate than traditional automations (38 of 75 automations). As mentioned in Section 5.1, the use cases for platforms vary, with Alexa being used heavily for built-in voice controls and Home Assistant used for user-configured automations. Since users do not set up the built-in automation features, they may not consider all the possible consequences, highlighting the need to analyze built-in automation features in addition to end-user programs.

Participants are more concerned about integrity violations than secrecy violations. Most participants were concerned about at least one scenario discussed with potential integrity violations (18 of 22 participants), but only 6 of the 14 participants we discussed secrecy violations with expressed concern for those scenarios. Out of the 11 participants where we discussed a scenario relating to guests accessing voice assistants from inside the home, 5 expressed concern, including P6, who said “I would not want messing around with the thermostat. That’d get weird. That would get terrible” (P6). 11 out of 12 participants were concerned about scenarios where someone, typically a stranger, activates a voice assistant from outside the home, such as through an open window. Participants said this would feel like a threat to their home security, including P18: “That’s not a good thing. I feel like that’s kinda invasion of privacy. ...if somebody is able to control something inside my house, I feel like a little bit scared actually.”

Concerns about activating voice assistants outside the house were often mitigated by some other belief: in 8 cases, participants either mentioned that it was an unlikely scenario (e.g., that they would not keep their windows open when not at home), or that voice match authentication would prevent strangers from being able to use their voice assistant.

Other integrity violations that raised fewer concerns include other users triggering notifications by walking in front of outdoor security cameras, like doorbells. Participants were unconcerned about these (0 of 3 concerned), unless it was done maliciously (4 of 5 concerned), in which case they said they would find it annoying.

Overall, participants did not find most secrecy violations concerning. The most common secrecy violation involves another user or stranger learning about their behavior by observing an automation. However, in many cases, these weren’t concerning, because

they didn't consider the information to be private. For example, P12 configured Home Assistant to notify both them and their spouse when the door was locked or unlocked. When we asked if they were concerned this automation would leak information about when they were leaving the house, they said: "We're very open with each other about, you know, where we're at, and what's going on. So there's no issue with that." This sentiment was also common with automations like motion- or timer-activated lights, which could reveal which room a person is in or whether the house is occupied.

Participants are more concerned about strangers than their guests or household members. 34 scenarios involved guests, 14 involved household members, and 41 involved strangers. Participants were most concerned about scenarios involving strangers (15 of 15 participants were concerned), and slightly less concerned with guests (9 of 15) or members of their household (4 of 8).

Participants who were concerned about guests triggering automations or voice assistants expressed that they had boundaries for their smart home devices and did not want guests changing the temperature, setting alarms, or playing music. In contrast, some participants said that this was not a strong concern because they trusted their guests to not engage in concerning behaviors.

5.3 Dataset Preparation

Ultimately, our goal was to use this home automation dataset to evaluate our information flow analysis and compare it to the coarse-grained information flow technique from prior work [10]. To do this, we formatted the automations so they could be understood by the tool. Following the method described in Section 4.3.4, we created an initial dataset of 373 automations for the 22 participants enrolled in our study. The tool expects automations in the "IF *trigger* THEN *action*" format for each row, so automations with multiple triggers or actions were split into multiple rows for each trigger-action pair. For example, the automation "IF someone says *Alexa turn off living room lights* or *Alexa turn off the living room lights* THEN turn off living room lamp and living room light bulb" would be split into four rows: one where the first voice command turns off light living room lamp, another where the first voice command turns off the light bulb, and similarly for the second voice command.

In cases where there is ambiguity about labeling, rows would be duplicated so that they could be assigned all possible labels. Each row contains eight fields: Trigger Device, Trigger, Trigger label (secrecy), Trigger label (integrity), Action Device, Action, Action label (secrecy), and Action label (integrity). A more detailed description of the dataset may be found in Appendix E.

6 Results of Information Flow Analysis

In this section, we discuss the results of our information flow analysis. Our tool is based on one from prior work [34] and uses Python 3.9 for preprocessing and Prolog for evaluating automations. The tool is general enough to use different information flow lattices, which allows us to compare the coarse-grained analysis from prior work [10] to our location-enhanced analysis.

In Section 6.1, we explain our tool implementation, then in Sections 6.2 and 6.3, we present the results of the coarse-grained [10] and our location-enhanced analyses. Finally, in Section 6.4, we

compare the accuracy of the two analyses. For both the coarse-grained and location-enhanced analysis, we report results on a per-automation basis (which reflects whether a user's automation would be flagged as a violation) as well as a per-row basis where each trigger-action pair in an automation is a separate row (which reflects the total size of the dataset and is more relevant to the runtime performance evaluation). Run time evaluation was conducted on a server with an Intel Xeon Gold 5222 CPU and 128 GB of RAM.

6.1 Information Flow Violation Detection

As described in Section 2.3, permissible information flows form a lattice structure. An information flow from trigger to action is permitted if the label of the trigger is at or above the label of the action in the lattice. Any information flow that is not permitted is a violation. The location-enhanced lattice (Section 3) is specific to an individual's home, while the coarse-grained analysis (Section 2.3) is the same for all users. To reduce the manual effort required to use our tool, the location-enhanced lattice is generated automatically based on a list of a user's rooms.

Following prior work [34], we convert the lattices into an *allowlist* of permitted information flows by traversing the lattice from the bottom to the top. All other combinations of labels form a *blocklist* of violations. The blocklist is then transformed into a Prolog violation rule database. Each rule is formatted as either `violation_s(trigger, action)`, which is a secrecy violation, or `violation_i(trigger, action)`, which is an integrity violation. When the labels on an automation's trigger and action match a rule in the violation rule database, the tool flags it as a violation.

6.2 Coarse-Grained Analysis and Results

Using the labeling procedure for the coarse-grained lattice [10], the resulting dataset had 729 rows across all participants (ranging from 4 to 51 rows, mean ≈ 33). The coarse-grained analysis required an average of 1.62 seconds per participant. The coarse-grained lattice is the same for every participant, so the Prolog rules for identifying violations can be generated once ahead of time. Thus, the performance results include only the time spent detecting violations.

This analysis discovered 53 automations (14%) with secrecy violations, 87 automations (22%) with integrity violations, and 34 automations (9%) exhibiting both secrecy and integrity violations. In addition to the automation-level analysis, we also looked at the rows created by the labeling process. In total, 105 rows (14%) exhibited secrecy violations, 150 rows (21%) had integrity violations, and 63 rows (9%) contained both types of violations. Detailed results for each participant may be found in Appendix E.

6.3 Location-Enhanced Analysis and Results

Because our location-enhanced lattice has more fine-grained labels, there was more labeling ambiguity, leading to more rows requiring different labeling. Incorporating location information expanded the dataset to 2,347 rows (ranging from 4 to 762 rows, mean ≈ 107). Our tool completed its analysis in less than 2.7 seconds for each participant, on average (excluding P19²), with less than 1.2 seconds devoted to generating Prolog violation rules and less than

²Participant P19 was excluded as an outlier due to having an unusually large number of rooms in their house.

Table 2: Run times of the analyses in seconds. For the location-enhanced analysis, we report the run time for Prolog rule generation and violation detection.

	Auto-Rooms	mations	Coarse-grained	Location-enhanced		Total
			Total Runtime	Rule generation	Violation detection	
P1	9	15	2.10	1.34	3.38	5.65
P2	8	8	2.19	1.12	0.89	2.50
P3	2	9	1.56	0.96	0.16	1.63
P4	2	10	1.46	0.99	0.15	1.70
P5	2	9	1.46	1.02	0.15	1.69
P6	7	9	1.45	0.95	0.33	1.81
P7	6	12	1.66	1.00	0.42	1.98
P8	3	13	1.69	0.93	0.13	1.55
P9	1	4	1.51	0.93	0.12	1.55
P10	2	5	1.52	0.93	0.14	1.55
P11	4	15	1.57	0.96	0.20	1.66
P12	9	39	1.60	1.14	2.84	4.49
P13	8	20	1.49	0.99	0.87	2.41
P14	5	27	1.45	0.93	0.20	1.62
P15	7	33	1.61	0.98	0.46	1.94
P16	6	51	1.56	0.97	0.35	1.84
P17	2	7	1.47	1.11	0.16	1.79
P18	2	12	1.71	0.97	0.16	1.72
P19	13	24	1.63	113.27	3053.12	3240.26
P20	3	17	1.51	1.15	0.22	1.92
P21	10	21	1.51	1.95	11.62	14.20
P22	4	13	1.68	1.00	0.17	1.72

1.1 seconds to violation detection. The analysis run time for each participant is shown in Table 2.

Our tool identified 156 automations (42%) with secrecy violations, 191 automations (51%) with integrity violations, and 115 automations (31%) with both types of violations. In the overall dataset, these violations were distributed across 1690 rows (72%) with secrecy violations, 1867 rows (80%) with integrity violations, and 1517 rows (65%) containing both types of violations. Table 9 presents detailed results for each participant.

6.4 Accuracy of the Analyses (RQ3)

To better understand the evaluation results and compare the two analyses, we collected ground-truth data in our interviews (see Section 5.2) to calculate the rate of False Positives (FP) and False Negatives (FN). Specifically, the ground-truth data are based on the participants’ *perception* of risk and if they care about the identified violations, rather than an attempt to define risk objectively. What one participant views as a violation, another participant might not. The FP rate could be interpreted as a measure of the (lack of) utility of the tool, since a higher FP rate means the tool will unnecessarily notify the user about violations they don’t care about. Meanwhile, the FN rate reflects how often the tool misses vulnerabilities, as it shows how many risky information flows remain undetected by

Table 3: Comparison between ground-truth data and information flow analysis

(a) **Ground-Truth Data:** In the interviews, we discussed information flow scenarios related to participants’ automations. This table shows the number of automations that these scenarios covered, across all participants. Note that seven automations had both secrecy and integrity scenarios, hence the discrepancy with the totals.

