

LLM-Powered Analysis of IoT User Reviews: Tracking and Ranking Security and Privacy Concerns

Taufiq Islam Protick¹, Sai Teja Peddinti², Nina Taft², Anupam Das¹

¹North Carolina State University, Raleigh, NC, USA

²Google Inc., Mountain View, CA, USA

tprotic@ncsu.edu, psajteja@google.com, ninataft@google.com, anupam.das@ncsu.edu

Abstract

Being able to understand the security and privacy (S&P) concerns of IoT users brings benefits to both developers and users. To learn about users' views, we examine Amazon IoT reviews - one of the biggest IoT markets. This work presents a state-of-the-art methodology to identify and categorize reviews in which users express S&P concerns. We developed an automated pipeline by fine-tuning GPT-3.5-Turbo to build two models: the Classifier-Rationalizer-Categorizer and the Thematic Mapper. By leveraging dynamic few-shot prompting and the model's large context size, our pipeline achieved over 97% precision and recall, significantly outperforming keyword-based and classical ML methods. We applied our pipeline to 91K Amazon reviews about fitness trackers, smart speakers and cameras, over multiple years. We found that on average 5% contained S&P concerns, while security camera exhibited the highest prevalence at 10%. Our method detected significantly more S&P-relevant reviews than prior works: 15x more for fitness trackers, 29% more for smart speakers, and 70% more for cameras. Our longitudinal analysis reveals that concerns like surveillance and data control have persisted for years, suggesting limited industry progress. We demonstrate that across all device types, users consistently demand more precise control over what data is collected and shared. We uncover challenges in multi-user and multi-device interactions, identifying two previously unreported themes concerning inadequate controls for account separation and data access. These findings, ranging from broad persistent trends to specific instances of customer loss, offer actionable insights for developers to improve user satisfaction and trust.

Artifacts — <https://github.com/tprotic/IoT-Products-SecurityAndPrivacy>

1 Introduction

The Internet of Things (IoT) market is estimated to have 29 billion connected devices by 2027. Popular IoT devices, such as fitness trackers, monitor health metrics; smart speakers manage media, calendars, messaging, and web queries; and security cameras enhance home safety. Despite their utility, these devices pose serious security and privacy (S&P) risks. Fitness apps may expose location data, always-listening microphones of smart speakers have led to unau-

thorized recordings, smart cameras often lack robust protection against unwanted data access. Notable breaches like hacks of Foscam and VTech baby monitors underscore these vulnerabilities (Shekyan and Hartutyunyan 2013).

S&P issues raise concerns over data misuse, online fraud, and loss of control which can erode user trust. These concerns are often amplified by media, peer discussions, or unexpected device behaviors (Emami-Naeini et al. 2019). A UK survey found that security is the third most important factor in IoT purchase decisions, after cost and functionality (Interactive 2019). However, S&P details are often hard to understand before purchase (Emami-Naeini et al. 2019). To address this, researchers and policymakers have proposed comprehensive S&P labels (Emami-Naeini et al. 2020). During 2021-2023, the U.S. government advanced this effort through NIST, proposing a Cyber Trust mark with a QR code linking to detailed information (Chen et al. 2024). Such labels are not yet widely adopted, and similar privacy labels in mobile apps remain insufficient (Zhang and Sadeh 2023).

Advances in natural language processing (NLP) now allow automated analysis of product reviews at scale (Harkous et al. 2022), surpassing traditional machine learning in quality. These techniques can extract S&P concerns across devices, offering value to consumers, developers, and researchers. For consumers, this information supports informed purchase decisions. Developers can use feedback to improve product security, prioritize issues, and address root causes of negative reviews. Notably, 60% of S&P-related reviews in one study led to relevant app updates (Nguyen et al. 2019). Retailers have started integrating AI-powered shopping assistants. For instance, Amazon's "Rufus," launched in September 2024, allows users to query product details and review summaries. However, it currently overgeneralizes feedback and lacks focus on S&P concerns. Our work aims to enhance such assistants by improving their ability to identify and summarize S&P issues effectively.

Prior research on users' S&P perceptions has relied on small-scale surveys (Emami-Naeini et al. 2021) or keyword-based analyses (Protick et al. 2024). The most advanced prior work (Harkous et al. 2022) used a T5 model and showed improved accuracy over manual and classical methods like LDA or SVMs (for apps, not IoT products). To our knowledge, we are the first to leverage LLMs for this task.

Our first research question asks **RQ1: *What level of qual-***

ity can be achieved for LLM-based detectors aimed at extracting S&P reviews from a large corpus of IoT reviews?

We aim to understand if one detector can uncover a wide range of S&P issues, while simultaneously reaching high accuracy. With the emergence of shopping assistants, it is important to understand the quality of such a detector and its ability to uncover a wide range of issues. While we expect LLMs to do well, our work offers a benchmark for future efforts since the labeled data we use is public. We fine-tune GPT-3.5-Turbo to build two models: the *Classifier-Rationalizer-Categorizer* (CRC), which detects and labels S&P issues in reviews, and the *Thematic Mapper* (TM), which groups these issues into higher-level themes based on established taxonomies. We see that our models identify substantially more S&P-relevant reviews than prior approaches.

With our models, we explore two more research questions. **RQ2: What is the prevalence of S&P concerns in IoT product reviews, and how does prevalence vary across device categories?** We analyze Amazon reviews of fitness trackers, smart speakers, and security cameras to measure how frequently users raise S&P concerns and how this varies by device type. **RQ3: What are users' S&P concerns and how do they rank relative to one another?** Our LLM-based pipeline supports issue extraction and ranking, allowing for trend analysis of S&P issues across device types over time.

In summary, we make the following contributions:

- We present a state-of-the-art approach by fine-tuning GPT-3.5-Turbo to develop two models. The first model, *CRC*, achieves over 97% precision and recall on a human-annotated validation set, and macro-averaged precision and recall of 88.7% and 93.4%, respectively, on real-world data. The second model, *TM*, achieves a micro-averaged precision of 96.8%, recall of 95.7%, and an F1-score of 96.3%. Comparative analysis with existing approaches demonstrates that our method outperforms traditional keyword-based techniques, classical ML/NLP models, and baseline LLMs.
- We applied our pipeline to 91K IoT product reviews and identified 4,896 reviews ($\approx 5\%$) with S&P concerns. Security cameras had the highest rate of such concerns (10.07%), significantly more than fitness trackers and smart speakers. We also observed a slight negative correlation between S&P mentions and star ratings, suggesting these concerns may contribute to lower satisfaction.
- The combination of more accurate methods (that extract more content than prior efforts), coupled with our categorizations of S&P concerns into 28 categories (permitting more distinction than prior work (Vetrivel et al. 2023) with 7 categories) enables simultaneously a broad view across issues (top issues, issue ranking, temporal trends), and also specific insights (e.g., trust, ethics, customer loss, and the variety of specific controls users seek). Surveillance dominates concerns for smart speakers and cameras. For both fitness trackers and cameras, the top three issues persist over time, indicating limited progress in addressing user concerns. A common, cross-device theme is users' desire for finer control over both data types and recipients, including household members, visitors, land-

lords, corporations, and government agencies.

- Our analysis further uncovers significant challenges in multi-user interactions with IoT devices, identifying two previously unreported themes that reveal distinctive privacy risks associated with cross-device interactions. Users frequently report inadequate and poorly designed controls for account separation and data access. Examining 52 speaker and camera device types over multiple years, our study highlights the widespread nature of these issues. We conclude by summarizing these challenges as actionable feedback for developers.

2 Related Work

Analysis of Mobile App Reviews. Multiple studies have analyzed Google Play reviews to identify S&P concerns in mobile apps. Nguyen et al. use an SVM classifier to detect privacy-relevant reviews and correlate them with app updates (Nguyen et al. 2019), while Mukherjee et al. apply a similar approach to study S&P concerns (Mukherjee et al. 2020). Cen et al. perform two-dimensional labeling of S&P issues using supervised ML (Cen et al. 2014). Besmer et al. combine ML and NLP to explore privacy perceptions based on reviews and ratings (Besmer, Watson, and Banks 2020). Harkous et al. use deep learning to summarize privacy feedback (Harkous et al. 2022), and Akgul et al. analyze 12 million privacy-relevant reviews using a large language model (Akgul et al. 2024). Our work builds on these efforts by leveraging a more advanced LLM. These latter 2 works focused on app reviews, while we study IoT reviews.

Analysis of IoT Product Reviews. Linden et al. analyzed 20K reviews of human wearables (e.g., fitness trackers) and found privacy concerns were rarely expressed (0.14%) (van der Linden Dirk et al. 2020). Fruchter et al. used a 53-keyword lexicon to identify S&P concerns in 109K smart speaker reviews, finding 2.04% raised issues related to data collection, creepy behavior, and privacy violations (Fruchter and Liccardi 2018). Vetrivel et al. examined 68K Amazon reviews of surveillance systems, hubs, set-top boxes, and routers, identifying S&P concerns in 9.8% of them (Vetrivel et al. 2023); for cameras specifically, they report only 5.9% compared to our 10.07%. Most recently, Protick et al. used a keyword-based approach to qualitatively analyze smart home IoT reviews, surfacing common and unique S&P issues, such as invasive customer support and user skepticism (Protick et al. 2024).

