Toward Transparent IoT Purchases: Understanding User Preferences for Privacy and Security Properties of IoT Devices

Lindrit Kqiku
Institute of Computer Science
Campus Institute Data Science (CIDAS)
University of Göttingen
kqiku@cs.uni-goettingen.de

Delphine Reinhardt
Institute of Computer Science
Campus Institute Data Science (CIDAS)
University of Göttingen
reinhardt@cs.uni-goettingen.de

Abstract

The rapid proliferation of *Internet-of-Things* (IoT) devices raises privacy concerns, as these devices collect extensive information about both users and bystanders. Yet users often remain uninformed of how such devices handle their data. In response, various privacy labels have been proposed, but although many designs have undergone user evaluation, we still lack a clear understanding of which privacy-related features users consider important to be informed about, and how these priorities differ across user profiles. Addressing this gap is essential for developing privacy communication mechanisms that genuinely align with users' preferences. This paper reports a study (N = 565) examining how users prioritize different privacy features in IoT devices and why. Our findings reveal which properties users value most, how these priorities differ across gender, age, technical affinity, and privacy concern, and which features are interpreted as baseline expectations rather than choices (e.g., encryption and updates). We also identify context-dependent features, such as multi-user management, that are rejected by numerous single-household users. Moreover, we show that users' privacy concerns predict preferences more strongly than demographics. Notably, open-source property elicits polarized reactions, serving as a sign of trust for some and a perceived security risk for others. Drawing on these insights, we advocate for adaptive privacy labels that automate security baseline privacy expectations (e.g., features like operational security, secure defaults, and updates) while foregrounding user sovereignty (e.g., features like data storage location, cloud provider, jurisdiction) and contextual factors (e.g., living alone), moving beyond one-size-fits-all privacy labels toward dynamic decision-support solutions.

Keywords

Internet of Things, privacy, security, labels, transparency, user preferences

1 Introduction

The increasing adoption of IoT devices in private households [40], from smart speakers and televisions to connected thermostats and security cameras, brings multiple benefits, including convenience, energy efficiency, and enhanced security. However, the vast amounts of personal data collected and processed by these devices also pose substantial privacy risks to users. These risks range from unauthorized access to sensitive information to the potential for user profiling and surveillance.

The consequences of inadequate privacy protections are multifaceted, ranging from precise location data and audio recordings to health metrics and usage patterns, often with limited user awareness. According to current statistics [32], 72% of smart home product owners express concern about the security of personal data collected by their devices, while nearly half of all U.S. internet households report experiencing at least one privacy or security incident in the past year.

Recent incidents intensify these concerns. The popular press has reported incidents such as Amazon Alexa sending private conversations to a random phone number [25] and baby monitors being hacked [9]. Even well-known IoT manufacturers have faced criticism for a lack of transparency about their data practices. For instance, recent news articles reported that Google forgot to mention that its Nest Secure hub has a microphone [26] and Amazon revealed that, in addition to *Artificial Intelligence* (AI) algorithms, human employees listen to a subset of audio recordings from Echo devices [33].

Drawing on the nutrition-label idea for privacy policies [21], existing IoT transparency mechanisms [7, 8, 28, 34–36, 47] attempt to address these risks by providing structured disclosures, often prominently through privacy labels. These privacy labels provide a comprehensive overview of what data a device collects, how it is used, with whom it is shared, and under which conditions it is retained, giving consumers a quick, standardized overview of a product's privacy practices. However, these labels currently suffer from two critical shortcomings. Firstly, they typically present all information at the same level, giving dozens of properties equal visual weight. As a result, consumers must mentally sort through complex technical terms, such as retention periods, third-party sharing, and data categories. This "cognitive overload" effect has been repeatedly observed in empirical work on app and IoT labels, and it ultimately increases uncertainty instead of reducing it. Secondly, the proposed solutions are mainly designer-driven (i.e., based on regulations or expert assumptions) rather than user-driven (i.e., based on users' own priorities) [12]. Labels, therefore, include many items that are unimportant to specific user groups (e.g., multi-user features for single households) while failing to foreground the strongest purchase-time differentiators (e.g., cloud provider, data jurisdiction).

Proceedings on Privacy Enhancing Technologies 2026(3), 505–522

https://doi.org/10.56553/popets-2026-0093

As a result, current transparency solutions often communicate more data than meaning, leaving consumers to infer which properties matter. To complement and ultimately strengthen these solutions, this paper aims to bridge this gap by investigating which privacy and security properties users prioritize, how these priorities differ across demographic and attitudinal factors, and which features users consider non-negotiable baselines rather than optional settings. Only with such user-centered content guidance can we move from one-size-fits-all labels to effective, personalized decision support solutions.

In summary, our contribution can be summarized as follows:

- **Empirically Derived Taxonomy of Privacy Priorities.** We first compile a comprehensive list of 57 privacy- and security-relevant IoT properties grounded in prior work and systematically quantify users' perceived importance for each. We then conduct a user study with 565 German-speaking participants to assess users' perceptions of the importance of different privacy properties and functionalities of IoT devices. Among others, we reveal a clear distinction between non-negotiable security baselines (e.g., encrypted communication, update visibility, data minimization) and purchase-critical sovereignty privacy properties (data storage location, cloud provider, jurisdiction).
- **Analysis of Influencing Factors.** We examine how users' general privacy concerns, their level of technical affinity, age, gender, and IoT device ownership relate to their perceived importance of the identified privacy and security properties. Using *Affinity for Technology Interaction* (ATI) [17], *Internet Users' Information Privacy Concerns* (IUIPC) [29], gender, age, and IoT device ownership factors, we show that privacy concern is the strongest predictor of property importance, while contextual factors (e.g., living alone) shape which features are irrelevant or burdensome.
- **Justifications and Qualitative Analysis.** We collect participants' justifications, both from predefined answer options and open-ended responses, regarding why they consider specific properties important or unimportant. Moreover, through inductive content analysis, we uncover key reasoning patterns, including (i) mistrust in manufacturer-provided transparency, (ii) the "single household effect", and (iii) the "Open-Source duality" where transparency is seen as both a benefit and a security risk.

These findings enable the design of privacy labels and purchasing platforms that are not only usable but aligned with what users genuinely need to make informed IoT purchase decisions, moving from designer-driven to user-centered transparency. Thus, indirectly contributing to better protecting their privacy.

2 Related Work

Research work on transparency, labeling, visualization, and controls for consumer IoT devices has evolved along several complementary directions: (1) physical and digital privacy labels and disclosures, (2) tools and interfaces that visualize device data flows and make risks tangible, (3) in-device or ecosystem controls that let users trade functionality for privacy, and (4) developer tooling and empirical audits that examine real-world practices and compliance. Below, we organize prior work around these themes and highlight how our contribution extends and differs from them.

Privacy Labels and Consumer-Facing Disclosures. Building on early proposals for standardized privacy labels in other domains [21], a substantial body of research has investigated what information should appear on IoT or app privacy labels and how consumers use those labels. Several studies test label complexity and format, finding that medium-to-high complexity labels, i.e., explicit properties and layered details, improve users' ability to compare devices and identify safer choices, whereas overly minimal labels are confusing. For example, Chen et al. [7] compare low-, medium-, and high-complexity label designs and show that consumers value access control, data sharing, and collection properties. Work comparing mobile app labels similarly documents widespread terminology confusion and missed information by users [28], while analyses of whether labels answer real user questions reveal large coverage gaps [47]. Railean et al. [34, 35] propose *General Data Protection Regulation* (GDPR)-grounded physical and online labels for IoT and show users want such information but struggle with some purpose/recipient fields and complex visualizations. Prior work on privacy-term updates proposes formal notations and tabular change-views to reduce consent fatigue [36].