	Secrecy	Integrity	Total
Automations discussed	36	85	114
Automations with concern	12	58	67

(b) **Analysis Accuracy:** The location-enhanced analysis had a lower false negative rate, where participants expressed concern about a scenario related to their automation, but the analysis tool did not detect a violation.

Metric	Analysis	Secrecy	Integrity	Total
False positive	Coarse-grained	2	9	11 (22%)
	Location-enhanced	6	24	30 (59%)
False negative	Coarse-grained	9	47	56 (80%)
	Location-enhanced	7	4	11 (16%)
True positive	Coarse-grained	3	11	14 (20%)
	Location-enhanced	5	54	59 (84%)
True negative	Coarse-grained	22	18	40 (78%)
	Location-enhanced	18	3	21 (41%)

the tool, despite them being concerning to the user. The results are summarized in Table 3. We discussed scenarios covering 114 automations out of the 373 in our dataset (seven automations were discussed twice, for both secrecy and integrity). Notably, 9 secrecy and 47 integrity violations were missed by the coarse-grained analysis (80% overall), while our location-enhanced analysis missed much fewer—7 secrecy and 4 integrity violations (16% overall).

Our location-enhanced analysis reduced FNs, missing fewer vulnerabilities than the coarse-grained analysis. This improvement can be attributed to the more nuanced and expressive labeling in the location-enhanced analysis. Most of the coarse-grained FNs, 49 of the 56, are due to both the trigger and action being labeled Restricted Physical. For example, when processing a voice command “Mute the TV,” labeling both the voice input and TV control Restricted Physical fails to capture essential spatial relationships and access control implications, like a stranger muting the TV through the kitchen window. A similar phenomenon was observed by Cobb et al. [10]. Our enhanced analysis addresses this limitation by incorporating fine-grained location information, allowing us to properly identify potential integrity violations that arise from spatial access control requirements.

On the other hand, the FPs in the location-enhanced approach increased. Notably, 16 of these 30 FPs are violations related to cross-location lighting controls that users deliberately set up for convenience (described in Section 5.2). For instance, if a user issued the voice command “turn on office light” from their living room,

our location-enhanced analysis flags this as a potential integrity violation since the command’s target location (bedroom) is different from the issuing location (living room). While cross-location lighting controls are sometimes intentionally designed for user convenience, our current approach flags these scenarios as violations due to the location mismatch.

This points to a limitation in our analysis framework: it cannot yet distinguish between intentional cross-location functionality and actual security concerns. If we were to adjust the analysis for just these users so that lighting controls were labeled physical trusted, then these automations would no longer be flagged as violations and the FP rate would decrease from 59% to 26%, bringing it in line with the coarse-grained approach. Importantly, while filtering cross-location lighting for just these participants would lead to no increase in FNs, a universal filter that dismisses all cross-location lighting controls as safe would lead to an *increase* in FNs from 16% to 49%. A complete list of the FPs can be found in Appendix F. We briefly discuss in Section 8 (under “Does device location improve automation analysis?”) how our tool can be improved in future work, allowing users to re-label automations to reduce the false positive rate while avoiding the high rate of false negatives that more coarse-grained strategies suffer from.

7 Visitor and Guest Interactions with Smart-Home Devices

With some participants, we had the opportunity to discuss other users interacting with their devices. We wanted to know if there were parts of their smart home they wanted to keep private from other people and if they were concerned about how their smart home would affect users who only interact with devices incidentally (like a mail carrier passing by a smart doorbell). We describe our findings in this section.

Participants expect others to encounter their devices. Almost all participants (19 out of 22 participants) expected visitors to enter their home, about half had one or more other household members who might interact with their devices, some of whom (6 participants) were children. Six participants mentioned that other household members had admin-level access. Three participants also mentioned other people, like neighbors, who might interact with their devices. The most common way participants expect visitors to learn about devices is by telling them (13 participants). Seven participants said visitors would see the devices without being told about them. P2 said there was “no way” someone would not learn about their smart devices “because Alexa controls so much in here.” Two participants said that they don’t tell guests about their devices, in one case for cultural reasons: “I wouldn’t feel comfortable only because we’re a [cultural identity] household, and because of that, like, when my parents and our cleaner were growing up, they didn’t really have as much technology” (P10).

Collecting guests’ data is typically considered okay. Few participants had concerns about collecting data about other people. Most commonly, people said that they did not consider the data sensitive (7 participants). Three people said the data collection is okay because they would not mind it if the roles were reversed, and two others said it is permissible because the devices are commonplace. Other justifications include keeping people safe (2 participants) and

that the front of their house is a public space (1 participant). While we heard many justifications for collecting data about others, these feelings were not universal. Three participants reported feeling uncomfortable collecting data about other people.

Most participants would adjust devices for guests. We also asked participants how they would react if they had guests who felt uncomfortable with their smart devices. Twelve participants said that they would adjust their devices by turning them off, unplugging them, muting them, or even removing them from sight after turning them off. Another common reaction to guests’ discomfort was to talk through their concerns (8 participants). On the other hand, four participants acknowledged that they might not do anything. P12 told us about recording people with their doorbell camera, “You’re in the public, like, outside your home is the public ... There’s no expectation of privacy there.”

8 Discussion

In this section, we revisit the original hypothesis from prior work that incorporating additional context about where devices are located could improve information flow analysis. We also make recommendations for building security analysis tools, provide additional commentary on the performance of our location-enhanced analysis, and make suggestions about where to head next.

Does device location improve automation analysis? To summarize our results from Section 6, we find that incorporating information about where smart devices are located *does* help identify more security and privacy vulnerabilities than a coarse-grained analysis. However, it can also lead to more false positives.

Unsurprisingly, the data in Section 5.2 showed that different people have different concerns. Therefore, in order for an analysis to identify all of the risks a user might care about, it must also risk alerting people to things they are not concerned about. As we discussed briefly in Section 6.4, though, we cannot adjust the analysis to ignore, for example, all cross-location lighting, because similar automations *did* concern some participants. Instead, we should allow users to give the analysis tool feedback about the types of violations they do (and do not) care about. These could be incorporated into information flow analyses by relaxing the strict criteria for what is considered a violation through declassification and endorsement. For example, in the case described in Section 6.4, 53% of the false positives would be eliminated by re-labeling (i.e., endorsing) the cross-location voice commands that trigger lights for the four participants who were unconcerned about these automations.

Recommendations for tool-building. Our evaluation (Section 6) and other interview results (Section 7) give us insight into what kinds of security analysis tools might be most relevant to smart-home users, some of which are counter to what the research community has focused on in the past. We make concrete recommendations based on our findings here.

(1) *Incorporate real-world context about where devices are located and how they are used.* To reiterate our point above, incorporating real-world context about where devices are located and how they are used into an information flow analysis can lead to an analysis that identifies more vulnerabilities than one that does not. However, more research is needed to understand how users can give tools feedback to account for the increased rate of false positives.

(2) *Analyze automations across platforms.* While most prior work focuses on specific automation platforms, like IFTTT [30, 34, 38, 39] or Amazon Alexa [3, 21], around half of the participants in our study (12 of 22 in Table 6) used two or more automation platforms. Moreover, the automations that led to concerning information flows (Section 5.2) came from multiple platforms. Ours is the first study to provide empirical evidence of the importance of supporting cross-platform analysis to identify all the security and privacy risks that users might care about.

(3) *Include support for built-in automations.* To our knowledge, prior work analyzing smart-home automations has exclusively considered end-user programming, such as IFTTT and Alexa Routines. While the research community has largely overlooked built-in automation features, our results demonstrate that built-in automations can be analyzed by the same information flow analysis and sometimes lead to undesirable behavior (Section 5.2). To avoid under-reporting security and privacy risks, the research community should consider both end-user and built-in automations.

Location-enhanced analysis run time. Although the run time of our location-enhanced analysis is generally longer than that of the coarse-grained analysis, the Prolog rule generation only needs to happen once for each participant, and the current performance is reasonably efficient overall. However, as noted in Appendix E, the number of detection rules the tool must check grows super-exponentially with the number of rooms. This is particularly evident in the run time for Participant P19, which was 54 minutes—significantly longer than for other participants—due to the unusually large number of rooms in their house.

More specifically, our lattice may contain labels that are not assigned to any automations. Because there is a Prolog rule for every possible information flow violation, removing the labels in the lattice that are not used in the dataset would speed up the violation-detection process by removing unused Prolog rules. We find that pruning unused labels in our lattice reduced the violation-detection time to under 1 second for every participant, including P19 (see Table 4 for the complete results). While the pruning process still took a non-trivial amount of time for large numbers of rooms, it was much faster than the original violation-detection time (pruning the lattice and detecting violations for P19 took around 8.5 minutes, while violation-detection without pruning took 50 minutes). Building a lattice out of the labels that are used (instead of pruning the unused labels) or implementing a custom tool instead of using Prolog could further improve the run time. We leave further optimizations to future work.

What is the best way to prevent violations? In this work, we investigated whether incorporating device location into information flow analysis helped identify security and privacy risks relevant to users. Broadly, we found that users are concerned about some types of risks that can be identified by a location-enhanced information flow analysis but less concerned about others. Participants were most concerned about situations where a person outside the home could control devices in their home unexpectedly, such as a person standing outside being able to trigger Alexa commands. People were much less concerned about guests or other people *inside* the home controlling devices, as there was more trust or norms governing their behaviors (as found in prior studies [42]).

Table 4: Run time of the location-enhanced analysis, in seconds, when unused labels are pruned from the lattice. We report the number of Prolog rules pre- and post-pruning, the time taken to prune unused labels, and the violation-detection time pre- and post-pruning.