Comparison with Prior Works. To our knowledge, we are the first to fine-tune a large language model on unstructured IoT product reviews. Prior studies (van der Linden Dirk et al. 2020; Fruchter and Liccardi 2018; Protick et al. 2024) typically rely on keyword-based filtering, which struggles to capture semantically similar content and is prone to high false positive and false negative rates (Harkous et al. 2022). While large open-source models like T5 have been used to analyze privacy-related mobile app reviews (Harkous et al. 2022; Akgul et al. 2024), our method leverages a (larger size ¹) closed-source API-based model that does not need

¹T5-11B has 11 billion parameters. The exact size of GPT-3.5-

Context	Reference	App/Device-Types (#)	Training Data	Model	% of S&P reviews (#)	S&P Themes
Mobile Apps Reviews	(Cen et al. 2014)	6,938 apps	keywords	Logistic Regression	0.05% (5M)	5
	(Nguyen et al. 2019)	2,583 apps	keywords	SVM	0.12% (4.5M)	N/A
	(Mukherjee et al. 2020)	539 apps	keyword	SVM	0.5% (2.2M)	N/A
	(Besmer, Watson, and Banks 2020)	5M apps	keyword	Logistic Regression	0.07% (4.9M)	9
	(Harkous et al. 2022)	1.3M apps	taxonomy+NLI	T5	0.9% (626M)	25 (P only)
	(Akgul et al. 2024)	160K apps (over 10 yrs)	taxonomy+NLI	T5	0.65% (1.9B)	89 (P only)
IoT Product Reviews	(Frucher and Liccardi 2018)	speakers (5)	keywords	Manual	2.0% (109K)	3
	(van der Linden Dirk et al. 2020)	activity & location trackers (14)	keywords	Manual	0.14% (20K)	8
	(Vetrivel et al. 2023)	cameras (881), routers (440), hubs (54), set-top boxes (41)	keywords	LDA	9.8% (68K)	22
	(Protick et al. 2024)	trackers (3), speakers (3), cameras (7)	keywords	Manual	N/A	14
	Our Work	trackers (23), speakers (24), cameras (28)	qualitative	GPT-3.5-Turbo	5% (91K)	28

Table 1: Comparison with existing works that elicit security and/or privacy issues from user generated reviews.

self-hosting for fine-tuning and is continuously updated. Also, we focus on the IoT ecosystem, such as fitness trackers, smart speakers, and cameras—that involve distinct data modalities and S&P risks. Moreover, previous methods often rely on hypothesis-based Natural Language Inference (NLI) to bootstrap data for manual labeling and later use it for fine-tuning large models. In contrast, we transform a pre-existing IoT labeled dataset (Protick et al. 2024), and fine-tune GPT-3.5-Turbo using dynamic few-shot prompting that effectively utilizes the larger context size of 16K tokens by including manually labeled examples. And lastly, while prior work has utilized four models/steps to organize reviews into thematic topics (one fine-tuned model for privacy-relevant review detection, another fine-tuned model for low-level issue tag generation, leader clustering technique to merge these low-level issues into higher level themes, and one fine-tuned generative model to generate theme titles), we compressed these into two fine-tuned models (CRC and TM) and focused the thematic analysis on 28 high-level themes derived from a unified S&P taxonomy merged from eight vetted IoT taxonomies. Our methodology requires fewer models and is anchored on vetted taxonomies. In summary, our study is not keyword-dependent, covers 28 nuanced S&P themes (e.g., privacy ethics, data accuracy), and enables automated analysis of IoT product reviews over time. This approach captures both the evolution and ranking of issues, offering actionable insights to help developers prioritize mitigations. Table 1 summarizes these key differences with existing related works.

3 Methodology

Data Description

Data Collection. We adopted the following steps for data collection. First, we go to `amazon.com` and search for device categories like “fitness trackers”, “smart home speakers”, and “security cameras”. We focus on these three device categories for two key reasons. First, they encompass diverse data modalities — such as motion sensors, heart rate, GPS, audio, and video — each presenting distinct security and privacy risks. Second, this selection allowed us to leverage thousands of existing human-annotated data points from existing work (Protick et al. 2024) for model training and evaluation. We next list products and their ASINs from the first search page. Then, with a crawler created with Sele-

Turbo is unknown, but is estimated to be larger than T5-11B.

nium and BeautifulSoup, we collect review texts and meta-data (including rating, user’s country information, date, and ASIN product IDs). The data was gathered in May 2023. Finally, we preprocess the crawled data into the wild review dataset (\mathcal{W}), described in more detail later in the section.

Ethical Considerations. We comply with Amazon’s terms and conditions of use during our data collection. Our data crawling process strictly follows the guidelines outlined in Amazon’s `robots.txt` file, particularly when accessing product review pages. Furthermore, we implement rate limiting for our crawlers to minimize server load and avoid any adverse impact on their infrastructure. To safeguard consumers’ privacy, we exclude any personally identifiable information, even if publicly accessible, such as usernames and posted images, as part of our data collection process. Also, fake reviews are a prominent problem on online marketplaces (including on Amazon) (He, Hollenbeck, and Proserpio 2022), and we rely on Amazon’s mechanisms for removing fake reviews and banning abusive accounts. Since we are collecting publicly available data, our IRB office, upon review, classified the study as an exempt.

Wild Review Dataset (\mathcal{W}). The crawled Amazon reviews are preprocessed to obtain a good quality wild review dataset \mathcal{W} . We collected a total of 127,821 reviews, which, after removing the empty entries, came down to 127,270 reviews. The number came down to 114,383 reviews after filtering out non-English content. We then eliminated duplicate reviews, resulting in 91,749 reviews. The distribution of reviews across the three device categories is as follows: trackers (24,046), speakers (32,179), and cameras (35,524). These reviews are for products from 75 vendors (trackers: 23; speakers: 24; cameras: 28), identified by ASINs. We will be making this wild review dataset public. With Amazon changing its policy towards the end of our data collection period and drastically reducing the number of reviews that can be collected (from 5,000 to just 100), we hope other researchers will find our dataset (Protick and Das 2025) useful.

Qualitative Labeled Data for Training. We fine-tuned our models using pre-labeled qualitative data from previous research eliciting S&P concerns from IoT product reviews (Protick et al. 2024). The dataset follows a three-step analysis of IoT products: keyword selection, review labeling, and thematic analysis, all conducted by multiple coders with high inter-rater reliability.

Model Choice & Rationale for Fine-tuning

To address **RQ1** in evaluating whether LLM-based detectors can effectively extract S&P reviews, we opted to use the GPT-3.5-Turbo model. We choose GPT-3.5-Turbo due to its high throughput limits (up to 1M tokens/min, 100M tokens/day for Tier-4 users), cost-efficiency, and strong performance across tasks. The computation was done on a machine that uses Linux with 16 GB of RAM. It uses an NVIDIA RTX-3070-M GPU for faster computation of BERT embedding. For deterministic outputs, we set the `temperature` parameter to 0, and to overcome hallucination-induced deviations in the LLM response, we set both `frequency_penalty` and `presence_penalty` values to 2.

We utilize the Supervised fine-tuning (SFT) method to adapt the GPT-3.5-Turbo model for our case. Rather than simply give input-output pairs during fine-tuning, we leveraged the larger context size of the model (16K tokens) to include additional few-shot examples that are dynamically chosen for each input text in the training data. We fine-tune the base model for multiple reasons, all of which are related to the downstream task of our interest. First, accurate identification of S&P concerns requires domain knowledge, which the base model may lack without fine-tuning (we provide a comparative analysis of vanilla vs. fine-tuned models in the next section). Second, the base model exhibits inconsistent responses for the same input, even with the `temperature` set to 0. Finally, it often fails to follow the structured response format required for our task.

Although GPT-3.5-Turbo is sometimes described as a “base model,” it is already instruction-tuned and not a raw pre-trained foundation model. Our work adds an additional round of task-specific supervised fine-tuning to adapt it to IoT security and privacy concerns. For fine-tuning, we use OpenAI’s API with 3 epochs and evaluate temperatures between 0.0 and 0.8 in increments of 0.2. The model’s learning rate is internal. However, the API gives a learning rate multiplier value to be used, which in our case was set to 2. The 16K-token context window allows dynamic few-shot examples. The best performance is achieved at `temperature = 0`.

Classifier-Rationalizer-Categorizer Model

Given a dataset of reviews, our first goal is to detect S&P concerns mentioned in these reviews. For this, we focus on three key aspects: first, determining whether a review text contains any concerns (*classify*); second, obtaining a justification/explanation for the decision (*rationalize*); and third, determining what specific concerns are present in the review text (*categorize* into low-level issues). We have modeled these three aspects as three tasks to be performed on a given review text ‘*t*’, as elaborated in this section. We have included the second step based on prior work showing that ‘chain-of-thought’ prompting improves performance on a range of reasoning tasks (Wei et al. 2022).

Dataset for Training and Validation. We leverage an existing dataset from (Protick et al. 2024) containing IoT product reviews labeled for S&P concerns. However, this dataset was created for a different purpose and has different seman-

```
—— System Prompt ——
You are a Large Language Model for classifying, rationalizing,
and categorizing security and privacy concerns in text.
—— User Prompt ——
Given the text t: t, a review to label
I want you to perform three tasks on t:
Task 1: Determine if text t explicitly mentions any security or
privacy concerns. Respond only with a ‘Yes’ or a ‘No’.
Task 2: Provide a brief explanation as to why your response is as
such for Task 1.
Task 3: If the response for Task 1 was a ‘Yes’, categorize the
explicit concerns mentioned in text T with themes related to
security and privacy.
[At this point, a random number x is generated between 0 and 1.
If x < 0.5, we assign Example A = t1, the text closest to t. Else,
we assign Example A = t2, the text second closest to t with the
opposing label as that of t1 from the training dataset]
Here are two examples that might be helpful:
Example text A: x < 0.5 ? t1: t2
Ideal response for A: x < 0.5 ? Label(t1): Label(t2)
Example text B: Show the other example
Ideal response for B: Show the other example’s label
```

Figure 1: The prompt for the CRC model

tics. So we first manually transform this dataset according to our task-set requirements, as shown in Table 6 in Appendix C. Since this pre-existing dataset was already labeled by two annotators and went through rigorous disagreement resolution to generate high-quality labels, one of the authors focused on manually transforming this data into our desired format. The labels for task-1 (*classify*) and task-3 (*categorize*) were directly derived from prior labels, where reviews containing S&P concerns were considered as positive samples and the rest as negative. The same author manually augmented these labels by generating explanation strings (i.e., task-2 *rationalize* output) to capture the reasons behind these labels based on prior annotations. As we focused on data transformation rather than new-label generation, multiple annotators were not involved and inter-annotator agreement metrics were not computed. Since there were not enough negative samples in this pre-existing dataset, we augmented it with additional negative samples from our crawled dataset, \mathcal{W} , through manual selection until approximately 600 were identified (this reduced our \mathcal{W} dataset from 91,749 to 91,149 reviews). Our final labeled dataset contains 2,454 reviews, *equally divided* between positive and negative samples.

Prompt Engineering. Given the complexity of the CRC task (containing 3 sub-tasks), we curate a good prompt for conducting Supervised fine-tuning (SFT) of the GPT-3.5-Turbo model that takes advantage of its longer context window (16K tokens). We utilize the knowledge gained in the manual data transformation process to create this prompt. For a given review text and its converted labels (right column of Table 6), we finalized the prompt shown in Figure 1, where underlined texts vary for each input review text.