While those papers focus primarily on how to present information, i.e., format, complexity, visualizations, and on whether labels answer users' questions, our work complements them by investigating which specific privacy and security properties users consider important at purchase time by proposing a ranked, perception-focused set of privacy properties. Instead of testing label formats, we compile a comprehensive list of candidate properties and measure their perceived importance across users with varying privacy concerns and technical affinity. This provides content-level guidance, i.e., what to include on labels, that complements prior formats and usability-focused findings.

Other Visualizations. Several other works build visualizations or *Augmented Reality* (AR) systems to make invisible network activity and privacy leaks visible. AR-based tools (Privacy Plumber, PARA, PARA-like studies) and non-AR visualizations increase awareness and motivate protective behavior by showing device location, data type, or live traffic; users often enable more protections after seeing such visualizations [4, 5, 11]. Privacy Plumber and related AR systems also provide interactive blocking and filtering, demonstrating that closed-loop interventions increase perceived control [4, 11]. Work such as [30] articulates user-centered transparency requirements and identifies concrete *User Interface* (UI) elements, e.g., logs, maps, status indicators, that help users diagnose incidents.

While visualization and AR work emphasize post-purchase awareness and control, on how to reveal device behavior and enable immediate mitigation, our study focuses on pre-purchase decision-making by examining which privacy and security properties drive purchase preferences, how people prioritize them, and how those priorities relate to individual profiles. Our findings can thus inform which visualizations should be foregrounded for different audiences.

Device-Side Controls and Configuration Ecosystems. Research on control mechanisms proposes practical ways to limit or tune device

connectivity and features for privacy. Thalhammer et al. [41] introduce a multi-stage connectivity model from offline to online with a physical slider and dynamic privacy labels to help users manage privacy–utility trade-offs. Chhetri et al. [8] demonstrate a high-fidelity UI for opt-in/out, data retention, sharing, and deletion controls, finding users appreciate granularity but can be overwhelmed and prefer tiered presets. These systems highlight the trade-offs between feature availability and privacy risk, and surface the tension between fine-grained controls and user burden.

The aforementioned efforts design and evaluate control architectures and interaction mechanisms. We contribute by empirically mapping which properties and functionalities users deem important before purchase, e.g., access control, local processing, data retention, and third-party sharing. Understanding these priorities helps designers of control systems, such as [8, 41], decide which controls to expose best and where to use presets vs. granular toggles.

Developer Tools, Audits, and Compliance. Other lines of work address the accuracy of developer-reported disclosures and the real-world transparency of IoT vendors. Li et al. [27] propose an IDE plugin for developers that uses static analysis to improve the accuracy of App Store data safety labels. Hudig et al. [20] investigate the state of data transparency in the consumer IoT domain by analyzing network traffic (e.g., during device setup and regular use) to identify destination IP addresses and their geolocation data for 43 IoT products across 11 categories (e.g., cameras, wearables, smart plugs). The study examines data flows, assesses vendor compliance with GDPR data-access and data-portability rights, and evaluates the practicality of blocking unwanted data transmissions.

Both studies [20, 27] diagnose and improve supply-side transparency (i.e., developer labels, vendor compliance). However, while these approaches examine transparency as implemented by developers and vendors, they do not directly look at how consumers themselves interpret or prioritize such information. Our study adopts a demand-side perspective (i.e., what consumers consider when evaluating IoT devices), explicitly linking these concerns to individuals' privacy profiles, thereby complementing consumer-focused evidence with audits and developer tools by indicating which transparency failures matter most to users and are therefore high-value targets for remediation.

Label Evaluation, Effectiveness, and Behavioral Consequences. Several studies combine qualitative interviews and larger surveys to measure how privacy and security information affects user behavior. For example, Emani et al. [14] find concerns often arise post-purchase, that privacy matters more for some device types (e.g., cameras), and that consumers would pay a premium for clear privacy and security information. Zhang et al. [46] show that iOS App privacy labels can be difficult for users to find and interpret, highlighting that such labels require careful integration and presentation to be usable. For IoT specifically, Shen et al. [39] propose IoT security and privacy label layouts. Other work shows that visualizations and AR-based transparency tools (e.g., network visibility) reduce undue comfort and increase protective choices [4, 5, 11].

These prior studies examine the existence of privacy concerns and the influence of specific interventions on behavior. Our study differs in three specific ways: (1) we propose a systematic list of privacy and security properties (across many candidate privacy properties) rather than testing a few isolated label items, (2) we quantify perceived importance for each property across participants and device classes, and (3) we explicitly analyze how those importance ratings correlate with general privacy concern and technical affinity, helping to explain heterogeneity in consumer priorities.

Summary. Prior work has extensively explored how to present privacy information (e.g., labels, visualizations, AR transparency tools), how to control device behavior (e.g., connectivity modes, dashboards), and how to audit developer compliance. However, what remains missing is a systematic, empirically grounded understanding of which privacy and security properties consumers actually prioritize during purchase decisions, and how these priorities vary across individuals and their contexts. Without this content-level evidence, even well-designed labels risk presenting irrelevant, overwhelming, or misaligned information. Our work fills this gap by (1) building a comprehensive taxonomy of candidate IoT privacy properties, (2) quantifying user importance ratings across 57 privacy properties, (3) modeling how privacy concern, technical affinity, gender, age, and the ownership of different IoT devices shape these preferences, as well as (4) analyzing open-ended explanations to reveal underlying rationales. This enables the design of privacy labels and transparency tools that are not only usable but also aligned with what users actually value, supporting informed IoT purchases.

3 Methodology

To investigate which privacy properties are relevant to users, which can serve as a basis for privacy label design, and which online platforms allow comparison of different IoT devices, we conducted an online questionnaire-based study. We first detail our privacy properties taxonomy in Sec. 3.1, our survey design and procedure in Sec. 3.2, before highlighting our analysis in Sec. 3.3.

3.1 Derivation of Privacy Properties Taxonomy

We derived an initial input space of candidate IoT privacy and security properties from (a) prior work proposing or evaluating consumer-facing privacy/security labels and disclosures, (b) legal and regulatory frameworks (e.g., GDPR, CRA) that motivate disclosure obligations and baseline safeguards, and (c) transparency tools and empirical studies that identify privacy-relevant device behaviors (e.g., data collection, transmission, storage, sharing, and related safeguards). These candidates were then reviewed, sorted, and refined through interdisciplinary expert consensus involving two university professors, one postdoc, two PhD researchers in usable security/privacy, one industry practitioner, and one legal practitioner, across an initial workshop and followed by multiple online discussion iterations. We stopped once conceptual saturation was reached, i.e., no new candidate properties emerged, and remaining disagreements could be resolved by consensus. We applied inclusion criteria requiring that a property (a) is attributable to an IoT product or its associated service, (b) can be operationalized to support consumer-facing IoT device purchase time decisions, and (c) is sufficiently specific to be rated or verified in principle. We excluded candidates that primarily described user-interface or label design guidelines (e.g., wording, accessibility colors), platform/tool

engineering requirements (e.g., installation effort, cross-platform deployment), or that were insufficiently clear.

We chose this approach over a systematic literature review because our goal was not exhaustive coverage of all possible privacy-related concepts but to ensure they are legally sound and applicable in practice. Expert input helped integrate regulatory and implementation constraints that are often underrepresented in the literature, while cross-checking ensured alignment with prior research. Tab. 1 summarizes the resulting 57 properties and their sources. After the final set was fixed, an expert subgroup proposed the category definitions (Appendix B.1) and initial mappings, which the full group reviewed and discussed in a single iteration before agreeing on a final version.