	Rules (Pre)	Rules (Post)	Prune Time	Time (Pre)	Time (Post)
P1	493669	281	0.42	3.38	0.20
P2	122345	137	0.09	0.89	0.15
P3	125	102	<0.01	0.16	0.17
P4	125	76	<0.01	0.15	0.14
P5	125	60	<0.01	0.15	0.17
P6	30613	150	0.02	0.33	0.14
P7	7865	218	0.01	0.42	0.16
P8	253	36	<0.01	0.13	0.14
P9	79	9	<0.01	0.12	0.14
P10	125	16	<0.01	0.14	0.14
P11	665	75	<0.01	0.20	0.14
P12	493669	353	0.40	2.84	0.16
P13	122345	288	0.08	0.87	0.14
P14	7865	152	0.01	0.20	0.14
P15	30613	337	0.03	0.46	0.14
P16	7865	265	0.01	0.35	0.16
P17	125	27	<0.01	0.16	0.14
P18	125	50	<0.01	0.16	0.18
P19	131168389	387	514.26	3053.12	0.15
P20	253	54	<0.01	0.22	0.14
P21	1996505	240	1.84	11.62	0.15
P22	665	60	<0.01	0.17	0.18

A natural way to follow up this work would be using information flow techniques to *prevent* these risks, for instance by stopping automations that might be risky rather than notifying users. However, to address the primary area of concern—unauthorized users issuing voice commands to devices—we suggest a more practical solution: fine-grained access control.

This is not to suggest that our information flow analysis is not useful. The information flow scenarios we discussed in our interviews *did* help participants in our study identify use cases for their devices that made them feel uncomfortable or unsafe. An information flow analysis, like our location-enhanced analysis in Section 3, would nicely complement robust access controls by helping people identify the types of unauthorized usage they want to prevent. We leave the investigation about how to incorporate an information flow analysis into access controls to future work.

Security and privacy concerns of other users. Our location-enhanced analysis focused on smart device owners. While this is a common choice for information flow analyses in smart homes [10, 34], real smart homes are multi-user environments [9, 31]. Indeed, many participants reported living with other people who might interact with their devices (see participant demographics in Appendix C). However, even if our analysis accounted for other users, it is unclear how those preferences should be communicated to the tool. One solution would be to have the device owner make those

choices on behalf of their guests, but our interviews show that their preferences for how their devices interact with other users can vary. While many participants agreed they would be willing to adjust their devices for others, several felt justified in their data collection practices (Section 7). More research is needed to understand how to account for the security and privacy preferences of multiple users.

Limitations. We collected data from a limited number of participants (22) due to resource constraints. As a result, certain use cases for home automations, devices, or platforms may be under-represented in the data, and there may be information flow violations that we did not observe. We recruited from Prolific, which has a population younger and more educated than the general US population; compared to other studies about security and privacy on Prolific, the participants in our study were older and more educated [33, 36]. However, we believe our dataset, which includes users of many smart-home platforms, provides data about smart-home usage not fully measured in prior work, which focuses primarily on publicly available IFTTT rules.

Our screening process prioritized participants who had a high number of automations, because our goal was to collect a dataset where we would likely find information flow violations to discuss. As a result, our dataset is biased towards “power users” of smart homes. Prior work has found that users who set up smart-home devices in their home are less likely to be concerned about their own privacy or the privacy of other users in the home, compared to other people who live in or visit the home [13, 41]. This may have resulted in a participant pool that was less concerned about information flow violations than smart-home users as a whole.

For our information flow analysis, the ground-truth labels for whether participants considered information flows to be actual violations is incomplete. This is because labels were collected during the interviews, based on real-time analysis done by interviewers, prior to running the analysis tool. The results reported in Section 6.4 are only based on the subset of automations that we had ground truth data for. A more-complete analysis would be possible if we could have invited participants back for a second interview. Because interviewers prioritized discussing automations that were likely to be identified as violations, it is possible that a more-complete analysis would result in more false positives. However, this would not change our main findings. A location-enhanced analysis would still identify vulnerabilities that are missed by the more coarse-grained analysis and the strategy described at the end of Section 6.4 would still help reduce the false positives, even if they were more prevalent in the more-complete analysis.

Sometimes, people configure automations (accidentally or on purpose) so that they interact to create new behaviors. Because our analysis is based on an existing tool [34], we do not currently reason about cross-automation interactions. That being said, cross-automation interactions were relatively uncommon in our dataset; two participants (P15 and P16) explicitly chained automations by having automations read and write to shared variables, and one participant (P19) implicitly chained automations using a room’s humidity and a fan. The tool could be modified to reason about cross-automation interactions by modeling the interactions as new automations, though this could be labor-intensive for users with a large number of automations and could underreport the interactions that users created unintentionally (users cannot identify

interactions if they do not know they exist). Ideally, there would be another phase of the analysis that helps users identify cross-automation interactions automatically, like what has been done by prior work [32, 39]. We leave tool extensions that account for cross-automation interactions to future work.

Finally, the goal of this work was to investigate the feasibility and relevance of a location-aware information flow analysis, using real users’ perceptions of risk as the ground truth for evaluation. However, we did not investigate how people feel about sharing fine-grained automation details with our tool, nor did we investigate how we can help users accurately evaluate risk. During our interviews, we observed that participants sometimes held inaccurate beliefs that led to them dismissing potential violations, such as the belief that *everybody* that enters their home will be someone they trust (which is possible, but unlikely) or that they could prevent the violations we identified by setting up access controls (which don’t exist in the devices they own). Ensuring users trust their tools and that the risks identified by the tools are accurately communicated are both necessary for security analysis tools to be usable in practice.

We conducted an interview study rather than a larger-scale data collection to allow for more discussion with the participants in our study. This allowed us to ask follow-up questions, for instance about why violations were or were not concerning (Section 5.2). However, we would need more data to investigate other things, like why the rate of concern varied across devices or if integrity violations are actually more prevalent (and concerning) than secrecy violations. We aim to investigate these aspects and other issues related to the user’s experience in future research.

9 Conclusion

In this paper, we developed a new information flow analysis tool that takes into account where smart devices are deployed and how people interact with them. We interviewed 22 smart-home users to learn how they automated their devices and what kind of potentially risky information flows they found concerning and why. Our conversations led to a dataset of 373 real home automations, which we used to evaluate our tool. We found that introducing device location did lead to an analysis that identified more violations users cared about, if at the cost of additional false positives. Perhaps more importantly, our information flow scenarios helped some participants think more critically about the potential security and privacy risks in their smart home. Overall, our study demonstrates that information flow analysis could play an important role in helping users understand the risks they may encounter with smart-home automations, even in situations where non-information-flow-based strategies may be best for preventing harm.

Acknowledgments

We would like to thank the people who took the time to participate in our study and make this research possible. We would also like to thank the anonymous reviewers and revision editor for their helpful feedback on our paper. This work was supported in part by Cisco Research through the Carnegie Mellon CyLab partnership program and the the National Science Foundation via grants 2114148, 2323105, and 2317184.