The first statement is the system prompt, which guides the model’s behavior. What follows is the user prompt, where we first give a review text *t* and specify the task set that al-

lows the model to know what to do with t . We show two helpful examples t_1 and t_2 with their ideal responses from the training dataset so that the generative model can understand the correctness and the desired response format for t . Instead of fixing these few-shot examples in the prompt, we dynamically choose these examples (dynamic few-shot prompting) based on t . We select review text t_1 that is closest to t in the embedding space, and t_2 that is second closest to t in the embedding space but has an opposite label for task-1. Here, the texts are represented with their BERT embeddings (Devlin et al. 2019) (which help represent text in a multi-dimensional space, where semantically related texts are closer together) and closeness is calculated with cosine similarity. To make sure we do not have label-specific bias, t_2 is chosen to have the opposite task-1 label from t_1 (Line 18 of Algorithm 1 in Appendix A), i.e., t_2 is a helpful review with no S&P concerns if t_1 is a review with S&P concerns, and vice versa. To further reduce positional bias between the examples, we randomly show the closest or the second closest example first based on a coin toss (text in square brackets in Figure 1). Note that $Label(t_i)$ includes the outputs for all three tasks for t_i (classification, rationale, and categorization outputs). An example of a CRC prompt with all values filled in is shown in Appendix B.

Model Training. During the training phase, we utilize the manually annotated task set responses for each review t .

The model is trained using 80% of the dataset (1964 reviews), with the remaining 20% (490 reviews) reserved for validation. The temperature setting ($temp$) is a hyperparameter during model training, and we vary it from 0 to 0.8 in increments of 0.2, resulting in five finetuned models. Each of these models is evaluated on the manually labeled validation set (balanced between positive and negative samples) to identify the best-performing one. We found that the model with $temp = 0$ yielded the highest accuracy, precision, recall and F1-score – all close to 97.8%.

Impact of Prompt Engineering. To quantitatively assess the utility of the dynamic few-shot prompting, we evaluated the performance of the Vanilla model (i.e., without finetuning) using all three prompting strategies: zero-shot (prompt in Figure 1 with no examples), fixed-shot (prompt with two fixed examples), and our dynamic few-shot prompting. We found that the dynamic few-shot prompting method achieved the highest performance (89%) on the validation dataset, outperforming fixed-shot (88%) and zero-shot (83%). Our fine-tuned CRC model employing dynamic few-shot prompting further boosted the accuracy to 98%.

We would like to point out that the closest neighbors identification in the dynamic few-shot prompt is necessary only during training, and we do not actually need to add few-shot examples in the prompt during inference – the high performance of the model is maintained even if we skip populating the few-shot examples during inference. Hence, the additional cost of computing BERT embeddings and cosine similarities for finding the nearest neighbors is a one-time training cost – where we compute and store embeddings for all the samples in the training data, and perform quick matrix multiplications to identify nearest neighbors (both positive

Model	Detail	Acc. %	P %	R %	F1 %
Keyword-based	(Cen et al. 2014)	71.0	92.6	45.7	61.2
	(Nguyen et al. 2019)	58.0	64.4	35.5	45.8
	(Mukherjee et al. 2020)	76.1	96.4	54.3	69.5
	(Fruchter and Liccardi 2018)	79.6	95.0	62.4	75.4
	(van der Linden Dirk et al. 2020)	55.5	93.5	11.8	21.0
	(Vetrivel et al. 2023)	78.8	86.9	67.8	76.1
	(Protick et al. 2024)	85.9	98.9	72.7	83.8
Classical Models	Logistic Regression	90.4	90.2	90.6	90.4
	SVM	84.3	80.9	89.8	85.1
	KNN (k = 37)	80.2	76.1	91.0	82.9
LLM	Vanilla GPT-3.5-Turbo	88.7	83.4	96.0	89.3
	CRC (our model)	97.8	97.6	98.0	97.8

Table 2: Performance comparison between our model and others. Model performance for Class 1 (S&P concerns), showing precision, recall, and F1-scores.

and negative cases) for each sample.

Comparison with Other Models. Using the validation dataset of 490 reviews (balanced 245 positive and 245 negative samples), our evaluation (as shown in Table 2) highlights the limitations of keyword-based and classical ML methods as well as the advantages of advanced models like CRC. Keyword-based approaches achieve high precision by relying on fixed terms but suffer from very low recall, missing concerns expressed in varied wording. While this validation dataset is balanced, real data is heavily imbalanced with very few S&P examples. To evaluate how keyword-based approaches work on real datasets, we curated a second validation dataset containing a random sample of 300 reviews that are manually labeled. This dataset contains eight S&P and the rest non S&P reviews. All the keyword-based models have a really high number of false positives, resulting in precisions less than 20% and recall values less than 50%.

Classical machine learning models outperform keyword-based approaches by capturing broader patterns and improving recall, but they still face challenges in balancing precision and recall effectively. Prior works using LLMs such as T5 (Harkous et al. 2022; Akgul et al. 2024) have not released their models publicly, preventing direct comparison on the same datasets. In contrast, general-purpose LLMs like Vanilla GPT-3.5 achieve high recall but exhibit lower precision, frequently misclassifying benign reviews as security and privacy concerns. In contrast, our CRC model achieves both precision and recall above 97%, effectively capturing nearly all true S&P concerns while minimizing false alarms—a balance critical for trustworthy S&P analysis. Moreover, vanilla GPT-3.5-Turbo failed to generate outputs for 9% of validation samples (44 of 490) and achieved only 88.7% accuracy on the remainder, with 52 misclassifications—five times higher than CRC. By comparison, our fine-tuned CRC model reached 97.8% accuracy, demonstrating the effectiveness of task-specific fine-tuning in achieving reliable performance.

Sampled Performance on Wild Dataset. As a sanity check we also randomly sampled 300 reviews from the wild dataset

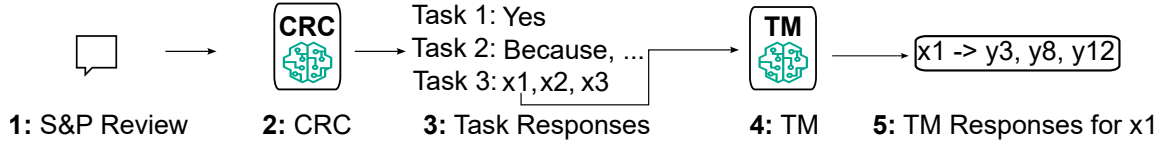


Figure 2: The CRC and TM models in tandem. The Task 3 responses of the CRC model are passed to the TM model as input.

\mathcal{W} (with 100 from each device category), manually labeled them, and used these to verify the model’s performance.

Despite significant class imbalance (only 8 of the 300 samples are positive examples), the model achieves a good performance: 99% accuracy with macro average precision of 88.7%, recall of 93.4%, and F1-score of 90.9% (weighted average precision, recall, and F1-score are all 99%). There were only 2 false positives and 1 false negative, indicating high accuracy and reliability in classifying S&P concerns. Table 4 in Appendix B shows the full evaluation of the CRC model across the three different tasks; different metrics are used for evaluating closed and open-ended responses.

Processing the Full Wild Dataset. Next, we processed the entire wild dataset \mathcal{W} through the model. It took approximately 30 hours to submit all the prompts to OpenAI’s API, yielding labels for 91,149 reviews.

Thematic Mapper Model

The output of the CRC model helps us quantify the number of S&P-related reviews and identify the low-level issues raised in these reviews. However, these identified concerns were often too review-specific and contained variability in their expression, yielding too many uniquely worded low-level issues that are similar in nature. This prevents us from succinctly summarizing the key issues users share in their reviews. So we propose to fine-tune a Thematic Mapper (TM) model that will map the low-level issues (i.e., task-3 responses) of the CRC model to one or more high-level themes related to S&P topics. For instance, both ‘concerns about eavesdropping’ and ‘concerns about listening’ can be mapped to high-level ‘surveillance’ S&P topic. Such high-level themes are more digestible for stakeholders.

Collection of High-level Themes. To ensure broad and comprehensive coverage of S&P topics, we merged eight vetted IoT taxonomies (Williams, Rojas, and Bayoumi 2019; Rizvi et al. 2018; Alsamani and Lahza 2018; Das, Zeadally, and He 2018; Gulzar and Abbas 2019; Akbar et al. 2021; Chen et al. 2018). Existing taxonomies for IoT products face two main limitations: (1) some include a limited number of subcategories, omitting important issues like anonymity, privacy ethics, or data accuracy; and (2) some categories are overly broad (e.g., “user fears,” could encompass issues like unintended data sharing, account management difficulties, or hacking), making them unhelpful for users to interpret and too ambiguous for developers to know what to fix. Three S&P experts from academia and industry collaboratively merged the taxonomies, removing redundancies and disaggregating overly broad topics. All dis-

agreements were discussed and resolved through consensus, consistent with Delphi methodology (Linstone and Turoff 1975), which emphasizes convergence over inter-rater reliability. The definitions of our 28 themes are available in our open-source repository (Protick and Das 2025).

Dataset for Training and Validation. To train the TM model, we leverage 1,200 *unique* low-level issues (from 1,989 total) stemming from the positive reviews in the training data for the CRC model, and manually map them to the 28 high-level themes. Annotators performed this manual mapping in multiple rounds and resolved any discrepancies through discussion until they arrived at a final agreement. This mapped dataset is used to fine-tune our TM model.

Prompt engineering. Figure 4 in Appendix C shows the prompt we use to fine-tune the GPT-3.5-Turbo model to be our thematic mapper (the underlined texts vary for each example). Like in the CRC prompt, the first statement is a system prompt, and this is followed by the user prompt. In the user prompt, we first give the low-level issue x to map. Next, we provide the list of taxonomy terms Y with their definitions, and later follow this with the task that the model should perform on x . Finally, we show five helpful examples t_1 through t_5 with their correct labels $Label(t_1)$ through $Label(t_5)$ from the training dataset, so that the generative model can learn from some ground truth. Similar to the CRC model, we follow a dynamic few-shot prompting technique where the shown examples are chosen based on how close they are to x in the embedding space using cosine similarity. For the TM model, we choose TF-IDF embeddings (Line 3 of Algorithm 2 in Appendix A) as opposed to BERT embeddings based on its performance on short phrases (like our low-level issues).

For example, when using TF-IDF embedding, the closest helpful example retrieved for ‘password sharing as a violation of basic IT security principles’ is ‘password security’, whereas it is ‘unauthorized listening or recording in a joking tone’ when using BERT embedding.