3.2 Survey Design and Procedure

To empirically investigate user preferences regarding the privacy and security properties of IoT devices, we designed an online questionnaire administered via the LimeSurvey platform. We used a LimeSurvey instance hosted on our institutional server. The study was conducted in full compliance with the GDPR. All data were collected pseudonymously, and the survey's landing page required participants to provide informed consent, confirming their legal capacity and understanding of the study's purpose before proceeding. Although our university does not require *Institutional Review Board* (IRB) approval for this study type, all procedures followed established ethical research principles. Participants were informed about the study procedure and their rights, including the ability to withdraw at any time without consequences.

We recruited 565 German-speaking participants via Norstat (ISO 9001 / ISO 20252:2019), using a controlled, invitation-only panel. Recruitment targeted a sample broadly representative of the German population in terms of gender and age, and reflecting the natural prevalence of IoT usage. All participants were compensated at a rate of 12 EUR per hour, corresponding to the legal minimum wage in Germany. The survey took an average of 18.53 minutes.

We used a scenario-based approach to elicit participants' stated preferences. Participants were informed that the study aimed to identify the privacy and security properties most relevant to them. We highlighted to them that their responses help inform the development of a platform that allows consumers to compare privacy and security features and functionalities of IoT devices and make more informed purchasing decisions.

The scenario-based tasks were supplemented with measures of IoT usage and frequency, standardized scales assessing individual differences in technology affinity and privacy concerns, and demographic information. The complete study protocol is provided in Appendix A.

3.2.1 Phase 1: Assessment of IoT Usage Context. The study started by establishing a baseline of the participants' engagement with the IoT devices. Respondents were presented with a list of 22 distinct smart device categories, ranging from common household items like smart speakers and thermostats to specialized health devices, such as insulin pumps and sleep trackers. For each device the participant reported owning or using, a conditional follow-up question was triggered to assess usage frequency, ranging from daily to annual. Furthermore, the survey captured the operational context of these devices, distinguishing between private and professional use, as well as their physical location, such as the home or workplace.

3.2.2 Phase 2: Evaluation of Privacy and Security Features. The core part of the study introduced participants to a scenario in which they intended to purchase a new smart device and were asked to rate the importance of various privacy and security functions.

Considering the context-dependent nature of privacy, we deliberately framed the survey around purchase-time decision-making to reduce ambiguity and improve internal validity. Otherwise, participants might have also considered usage-time aspects. This ensures that responses reflect the intended context and align with prior work on privacy labels. At the same time, this framing may prime respondents toward properties perceived as immediately decision-relevant and may under-represent post-purchase concerns, which we reflect in our interpretation (see Limitations).

To cover a comprehensive design space of 57 distinct privacy properties, spanning categories such as technical implementation, transparency mechanisms, and user control, without inducing survey fatigue, we employed a randomization function. This ensured that while the final dataset covered the full spectrum of rated features, individual participants were presented with about half of the features each.

For each displayed feature, participants rated its importance on a 5-point Likert scale, ranging from "not at all important" to "very important", including a "neutral" option. The inclusion of the latter allowed respondents to express their uncertainty or indifference without being forced into an evaluative stance that might distort the measured importance of each privacy property. To gain deeper insight into the rationale behind these ratings, a conditional branching logic was applied. If a participant rated a feature as important or very important, they were prompted to explain their reason, such as (a) specific risk awareness, (b) a general desire for data protection, or (c) the feature's impact on user experience. Conversely, if a feature was rated as unimportant, respondents were asked to clarify whether they viewed the aspect as (a) unproblematic, (b) lacked sufficient knowledge to evaluate it, or (c) considered it irrelevant to their usage. In both the important and unimportant ratings, participants could also use the free-text field for other answers. We did not include follow-up questions for "neutral" ratings to reduce respondent burden, as these responses reflect indifference or uncertainty and were expected to provide limited additional insight.

3.2.3 Phase 3: Data Quality Control. To ensure the reliability and validity of the collected data, *Instructional Manipulation Checks* (IMCs) were included within the feature evaluation blocks. These attention checks explicitly instructed participants to select a specific response option, such as "please select very important to show that you are paying attention". This allowed for the identification and subsequent exclusion of careless responders who may have skimmed the questions without cognitive engagement, thereby maintaining the integrity of the dataset.

3.2.4 Phase 4: Technical Affinity and Privacy Concerns Scales. To account for individual differences in technical affinity and privacy concerns that might influence privacy preferences, two standardized scales were adopted following the feature evaluation step.

Table 1: IoT privacy properties derived taxonomy grouped by category.