References

- [1] Desiree Abrokwa, Shruti Das, Omer Akgul, and Michelle L. Mazurek. 2021. Comparing Security and Privacy Attitudes Among U.S. Users of Different Smartphone and Smart-Speaker Platforms. In *Seventeenth Symposium on Usable Privacy and Security (SOUPS)*.
- [2] Amazon. 2025. Alexa Routines. <https://www.amazon.com/alexa-routines>.
- [3] Tawfiq Ammari, Jofish Kaye, Janice Y. Tsai, and Frank Bentley. 2019. Music, Search, and IoT: How People (Really) Use Voice Assistants. *ACM Transactions on Computer-Human Interaction* (2019).
- [4] Apple. 2025. Home app. <https://www.apple.com/home-app/>.
- [5] Z. Berkay Celik, Patrick McDaniel, and Gang Tan. 2018. SOTERIA: Automated IoT Safety and Security Analysis. In *2018 USENIX Annual Technical Conference*.
- [6] Claire C Chen, Dillon Shu, Hamsini Ravishankar, Xinran Li, Yuvraj Agarwal, and Lorrie Faith Cranor. 2024. Is a Trustmark and QR Code Enough? The Effect of IoT Security and Privacy Label Information Complexity on Consumer Comprehension and Behavior. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*.
- [7] Yunang Chen, Mohannad Alhanahnah, Andrei Sabelfeld, Rahul Chatterjee, and Earlene Fernandes. 2022. Practical Data Access Minimization in Trigger-Action Platforms. In *31st USENIX Security Symposium*.
- [8] Yunang Chen, Amrita Roy Chowdhury, Ruizhe Wang, Andrei Sabelfeld, Rahul Chatterjee, and Earlene Fernandes. 2021. Data Privacy in Trigger-Action Systems. In *Proceedings of the 41st IEEE Symposium on Security and Privacy*.
- [9] Camille Cobb, Sruti Bhagavatula, Kalil Anderson Garrett, Alison Hoffman, Varun Rao, and Lujo Bauer. 2021. "I would have to evaluate their objections": Privacy tensions between smart home device owners and incidental users. *Proceedings on Privacy Enhancing Technologies* (2021).
- [10] Camille Cobb, Milijana Surbatovich, Anna Kawakami, Mahmood Sharif, Lujo Bauer, Anupam Das, and Limin Jia. 2020. How Risky Are Real Users' IFTTT Applets?. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*.
- [11] Dorothy E. Denning. 1976. A Lattice Model of Secure Information Flow. *Commun. ACM* 19, 5 (1976).
- [12] Youngwook Do, Nivedita Arora, Ali Mirzazadeh, Injoo Moon, Eryue Xu, Zhihan Zhang, Gregory D. Abowd, and Sauvik Das. 2023. Powering for Privacy: Improving User Trust in Smart Speaker Microphones with Intentional Powering and Perceptible Assurance. In *32nd USENIX Security Symposium*.
- [13] Christine Geeng and Franziska Roesner. 2019. Who's In Control? Interactions in Multi-User Smart Homes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*.
- [14] Google. 2025. Manage Your Smart Home with Google Home. <https://home.google.com/welcome/>.
- [15] Weijia He, Maximilian Golla, Roshni Padhi, Jordan Ofek, Markus Dürmuth, Earlene Fernandes, and Blase Ur. 2018. Rethinking Access Control and Authentication for the Home Internet of Things (IoT). In *27th USENIX Security Symposium*.
- [16] Home Assistant. 2025. Home Assistant. <https://www.home-assistant.io>.
- [17] Danny Yuxing Huang, Noah Apthorpe, Gunes Acar, Frank Li, and Nick Feamster. 2020. IoT Inspector: Crowdsourcing Labeled Network Traffic from Smart Home Devices at Scale. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2020).
- [18] Justin Huang and Maya Cakmak. 2015. Supporting Mental Model Accuracy in Trigger-Action Programming. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*.
- [19] IFTTT. 2025. IFTTT: Every thing works better together. <https://ifttt.com>.
- [20] Deepak Kumar, Kelly Shen, Benton Case, Deepali Garg, Galina Alperovich, Dmitry Kuznetsov, Rajarshi Gupta, and Zakir Durumeric. 2019. All Things Considered: An Analysis of IoT Devices on Home Networks. In *28th USENIX Security Symposium*.
- [21] Josephine Lau, Benjamin Zimmerman, and Florian Schaub. 2018. Alexa, Are You Listening?: Privacy Perceptions, Concerns and Privacy-seeking Behaviors with Smart Speakers. *Proceedings of the ACM on Human-Computer Interaction* 2 (2018). Issue CSCW.
- [22] Chieh-Jan Mike Liang, Börje F. Karlsson, Nicholas D. Lane, Feng Zhao, Junbei Zhang, Zheyi Pan, Zhao Li, and Yong Yu. 2015. SIFT: Building an Internet of Safe Things. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*.
- [23] Nathan Malkin, Alan F. Luo, Julio Poveda, and Michelle L. Mazurek. 2023. Optimistic Access Control for the Smart Home. In *Proceedings of the 44th IEEE Symposium on Security and Privacy*.
- [24] Sunil Manandhar, Kevin Moran, Kaushal Kafle, Ruhao Tang, Denys Poshyvanyk, and Adwait Nadkarni. 2020. Towards a Natural Perspective of Smart Homes for Practical Security and Safety Analyses. In *2020 IEEE Symposium on Security and Privacy (SP)*.
- [25] Prianka Mandal, Sunil Manandhar, Kaushal Kafle, Kevin Moran, Denys Poshyvanyk, and Adwait Nadkarni. 2023. Helion: Enabling Natural Testing of Smart Homes. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*.
- [26] McKenna McCall, Ben Weinschel, Kunlin Cai, Ying Li, Eric Zeng, Devika Manohar, Lujo Bauer, Limin Jia, and Yuan Tian. 2025. Location-Enhanced Information Flow Analysis for Smart Home Automations (Artifact). <https://osf.io/nphrv/>.
- [27] McKenna McCall, Eric Zeng, Faysal Hossain Shezan, Mitchell Yang, Lujo Bauer, Abhishek Bichhawat, Camille Cobb, Limin Jia, and Yuan Tian. 2023. Towards Usable Security Analysis Tools for Trigger-Action Programming. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*.
- [28] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019).
- [29] Abraham Mhaidli, Manikandan Kandadai Venkatesh, Yixin Zou, and Florian Schaub. 2020. Listen Only When Spoken To: Interpersonal Communication Cues as Smart Speaker Privacy Controls. *Proceedings on Privacy Enhancing Technologies* 2020 (2020).
- [30] Xianghang Mi, Feng Qian, Ying Zhang, and XiaoFeng Wang. 2017. An Empirical Characterization of IFTTT: Ecosystem, Usage, and Performance. In *Proceedings of the 2017 Internet Measurement Conference*.
- [31] Phoebe Moh, Pubali Datta, Noel Warford, Adam Bates, Nathan Malkin, and Michelle L. Mazurek. 2023. Characterizing Everyday Misuse of Smart Home Devices. In *Proceedings of the 44th IEEE Symposium on Security and Privacy*.
- [32] Muslum Ozgur Ozmen, Xuansong Li, Andrew Chu, Z. Berkay Celik, Bardh Hoxha, and Xiangyu Zhang. 2022. Discovering IoT Physical Channel Vulnerabilities. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*.
- [33] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. 2019. How Well Do My Results Generalize? Comparing Security and Privacy Survey Results from MTurk, Web, and Telephone Samples. In *Proceedings of the 40th IEEE Symposium on Security and Privacy*.
- [34] Milijana Surbatovich, Jassim Aljuraidan, Lujo Bauer, Anupam Das, and Limin Jia. 2017. Some recipes can do more than spoil your appetite: Analyzing the security and privacy risks of IFTTT recipes. In *Proceedings of the 26th International Conference on World Wide Web*.
- [35] Madiha Tabassum, Tomasz Kosinski, and Heather Richter Lipford. 2019. "I Don't Own the Data": End User Perceptions of Smart Home Device Data Practices and Risks. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*.
- [36] Jenny Tang, Eleanor Birrell, and Ada Lerner. 2022. Replication: How Well Do My Results Generalize Now? The External Validity of Online Privacy and Security Surveys. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS)*.
- [37] Blase Ur, Jaeyeon Jung, and Stuart Schechter. 2014. Intruders versus intrusiveness: teens' and parents' perspectives on home-entryway surveillance. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*.
- [38] Blase Ur, Melwyn Pak Yong Ho, Stephen Brawner, Jiyun Lee, Sarah Mennicken, Noah Picard, Diane Schulze, and Michael L. Littman. 2016. Trigger-Action Programming in the Wild: An Analysis of 200,000 IFTTT Recipes. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*.
- [39] Qi Wang, Pubali Datta, Wei Yang, Si Liu, Adam Bates, and Carl A. Gunter. 2019. Charting the Attack Surface of Trigger-Action IoT Platforms. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*.
- [40] Svetlana Yarosh and Pamela Zave. 2017. Locked or Not? Mental Models of IoT Feature Interaction. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*.
- [41] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End User Security and Privacy Concerns with Smart Homes. In *Thirteenth Symposium on Usable Privacy and Security (SOUPS)*.
- [42] Eric Zeng and Franziska Roesner. 2019. Understanding and Improving Security and Privacy in Multi-User Smart Homes: A Design Exploration and In-Home User Study. In *28th USENIX Security Symposium*.
- [43] Lefan Zhang, Weijia He, Jesse Martinez, Noah Brackenburg, Shan Lu, and Blase Ur. 2019. AutoTap: Synthesizing and Repairing Trigger-Action Programs Using LTL Properties. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*.
- [44] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proceedings of the ACM on Human-Computer Interaction* 2 (2018). Issue CSCW.

A Screening Survey Questions

A.1 Initial Screening Survey

Platforms like IFTTT, Samsung SmartThings, Apple HomeKit, or openHAB allow people to automate behaviors and connect their smart home devices so they interact in new ways.

SmartThings and Amazon Alexa allow you to set up **routines** which automatically perform some action when they are triggered. SmartThings and Apple Home support **scenes** which allow multiple devices to be controlled simultaneously. Philips Hue can operate automatically according to a **schedule**.

In the following question, we collectively refer to all types of automations as **home automations**.

Do you currently have home automations enabled in your home like the ones described above? [Yes, No]

If you are selected to participate in our interview, one of our researchers will contact you via Prolific to schedule a call over Zoom or Google Meet.

As a reminder, during the call we will ask you questions about your smart home devices, including where in your home you have them set up (e.g., “living room” or “a child’s bedroom”), which automations you have installed, and who can interact with your devices (e.g., “only people living in the home” or “family members and guests”). We will ask you to share your screen to show us your automations while we discuss them.

We will record audio and video from interview so we can review the automations you show us later, but you do not have to turn on your camera if you don’t want to. The recordings will not be shared with anyone outside of our team. We may share pieces of what you tell us, or describe your automations, but not in a way that could be connected to you or anyone else individually (e.g., we may say you have an automation which “sends a private ‘good morning’ email every morning at 7AM” but we would not say who receives the email, even if the email address is visible to us during our interview).

We expect the interview will take between 40 and 60 minutes, though it may be longer or shorter depending on how many devices and automations we have to discuss. For participating in the interview, you would be compensated with \$20.

Are you interested in scheduling an interview with someone from our team?

To ensure we have a diverse set of participants for our interviews, we may follow-up with another 3-5 minute screening survey to learn more about you and your home automation habits before attempting to schedule an interview with you. [Yes, No]

During the interview, we will ask you to share your screen so we can look at your home automations together.

Do you (1) have access to a device where you can easily view the automations you have enabled (a smartphone to view mobile apps, for example) and (2) know how to share your screen on Zoom or Google Meet?

If you can view most, but not all, automations, or would need to use multiple devices, like a mobile phone and desktop computer, please select “Yes”. If you do not feel comfortable sharing your screen, please select “No”. [Yes, No]

How many automations do you currently have running in your home? If you use multiple platforms for automation, you can report the total number of automations. If you are not sure how to count your automations, give us your best approximation. [1-2 automations, 3-4 automations, 5-6 automations, 7-10 automations, 11-15 automations, More than 15 automations]

[Included in later rounds of screening] On the next page, we will ask you to upload a screenshot of at least one of your automations. Is this something you are willing to do? If you are not sure how to take a screenshot on your device, you can find instructions for Android phones, iPhone, Windows, and Mac. (*Links for each platform*) (*Example screenshots*) Please try to ensure your screenshot only includes your automations and avoid capturing any personal information in your screenshot. [*File upload*]

A.2 Secondary Screening Survey

If you are selected to participate in our interview, one of our researchers will contact you via Prolific to schedule a call over Zoom or Google Meet.

As a reminder, during the call we will ask you questions about your smart home devices, including where in your home you have them set up (e.g., “living room” or “a child’s bedroom”), which automations you have installed, and who can interact with your devices (e.g., “only people living in the home” or “family members and guests”). We will ask you to share your screen to show us your automations while we discuss them.

We will record audio and video from interview so we can review the automations you show us later, but you do not have to turn on your camera if you don’t want to. The recordings will not be shared with anyone outside of our team. We may share pieces of what you tell us, or describe your automations, but not in a way that could be connected to you or anyone else individually (e.g., we may say you have an automation which “sends a private ‘good morning’ email every morning at 7AM” but we would not say who receives the email, even if the email address is visible to us during our interview).