Model Training. Each low-level issue from Task-3 responses from the CRC model is passed to the TM model (see Figure 2). We utilize 70% of the 1,200 curated low-level to high-level theme mappings as our training dataset, and the rest (i.e., 360 or 30%) as a validation dataset to train our TM model. Based on our experience with the CRC model, we set the `temperature` parameter to 0. To assess the performance of the TM model, we verified each high-level theme prediction made for the 360 low-level issues in our validation dataset. Note that a single low-level theme can be mapped to multiple high-level themes. For instance,

‘concerns about sharing login details’ was mapped to ‘authentication’ and ‘data sharing’. Two contingency tables of dimension $rows \times terms = 360 \times 28$, one for TM predictions and the other for ground truths, were created for comparison. The model achieves a macro-averaged precision of 90.99%, recall of 88.14%, and F1-score of 89.33 (micro-averaged precision is 96.88%, recall is 95.76%, and F1-score is 96.32%) (see Table 7 in Appendix C for full results).

In summary, the answer to **RQ1** is that our CRC and TM models effectively detect S&P issues in user reviews with a micro-averaged F1 score of 98% and 96%, respectively.

4 Results

Prevalence of S&P Concerns

We now address our second research question, **RQ2**, regarding the prevalence (indicated by the proportion of reviews) of S&P concerns in IoT product reviews. Task 1 of our CRC model classifies 4,896 out of 91,149 reviews (5.37%) from the wild dataset to have S&P concerns. To compare these concerns across device category c , we calculated the *concern ratio* ($|C_c|/|R_c|$), where $|C_c|$ is the count of reviews with concerns and $|R_c|$ is the total reviews in category c . We check for significant differences in concern ratio with the Chi-squared test for proportions, followed by pairwise comparison with Bonferroni correction to control family-wide error rate. The Chi-squared test for proportion (Table 8 in Appendix D, middle) assesses whether the concern ratios are consistent across the device categories. The test yielded a $p < 2.2e^{-16}$, indicating a highly significant difference in concern ratios across the categories. Pairwise comparisons (Table 8, right) showed significant differences between each pair of categories (all $p < 0.001$), rejecting the null hypothesis and confirming that the concern ratios are not the same across trackers, speakers, and cameras. While the pairwise comparison shows significant inequality between the concern ratios of trackers and speakers, their numerical values are relatively similar (trackers = 2.11% vs. speakers = 2.64%). In contrast, cameras have a significantly higher concern ratio of 10.07%, surpassing that of trackers and speakers. Prior work found 0.14% S&P reviews for fitness trackers (van der Linden Dirk et al. 2020), 2.04% for smart speakers (Fruchter and Liccardi 2018), and 5.9% for cameras (Vetrivel et al. 2023). Thus our classifier identifies 15× more S&P reviews for fitness trackers, 29% more for speakers, and 70% more for cameras. This is likely due to higher recall compared to prior solutions (Table 2).

We checked for correlations between the S&P reviews and the star ratings by considering the counts of S&P concerned and non-concerned reviews across different star ratings, and conducting a point-biserial correlation test. Results are shown in Table 9 in Appendix E. Comparing correlation coefficients (r_{pb}) category-wise and overall, we see no correlation in speakers, however we do observe low negative correlation otherwise. This confirms a negative correlation between S&P concerns and review ratings, indicating that such concerns likely contribute to lower user satisfaction, as noted in prior work (Vetrivel et al. 2023).

Types of S&P Concerns

We now look at **RQ3**, which asks what the S&P concerns are and how they rank. In Table 3 we rank all the concerns for each device, showing the top concerns as well as lesser visible but important issues, and the relative level of concerns across these issues. Our longitudinal data allows us to view the evolution of reviews for top themes per device over multiple years, as shown in Figure 3. The plot displays absolute values rather than normalized proportions (which is around 5% per year) since our analysis focuses on trends in total user contributions rather than relative growth (see discussion below). We include total unique S&P review counts as separate bars to reflect likely increases in product adoption, noting that the historical data per device type is influenced by our process for identifying top devices in each category.

Fitness Trackers’ Top Issues. Table 3 shows that the top three issues for fitness trackers are Data Security and Data Theft, Privacy Controls, and Data Collection. In Data Security and Data Theft reviews, users express significant worry about apps that require credit card information for ‘free’ trials, as they fear it can lead to reports of unauthorized charges and account breaches. One user writes:

“Shortly after using this and giving it all the permissions it asked for, there were mysterious charges on my debit card...”

In the reviews about Privacy Controls, users express frustration with the limited privacy controls to customize the apps or notifications (especially when they display sensitive data), and to restrict third-party data sharing. For example:

“I hate that it shows what day of my menstrual cycle I’m on. How embarrassing ... no way to remove that”.

We found a sizeable number of reviews about embarrassing notifications and an inability to control them. This indicates users are worried about their privacy in front of others in the same room, possibly due to shoulder surfing. In reviews labeled Data Collection, users express concern that fitness trackers and their apps collect excessive personal data. These concerns have been reported in prior work, and we expected them to appear in our corpus.

Our ranking of themes reveals that a number of issues are more frequent than location tracking. This contrasts prior work in which location tracking for fitness trackers has been identified as a top 3 theme (Gabriele and Chiasson 2020). This contrast may arise because prior studies are smaller in scale (e.g., 200 participants), while our automated analysis examines thousands of reviews and tens of topics.

Our temporal measurement in Figure 3 shows steady growth in the absolute number of reviews for the top 3 issues, suggesting that as device adoption increases, related concerns are rising too—indicating these issues remain inadequately addressed by the industry.

Smart Speakers’ Top Issues. The top issue for Smart Speakers is Surveillance, comprising 17% of the reviews. In these reviews about privacy, we see that users are deeply anxious about the ‘always-on’ nature of smart speakers, fearful of continuous conversation recording and data sharing with manufacturers, the government, or third parties. Concerns are heightened by unexpected activations without a clear wake word. An example review is:

Trackers			Speakers			Cameras		
High-level Theme	#	%	High-level Theme	#	%	High-level Theme	#	%
1. data security and data theft	215	12.47	1. surveillance	501	17.32	1. surveillance	1914	16.56
2. privacy controls	175	10.15	2. authentication	302	10.44	2. privacy controls	1347	11.66
3. data collection	155	8.99	3. authorization	288	9.96	3. usability	1116	9.66
4. usability	116	6.73	4. access control	232	8.02	4. data security and data theft	992	8.58
5. authentication	109	6.32	5. privacy controls	218	7.54	5. access control	799	6.91
6. policies and law	99	5.74	6. data security and data theft	212	7.33	6. authentication	701	6.07
7. access control	98	5.68	7. usability	192	6.64	7. authorization	655	5.67
8. authorization	92	5.34	8. data collection	137	4.74	8. trust and transparency	642	5.56
9. location tracking	88	5.1	9. secure communication	119	4.11	9. policies and law	424	3.67
10. surveillance	76	4.41	10. trust and transparency	117	4.05	10. data management and storage	400	3.46
11. trust and transparency	64	3.71	11. privacy ethics	75	2.59	11. general comments related to S&P	374	3.24
12. data accuracy	54	3.13	12. general comments related to S&P	73	2.52	12. secure communication	321	2.78
13. general comments related to S&P	45	2.61	13. policies and law	70	2.42	13. data collection	316	2.73
14. consent	44	2.55	14. personalized advertising	64	2.21	14. security vulnerabilities	233	2.02
15. data sharing	41	2.38	15. data sharing	49	1.69	15. software and firmware updates	207	1.79
16. personalized advertising	41	2.38	16. security vulnerabilities	47	1.63	16. personalized advertising	188	1.63
17. software and firmware updates	39	2.26	17. software and firmware updates	43	1.49	17. privacy ethics	146	1.26
18. secure communication	37	2.15	18. consent	33	1.14	18. data deletion	145	1.25
19. data management and storage	27	1.57	19. data harms	23	0.8	19. availability	132	1.14
20. data exposure	22	1.28	20. data deletion	22	0.76	20. data sharing	116	1
21. security vulnerabilities	21	1.22	21. data management and storage	21	0.73	21. data accuracy	99	0.86
22. data deletion	15	0.87	22. confidentiality	16	0.55	22. consent	98	0.85
23. data harms	15	0.87	23. location tracking	12	0.41	23. data harms	51	0.44
24. privacy ethics	15	0.87	24. data accuracy	8	0.28	24. location tracking	47	0.41
25. availability	11	0.64	25. availability	7	0.24	25. data exposure	34	0.29
26. confidentiality	10	0.58	26. data exposure	6	0.21	26. confidentiality	31	0.27
27. anonymity	0	0	27. anonymity	4	0.14	27. anonymity	22	0.19
28. data hiding	0	0	28. data hiding	1	0.03	28. data hiding	7	0.06

Table 3: Distribution of high-level S&P themes (28 in total) across different device categories. The column # denotes the number of reviews tagged with a high-level theme. The column % denotes the overall percentage of their prevalence.

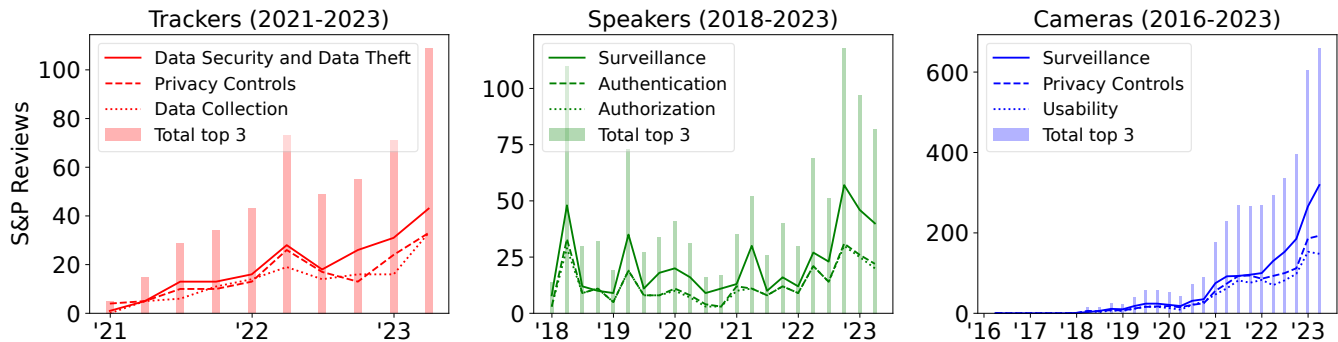


Figure 3: Quarterly trends of the top three S&P themes and unique review counts with S&P concerns across trackers, speakers, and cameras. Each review is counted once, even if it covers multiple themes.

“You can believe the stories you hear about [PRODUCT] listening, ... talk about something and it would appear in your news feed advertisements. Scary. Unplugged it.”