Label	Property Name	Description	Source
Transparency Information (TI)			
TI-LGL	Legal Grounds	IoT products process personal data in accordance with a legal basis and make this transparent to the user.	[38]
TI-ENC	Encrypted Data Storage Info	Get informed that your own files are encrypted on the device, during transport, and in cloud components.	[1, 8, 27, 28, 35]
TI-SPC	Special Categories of Data	An IoT product processes personal data of the special category.	[15, 38]
TI-ACT	Actors Involved in Processing	Which actors (controller, joint controller, processor) are involved.	[38]
TI-RCP	Recipient of Information	Which companies receive data when you use the device.	[47]
TI-APP	App Permissions Required	Which access permissions must be granted for the use of an app belonging to an IoT product.	[47]
TI-CAT	Data Categories Collected	The system specifies which categories of data are collected when using the device.	[47]
TI-SHR	Data Sharing	Whether data collected from the use of an IoT product is shared with third parties.	[47]
TI-SEC	Optional Security Review	If an additional security review has taken place, it should be clear exactly which aspects were examined and who performed the tests.	[47]
TI-PRF	Profiling	The IoT-Product features functions that involve automated decision-making (profiling).	[38]
TI-NPI	Non-Personally Identifiable Info	The IoT-Product provides functionality to list *Non-Personally Identifiable Information* (NPII) collected from IoT sensors or other sources.	[36]
TI-SND	Informed About Sent Data	Informed about which types of data are transferred from the device to other recipients.	[38]
TI-EXR	Exact Recipients Information	Get informed about the exact recipients to whom the device sends data.	[28, 38]
TI-LKG	History of Data Leakage	Documented data leaks in connection with the IoT product.	[14]
TI-BRC	Privacy Breaches	There are currently known data breaches associated with the IoT product.	[1, 38]
TI-PSI	IoT-Product Specific Info	The IoT-Product informs you whether the full product name, firmware version, and country of manufacture are specified.	[7, 10]
TI-DUR	Data Storage Duration & Entity	The IoT-Product should clearly specify the duration for which the data is stored and identify the company or entity responsible for storing it.	[34, 35, 38]
TI-PUR	Purpose of Data Collection	The IoT-Product provides a detailed explanation of the reasons for collecting the data.	[7, 10, 34, 35, 38]
TI-FRQ	Frequency of Data Sent	The IoT-Product indicates how often the data is transmitted.	[34, 35]
TI-OEU	Data Processing Outside the EU	The IoT-Product indicates whether data is being transmitted outside the EU.	[35, 38]
TI-LOC	Location of Data Sent To	The IoT-Product indicates to which location (company & country) data is being transmitted.	[35, 38]
Implementation Detail (ID)			
ID-DEF	Secure Default Settings	IoT-Products are shipped with maximized privacy protection settings that allow processing only of personal data relevant to the device's purpose.	[1, 38]
ID-PWD	Encryption Status for Passwords	Whether access data is stored in encrypted form on the IoT products or the associated apps.	[47]
ID-AUT	Authentication Information	The type of user authentication of an IoT device, e.g., single-factor authentication, user-generated password, no password recovery.	[7]
ID-DSC	Controllable Discovery Function	Decide for yourself whether the device searches for other devices in your home network.	[30]
ID-PET	Privacy Enhancing Technologies	The IoT-Product features PETs to safeguard users' privacy.	[4, 20]
ID-BKD	Backend Service Demands	The IoT-Product requires a connection to a specific backend service.	[1, 30]
ID-UPT	Updateability	The IoT-Product features an update functionality and specifies the planned lifetime.	[1, 7, 34, 35]
ID-OSS	Open Source	Which part of the IoT product is based on open-source software.	[1]
ID-TLS	TLS Certificate Validation	The IoT-Product uses secure communication channels and validates server certificates.	[1]
ID-CAD	Controlled Automated Discovery	Specified whether the IoT product gives the user control over whether or not to search for other devices on the network automatically.	[30]
ID-ODP	On-device Data Processing	The data is preferably processed locally.	[30]
ID-BST	Account Required for Bootstrapping	An account is required to set up the IoT device.	[20]
Control Function (CF)			
CF-RST	Resetting to Secure Defaults	IoT-Products are able to be easily reset to data protection-preserving and secure default settings.	[1, 38]
CF-DAT	Control of Sent Data	IoT-Products allow users to control the type of data that is sent.	[38]
CF-RCP	Control of Data Recipient	IoT-Products allow users to control the recipient of sent data.	[38]
CF-LOC	User-controlled Processing Location	Decide for yourself which cloud provider the device uses (e.g., iCloud, OneDrive, or Google Drive).	[30]
CF-ACC	User Control Accounts	The IoT-Product offers to manage multiple authorized users.	[30]
CF-UIE	User Interaction and Engagement	The IoT-Product offers an appropriate user interface (e.g., through an app or a web application).	[30]
CF-PRT	Data Portability	The IoT-Product / related service features portability functions to migrate from/to other products.	[38]
Transparency Function (TF)			
TF-CTX	Contextual Data Disclosure	The IoT-Product explains to you clearly which of your data is to be processed before you give your consent.	[4]
TF-ADD	Additional Transparency Mechanisms	The IoT-Product regularly displays information about protecting your privacy.	[8]
TF-UPD	(Security) Update Status Display	The IoT-Product provides functionality for users to gain an overview of the current system's update status.	[30]
TF-EXP	Exploration Functionality	The IoT-Product provides functionality for users to browse specific system interactions, e.g., a list of communications within the system or visualizations of activities.	[30]
TF-INV	Investigation Functionality	The IoT-Product features a log of events, including errors, updates, etc. This feature would allow for in-depth tracking and thorough investigation of occurring issues.	[30]
TF-NTF	Notification When Attention Needed	The IoT-Product triggers notifications when users' attention is needed (e.g., in case of information being sent over the internet, errors, or suspicious occurrences).	[30]
TF-PPU	Privacy Policy Highlighted Updates	The IoT product highlights the changes between the old and new versions when changes are made to the privacy policy.	[36]
Usability of Transparency (UT)			
UT-SEP	Separate Data Protection Info from T&C	Data protection information should be separated from common Terms & Conditions.	[38]
UT-CLR	Clear & Transparent Processing Info	Information on the processing of personal data is provided in a concise, transparent, intelligible, and easily accessible form.	[14]
UT-VIS	Visual Representation of Privacy Flow	The IoT-Product features graphical representations of data flows or visualizations of privacy settings.	[11]
UT-SBJ	Subject Rights	The IoT-Product (service) supports users to make use of their subject rights.	[38]
Compliance/Support (CS)			
CS-CRA	Compliance with Cyber Resilience Act	IoT-Products provide information conforming to the Cyber Resilience Act.	[1]
CS-SUP	Compliance Support	IoT-Products support informing other persons (e.g., guests) about data protection.	[35]
CS-CNT	Contracts for Joint Controllers	Contract templates are available for joint controllers/processors.	[38]
Usability of Control (UC)			
UC-INT	Intuitive Privacy Control Interface	The IoT-Product features an intuitive and user-friendly interface that allows users to easily control their privacy settings.	[11]
UC-GRN	Granularity of Privacy Settings	The IoT-Product presents options that define a combination of privacy settings that users can easily select to achieve their desired level of w/o extensive manual configuration.	[8]
Data Handling (DH)			
DH-MIN	Data Minimisation	IoT-Products only process personal data that is necessary in relation to the purpose of the device.	[1, 38]

Firstly, to assess participants' technical affinity, we used the ATI scale [17] (shown in Appendix A.5.1). The ATI scale is a 9-item questionnaire designed to measure an individual's tendency to actively engage with and explore technology. This scale helps to characterize the user sample in terms of their comfort and willingness to interact with technical systems.

Secondly, to measure participants' general online privacy concerns, we utilized the IUIPC scale [29]. The IUIPC is a well-established instrument that captures three key dimensions of privacy concerns: collection, control, and awareness. Given recent research suggesting the advantage of using a shorter version [18], we opted for the 8-item IUIPC scale.

3.2.5 Phase 5: Demographics. We finally collected demographic data to characterize the sample.

3.3 Analysis

Importance Rating of Privacy Properties. We first assess quantitatively which properties are relevant by grouping the answers according to participants' perceived importance ratings for the privacy properties.

To assess whether observed differences in importance ratings between privacy properties were statistically significant, we conducted pairwise Mann-Whitney U tests for all 1,596 possible property pairs (57 × 56 / 2). The Mann-Whitney U test is appropriate for comparing independent samples with ordinal data, as participants were randomly assigned to rate approximately half of the properties.

We further examine how ATI, IUIPC, age groups, gender, and device ownership influence the perceived importance of 57 privacy-related IoT properties. ATI and IUIPC scores were grouped into low and high groups using median splits. Because only a small number of participants identified as genders other than male or female (0.4% as diverse resp. 0.5% preferred not to say), our gender-based analyses focus solely on male and female participants.

We assessed differences in importance ratings across age groups, gender, and device ownership using χ^2 tests, treating each privacy property as a binary outcome (important vs. not important). In addition, we conducted logistic regression analyses to model the effects of ATI and IUIPC (each entered as a binary predictor) on the likelihood that participants rated a given property as important.

Privacy Properties Rationale Analysis. We analyzed the participants' rationale by first aggregating the fixed-choice responses to provide a descriptive overview. Additionally, we analyzed the free-text responses using an inductive coding approach. One author reviewed all open-ended answers to develop a codebook, resulting in 109 coding categories. A second independent rater then classified the responses using this codebook. Finally, the two raters resolved any discrepancies through a consensus meeting to determine the final category for each response. Additionally, noting opportunities to consolidate thematically similar categories, the 109 coding categories were dimensionally reduced to 106 distinct categories.

4 Results

We begin by detailing the participants' demographics, including their IoT device ownership and usage information in Sec. 4.1. We then present the prioritization of privacy properties in Sec. 4.2, followed by an analysis of how technical affinity, privacy concerns, gender, age, and device ownership influence these preferences (Sec. 4.3). We further examine the qualitative rationale behind the perceived importance ratings in Sec. 4.4.

Table 2: Sample demographics and usage type distribution.