We expect the interview will take between 40 and 60 minutes, though it may be longer or shorter depending on how many devices and automations we have to discuss. For participating in the interview, you would be compensated with \$20.

In our last screening survey, you said you are interested in being interviewed. **Are you still interested in scheduling an interview with someone from our team?** [Yes, No]

Platforms like IFTTT, Samsung SmartThings, Apple HomeKit, or openHAB allow devices to be configured to behave **automatically** on behalf of users and connect their smart home devices so they interact in new ways.

Automations are sometimes written in the format “IF trigger THEN action” where the *trigger* decides when the automation runs and the *action* decides what the automation does.

Depending on the platform you use, these might be called **automations**, **routines**, or **scenes**. SmartThings and Amazon Alexa allow you to set up **routines** which automatically perform some action when they are triggered. SmartThings and Apple Home support **scenes** which allow multiple devices to be controlled simultaneously. Philips Hue can operate automatically according to a **schedule**.

For example, this is what *automations* look like in Apple Home: (screenshot)

You can also create *routines* in Amazon Alexa, which look like this: (screenshot)

Some devices already do things automatically, like sending notifications. We are interested in the automations, routines, scenes, and schedules that *you or someone else in your home have set up for your devices* rather than what is set up in the device by default. In the following question, we collectively refer to automations, routines, scenes, and schedules as **home automations**.

We want to learn **which platforms** people use to automate their smart home and how many home automations they have on each platform. **How many** home automations do you currently have running on each of the following platforms? If you are not sure how to count your home automations, give us your best approximation. [*Participants select number of automations they have on each platform*]

If you use another platform, not listed above, please describe the platform and how many home automations you have running on it. If you are not sure how to count your home automations, give us your best approximation. [*Free response*]

For each of the following categories, tell us if you **use the device for your home automations**. If you own or have set up a device not listed here, you can use the box at the bottom to describe the device. [*Participants select from list of device types*]

How many home automations do you have currently running, total? If you are not sure how to count your home automations, give us your best approximation. [*Free response*]

How are your home automations triggered? Please select all that apply. (Recall: the trigger determines when a home automation runs)

- Manually (e.g., when someone clicks a button in an app or gives a voice command to a smart assistant)
- By location (e.g., when someone enters/leaves a room)
- When a device senses something (e.g., motion or a change in temperature is detected)
- At a specific time of day (e.g., 9AM or at sunset)
- When someone interacts with a device (e.g., turning on a light)
- Another way (please describe)

Why do you use home automation? Please tell us more about what your home automations do, in a few sentences. [*Free response*]

We will use the following information about your home during our interview.

Which of the following describes where you currently live?

(If you have multiple residences, pick the one you had in mind when you answered the previous questions about smart home devices and home automation.)

- Single-family (detached) house
- Townhouse (attached house)
- Shared building with 3 or fewer apartments
- Shared building with 4 or more apartments
- Mobile home/trailer
- Not listed here [*Free response*]
- Prefer not to answer

Does your home have any of the following features? (Select all that apply.)

- Multiple floors
- Private space outside my home that only the people in my household and my guests can enter (e.g., a gated backyard)
- Shared space outside my home that only some people outside my household can enter (e.g., a shared hallway or lobby in a locked apartment building)
- Public outdoor space that anybody can enter (e.g., a courtyard)
- My home has none of the features listed above
- Prefer not to answer

Who else lives in your home? (Select all that apply, even if they only live your home part-time.)

- Children under the age of 1
- Roommates who are not family (over the age of 18)
- Family (over the age of 18)
- Pets
- None of the above

We will use the following information to select a diverse group of participants for our interviews.

What is your age?

- 18-24 years
- 25-29 years
- 30-34 years
- 35-39 years
- 40-44 years
- 45-49 years
- 50-59 years
- 60 or older
- Prefer not to answer

How would you describe your gender?

- Male
- Female
- Non-binary

- Prefer to self-describe [*Free response*]
- Prefer not to answer

What is the highest level of education that you've completed?

- Some high school
- High school or equivalent
- Some college or associate's degree
- Bachelor's degree
- Master's degree
- PhD or higher
- Prefer to self-describe [*Free response*]
- Prefer not to answer

Are you majoring in, or do you hold a degree in a computing field (e.g., computer science, computer engineering, information technology)? [*Yes, No*]

Are you currently, or have you in the past been, employed in a computing field (e.g., software engineering, IT)? [*Yes, No*]

Is there any other feedback you'd like to give us? [*Free response*]

B Interview Script

B.1 Introduction

My name is [Interviewer 1] and this is [Interviewer 2]. I'm a [Title 1] and [Interviewer] is a [Title 2] in [Department] at [Institution]. I'll be the one asking questions during the interview while [Interviewer 2] takes notes to help me organize questions for the second half of the interview.

[*Open consent form and prepare to share screen, if needed*]

Thank you for your interest in our research! We shared a consent form with you when you scheduled this interview. Do you have any questions about that form? [*Answer any questions*]

B.2 Consent

Just to make sure everything is fresh, I'm going to review some of the information covered in that form.

Our conversation today will last about an hour. During this interview, I will ask you to share your smart home automations with me and describe their purpose. We'll also talk about where in your home your smart devices are located, and who can operate them. Finally, I'll read some hypothetical scenarios about how your devices might behave and ask how you feel about those situations.

When the interview begins, I will record both audio and video. In our research, we are interested in comparing what you see in the app versus what you tell me during our interview, so I will record video so our team can review the automations you show me today. I will keep my camera on during the interview, but you don't need to do the same if you don't want to. If you [*turn/leave*] it on, please be aware that your face will appear in the recording.

We will use Otter.ai to transcribe what is said during the interview today. Pieces of the transcription may appear in our research, but we will not share anything that could identify you.

To the extent that you are able, try not to reveal anything in your responses that might be used to identify you or anyone else. If you accidentally disclose something, whatever you share will be removed from the transcript of our conversation. Your participation is completely voluntary, so if you wish to stop the interview or skip a question, you can tell me that at any time.

When we're done, you'll receive \$20 via Prolific.

Do you have any questions before we get started? [*Answer questions*]

Are you still interested in proceeding with the interview? [*If yes, continue*]

Are you still okay with the interview being recorded? [*If yes, continue*]

I'm going to start recording now. As a reminder, please do not share anything private or identifiable about yourself or anyone else. Before we begin the interview, I'm going to ask you to repeat your consent so that we have it on the recording. [*Start recording*]

Alright, we are recording now. Can you confirm that we reviewed the consent information? [*If yes, continue. If no, turn off recording.*]

Do you agree to participate in our research? [*If yes, continue. If no, turn off recording.*]

Do you agree to have the interview recorded today? [*If yes, continue. If no, turn off recording.*]

B.2.1 Prolific ID. Before we get started, we're going to confirm your Prolific ID with you so that we can ensure you are paid for your time here.

Our notes indicate that your Prolific ID begins with <first 4-8 digits of ID>. Is that correct?

B.3 Collecting Automations

For the first part of the interview, I'm going to ask you to tell me about your experiences with home automation. Then I'm going to ask you to show me the automations you have installed and describe their purpose.

B.3.1 Automations Background. For the automations running in your home, did you write all of them? Did someone else in your home write some of the automations? Are your automations custom or do you find them online?

How often do you update or remove automations?

Have you ever turned off an automation because you thought it was not working correctly? [*If yes*] What happened that made you turn off the automation? [*If no specific risk mentioned*] How did that make you feel? [*Depending on automation*] Did it lead to annoyance? Did you worry it would increase your electricity bill?

Do you often have guests in your home, like friends, family, or even a cleaner?

Do you tell these guests about your smart devices? Would you feel comfortable telling them about your automations? Do you feel more comfortable telling guests about some automations than others? [Why?] Do you have devices that you would feel more comfortable with people interacting with than other devices? [Why?] Do you have automations that you would feel more comfortable with people triggering than other automations? [Why?]

Have you ever tried to create automations that trigger other automations? [If yes] Did it work the way you wanted? Have you ever accidentally had automations trigger other automations?

Now I'd like to learn more about the automations you have installed. In the pre-interview survey, you mentioned using <list platforms> for automation. Is that accurate? And the devices you use for automation include <list devices>. Is that accurate? [If yes, continue. Otherwise, update lists]

The next thing I'm going to do is ask you to share your screen with your home automations pulled up so we can discuss them. Before you begin sharing your screen, please make sure you close any other apps or webpages so that we are sure you don't share anything else while I'm recording. I'd also like to ask you to turn on "Do Not Disturb" mode so that we aren't interrupted by any notifications while you are sharing your screen. [Wait until participant indicates they have closed other apps]

I'll wait for you to pull up your automations. Are you going to share your automations on this device that you're using to speak with me now, or do you need to get another device?

[If the participant needs another device] You can add that device to our call, using the same link we're using now in order to share your screen. Just make sure you don't connect to audio. Let me know if you need any help getting set up in Zoom.

B.3.2 Automations & Devices. Now I'd like you to read through your automations and describe each of the automations. I'm going to take some notes as you're talking so you may hear me typing. I'll let you know if I need you to pause or slow down.

Trigger/Action Questions [If the participant does not mention a specific trigger/action and it is unclear from screen-share] Could you tell me specifically what triggers the automation to run? Could you tell me specifically what happens when the automation runs?

Device Questions [If participant does not mention specific device for trigger/action] Could you tell me which device <triggers the automation/performs <action>?

Old Device [If participant mentions a device that was involved in an earlier automation] Is this the same device as <other automation>?

New Device [If participant mentions a device that was not mentioned previously] Where in the home is this device located? Can you operate the device remotely via an app? Do you share that <app> account with anyone else in the home?

Old Online Service [If this is not the first shared account] Is this account shared with the same people as <other service>?