This review is interesting as it indicates that the user did more than complain — they stopped using the device.

The privacy challenges of this sort have been identified before through surveys (Lutz and Newlands 2021) and via experimentation (Dubois et al. 2020). One benefit of automated review analysis is that it enables longitudinal trend tracking. In Figure 3, we see that the number of reviews about surveillance issues for smart speakers has remained fairly stable over a 5-year period. We examined the spikes in this plot and found that they mostly happen in January. We found many mentions of words like “gift”, “holiday”, and “Christmas”, suggesting these spikes are occurring after

people have received holiday gifts.

The second and third issues are Authentication and Authorization. In these reviews, users tell stories about unauthorized access to user accounts and devices by people they view as scammers who impersonate others to gain remote access and steal personal information. For example:

“I used the pop up virtual assistant during setup who re-moted into my computer, and installed a network security program. Two days later I got 8 calls about fraudulent charges on my account.”

The temporal trends for smart speakers (from Figure 3) appear stable. Since the adoption is likely growing, this indicates a relative decrease in complaints about smart speakers. We can only speculate about this, either being due to growing learned helplessness or device improvement; this would

be interesting to explore in future work.

Camera’s Top Issues. As with smart speakers, Surveillance is the top issue for cameras (17% of S&P reviews). Users fear unauthorized surveillance by strangers, landlords, or even family members without their knowledge or consent. The following review implies the user worries their landlord may have access to their video and knows when no one is home:

“I surely hope to never find our landlord sneaking in ... he is nosy and goes in tenants apartments.”

Park et al. (Park et al. 2024) have interviewed short term rental hosts about their IoT device usage, and identified that hosts/landlords often have competing needs and goals (e.g., protecting property) from protecting guests’ privacy. Beyond landlords, we also saw reviews stating concerns about neighbors, friends, and family members accessing footage.

In addition, many users worry that recordings are shared with governments and or third-party corporations.

“I own a bunch of these cameras but will replace them and never buy a [COMPANY] products as long as they give MY FOOTAGE to authorities without consent or warrant. Do not enable this behavior by supporting [COMPANY].”

In the Privacy Controls theme, users complain about not having the right level of control for data sharing as described in the review below:

“I purchased 4 minicams but I’m not happy. The cameras lack basic features like not being able to share access ... without supplying them your user credentials.”

In Usability, the third top theme, users complain about the lagging performance of cameras, causing missed events and delayed notifications, as well as difficulties with the user interfaces in accessing live feeds and challenges managing multiple cameras. One user writes:

“When you get alerts on your phone about motion detection, there’s no way to disable individual cameras - you can disable all or enable all.”

In Figure 3, we see a stark growth of reviews discussing these three topics. The growth of surveillance is noteworthy, as this issue has been known for a while, and thus, we hypothesize that the problem is not really getting addressed.

Many of the remaining dominant issues for speakers and cameras are related to multi-user households and shared account management difficulties. We thus discuss multi-user environments separately in a later section.

Cross Device Insights

Customer Loss. Customer retention is a critical challenge for developers. Having found some reviews that discuss replacing devices or stopping to use them, we investigated more broadly if security and privacy issues were causing users to abandon apps or devices. We utilized GPT-5 (see the prompt in Appendix F) to label the reviews. This analysis identified 321 out of 4896 concern reviews, representing 6.6% of the total across all device types. We then manually verified these reviews to confirm that users had indeed stopped using the device or service. Our straightforward prompt and manual review likely led to an underestimation of this behavior. For example, one user says:

“Because it wants complete access to everything on my phone, such as phone numbers, messages, photos, contacts and “other files”, I returned this device.”

From these reviews, we extracted 15 recurring themes, indicating that a variety of issues contributed to customer loss. The most prominent reasons included Surveillance, Data Security and Theft and Privacy Controls. However, we also found that reviews highlighting Trust and Transparency, Policy and Law, and Privacy Ethics showcase when user dissatisfaction can lead to abandonment.

Seeking More Privacy Control. We found users writing about wanting more fine-grained controls across all device types and many themes, including Privacy Controls, Authentication, Data Sharing, Access Control, Surveillance. For example, consider the Privacy Controls issue. With fitness trackers, users want more control over notifications that reveal sensitive information, while speaker users are concerned about the inability to control the “always-listening” feature, and camera users focus on limited control over motion detection zones and who has access to the stored video.

Ethics and the Law. Because our taxonomy enables the separation of reviews into many categories, we are able to look at issues beyond the top few to gain a deeper understanding of users’ perspectives. For example, in the theme of Privacy Ethics, we see comments indicating a loss of trust. One speaker user writes

“Your internet will be automatically shared through [BRAND-NAME] devices with people near ur home. This policy is not customer centric... Really degrades trust. [COMPANY] should have notified us.”

Similarly, one camera user states

“Let’s be clear, this is not a technology limitation, it’s a purposeful limitation to force you into a [COMPANY] subscription. This kind of crap is why we can’t have nice things.”

The theme Policy and Law is a top 10 issue for trackers and cameras, and top 13 for speakers, and addresses a more specific aspect of data collection. For both speakers and cameras, we see reviews in which users discuss their data being given to government or law enforcement agencies. This illustrates an awareness of this data-sharing practice. Interesting, we see a divergent set of views about this. One upset reviewer says about their speaker product:

“A clause in the fine print allows law-enforcement agencies to simply fill out a form with Amazon to invade your privacy ... did not get any notification ... do I see a way to opt-out to protect my right to privacy ... buyer beware!”

Another review for a camera product, however, espouses an opposing view, showing they see validity for preventing physical harm.

“If we reasonably believe that we can prevent someone from dying or from suffering serious physical harm, we may provide information to a government agency — for example, in the case of bomb threats, school shootings, kidnappings, suicide prevention, and missing persons cases... Why not?”

Sharing these reviews with developers can aid them in improving product communications to clarify data sharing with governments or law enforcement.

5 Multi-User/Device Interactions Use Case

The previous section shared that many users write about challenges occurring due to multi-user and multi-device interactions. These issues can surface under the Authorization (one user can access another's account), Privacy controls (one user changes another's settings), or Usability (cannot figure out how to manage multiple users) themes. In a collaborative space, proper multi-user interaction is of paramount importance to ensure that the privacy of each user is not violated. In this section, we highlight four different types of multi-user and multi-device interaction issues identified in our review analysis. We focus on reviews for speakers and cameras, as very few such reviews exist for fitness trackers.

One problematic scenario is when a single device is shared by many people, and the device supports only one account. Many home security camera systems lack granular access controls, preventing users from assigning different permission levels to multiple users (such as controlling which stored videos they can see). Device owners are often forced to share full administrative credentials with individuals who only need limited viewing. For example, *"The biggest issue is that you can only have one login per setup ... if you want someone else to have access to the camera, you have to share your login details. That person then also has full admin rights ..., and there's no way to see what anyone might do..."*

This situation also limits notification capability since the camera cannot tailor alerts for specific users.

A second problematic interaction occurs when different user accounts and their personal devices are unintentionally linked due to using the same app/service within the same household. In one incident, one user had full visibility of another user's request to the voice assistant.

"BEWARE!!! Your alexa app may be linked to other people's apps even if they are using their own account ... bought a second echo as a gift for my mother. I could see everything she requested on her echo and she could see what I requested on mine. She was logged into alexa through her own account and the same was true for me."

In another example, multiple speakers are present, and the notifications from one device get announced on another device, creating unwanted privacy disturbances. This creates potential for unauthorized access and data sharing.

"Thought my unit would be some what independent of my roommates unit...registered units several times with little success...Could not drop in without problems of a contact being added to both units or their friend coming in on me."

In a third scenario, one user has multiple devices in multiple locations and wishes to manage them via a single account. For some devices, this requires multiple log-ins and separate accounts, creating inconvenience and security risks. *"And when I wanted to install cameras on my property and my rental I was told I would have to start a whole new [COMPANY] account AND get another subscription. Then it requires me to log in and out all day to keep a consistent eye on all of the properties."*

A fourth scenario is when camera usage generates concerns about bystander privacy, particularly in shared spaces

where they can capture multiple users without their consent. We uncovered such practices among some consumers who use cameras to secretly monitor specific family members or tenants without their knowledge.

"I purchased these to spy on my stepson and his lazy girlfriend, they gonna stop getting into my stuff now because I can see everything they do even when we're gone."

This underscores the need for more transparent and ethically designed camera systems to protect the privacy of all individuals in shared environments.

The challenges due to multi-user and multi-device interactions have been partly explored before. Lee et al. have focused on designing privacy-preserving multi-user access control systems for IoT environments (Lee et al. 2024), but did not survey users to identify their concerns or needs. Of the four scenarios we shared earlier, scenarios 2 and 3 are reported for the first time, while scenarios 1 (Park et al. 2023; Lee et al. 2024) and 4 (Ahmad et al. 2020; Park et al. 2023) have been reported previously. The study in (Park et al. 2023) is based on interviews with 32 smart home power users. In contrast, our findings draw on thousands of reviews across 52 speaker and camera devices collected over multiple years, enabling us to both corroborate prior results and demonstrate them at a much larger scale.

6 Discussion

Generalizability on Unseen Devices

To demonstrate the generalizability of the CRC and TM models beyond the training data, we collected Reddit posts related to device types not included in training (i.e., virtual reality devices). We collected 11,233 posts from three VR-related subreddits: r/VirtualReality (4,124 posts), r/Oculus (3,944 posts), and r/VisionPro (3,165 posts). Applying the CRC model, we identified subsets of posts discussing S&P issues (200 from r/VirtualReality, 252 from r/Oculus, and 114 from r/VisionPro). To evaluate accuracy, we randomly sampled 50 posts (17 from r/VirtualReality, 17 from r/Oculus, and 16 from r/VisionPro), which two coders manually labeled to create ground-truth labels. Inter-rater reliability was high, Cohen's κ was 0.96, which is almost perfect, according to (Landis and Koch 1977). Using these finalized labels, our model achieved an accuracy of 98%, indicating strong performance in unseen contexts. To further assess the generalizability of the TM model, we evaluated it on 100 randomly selected VR-related posts flagged for S&P issues. Again, two coders independently annotated the posts, achieving high initial agreement; Cohen's κ was 0.71, which is substantial, according to (Landis and Koch 1977); any remaining conflicting labels were resolved and finalized through collective discussion. The finalized ground-truth labels were then compared against the model's outputs. The TM model achieved an accuracy of 88%, providing strong evidence that it generalizes effectively across contexts.