Gender	n	%	IoT Usage Type	n	%
Male	286	50.6	Private only	380	67.3
Female	274	48.5	Neither private nor professional	152	26.9
Diverse	2	0.4	Both private & professional	27	4.8
Unspecified	3	0.5	Professional only	6	1.1

Age Group	18–24	25–29	30–34	35–39	40–44	45–49	50–54	55–59	60–64	65–69	70+
n	39	36	43	43	65	45	41	74	38	59	82
%	6.9	6.4	7.6	7.6	11.5	8.0	7.3	13.1	6.7	10.4	14.5

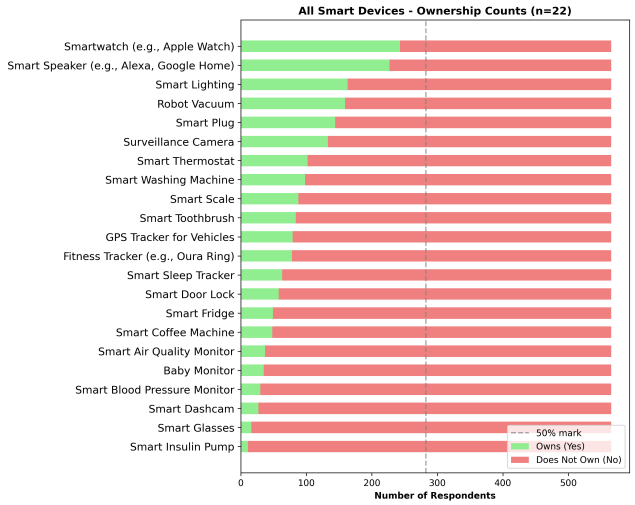

Figure 1: Smart device ownership.

4.1 Participant Sample, Device Ownership, and Usage Context

After excluding 155 participants who failed our attention check, the final dataset, detailed in Tab. 2 includes $N = 565$ participants with a nearly balanced gender distribution (50.6% male, 48.5% female). The sample covers a broad age range, spanning from young adults (18–24) to seniors (70+). Our sample size thus ensures a wide representation across both genders and multiple age groups.

As introduced in Sec. 4.1, to contextualize the privacy priorities, we analyzed device ownership and usage type among $N = 565$ participants. The ownership data shows a clear preference for wearable and general household smart devices, as shown in Fig. 1. The most commonly owned devices are smartwatches (43.01%, $n = 243$) and smart speakers (40.18%, $n = 227$). Ownership significantly drops for highly specialized or niche products, such as smart insulin pumps (1.95%, $n = 11$), suggesting that participants' experience is concentrated on more common consumer IoT devices. Regarding usage context, Tab. 2 confirms that IoT devices are predominantly used for personal purposes. Approximately 67.26% ($n = 380$) of respondents indicated usage as only "private", with only 1.06% ($n = 6$) reporting only professional use. This strong skew toward private usage confirms that the articulated privacy concerns are predominantly rooted in the personal consumer experience and the data generated within the domestic sphere.

4.2 Privacy Properties Prioritisation

To assess the perceived significance of various IoT privacy properties, we asked participants to rate a random half of a set of 57 distinct properties using a five-level importance scale. This prioritization allows us to identify which of the 57 privacy features users consider most critical.

Fig. 2 shows that respondents in our sample rated most properties as important, reflecting a high baseline expectation. Nevertheless, clear differences between properties (e.g., CS-CRA vs. CS-SUP or CF-RST vs. CF-ACC) indicate meaningful prioritization rather than uniformly high ratings across all properties.

The most important properties center on immediate operational security and reliability, i.e., being alerted to device malfunctions (TF-NTF, $Mean = 4.38$ where 1 mapped to "not at all important" and 5 to "very important"), ensuring encrypted storage of authentication data (ID-PWD, $M = 4.35$), and guaranteeing secure, encrypted communication over the internet (ID-TLS, $M = 4.35$). Furthermore, data minimization, meaning that only data necessary for the product's purpose are processed (DH-MIN, $M = 4.32$), and the availability of simple and intuitive privacy controls (UC-INT, $M = 4.29$) also emerge as top priorities. The consistently high scoring of these properties suggests that users prioritize protective, easily manageable security features.

As the properties move down the hierarchy, a cluster of features falls within the mid-to-high mean importance range. These often relate to transparency, data flow control, and legal compliance. Properties concerning information on data recipients (TI-RCP, $M = 4.20$; TI-SHR, $M = 4.20$; TI-EXR, $M = 4.11$), purpose of data collection (TI-PUR, $M = 4.10$), and the provision of clear, transparent privacy information (UT-CLR, $M = 4.16$) are viewed as highly important, although slightly less critical than the top operational security features.

Another key observation concerning the distribution of perceived importance is the tendency of participants to rate certain properties as neutral. Specifically, these ratings were most frequently applied to properties that, across the full spectrum of responses, were skewed toward the non-important end. This suggests that while participants did not explicitly dismiss these properties as irrelevant, they did not perceive them with the same level of criticality. This differentiation highlights a hierarchy of privacy concerns, indicating where developers and privacy label designers should primarily focus their resources to satisfy users' preferences.

The least important properties, with mean scores falling below 3.30, include features like the ability to use the device without a user account (ID-BST, $M = 3.22$), the use of Open-Source software (ID-OSS, $M = 3.29$), and properties related to novel or indirect concerns. Notably, the ability to provide graphical representations of data flows (UT-VIS, $M = 3.11$), the ability to set up multiple user accounts (CF-ACC, $M = 2.89$), and the feature of assisting users in informing guests about data protection (CS-SUP, $M = 2.76$) received the lowest scores. This suggests that while participants did not explicitly dismiss these properties as irrelevant, they did not perceive them with the same level of urgency as the more core privacy properties. This differentiation highlights a hierarchy of privacy concerns, indicating that practical, self-protective measures are preferred over features dedicated to guest management or abstract data visualization.

Fig. 3 presents the pairwise significance matrix comparing all 57 privacy properties. While non-significant differences should not be interpreted as conceptual equivalence, the prevalence of statistically indistinguishable pairs suggests that several properties form tiers of similar perceived importance, which can inform a layered content structure to reduce user overload.

We later expand on this prioritization in Sec. 4.4, where we discuss the underlying rationale behind participants' choices and reveal additional nuances in users' interpretations of these privacy properties.

4.3 The Effect of Different Factors on Privacy Properties Prioritisation

Before examining each user's rationale in detail in Sec. 4.4, we conducted a set of analyses to explore how different factors shape the perceived importance of IoT privacy properties. Specifically, we investigated the effects of technical affinity (using the ATI scale), privacy concerns (using the IUIPC scale), gender, age, and device ownership on participants' prioritization patterns. These analyses reveal both shared baseline expectations and distinct subgroup-specific preferences, which we elaborate further.

ATI Effect. ATI produces a nuanced pattern as shown in Fig. 4 and Fig. 5. While higher ATI is often assumed to correlate with higher privacy expectations, our results show a more differentiated effect. Large ATI-related differences are observed in properties that require active configuration. These include privacy properties such as multiple-user management (CF-ACC), visual representation of data flows (UT-VIS), control of software updates (ID-UPT), granular privacy settings (UC-GRN), and exploration functionality (TF-EXP). High-ATI users place greater value on hands-on privacy management and interactive control interfaces.

Low-ATI users value strong defaults and passive protections. Some features are more valued by low-ATI users, particularly those that automatically protect them, such as secure default settings (ID-DEF), account requirements (ID-BST), resetting to defaults (CF-RST), and controllable discovery functions (ID-CAD). This indicates a preference for passive, default, and low-effort privacy protections.

Regardless of technical affinity, some privacy properties remain nearly always valued, including secure communication (ID-TLS), encrypted storage (ID-PWD), update status (TF-UPD), data minimization (DH-MIN), Privacy Enhancing Technologies (ID-PET), and alerts (TF-NTF). These form a shared baseline expectation independent of technical skill.

IUIPC Effect. IUIPC exhibits the strongest and most consistent segmentation of all dimensions as illustrated in Fig. 4 and Fig. 5.