New Online Service (Non-Social) [If participant mentions an unfamiliar online service] I'm not familiar with <online service>—is that a public site or do you need to make an account to access that? Do you share that account with anyone? [If yes] What is their relationship to you? [Might need a follow-up to determine: Spouse, sibling, child, other family, friend, coworker]

New Online Service (Social Media) Do you share that account with anyone? Who can see what you post on that account? [If yes] What is their relationship to you? [Might need a follow-up to determine: Spouse, sibling, child, other family, friend, coworker] Are these the same people who see your posts on <other social media>? [Might need a follow-up to confirm what overlap exists between the two accounts, if any]

[If there are other platforms left to discuss] Which platform would you like to move to next?

Break You can go ahead and stop sharing your screen now. Before we move on to the next section of the interview, do you want to take a break for a few minutes?

B.4 Additional Device/Home Information

Next, I'd like to ask some follow-up questions about the smart devices you have in your home. I'll be asking about where the devices are located in your home and who can interact with the devices, but I'll only be asking these questions about the devices you use for automation. The last thing I'm going to do is ask you to imagine a few scenarios related to your automations and tell me how comfortable you would be.

Just to confirm, in the pre-interview survey you said that your home has <list features>. Is that accurate? [If yes, continue. Otherwise, update list]

I'm going to ask you some follow-up questions about your home and devices. [Questions depend on which types of devices/house features the participant has]

Devices playing/recording audio Can you generally hear things between rooms in your home? For instance, can <device> be heard from other rooms? For instance, can <device> be triggered from other rooms? What about from between floors? What about from the hallway outside your apartment? What about outside your <home/building>?

Devices producing light Can you see light between rooms? What about outside your <home/building>?

Devices changing/sensing temperature/air quality/etc. Can the <temperature/air quality/...> of one room affect other rooms? For instance, if <device> turns on? For instance, can <device> be affected by the <temperature/...> in other rooms?

Break Before we move on to the last section of the interview, do you want to take a break for a few minutes?

B.5 IFC Comfort Scenarios

[Note taker has made a list of automations which could lead to undesirable information flows]

Next I'm going to describe some hypothetical scenarios that could arise with a few of the automations you told me about earlier. I want to know how you would feel about each scenario and if anything like these scenarios have ever happened to you.

Confidentiality: For <automation>, how would you feel if someone in the same room as <action device> learned about <trigger>?

Integrity: For <automation>, how would you feel if someone in the same room as <trigger device> caused <action>?

[If participant responds negatively (discomfort, etc.)] **Specific risk:** What about that makes you feel uncomfortable? Is there a specific risk that comes to mind?

[If participant mentions specific risk] Have you ever thought about how this automation might cause <risk> before our conversation today?

Would you want to be notified that this automation could lead to <risk or discomfort>?

[After going through all potential violations, repeat the process again for violations which weren't concerning; this time mentioning specific harms like annoyance, embarrassment, financial harms, or safety risks]

[If participant responds negatively (discomfort, etc.)] Have you ever thought about how this automation might cause <risk> before our conversation today?

[If the participant does not respond negatively] Can you think of a reason **someone else** would be worried about this automation?

[Regardless of whether the participant responds negatively] What specifically [is/is not] worrying about this scenario?

[Depending on what they say is/is not worrying] Do you think you would always feel this way about <trigger/action>?

[Even though you aren't worried about it/Given that you are worried about it,/Given that others might be worried about it,] would you want to be notified when you are setting up the automation that it could lead to <risk or discomfort>?

Thank You Thanks so much for your time! Do you have any questions or feedback about our study that you'd like to share with us before we wrap up the interview?

B.6 Follow-ups if Incidental Users are Mentioned

How does it make you feel to collect data about [guests in your home/other incidental users]?

Do the [guests in your home/other incidental users] know about [automation]?

Table 5: Participant demographics

		Count
Gender	Male	15
	Female	7
Age	18-29	2
	30-39	4
	40-49	4
	50-59	6
	60 or older	6
Education	Some college or associate's degree	6
	Bachelor's degree	8
	Master's degree	7
	PhD or higher	1
Lives with...	Alone	3
	Family (over the age of 18)	16
	Pets	13
	Family (under the age of 18)	9
	Roommates	2
CS Education	Yes	7
	No	15
CS Employment	Yes	9
	No	13

If the [guests in your home/other incidental users] told you they were uncomfortable with the data collection, how would you react? Would you feel differently? Would you change the automation?

Before our conversation today, have you thought about the data you're collecting about [guests in your home/other incidental users]?

Would you want [device] to remind you that it might collect data about other people?

C Participant Information

The demographics of the 22 participants are summarized in Table 5. Note that participants were allowed to select multiple options when asked who they live with. The participants in our study are not representative of the United States: most notably, they are disproportionately male and educated. In Table 6, we summarize the smart home devices, automation platforms, and how automations were created by each participant.

D Location Analysis Heuristics

In this section, we describe the heuristics we applied for the information flow labeling for the location analysis.

Voice command (trigger) or smart speaker (action).

Secrecy Each room has its own row (i.e., you would only know the voice command if you're in the same room as the person speaking) unless the participant specifies that you can hear people speaking between rooms in which case the label is the list of rooms you can hear from the room with the device

Integrity List of rooms where people can talk to the device from. If

Table 6: Summary of participants’ smart home environments. Automations created using traditional end-user programming are listed separately from built-in automation features. The total number of automations are shown in bold text.

Automation Platforms	Smart Device Types	Who Writes Automations	Automations		
			End-User	Built-In	Total
P1 Google	Voice Assistant, Lighting, Outlet	Participant, Other person	15	0	15
P2 Alexa	Voice Assistant, Lighting, Outlet	Participant	6	2	8
P3 macroDroid, HBN, eHomeLive, Magic Lantern	Lighting, Outlet, Car	Other person	9	0	9
P4 Alexa, IFTTT	Voice Assistant, Lighting, Outlet, Camera, Printer	Other person, Platform	9	1	10
P5 Alexa, Kasa, Rain Bird, Schlage	Voice Assistant, Outlet, Lock, Oven, Sprinkler	Participant, Other person	8	1	9
P6 Alexa	Voice Assistant, Camera, Doorbell, Thermostat, TV	Other person	7	2	9
P7 Google	Voice Assistant, Lighting, Outlet, Vacuum, TV	Participant	6	6	12
P8 Alexa	Voice Assistant, Lighting	Platform	8	5	13
P9 Alexa	Voice Assistant, Doorbell	—	0	4	4
P10 Alexa, Ring	Voice Assistant, Doorbell	Participant	0	5	5
P11 Google, IFTTT, Hatch, Ecobee	Voice Assistant, Lighting, TV, Thermostat, Alarm Clock	Participant	10	5	15
P12 Google, Home Assistant	Voice Assistant, Lighting, Camera, Doorbell, TV, Lock, Thermostat, Contact Sensor, Vacuum, Nanny camera	Participant, Found online	37	2	39
P13 Alexa, Flume, IFTTT, SmartLife, Sense	Voice Assistant, Lighting, Doorbell, TV, Water Sensor, Water Pump, Smoke Alarm, Electric Sensor, Garage	Participant	16	4	20
P14 Alexa, Ring, MyQ, Lutron, ScreenLogic	Voice Assistant, Lighting, Doorbell, Camera, Fan, Lock, Garage, Water Pump	Participant, Other person, Platform	27	0	27
P15 Google, Home Assistant	Voice Assistant, Lighting, Button, Thermostat, Smoke Alarm, Environmental Sensors, Activity Sensors, Lock	Participant, Found online	28	5	33
P16 Alexa, Home Assistant	Voice Assistant, Lighting, Outlet, Button, Camera, Activity Sensors, Thermostat, Security System, TV, Game Console	Found online, Generative AI	51	0	51
P17 Alexa	Voice Assistant, Doorbell	Participant, Other person	1	6	7
P18 Alexa, Ring	Voice Assistant, Lighting, Doorbell, Lock, Sprinkler, Blinds	Participant	5	7	12
P19 Home Assistant	Lighting, Outlet, Button, Fan, Environmental Sensors, Activity Sensors, Garage, Fan, Speaker	Participant, Found online	24	0	24
P20 Google	Voice Assistant, Lighting, Outlet, TV, Vacuum	Participant, Platform	17	0	17
P21 Apple, Home Assistant	Lighting, TV, Lock, Activity Sensors, Security System, Garage, Speaker	Participant, Found online	21	0	21
P22 Alexa	Voice Assistant, Lighting, TV, Lock, Security system	Participant, Platform	13	0	13

people can talk to the device from outside, use the label “Physical untrusted”.

Date/Time.

Secrecy Always labeled public

Integrity Always labeled most trusted

Date/Time and [something else]. For example, “Motion detected while it is 8PM-8AM”

Secrecy or Integrity Label of (something else)

Change setting on physical device (special cases listed below). For example, turning a light on or off

Secrecy Initially label [room], then check transcript for other rooms that can see the light/hear things from this room

Integrity Room where the device is located

Holiday lights.

Secrecy or Integrity If room is unspecified, use LR

Button.

Secrecy or Integrity If unspecified, assume location is same as action **TV**.

Secrecy or Integrity If location unspecified, default LR

Presence/camera detected.

Secrecy Each room has its own row (i.e., you would only know the person is there if you were in the same room), physical public if the device is outdoors

Integrity Room where the device is located, physical untrusted if the device is outdoors

Monitors. For example, air quality monitor or smoke detector

Secrecy If indoor air monitoring, same label as room. If outdoor air monitoring, use “Public” (default: “public” unless the device is specified as an air quality monitor, then default is all rooms; if the device is a smoke detector default is “K”)

Integrity If indoor air monitoring, same label as the room. If outdoor air monitoring, use “Most trusted” (default: “most trusted” unless the device is specified as an air quality monitor, then default is all rooms; if the device is a smoke detector default is “K”)

Battery level.

Secrecy Private shared

Integrity Most trusted

Lock/unlock door.