Recommendation

Q&A Agent for E-commerce Consumers. This work offers valuable support for consumer decision-making. Since S&P reviews make up less than 10% of total reviews for

many IoT products, manually browsing or keyword searches often fail to surface relevant concerns. While early stage chatbots like Amazon’s Rufus can highlight reviews on specific features, they currently lack the ability to accurately capture S&P issues. Our approach fills this gap by automatically identifying S&P-relevant reviews and generating both fine-grained and high-level thematic tags. These insights help surface meaningful comparisons across devices, enabling more informed purchases. A promising direction for future is to explore how to integrate S&P summaries into GenAI assistants (e.g., Rufus) in a user-centric way that educates consumers while preserving review sentiment.

Guidance for Product Vendors. Stakeholders such as product vendors and developers can leverage our methodology to surface and summarize S&P concerns from user reviews within development tools. Amazon sellers, for instance, already access customer feedback through the “Feedback Manager” portal. Our approach can enhance this by extracting and quantifying S&P issues, helping prioritize concerns during feature planning.

Two key categories of developer takeaways emerged. First, users consistently seek more granular control—both over recipients (e.g., household members, landlords, third parties, law enforcement) and over specific content (e.g., selectively sharing audio/video footage or organizing it by day/device). Lack of such control often leads to user frustration and device abandonment.

Second, there is a clear need for better support in multi-user and multi-device scenarios. We identify four key challenges: (1) single devices that are shared but are lacking in multi-account support, (2) shared devices enabling unintended visibility on non-shared devices, (3) demand for unified control of devices across multiple homes, and (4) camera misuse for surveillance. Addressing these scenarios is crucial to improving user trust and retention.

Furthermore, across all three device categories, we examined the manufacturers of reviewed devices and consistently found that each device type is dominated by a small set of manufacturers (2–4). As shown in (Peddinti et al. 2019), highlighting to developers that competitors avoid practices disliked by users can serve as an effective nudge for change.

Ethical Considerations of Using User Reviews. Amazon currently generates AI-based summaries of user reviews without requiring explicit user opt-in or consent—submitting a review is considered implicit consent. Our envisioned use of assistive agents for both consumers and sellers aligns with Amazon’s existing terms of service. However, we strongly advocate for greater transparency: a clear and explicit statement on the web portal regarding the use of user-generated data by AI is essential.

Limitations/Biases. Our work has a few limitations. First, we focus our analysis on reviews from only three device categories. However, since Amazon is a leading IoT retailer, we analyzed over 70 device types, and because our models and taxonomy are device-agnostic, we expect our approach to generalize to other products. Second, our analysis is limited to reviews in English. However, given that state-of-the-art LLMs are inherently multilingual, our approach can be

readily extended to non-English reviews. Third, we focused on expressed user concerns without assessing the technical accuracy of users’ perceptions. Fourth, our analysis may be affected by fake reviews on Amazon. Since we do not collect user-specific metadata (e.g., usernames), identifying fake reviews from text alone is challenging, so we rely on Amazon’s detection mechanisms² for removing these fake reviews and banning abusive accounts, but acknowledge that some fake reviews may have trickled into our dataset. Lastly, sociodemographic characteristics of users and their biases may influence who submits reviews on Amazon, and these biases also impact our results.

7 Conclusion

We propose a modern approach to identifying security and privacy concerns related to IoT devices by prompt engineering and fine-tuning GPT-3.5-Turbo. Compared to prior work, this approach detects more issues and groups them into clear themes. Analyzing over 70 products across years, we find persistent issues, highlighting limited industry progress. The study particularly highlights two major challenges: users’ demand for more precise control over their data, and the complexity of managing multi-user and multi-device setups, both of which must be carefully addressed in future design and development efforts.

Acknowledgments

We thank the anonymous reviewers for their feedback. The views expressed are those of the authors and do not necessarily reflect those of the participating organizations.

References

- Ahmad, I.; Farzan, R.; Kapadia, A.; and Lee, A. J. 2020. Tangible privacy: Towards user-centric sensor designs for bystander privacy. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2): 1–28.
- Akbar, M. A.; Alsanad, A.; Mahmood, S.; and Alothaim, A. 2021. A multicriteria decision making taxonomy of IoT security challenging factors. *IEEE Access*, 9: 128841–128861.
- Akgul, O.; Peddinti, S. T.; Taft, N.; Mazurek, M. L.; Harkous, H.; Srivastava, A.; and Seguin, B. 2024. A Decade of Privacy-Relevant Android App Reviews: Large Scale Trends. *arXiv preprint arXiv:2403.02292*.
- Alsamani, B.; and Lahza, H. 2018. A taxonomy of IoT: Security and privacy threats. In *2018 IEEE International Conference on Information and Computer Technologies (ICICT)*, 72–77.
- Besmer, A. R.; Watson, J.; and Banks, M. S. 2020. Investigating user perceptions of mobile app privacy: An analysis of user-submitted app reviews. *International Journal of Information Security and Privacy (IJISP)*, 14(4): 74–91.
- Cen, L.; Si, L.; Li, N.; and Jin, H. 2014. User Comment Analysis for Android apps and CSPI Detection with Comment Expansion. In *PIR@ SIGIR*, 25–30.
- Chen, C. C.; Shu, D.; Ravishankar, H.; Li, X.; Agarwal, Y.; and Cranor, L. F. 2024. Is a Trustmark and QR code enough? The effect of IoT security and privacy label information complexity on

²<https://www.amazon.com/gp/help/customer/display.html?nodeId=G8CXDF9GLRRSV3G>

- consumer comprehension and behavior. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–32.
- Chen, K.; Zhang, S.; Li, Z.; Zhang, Y.; Deng, Q.; Ray, S.; and Jin, Y. 2018. Internet-of-things security and vulnerabilities: Taxonomy, challenges, and practice. *Journal of Hardware and Systems Security*, 2: 97–110.
- Das, A. K.; Zeadally, S.; and He, D. 2018. Taxonomy and analysis of security protocols for Internet of Things. *Future Generation Computer Systems*, 89: 110–125.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics (ACL)*, 4171–4186.
- Dubois, D. J.; Kolcun, R.; Mandalari, A. M.; Paracha, M. T.; Choffnes, D.; and Haddadi, H. 2020. When speakers are all ears: Characterizing misactivations of IoT smart speakers. *Proceedings on Privacy Enhancing Technologies*.
- Emami-Naeini, P.; Agarwal, Y.; Cranor, L. F.; and Hibshi, H. 2020. Ask the experts: What should be on an IoT privacy and security label? In *2020 IEEE Symposium on Security and Privacy (SP)*, 447–464.
- Emami-Naeini, P.; Dheenadhayalan, J.; Agarwal, Y.; and Cranor, L. F. 2021. Which privacy and security attributes most impact consumers’ risk perception and willingness to purchase IoT devices? In *2021 IEEE Symposium on Security and Privacy (SP)*, 519–536.
- Emami-Naeini, P.; Dixon, H.; Agarwal, Y.; and Cranor, L. F. 2019. Exploring how privacy and security factor into IoT device purchase behavior. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–12.
- FORCE11. 2020. The FAIR Data principles. <https://force11.org/info/the-fair-data-principles/>.
- Fruchter, N.; and Liccardi, I. 2018. Consumer attitudes towards privacy and security in home assistants. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*.
- Gabriele, S.; and Chiasson, S. 2020. Understanding fitness tracker users’ security and privacy knowledge, attitudes and behaviours. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–12.
- Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- Gulzar, M.; and Abbas, G. 2019. Internet of things security: a survey and taxonomy. In *2019 IEEE International Conference on Engineering and Emerging Technologies (ICEET)*, 1–6.
- Harkous, H.; Peddinti, S. T.; Khandelwal, R.; Srivastava, A.; and Taft, N. 2022. Hark: A Deep Learning System for Navigating Privacy Feedback at Scale. In *2022 IEEE Symposium on Security and Privacy (SP)*.
- He, S.; Hollenbeck, B.; and Proserpio, D. 2022. The market for fake reviews. *Marketing Science*, 41(5): 896–921.
- Interactive, H. 2019. Consumer Internet of Things Security Labelling Survey Research Findings. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/950429/Harris_Interactive_Consumer_IoT_Security_Labelling_Survey_Report_V2.pdf.
- Kuckartz, U.; Rädiker, S.; Ebert, T.; and Schehl, J. 2013. *Statistik: eine verständliche Einführung*. Springer-Verlag.
- Landis, J. R.; and Koch, G. G. 1977. The measurement of observer agreement for categorical data. *biometrics*, 159–174.
- Lee, S.; Jeong, M.; Song, D.; Choi, J.; Son, S.; Song, J. Y.; and Shin, I. 2024. FLUID-IoT: Flexible and Fine-Grained Access Control in Shared IoT Environments via Multi-user UI Distribution. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–16.
- Linstone, H. A.; and Turoff, M. 1975. *The delphi method*, volume 1975. Addison-Wesley Reading, MA.
- Lutz, C.; and Newlands, G. 2021. Privacy and smart speakers: A multi-dimensional approach. *The Information Society*, 37(3): 147–162.
- Mukherjee, D.; Ahmadi, A.; Pour, M. V.; and Reardon, J. 2020. An empirical study on user reviews targeting mobile apps’ security & privacy. *arXiv preprint arXiv:2010.06371*.
- Nguyen, D. C.; Derr, E.; Backes, M.; and Bugiel, S. 2019. Short text, large effect: Measuring the impact of user reviews on android app security & privacy. In *2019 IEEE symposium on Security and Privacy (SP)*, 555–569.
- Park, S.; He, W.; Deldari, E.; Emami-Naeini, P.; Huang, D. Y.; Vitak, J.; Yao, Y.; and Zimmer, M. 2024. Well-intended but half-hearted: {Hosts’} consideration of {guests’} privacy using smart devices on rental properties. In *Twentieth Symposium on Usable Privacy and Security (SOUPS 2024)*, 179–198.
- Park, S.; Lenhart, A.; Zimmer, M.; and Vitak, J. 2023. “Nobody’s Happy”: Design Insights from Privacy-Conscious Smart Home Power Users on Enhancing Data Transparency, Visibility, and Control. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS)*.
- Peddinti, S. T.; Bilogrevic, I.; Taft, N.; Pelikan, M.; Erlingsson, U.; Anthonysamy, P.; and Hogben, G. 2019. Reducing Permission Requests in Mobile Apps. In *Proceedings of the Internet Measurement Conference, IMC ’19*, 259–266.
- Protick, T. I.; and Das, A. 2025. Data Artifacts. <https://github.com/protick/IoT-Products-SecurityAndPrivacy>.
- Protick, T. I.; Sabir, A.; Abhinaya, S.; Bartlett, A.; and Das, A. 2024. Unveiling Users’ Security and Privacy Concerns Regarding Smart Home IoT Products from Online Reviews. *ACM Journal on Computing and Sustainable Societies*.
- Rizvi, S.; Kurtz, A.; Pfeffer, J.; and Rizvi, M. 2018. Securing the internet of things (IoT): A security taxonomy for IoT. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications (TrustCom)*, 163–168.
- Shekhan, S.; and Hartutyunyan, A. 2013. Watching the watchers: hacking wireless IP security cameras. *Shape Security and Qualys Inc*.
- van der Linden Dirk; Edwards, M.; Hadar, I.; and Zamansky, A. 2020. Pets without PETs: on pet owners’ under-estimation of privacy concerns in pet wearables. *Proc. Priv. Enhancing Technol.*, 2020(1): 143–164.
- Vetrivel, S.; van Harten, V.; Gañán, C. H.; van Eeten, M.; and Parkin, S. 2023. Examining Consumer Reviews to Understand Security and Privacy Issues in the Market of Smart Home Devices. In *32nd USENIX Security Symposium (USENIX Security 23)*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Williams, P.; Rojas, P.; and Bayoumi, M. 2019. Security taxonomy in IoT—a survey. In *2019 IEEE 62nd International Midwest Symposium on Circuits and Systems (MWSCAS)*, 560–565.
- Zhang, S.; and Sadeh, N. 2023. Do privacy labels answer users’ privacy questions. In *Network and Distributed System Security Symposium*.