High-IUIPC users value detailed transparency and data-flow information. These participants assign significantly greater importance to properties such as getting informed on which data categories are transferred (TI-SND), which exact recipients the device sends data to (TI-EXR), which categories of data are collected (TI-CAT), why and for what purpose the data is collected (TI-PUR), in which countries data is stored (TI-LOC), functions to allow for

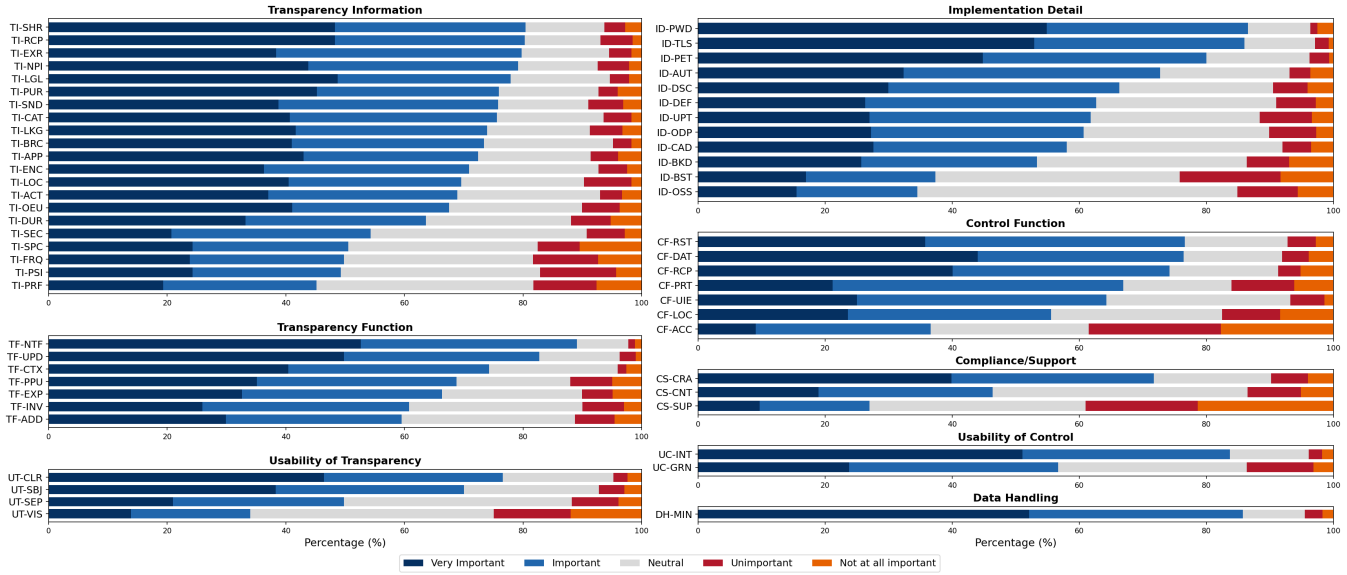

Figure 2: Perceived importance of privacy properties grouped by their corresponding categories (in %), ranked from most to least important.

controlling software updates (ID-UPT), and whether the IoT product requires connection to a specific backend service (ID-BKD). High-IUIPC participants want fine-grained transparency regarding data flows, recipients, storage locations, and purposes.

Even low-IUIPC users place a high value on secure communication (ID47), encrypted storage (ID-PWD), alerts (TF-NTF), update status (TF-UPD), and data minimization (DH-MIN). Thus, security essentials are consistently appreciated across all levels of privacy concern.

We observe low-importance features (even among IUIPC-high). Certain non-critical features, such as visual representation of flows (UT-VIS), guest information support (CS-SUP), special categories of data (TI-SPC), and open source software (ID-OSS), remain comparatively low. These features appear to be optional enhancements rather than core expectations.

Gender Differences. Women consistently rate most privacy features as more important (Fig. 4). The strongest gender gaps appear in cloud provider choice (CF-LOC), requiring connection to a specific backend service (ID-BKD), availability of joint controller/processor contracts (CS-CNT), which companies' data is passed on to (TI-RCP), which categories of data are collected (TI-CAT), and for what purposes data is collected (TI-PUR). These properties relate to transparency around data flows and institutional data handling, which women value more strongly.

Men place slightly greater importance on hands-off protections. Two features, i.e., secure default settings (ID-DEF) and account requirements (ID-BST), frequently appear higher for men. This mirrors ATI patterns and suggests that male participants show a slight preference for automatic, low-effort privacy protections.

We also observe gender-invariant properties. These properties, i.e., secure communication (ID-TLS), update status (TF-UPD), and alerts (TF-NTF) are equally valued across genders (>90%). This suggests that gender does not meaningfully affect the evaluation of essential security mechanisms.

However, the statistical analysis narrows these differences to two properties, i.e., IoT products to be shipped with secure default settings (ID-DEF) and information on personal data to be provided clearly and transparently (UT-CLR), both with small effect sizes (V=0.19 for ID-DEF, resp. V=0.14 for UT-CLR). This confirms that, with the exception of these two distinct priorities, privacy expectations are largely consistent across genders. Moreover, neither property remains significant after the Bonferroni correction (controlling the family-wise error rate), nor *Benjamini-Hochberg* (BH) (controlling the false discovery rate), as highlighted in Fig. 5, further suggesting that the perceived importance of privacy properties is stable across genders.

Age-Related Differences. Age introduces a non-linear but interpretable pattern (Fig. 4). Older participants value most privacy features more strongly. The importance of many properties increases with age, especially how frequently data is transferred (TI-FRQ), in which countries data is stored (TI-LOC), data transfer to a new device (CF-PRT), different user accounts (CF-ACC), and control software updates (ID-UPT). This suggests that older participants appear more attentive to data flow, data storage, and device life-cycle issues.

Younger users prioritize fewer properties. Participants aged 18–29 consistently give lower ratings to features that appear less relevant for digitally acclimated younger respondents who tend to rely on default configurations: visual representations (UT-VIS), open source (ID-OSS), controllable discovery (ID-CAD), and account requirements (ID-BST).

Some properties maintain very high importance across all age groups, e.g., secure communication (ID-TLS), update status (TF-UPD), encrypted storage (ID-PWD), data minimization (DH-MIN),