Secrecy Physical public if it is an exterior door as it could be determined by anyone passing the house

Integrity Entry” or similar (e.g., “Living room”) if it is an exterior door as it could be locked by anyone in the house (this is less trusted than “physical trusted” which would be the label if the door were locked from the outside with a key)

Thermostats.

Secrecy (if trigger) Same room label as action if the thermostat controls the entire house. If room is unspecified label is all rooms. If there is one control or device per room affecting temperature (e.g., window A/C) there should be one row per room

Integrity (if trigger) Same room label as action If thermostat controls the entire house, or if it is unspecified label is a single room where the thermostat is located (default: LR). If there is one control or device per room affecting temperature (e.g., window A/C) there should be one row per room

Secrecy (if action) If it’s change temperature, all rooms in the house should be able to feel it (unless it is a window A/C). If it’s report temperature, only the room the thermostat is located in (default: LR).

Integrity (if action) 2 rows, one for online and one for the room the thermostat is located in (default: LR)

Scenes.

Secrecy Labels should reflect the ability of people who can view that the scene is active in the app: Private (shared)

Integrity Labels should reflect the ability of people who can set the scene to be active in the app: Trusted (shared)

PC/smartphone/tablet notifications.

Secrecy Unless otherwise specified, (Most) private

Integrity Most trusted

Open front door.

Secrecy Physical public if it is an exterior door as it could be determined by anyone passing the house Physical untrusted

Integrity Depends on if door is unlocked or locked, have one row for each case. If unlocked, physical untrusted If locked, use LR or E (someone inside could unlock door and open it)

Participant location.

Secrecy Anyone in the same physical space (physical public if outside, same label as room if inside)

Integrity Most trusted

E Additional Implementation Details and Evaluation Results

Additional evaluation results can be found in Tables 7–9.

Complexity of the tool: Given a list of rooms, our tool automatically generates a violation rule database as discussed in Section 6.1. The number of room combinations grows quadratically with the number of rooms, causing the violation rules to grow

super-exponentially. Specifically, we consider all possible room combinations using the power set $P(R)$, where R represents the set of all rooms. For a total of n rooms, there are $\sum_{k=1}^n \binom{n}{k} = 2^n - 1$ possible room combinations. Each combination $\binom{n}{k}$ is positioned in an upper layer above the combinations from $\binom{n}{k+1}$ in the lattice hierarchy. Since we construct a lattice to represent the physical environment for each participant, and each participant may have home IoT devices distributed across various rooms, the lattice size grows rapidly. This expansion results in a significant increase in the number of rules describing permitted information flows with complexity $O(n^n)$ and rules describing disallowed information flows with complexity $O(n^{n^2})$.

Table 7: Dataset Descriptions: Participant-level statistics, including the number of rows after labeling the data with labels from the coarse-grained [10] and location-enhanced lattices.

	Automations	Rows Coarse-Grained	Rows Location-Enhanced	Rooms
P1	15	106	762	9
P2	8	16	80	8
P3	9	10	13	2
P4	10	13	24	2
P5	9	9	13	2
P6	9	9	29	7
P7	12	48	298	6
P8	13	13	19	3
P9	4	4	4	1
P10	5	5	5	2
P11	15	29	80	4
P12	39	53	110	9
P13	20	35	83	8
P14	27	38	50	5
P15	33	69	222	7
P16	51	69	245	6
P17	7	7	12	2
P18	12	12	17	2
P19	24	80	80	13
P20	17	38	102	3
P21	21	40	40	10
P22	13	25	42	4

Table 8: Lattice Description: Participant-level statistics on the number of nodes in the lattices. The size of the coarse-grained lattice remains fixed, while the location-enhanced lattice grows significantly with an increasing number of rooms.

	Rooms	Coarse-Grained Lattice Size (Nodes)		Location-Enhanced Lattice Size (Nodes)
		Secrecy	Integrity	Secrecy and Integrity
P1	9	6	8	518
P2	8	6	8	262
P3	2	6	8	10
P4	2	6	8	10
P5	2	6	8	10
P6	7	6	8	134
P7	6	6	8	70
P8	3	6	8	14
P9	1	6	8	8
P10	2	6	8	10
P11	4	6	8	22
P12	9	6	8	518
P13	8	6	8	262
P14	5	6	8	38
P15	7	6	8	134
P16	6	6	8	70
P17	2	6	8	10
P18	2	6	8	10
P19	13	6	8	8198
P20	3	6	8	14
P21	10	6	8	1030
P22	4	6	8	22

Table 9: Violation Descriptions: Participant-level statistics on the number of automations and rows in the labeled dataset with secrecy and integrity violations.

	Coarse-Grained Analysis						Location-Enhanced Analysis					
	Secrecy		Integrity		Both		Secrecy		Integrity		Both	
	Auto-mations	Rows	Auto-mations	Rows	Auto-mations	Rows	Auto-mations	Rows	Auto-mations	Rows	Auto-mations	Rows
P1	0	0	0	0	0	0	8	679	8	738	8	679
P2	0	0	0	0	0	0	6	61	8	80	6	61
P3	1	2	1	2	1	1	1	4	2	6	1	4
P4	3	14	1	14	0	0	5	25	8	21	3	7
P5	0	0	1	0	0	0	5	5	4	8	4	4
P6	0	0	0	0	0	0	5	22	4	24	4	21
P7	0	0	0	0	0	0	10	291	10	294	10	291
P8	1	2	0	2	0	0	0	0	0	0	0	0
P9	0	0	1	0	0	0	0	0	4	4	0	0
P10	0	0	1	0	0	0	0	0	5	5	0	0
P11	0	0	0	0	0	0	8	53	9	68	8	53
P12	5	10	14	10	1	2	11	47	18	57	6	23
P13	1	1	3	1	1	1	8	54	13	68	8	53
P14	0	0	8	0	0	0	8	20	14	26	5	14
P15	17	34	17	34	14	25	25	188	22	186	22	166
P16	22	33	25	33	15	26	19	126	21	104	6	58
P17	0	0	2	0	0	0	1	1	7	12	1	1
P18	0	0	5	0	0	0	4	5	5	8	3	3
P19	1	1	5	1	0	0	5	15	7	31	1	1
P20	0	0	0	0	0	0	14	58	15	96	14	58
P21	2	8	2	8	2	8	8	15	3	12	1	1
P22	0	0	1	0	0	0	5	21	4	19	4	19

F False Positives

A list of the false positives for the location-enhanced analysis from Section 6.4 are shown in Table 10. Many of the false positives (labeled with an asterisk in the table) relate to cross-location lighting controls. The other false positives vary. While these automations were not concerning to the participants in our study, they might be concerning to other people. For instance, P16 A18 plays a custom audio notification when an email with a specific subject line is received. For some, this could be sensitive information that they don't want broadcast if there are guests in the home. A more detailed, complete list of all automations may be found in our online dataset [26].

Table 10: False positives for the location-enhanced analysis. Each row describes one automation, identified by participant and automation ID, including: the triggering device(s), a description of the automation trigger(s), the action device(s), and a description of the automation action(s). When a row contains multiple triggers, it means that the automation runs when any one of the triggering event happens. When a row contains multiple actions, the automation triggers multiple actions every time the automation runs. Many automations triggered by voice commands have other voice triggers not shown here. Automations labeled with a * relate to cross-location lighting controls.

	Trigger Devices	Automation Triggers (OR)	Action Devices	Automation Actions (AND)
Automations incorrectly identified as secrecy violations				
P11 A1*	Date/Time Voice assistant (any)	6:30 PM Voice: "evening routine"	Dining room light Kitchen lights Living room light Bedroom light	Turn on light
P12 A24*	Driveway camera	Presence detected while nighttime 30 seconds elapsed while presence no longer detected	Outdoor light	Brightness 100% Brightness 25%
P12 A25*	Front yard camera	[same as above]	[same as above]	
P12 A26*	Porch camera	[same as above]	[same as above]	
P14 A9	Voice assistant (any)	Dog barking between 8AM-11PM	Voice assistant (front door)	Say "[dog name] hush"
P16 A18	Email	Receive [email]	Voice assistant (same room as participant)	Play sound
Automations incorrectly identified as integrity violations				
P1 A1*	Voice assistant (any)	Voice: "Halloween Mode" Voice: "October Mode" Voice: "Spooky Mode"	Living room lamp Living room light Dining room light Kitchen light	Light orange, brightness 30 Light purple, brightness 5 Turn off Light green, brightness 14
P1 A4*	Voice assistant (any)	Voice: "Reset the upstairs lights" Voice: "Reset the lights upstairs"	Living room light Dining room light Kitchen light Living room lamp	Set light color to candlelight brightness 35 Set light color to candlelight brightness 50
P1 A9*	Voice assistant (any)	Voice: "Turn off [name] lamp" Voice: "Turn off [name] light"	Teen room lamp	Turn off light
P1 A10*	Voice assistant (any)	Voice: "Turn off the upstairs lights" Voice: "Turn the upstairs lights off"	Dining room light Kitchen light Living room light Living room lamp	Turn off light
P1 A11*	Voice assistant (any)	Voice: "Turn [name] lamp on" 3 other variations	Teen room lamp	Turn on light
P1 A12*	Voice assistant (any)	Voice: "Turn off downstairs lights" 10 other variations	Dining room light Kitchen light Downstairs hallway light Teen room light Teen room lamp Utility room lamp	Turn off light