Paper Checklist

1. For most authors...
 - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes, to stop violating privacy norms and unfair profiling, the product reviews have been presented anonymously. Any mention of socio-economic divide or implication of disrespect to societies or cultures is out of scope.**
 - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes, claims made in the abstract and introduction are a reflection of the methods and results.**
 - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes, the appropriateness of the proposed approach is mentioned in our Methods section.**
 - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes, we clarify it in the Methods section.**
 - (e) Did you describe the limitations of your work? **Yes, in the Discussion section.**
 - (f) Did you discuss any potential negative societal impacts of your work? **No, there are no specific negative societal impacts due to our research. In fact, we help increase user awareness towards security and privacy concerns with IoT devices to alleviate any negative impact caused by these devices. Also, to protect specific devices against negative societal impact, we anonymized any named entities in product reviews.**
 - (g) Did you discuss any potential misuse of your work? **No, because our dataset or the models are not made public. We will anonymize the dataset before publishing it to avoid any misuse.**
 - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes, as outlined in the ethics subsection when describing the methodology, we refrain from collecting any personally identifiable information of users who submit reviews. Also, we comply with Amazon's terms when crawling to collect the reviews. We acknowledge that we will anonymize the review dataset before publishing it.**
 - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes, we acknowledge that we read the guidelines and the paper conforms to them.**
2. Additionally, if your study involves hypotheses testing...
 - (a) Did you clearly state the assumptions underlying all theoretical results? **Yes, in the results section.**
 - (b) Have you provided justifications for all theoretical results? **Yes, we shared interpretations for the statistical analysis in the paper.**
3. Additionally, if you are including theoretical proofs...
 - (a) Did you state the full set of assumptions of all theoretical results? **NA**
 - (b) Did you include complete proofs of all theoretical results? **NA**
4. Additionally, if you ran machine learning experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes, for all the machine learning experiments, we list the model choice, prompt text, algorithms for choosing helpful examples, and the chosen parameters for reproducibility. However, we haven't published our crawled dataset used for model finetuning. We plan to anonymize our dataset before publishing it.**
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes, in the methods section.**
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **No, because ours is a large dataset of 91K reviews, and each API call incurs cost. Although we did multiple experiments (five times on the validation dataset) to see what temperature parameter would give us the best results.**
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **Yes, in the model choice part of the methods section.**
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes, in the methods and results section.**
5. Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **Yes, where possible, we complemented our results with findings from the literature.**
6. Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **No, because our statistical tests were simple correlation tests chosen based on the variables, and there are no alternative explanations for the observed outcomes.**
7. Did you address potential biases or limitations in your theoretical framework? **No, because the samples are statistically large, and for all tests the entire population was considered. In the LLM prompts, we do take measures to reduce bias due to our dynamic few-shot prompting.**
8. Have you related your theoretical results to the existing literature in social science? **Yes, we did this where possible. For instance, the point-biserial test which implies a negative correlation between concerns and star ratings agrees with the findings of Vetrivel et al.**
9. Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **Yes, in the discussion section.**

- (f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? [Yes, we perform false positive and negative evaluation when fine tuning our models to ensure our models has fewer misclassifications. There is ‘cost’ for misclassifications, except for missing out on any related S&P reviews.](#)
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity...**
- (a) If your work uses existing assets, did you cite the creators? [Yes, we cited the assets.](#)
- (b) Did you mention the license of the assets? [Yes, the assets used in this study are publicly available.](#)
- (c) Did you include any new assets in the supplemental material or as a URL? [Yes, we include the algorithms used, the prompts with examples, and the taxonomy with their definitions.](#)
- (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [Yes, according to their paper, the data has been made public and is found in an OSF repository.](#)
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [NA](#)
- (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? [NA](#)
- (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? [NA](#)
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity...**
- (a) Did you include the full text of instructions given to participants and screenshots? [NA](#)
- (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? [NA](#)
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [NA](#)
- (d) Did you discuss how data is stored, shared, and de-identified? [NA](#)

Appendix

A Algorithms to Generate Helpful Examples for The CRC and The TM Models

Algorithm 1: Algorithm for generating two helpful examples to be shown at the prompt for the CRC model

Input: t and T , where

t = An unlabeled text passed to the CRC model

T = The training dataset

Output: $t_1, t_2, Label(t_1)$ and $Label(t_2)$

```
1:  $T \leftarrow t \cup T$ 
2:  $L \leftarrow Label(t) \cup Label(T)$ 
3:  $B \leftarrow BertEmbed(T)$ 
4:  $C \leftarrow CosineSimilarityMatrix(B)$ 
5:  $c \leftarrow C[t]$ 
6:  $maxSim = -\infty$ 
7:  $maxSimIdx = -1$ 
8: for  $i = 1$  to  $|c|$  do
9:   if  $c[i] \neq 1$  and  $c[i] > maxSim$  then
10:     $maxSim \leftarrow c[i]$ 
11:     $maxSimIdx \leftarrow i$ 
12:   end if
13: end for
14:  $secondMaxSim = -\infty$ 
15:  $secondMaxSimIdx = -1$ 
16: for  $i = 1$  to  $|c|$  do
17:   if  $c[i] \neq 1$  and  $c[i] < maxSim$  then
18:     if  $c[i] > secondMaxSim$  and  $Label(T[i]) \neq$ 
19:        $Label(T[maxSimIdx])$  then
20:        $secondMaxSim \leftarrow c[i]$ 
21:        $secondMaxSimIdx \leftarrow i$ 
22:     end if
23:   end if
24: end for
25:  $t_1 \leftarrow T[maxSimIdx]$ 
26:  $Label(t_1) \leftarrow L[maxSimIdx]$ 
27:  $t_2 \leftarrow T[secondMaxSimIdx]$ 
28:  $Label(t_2) \leftarrow L[secondMaxSimIdx]$ 
29: return  $t_1, t_2, Label(t_1), Label(t_2)$ .
```

Algorithm 2: Algorithm for generating five helpful examples to be shown at the prompt for the TM model

Input: t and T , where

t = An unlabeled text passed to the TM model

T = The training dataset

Output: $t_1 - t_5, Label(t_1) - Label(t_5)$

```
1:  $T \leftarrow t \cup T$ 
2:  $L \leftarrow Label(t) \cup Label(T)$ 
3:  $E \leftarrow TFIDF(T)$ 
4:  $C \leftarrow CosineSimilarityMatrix(E)$ 
5:  $c \leftarrow C[t]$ 
6:  $Sort(c, decending)$ 
7:  $ArgSort(T)w.r.t c$ 
8:  $ArgSort(L)w.r.t c$ 
9:  $t_1 \leftarrow null$ 
10:  $t_2 \leftarrow null$ 
11:  $t_3 \leftarrow null$ 
12:  $t_4 \leftarrow null$ 
13:  $t_5 \leftarrow null$ 
14: for  $i = 2$  to  $6$  do
15:    $t_{i-1} \leftarrow T[i-1]$ 
16: end for
17:  $Label(t_1) \leftarrow Label[1]$ 
18:  $Label(t_2) \leftarrow Label[2]$ 
19:  $Label(t_3) \leftarrow Label[3]$ 
20:  $Label(t_4) \leftarrow Label[4]$ 
21:  $Label(t_5) \leftarrow Label[5]$ 
22: return  $t_1, t_2, t_3, t_4, t_5,$ 
23:    $Label(t_1), Label(t_2), Label(t_3), Label(t_4), Label(t_5)$ .
```

B The CRC Model

The CRC prompt with an actual example from the training dataset

- You are a Large Language Model for classifying, rationalizing, and categorizing security and privacy concerns in text.
- Given the text T: *I bought this with no intention of utilizing the subscription and doing only local storage. When I first bought the product and set it all up, everything worked great and it would allow me to get alerts when there was motion and would shoot a little 10 second video so I could review it later. About 3 months or so later, that got moved behind the subscription paywall as well. Now I can only use them as live feed as any other service is blocked without paying them extra. Only buy Blink if you're planning to buy the subscription as well*
- I want you to perform three tasks on T:
 - Task 1: Determine if text T explicitly mentions any security or privacy concerns. Respond only with a Yes or a No.
 - Task 2: Provide a brief explanation as to why your response is as such for Task 1.
 - Task 3: If the response for Task 1 was a Yes, categorize the explicit concerns mentioned in text T with themes related to security and privacy.
- Here are two examples that might be helpful:
 - Example text A: *cannot use their app now, without giving them your phone number. useless cameras sitting on my pile of garage sale items. they work okay, if you can get the app to work. app and cameras worked fine for a few months, now i cannot use them as the app has an apparently new requirement that i give them a phone number so they can send me a text - on a landline. email was fine for this, now the app won't allow getting past the page that requires this. customer service says to tap on an icon that does not appear on this page - just "give us your phone number". done with blink, will find another camera company that is more oriented toward the customer*
 - Ideal response for A:
 - Task 1: Yes*
 - Task 2: The text raises security and privacy concerns because the app is now requiring the user to provide their phone number, which may be unnecessary for the functionality of the app. This raises questions about why the company needs access to the user's phone number and how they will use it. Additionally, the mention of the app not allowing the user to proceed without providing the phone number suggests a lack of transparency and control over personal information.*
 - Task 3: requirement of a phone number to get access to the app, lack of transparency in why a phone number is required to use the app*
 - Example text B: *my biggest complaint with these cameras having to tap the screenshot to update the picture. i also hate that you have to wait before loading into a view of another camera. it will prompt "live view in progress" never had that on my zmodo.*
 - Ideal response for B:
 - Task 1: No*
 - Task 2: The text provided does not explicitly mention any security or privacy concerns. The complaints expressed are primarily related to the functionality and user experience of the cameras, such as the need to tap the screenshot to update the picture and the delay in loading into a view of another camera. These complaints do not directly pertain to security or privacy.*
 - Task 3: N/A*
- The following goes in the assistant prompt during training:
 - Task 1: Yes*
 - Task 2: The text raises concerns about the subscription-based model of the product. The user mentions that certain features, such as motion alerts and video recording, were initially available without a subscription but later moved behind a paywall. This raises questions about the security and privacy of the user's data. It suggests that the company may be limiting access to essential security features unless the user pays for a subscription, which could potentially compromise the user's security and privacy.*
 - Task 3: Concerns about the subscription model, Limited access to security features, Potential data collection and privacy implications*