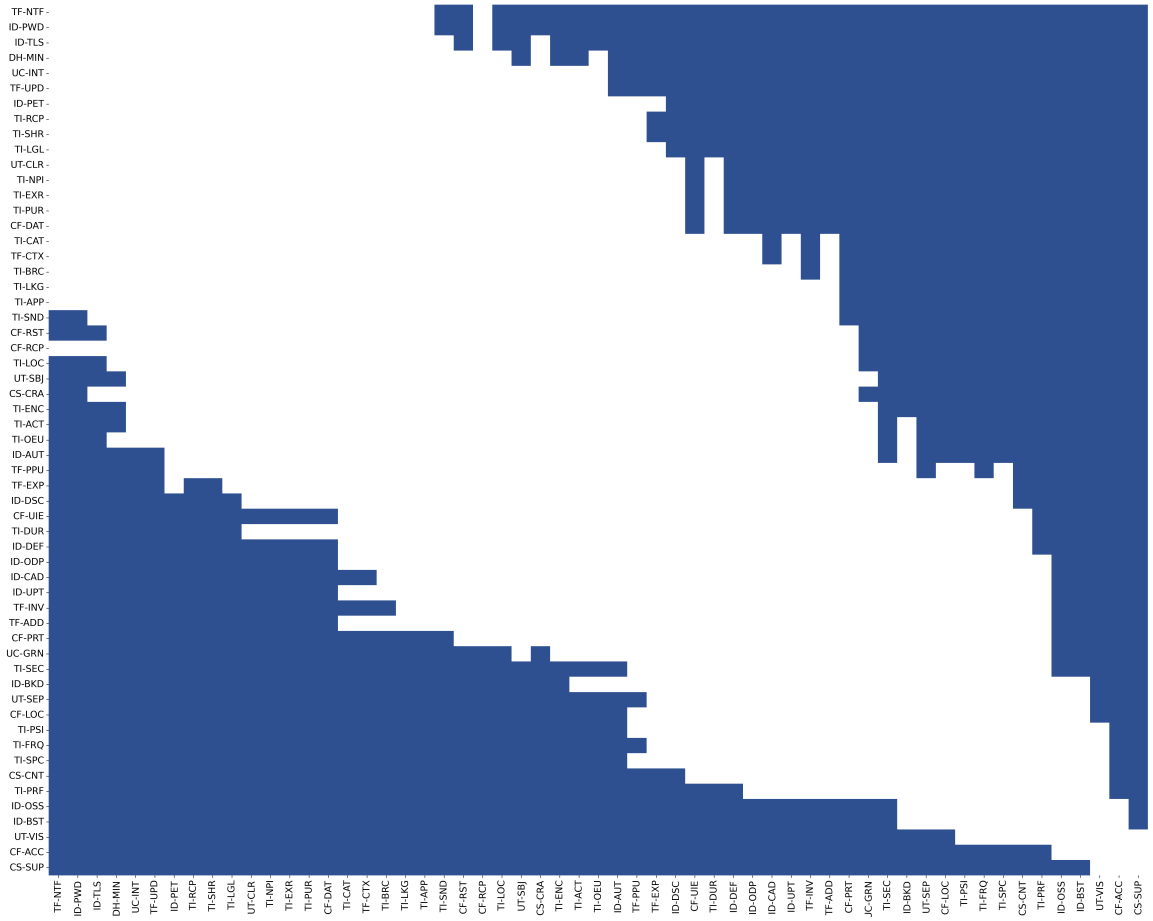

Figure 3: Pairwise significance matrix for privacy property importance ratings using Mann-Whitney U test, Bonferroni-corrected. Blue cells indicate significant differences in importance responses across corresponding properties.

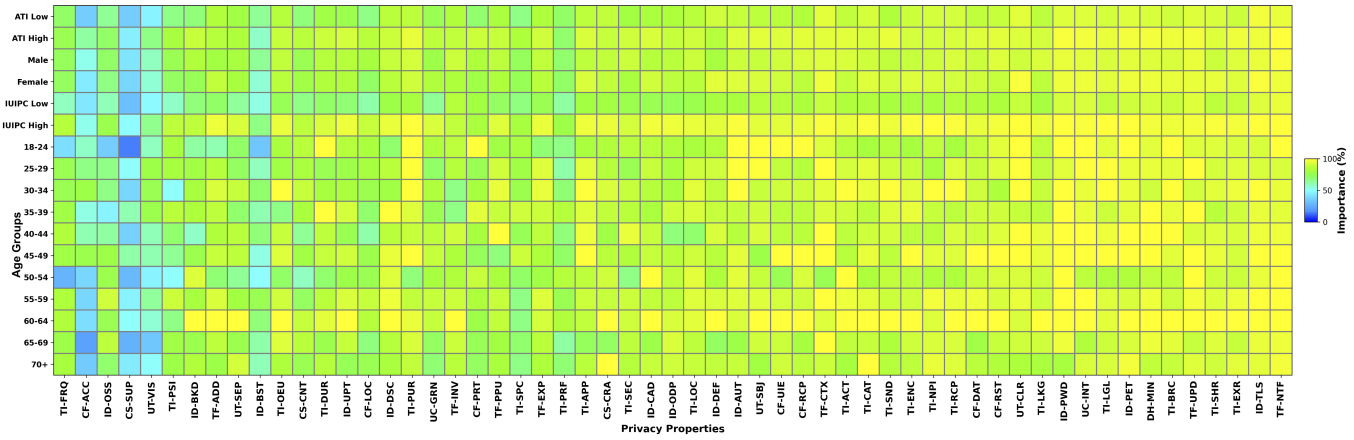

Figure 4: Heatmaps illustrating the perceived importance of 57 IoT privacy properties segmented across four user dimensions. The color intensity represents the proportion (%) of participants in each segment rating the property as important. Properties are sorted by variance across all groups.

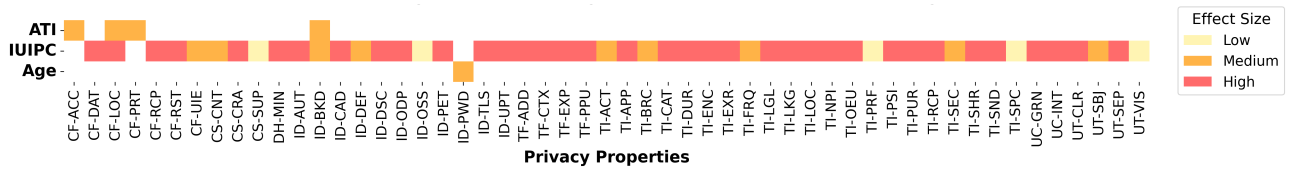

Figure 5: Significance overview for ATI, IUIPC, age, and gender effects on privacy property preferences after BH correction. Cell colors indicate effect-size ranges (odds ratios for ATI/IUIPC; Cramér's V for age). Gender shows no significance.

alerts (TF-NTF), and PETs (ID-PET). This again confirms a general baseline expectation. However, strict statistical validation yields a single significant result, driven by the 70+ age group, which rated encrypted password storage (ID-PWD) as less important than other age groups.

This suggests that while age groups slightly influence general rating tendencies, it rarely dictates the prioritization of specific individual features to a statistically significant degree.

Device Ownership Effect. To assess whether property importance depends on the type of IoT devices participants own, we compared owners and non-owners for each device category across the rated properties (important/highly important vs. otherwise, while excluding neutral responses). We restricted tests to comparisons with at least 20 owners and 20 non-owners among the participants who rated the respective property and applied χ^2 tests with effect sizes and multiple-testing correction. When the resulting 2×2 contingency table contained a zero cell or any expected cell count fell below 5, we instead used Fisher's exact test, as the χ^2 approximation is unreliable for small expected frequencies. Due to low sample size under 20 thresholds (Fig. 1), smart glasses and smart insulin pumps were excluded from this analysis. 37 were nominally significant at $p < .05$ without correction, but none remained significant after global multiple-testing BH correction.

To address device-conditional differences (i.e., which properties shift across device types), we applied BH correction within each device category, treating the set of properties tested for that device as a single hypothesis family and controlling the false discovery rate at the device level. We identify three device-conditional associations with small effect sizes (Appendix C.2). Owners of GPS trackers for vehicles and smart door locks prioritize multi-user account management (CF-ACC), and owners of smart washing machines prioritize accounts required for bootstrapping (ID-BST). This can be attributed to locks/trackers being shared/managed across household members (CF-ACC), and washing machines being commonly app/account-centric (ID-BST). Overall, device ownership yields only limited shifts in prioritization rather than broad device-specific preference patterns.

Cross-Dimensional Synthesis. Several cross-cutting insights emerge. Across ATI, IUIPC, gender, and age, the following properties consistently exceeded 90% importance: secure/encrypted connections (ID-TLS), encrypted storage of authentication data (ID-PWD), indication of a new security update (TF-UPD), data minimization (DH-MIN), easy and intuitive privacy settings adjustment (UC-INT), alerts for unusual events (TF-NTF), and use of technical techniques (ID-PET). These represent foundational IoT privacy expectations, regardless of demographics or attitudes.