P1 A13*	Voice assistant (any)	Voice: "Cozy mode"	Stairwell light Living room lamp Living room light Dining room light Hallway light Kitchen light	Turn on/off light
P1 A14*	Voice assistant (any)	Voice: "Night time" Voice: "Night mode"	Kitchen light Dining room light Living room light Downstairs hallway light Stairwell light Living room lamp Utility room lamp	Turn off light Turn on light
P4 A7	Voice assistant (any)	Barking for 45s	Alexa app	Send notification
P7 A1	Voice assistant (any)	Voice: "let's get ready for bed"	Smart plug (Living room AC) Living room light Living room TV Office light Bedroom light 1 Bedroom light 2 Smart plug (Bedroom AC) Google home speaker (bedroom)	Turn off plug Turn off Turn on light Brightness 20 Turn on plug Volume 51
P7 A2	Voice assistant (any)	Voice: "bedtime" Voice: "goodnight"	Bedroom light 1 Bedroom light 2 Google home speaker (bedroom)	Turn off "Tell me tomorrow's weather" Play rain sounds
P7 A3*	Voice assistant (any)	Voice: "someone's home" Presence detected	Hallway light 1 Hallway light 2 Living room light	Turn on light
P7 A4*	Voice assistant (any)	Voice: "Everyone's away" Voice: "Everyone's leaving" No presence detected	Bedroom light 1 Bedroom light 2 Hallway light 1 Hallway light 2 Living room light Office light	Turn off light
P7 A7*	Voice assistant (any)	Voice: "turn on office light"	Office light	Turn on light
P7 A8	Voice assistant (any)	Voice: "turn off living room AC"	Smart plug (Living room AC)	Turn off smart plug
P7 A9	Voice assistant (any)	Voice: "play [music]"	(Same) Google Home Speaker	Play [music]
P12 A10	Back yard camera	Presence detected	Home Assistant app (spouse)	Send notification
P12 A5	Doorbell camera	Presence detected	Home Assistant app	Send notification
P12 A6	Garage camera	Presence detected	Home Assistant app	Send notification
P12 A7	Back yard camera	Presence detected	Home Assistant app	Send notification
P12 A8	Doorbell camera	Presence detected	Home Assistant app (spouse)	Send notification
P12 A9	Garage camera	Presence detected	Home Assistant app (spouse)	Send notification
P13 A15*	Camera (front steps) Camera (driveway)	Motion 45s after motion Motion 8PM-6:30AM 30s after motion 8PM-6:30AM	Front porch light Spot lights (front corner)	Turn on light Brightness 100 set color cool white Turn on light Brightness 100 Turn off Turn off light
P16 A40	Motion sensor (mailbox)	Motion detected while mail=False	Home Assistant app Voice assistant (same room as participant)	mail=True Play audio notification

G Interview Codebook

The codebook used for analyzing our interviews is shown in Tables 11–12. The table includes codes, definitions, and counts.

Count	Code	Description
—	How automations are created	
17	wrote myself	Participant wrote the automations
6	another person wrote	Another person the participant knows wrote automations
5	found online	Participant indicates found automations online
5	platform suggested automation	Participant used automations suggested to them by the platform (e.g., Alexa recommended automating common household behaviors)
1	gen AI wrote	Participant used generative AI to write automations
—	How often automations are changed	
11	monthly/seasonal changes	Change automations infrequently, every month, few months, or “seasonally”
4	1-2 changes a year/don’t change	Don’t regularly change or tinker with them; or change/update them infrequently, on the scale of many months
4	changes every 1-2 weeks	Change automations often; every week or two
2	more than weekly changes	Change automations frequently; more than one every week or “constantly”
—	Who interacts with devices/automations	
19	guest user	People not living in the home interact with devices/automations
10	other home member user	Other people living in the home interact with devices, but may not have access to the app to change device/automation settings
6	other admin user	Other people living in the home have access to the automation apps to change device/automation settings
6	home member kid user	Children under 18 living in the home interact with the devices/automations
3	other user	People not living in the home interact with devices who may not even realize that they are interacting with them (e.g., someone walking by the home triggering the doorbell camera)
—	Guest reactions to devices/automations	
8	interested	Interested in the devices/automations and ask questions/play with them
7	afraid/uncomfortable	Fear/discomfort around their devices
4	uninterested	Uninterested in the devices/automations
2	prank or joke	Play a prank/do something mischievous (though not necessarily malicious) with the devices
2	unaware	Guests are probably unaware of the smart devices
—	Response to guests’ discomfort with devices	
9	turn off	Turn off the device
8	talk	Talk through their guests’ feelings if they express discomfort around smart devices (note: they may follow this up with other actions, so other codes may also be applied)
5	unplug	Unplug the device
4	no action	Not take any action (e.g., “my house my rules”)
3	mute	Mute the device
1	hide device	Hide the device (e.g., in a cabinet)
—	Guest data okay to collect	
7	not sensitive	Participant believes data being collected is not sensitive
3	I’d be fine with it	Participant would be comfortable with the same data being collected about them
2	safety issue	Data is collected for safety reasons (e.g., making sure their kids get home safely/aren’t sneaking out)
2	everyone has similar devices	It is a social norm/”everyone’s doing it”

1	public space	Participant believes data is being collected in a public space
—	Guest data other opinions	
3	uncomfortable collecting data	Participant acknowledges they are uncomfortable collecting other people's data
2	never thought about collecting data	Participant says that they hadn't thought about collecting other people's data before being asked in the interview
—	How guests learn of devices	
13	show/tell	Participant tells them about the devices and/or automations OR guests learn of the devices because they see the participant interacting with them
7	see	Guests see the devices (e.g., participant implies guests will notice the device)
2	don't	Participant doesn't talk about their devices with guests (e.g., because they won't understand them anyway)
—	Purpose of devices/automations	
22	repetitive tasks/convenience/comfort normally do	Repetitive tasks the participant normally does/make tasks they hands-free/make their lives more convenient/keep the home comfortable
16	security (external threats)	Keep participant's home secure from external threats (e.g., prevent unauthorized entry, lights to indicate when doors are locked/unlocked)
13	seasonal decor/jokes/fun	Fun, such as telling jokes or seasonal lighting displays (e.g., Halloween, Christmas)
11	reminders/tracking	Scheduled reminders/tracking things in the home
7	safety (environmental threats)	Keep participant's home safe from internal threats (e.g., water leak, smoke detector, carbon monoxide detection)
7	fixes house/device problem	Fix a problem they encountered (Note: this is only used when people specify that the device/automation fixes a problem with a device or something else in their home, not a problem with another home automation)
6	help with pets	Helps monitor/care for their pets (Likely to be paired with another code)
5	add voice trigger	Add trigger words to fix similar/forgotten voice commands
4	accessibility	Participant has accessibility needs that home automation helps with
4	help with kids	Help monitor/care for children (Likely to be paired with another code)
4	save money, energy	Save money and/or energy/other utilities (Note: this is only used when people specifically mention energy/money, if people are talking generally about turning lights on/off automatically "purpose:repetitive tasks/convenience/comfort" code is used)
3	not wake others	Minimize disturbances to other people living in the house (Note: this is only used when people specifically mention not waking people up, e.g., automatically turning on dim lights to go to the bathroom at night without waking others)
3	health	Health-related reasons (e.g., air quality monitoring for asthma)
2	privacy	Signal to others in the home that they shouldn't be disturbed (e.g., VR headset, on zoom call)
2	play with tech	Participant likes learning about/playing with tech
—	Device sensitivity	
18	devices public	All devices and/or automations are public knowledge
4	discomfort with some interactions	Less comfortable with other people interacting with some devices/interacting with devices in some ways (e.g., they don't want people to change light settings) (Note: this code was used when people bring up these preferences on their own, not for use during the information flow analysis section of the interview)
2	uncomfortable discussing devices	All devices/automations are public knowledge, but some would be easier than others to discuss because they are likely to make people uncomfortable

— Problems with devices/automations		
11	platform changes/issues	Issues getting their devices from different manufacturers/platforms connected (e.g., API being removed, IFTTT becoming more expensive, Google Home not working well with Philips Hue)
10	unanticipated behavior/seasonal change	Interacted with the device differently than they anticipated (e.g., coming home later than usual, kids leaving a door open longer than usual) (Note: this is where the human behaves differently than anticipated not where the device is triggering when it shouldn't) OR a seasonal change (e.g., sun setting earlier/later than usual) caused a problem
10	source unclear	Source of the problem (e.g., which automation was responsible) was not clear to them
8	false trigger	False or incorrect triggers on a voice assistant, e.g., non-trigger word triggers something (note: the problem here is with the voice recognition not with people using the wrong trigger)
6	voice not understood	Voice assistants not understanding/hearing them (Note: if someone uses home automation to add new voice commands because they keep forgetting the correct command, "purpose:add voice trigger" code is used)
5	moving homes/accounts	Moving homes/accounts made devices/automations stop working
4	wrong device activated	Automation activates the wrong device (e.g., music plays the wrong smart speaker)
3	location/geo fencing issues	Inconsistent detection of the user's location
2	automation conflict	One automation somehow interfered with another automation (e.g., one automation works fine until another automation changes something)
<hr/>		
— Pre-existing opinions on automation/device risks		
6	concerns about microphone privacy	Smart voice assistants record everything they hear
1	concerns about others using online services	Concerns about guests using participants' devices to access online services

Table 11: Interview codebook

Table 12: Codebook for participants' reactions to information flow violation scenarios.

Count	Code	Description
<i>Concerned</i>		
11	Annoying	Information flow would be annoying
9	General	General unspecified concern
9	Feelings of privacy violation	Information flow leads to feelings of privacy violation (e.g., creepiness, embarrassment)
8	Security threat	Information flow would lead to a security threat (from outside, e.g., burglars, harassment)
5	Feelings of fear	Information flow leads to a general feeling of fear
4	Not permitted to use device	Trigger device isn't supposed to be used by the other person in question
1	Leaks specific information	Information flow leads to unintentional exposure of information to others
<hr/>		
<i>Not concerned</i>		
14	Indifference	General indifference to the scenario
13	Information isn't private	They don't mind that other people learn information from the flow
10	Intended purpose of automation	Not worried because that was the purpose of the automation anyway
3	Not technically possible	Device features don't actually permit the scenario from happening
2	Funny	Wouldn't be bothered, would find it funny
<hr/>		
<i>Mitigations</i>		
9	Scenario unlikely	Scenario is unlikely to happen, low probability, requires very specific circumstances
5	Technical mitigations	Device features make it hard for the scenario to happen
3	Trust people to not do this	People mentioned in the scenario just wouldn't do it