Metric to Evaluate Classification			Metric to Evaluate Rationale and Low-level Issue Generation							
Metric	Task 1		Metric	Objective	Task 2		Task 3 = "N/A"		Task 3 = Other	
	Valid	Test			Valid	Test	Valid	Test	Valid	Test
Precision	97.8%	88.7%	ROUGE-L	Overlapping Similarity	0.531	0.739	0.9793	0.9947	0.508	0.6732
Recall	97.8%	93.4%	METEOR		0.5953	0.8119	0.9805	0.9893	0.6169	0.7462
F1-score	97.8%	91.0%	BERTScore	Semantic Similarity	0.9395	0.9731	0.996	0.999	0.936	0.9416

Table 4: Performance of the CRC model for the three tasks. Task 3 = "N/A", when Task 1 = "No".

C The TM Model

The Prompt for the TM Model

<p>—— System Prompt ——</p> <p>You are a Large Language Model for mapping a low-level issue to one or more high-level themes.</p> <p>—— User Prompt ——</p> <p>I have x, a low-level issue related to security and privacy: x, <u>an individual Task 3 response</u></p> <p>I also have Y, a list of high-level themes related to security and privacy: Y, <u>taxonomies with their definitions</u></p> <p>Your task is to map x to one or more y's from Y.</p> <p>Respond in the form: $x \rightarrow y_1$ if x should be mapped to just one y.</p> <p>Respond in the form: $x \rightarrow y_1, y_2, \dots, y_n$ if x should be mapped to more than one y.</p> <p>Here are some examples that might be helpful.</p> <p>E.g., low-level theme A: t_1; Response for A: <u>Label(t_1)</u></p> <p>E.g., low-level theme B: t_2; Response for B: <u>Label(t_2)</u></p> <p>E.g., low-level theme C: t_3; Response for C: <u>Label(t_3)</u></p> <p>E.g., low-level theme D: t_4; Response for D: <u>Label(t_4)</u></p> <p>E.g., low-level theme E: t_5; Response for E: <u>Label(t_5)</u></p>

Figure 4: The prompt for the TM model

The TM prompt with an actual example

- You are a Large Language Model for mapping a low-level theme to one or more high-level themes.
- I have x , a low-level theme related to security and privacy: *password sharing as a violation of basic it security principles*
- I also have Y , a list of high-level themes related to security and privacy:

Y, the list of high-level themes with their definitions

- Your task is to map x to one or more y -s.
- This means map x to one or more appropriate y -s from Y .
 - Respond in the form: $x \rightarrow y_1$ if x should be mapped to just one y .
 - Respond in the form: $x \rightarrow y_1, y_2, \dots, y_n$ if x should be mapped to more than one y .
- Here are some examples that might be helpful.
 - Example low-level theme A: *password security*
 - Ideal response for A: *password security* – \rightarrow *authentication*
 - Example low-level theme B: *concerns about password security*
 - Ideal response for B: *concerns about password security* – \rightarrow *authentication*
 - Example low-level theme C: *privacy violation*
 - Ideal response for C: *privacy violation* – \rightarrow *surveillance*
 - Example low-level theme D: *concerns about privacy violation*
 - Ideal response for D: *concerns about privacy violation* – \rightarrow *surveillance*
 - Example low-level theme E: *potential violation of privacy zones*

- Ideal response for E: *potential violation of privacy zones* – > *privacy controls*
- The following goes in the assistant prompt during training:
password sharing as a violation of basic it security principles – > *authentication, data sharing*

The labeled data used to create the training dataset for this work

Text	Label from (Protick et al. 2024)	Converted label based on our need
GOT HACKED<-tc-> I had one in my sons room because he has no sense of danger and it got hacked and they kept moving the camera and talking through the camera to my son when he was asleep	Label from the review labeling step Coder 1= 1 Coder 2= 1 Label from the thematic coding step Coder 1= Vignette of being hacked or someone taking control over the camera Coder 2= Hacking concerns	Task 1: Yes Task 2: The text mentions that a camera in the son’s room was hacked, indicating a security breach. The hackers were able to control the camera, move it, and even communicate through it, which raises concerns about privacy invasion and potential harm to the son’s well-being. Task 3: hacking incident, breach of security

Table 6: Conversion of the label from existing work (Protick et al. 2024) for our model training

Performance of the TM model on the validation dataset

Taxonomy Term	Rows Annotated	Rows Predicted	Precision	Recall	F1-score
access control	23	23	1	1	1
anonymity	0	0	0	0	0
authentication	41	41	1	1	1
authorization	23	23	1	1	1
availability	3	2	1	0.67	0.8
confidentiality	3	3	1	1	1
consent	13	13	1	1	1
data accuracy	4	4	1	1	1
data collection	32	33	0.94	0.97	0.95
data deletion	4	4	1	1	1
data exposure	6	5	1	0.83	0.91
data harms	3	3	1	1	1
data hiding	2	0	0	0	0
data management and storage	14	14	1	1	1
data security and data theft	35	41	0.85	1	0.92
data sharing	17	18	0.94	1	0.97
general comments related to security and privacy	21	21	0.95	0.95	0.95
location tracking	7	7	1	1	1
personalized advertising	6	5	1	0.83	0.91
policies and law	23	23	0.96	0.96	0.96
privacy controls	56	58	0.95	0.98	0.96
privacy ethics	35	31	1	0.89	0.94
secure communication	5	5	1	1	1
security vulnerabilities	14	12	1	0.86	0.92
software and firmware updates	1	1	1	1	1
surveillance	92	89	1	0.97	0.98
trust and transparency	19	18	0.94	0.89	0.92
usability	17	16	0.94	0.88	0.91
Macro-averaged	-	-	0.91	0.88	0.89
Micro-averaged	-	-	0.97	0.96	0.96

Table 7: Performance of the TM model for each taxonomy terms across 360 validation examples. ‘Rows Annotated’ are rows of examples mapped to that term as ground truth ‘Rows Predicted’ are rows of examples for which the TM model mapped that example with that taxonomy term. The macro-averaged calculates the precision, recall, and F1-score for each taxonomy term independently and then averages them. The micro-averaged aggregates the contribution of all classes before calculating the metric.

D Pairwise Chi-square Test

Category	Concern ratio			Test for equality of concern ratio				Pairwise test		
	$ C_c $	$ R_c $	$ C_c/R_c $	Test	Test Statistic	df	p-value	c_i	c_j	p-value
Trackers	505	23894	0.0211	Chi-squared	$\chi^2 = 2498.9$	2	$<2.2e-16^{***}$	Trackers	Speakers	$2e-04^{***}$
Speakers	847	32069	0.0264					Trackers	Cameras	$<2e-16^{***}$
Cameras	3544	35186	0.1007					Speakers	Cameras	$<2e-16^{***}$

Table 8: At the left: Calculation of the concern ratio. In the middle: test for equality of proportions. $H_0 : |C_c|/|R_c|$ is the same in each c . H_A : At least one $|C_c|/|R_c|$ for a c is different from the others. At the right: Posthoc pairwise comparison of the proportions of all pairs of $|C_c|/|R_c|$'s.

E Point Biserial Test of Correlation between S&P concerns in a review and rating given by a user

Category	Concern	Total	Rating					Point-biserial test				Interpretation
			1	2	3	4	5	t	df	p-value	r_{pb}	
Trackers	No	23389	4684	2263	2676	3364	10402	-16.603	23892	$<2.2e-16^{***}$	-.11	Low corr. (Kuckartz et al. 2013)
	Yes	505	235	75	55	64	76					
Speakers	No	31222	1716	1176	2437	6421	19472	-10.801	32067	$<2.2e-16^{***}$	-.06	No corr. (Kuckartz et al. 2013)
	Yes	847	92	51	100	226	378					
Cameras	No	31642	4057	1912	2667	4408	18598	-35.504	35184	$<2.2e-16^{***}$	-.19	Low corr. (Kuckartz et al. 2013)
	Yes	3544	970	469	537	438	1130					
Overall	No	86253	10457	5351	7780	14193	48972	-40.168	91147	$<2.2e-16^{***}$	-.13	Low corr. (Kuckartz et al. 2013)
	Yes	4896	1297	595	692	728	1584					

Table 9: Point-biserial test between S&P concerns in a review and the rating given by a user. $H_0 : \rho_{pb} = 0$ and $H_A : \rho_{pb} \neq 0$. The significance code *** means $p < .001$.

F Prompt used in GPT-5 to detect cases of app/device ‘Stop Using,’ ‘Replacing,’ or ‘Uninstalling.’

- You are a text classification assistant. Your task is to carefully read a piece of text (an Amazon product review) and do a classification task.
- Given the text T: [USER REVIEW]
- Classify whether the reviewer explicitly indicates one of the following actions:
 1. The user uninstalled the product.
 2. The user replaced the product.
 3. The user stopped using the product.
 4. None of the above three actions.
- Always output exactly one of: ‘1. Uninstalled’, ‘2. Replaced’, ‘3. Stopped Using’, or ‘4. None of them’. Do not add explanations or extra text.
- If a user uninstalled, replaced, or stopped using the product due solely to configuration issues rather than both confirmation and privacy concerns, label it 4. None of them.