Features with strong variance across the addressed factors include: cloud provider choice (CF-LOC), backend service demands (ID-BKD), data transfer categories (TI-SND), data recipients (TI-EXR), update control (ID-UPT), granular settings (UC-GRN), and visual flows (UT-VIS). These are prime candidates for adaptive privacy interfaces that tailor controls and information based on user profiles. Consistently less valued properties include guest information (CS-SUP), visual flows (UT-VIS), open source (ID-OSS), and profiling (TI-PRF).

Summary. Our analysis identifies both general privacy expectations and group-specific preferences across 57 IoT privacy properties. While security fundamentals such as encryption (ID-TLS, ID-PWD), update displays (TF-UPD), and data minimization (DH-MIN) are valued consistently across all factors, preferences for transparency, control, and configurability vary significantly across ATI, IUIPC, gender, and age. These findings provide a foundation for user-centered IoT privacy design, adaptive privacy interfaces, and evidence-based regulatory standards, and directly inform the development of a privacy comparison platform tailored to user needs.

4.4 The Rationale Behind Privacy Properties Perceived Importance

To understand the motivations behind the importance ratings, we analyze the distribution of standardized reasons (see Fig. 6) and conduct a qualitative content analysis of open-ended justifications (as introduced in Sec. 3.3). We further highlight the main findings.

4.4.1 Drivers of Importance: Security, Sovereignty, and Lifecycle. For properties rated as "important" or "very important", the quantitative data shows a dominance of ensuring that users' "data is protected" (see Fig. 6a). However, the qualitative responses reveal three distinct drivers for participants, as follows.

The Baseline Expectation: Users view operational security as a baseline requirement. For secure defaults (ID-DEF) and updates (TF-UPD), participants frequently noted "This is minimum standard" (P504) or "guaranteed security" (P558).

Data Sovereignty and "Mastery of the House": High importance ratings for cloud provider choice (CF-LOC) and storage location (TI-LOC) were driven by a strong desire for control and distrust of foreign jurisdictions. P309 stated "I want to remain master of the house", while P481 explicitly noted "I do not want to see my data stored in, for example, the USA". Similarly, P1582 emphasized ensuring data is not sent abroad.

The Product Lifecycle Perspective. Users look beyond immediate usage. Properties like reset to default (CF-RST) and data portability (CF-PRT) were rated highly due to future concerns. P283

Table 3: Descriptive statistics for ATI and IUIPC (N=565).

	Scale	M	SD	Mdn	Min	Max
	ATI	3.75	1.03	3.78	1.00	6.00
IUIPC Subscale	Collection	5.84	0.94	6.00	1.00	7.00
	Control	5.36	1.32	5.67	1.00	7.00
	Awareness	5.35	1.44	5.50	1.00	7.00
	IUIPC Total	5.59	0.95	5.75	1.00	7.00

mentioned "Deleting data when selling the device" as a key reason, and P504 noted that open source (ID-OSS) is critical "in case the manufacturer goes bankrupt", ensuring long-term usability.

4.4.2 Drivers of Non-Importance: Context, Distrust, and Duality. The "Not Important" ratings were driven by three primary factors, illustrated in Fig. 6b and explained by user comments as below.

Contextual Irrelevance (The "Single Household" Effect): The quantitative data shows that for user accounts (CF-ACC) and guest info (CS-SUP), the primary reason for non-importance was "no impact on my usage". The qualitative data confirms this is largely due to household composition. Numerous participants (e.g., Participants 309, 387, 481, 922, 1006) stated "I live alone" or "single household", considering multi-user features "superfluous" (P1645).

Distrust. A recurring theme for transparency properties, i.e., clear processing info (UT-CLR), subject rights (UT-SBJ), history of data leakage (TI-LKG) was resignation. Even when features were theoretically good, users rated them low because they believed manufacturers "lie anyway" or "one cannot verify it" (P309, P144). Regarding Cyber Resilience Act compliance (CS-CRA), P309 bluntly stated, "I don't believe in seals", and P398 feared that "security holes are built in" to certified devices.

The Open-Source Duality. While some users valued open source (ID-OSS) for transparency (P1700), others rated it low due to security fears. P775 described it as a "double-edged sword" because it could be used by "hostile hackers", illustrating a specific mental model where "open" implies also "insecure."

5 Discussion

Our findings provide a granular map of consumer privacy priorities for IoT devices. By synthesizing the importance ratings with the provided rationale, we identify critical implications for privacy research in IoT more generally and for privacy labels in particular.

5.1 Implications of Privacy Property Prioritization and Importance Rationale

Baseline vs. Configurable Properties. Users distinguish between non-negotiable safety features and configurable privacy options. Encryption (ID-PWD, ID-TLS) and updates (TF-UPD) are viewed as "minimum standards" (P504) that should be enabled by default, suggesting they function as baseline requirements rather than differentiators. In privacy labels, these can be represented as binary "pass/fail" indicators. In contrast, data-sovereignty properties (TI-LOC, CF-LOC) are perceived as choices users want to control, with a strong "EU vs. non-EU" framing in the qualitative data (e.g., P481, P1582). Accordingly, a specific label or comparison tools (e.g., derived from privacy policies) could highlight the data storage location and cloud provider to support users' desire to remain "master of the house" (P309), by purchasing IoT devices that conform to this.

Context-aware Prioritization. Multi-user and guest management features (CF-ACC, CS-SUP) were often rated as having no impact by single-household participants, a pattern reflected in our qualitative coding (e.g., Participants 481, 922, 1006), indicating that one-size-fits-all disclosures can overemphasize irrelevant properties. This motivates the disclosure of layered or filtered content.

Clarity Over Visual Complexity. Although respondents value transparency, they prefer semantic clarity over complex visuals. UT-VIS was rated lower than text-based purpose information (TI-PUR), and participants noted that visualizations can be unclear or exhausting (e.g., P406, P922). This suggests that deployments should prioritize clear, interpretable summaries over detailed topologies.

Trust, Verifiability, and Duality. Trust remains a barrier, reflected in comments such as "manufacturers lie anyway" (P144) and "I can't verify it" (P309). We also observe a polarized perception of open-source (ID-OSS). Some interpret it as transparency, others as increased vulnerability. When considering such properties, adding contextual qualifiers such as maintenance status or last audit date could help reduce misinterpretation.

5.2 The Effect of Different Factors on Privacy Properties Preferences.

Technology Affinity. The ATI shape shows users want privacy to be realized rather than uniformly increasing privacy concern. High-ATI participants favor interactive and granular control (e.g., CF-ACC, UC-GRN, and ID-UPT), whereas low-ATI participants prefer protective defaults and reduced configuration burden (e.g., ID-DEF, and ID-BST). This aligns with prior work suggesting that higher digital literacy lowers the perceived cost of managing privacy settings [31].

General Privacy Concerns. IUIPC is the most consistent predictor of higher importance ratings across the taxonomy, especially for transparency and control-oriented properties, while baseline security safeguards remain highly valued across groups. This suggests that, in an IoT purchase context, our IoT-specific property set reflects meaningful differences associated with general privacy concerns. While the privacy paradox [2] cautions that stated concerns may not translate into behavior, a next step would be to examine whether these preferences manifest in real decisions once such properties are implemented.

Gender. Women tend to place higher importance on transparency and data flow information (TI-CAT, TI-EXR, ID-CAD), consistent with findings on gendered risk perception [16, 19, 24].

Age. Older users exhibit slightly more comprehensive privacy assessments, contradicting stereotypes of low digital literacy [23, 37]. Their preferences reflect accumulated digital experience and higher exposure to risk.

5.3 Operationalizing the Taxonomy.

The proposed properties can be operationalized using different measurement and analysis tools. Network measurements (e.g., packet capture and flow analysis) can quantify recipients, destinations,

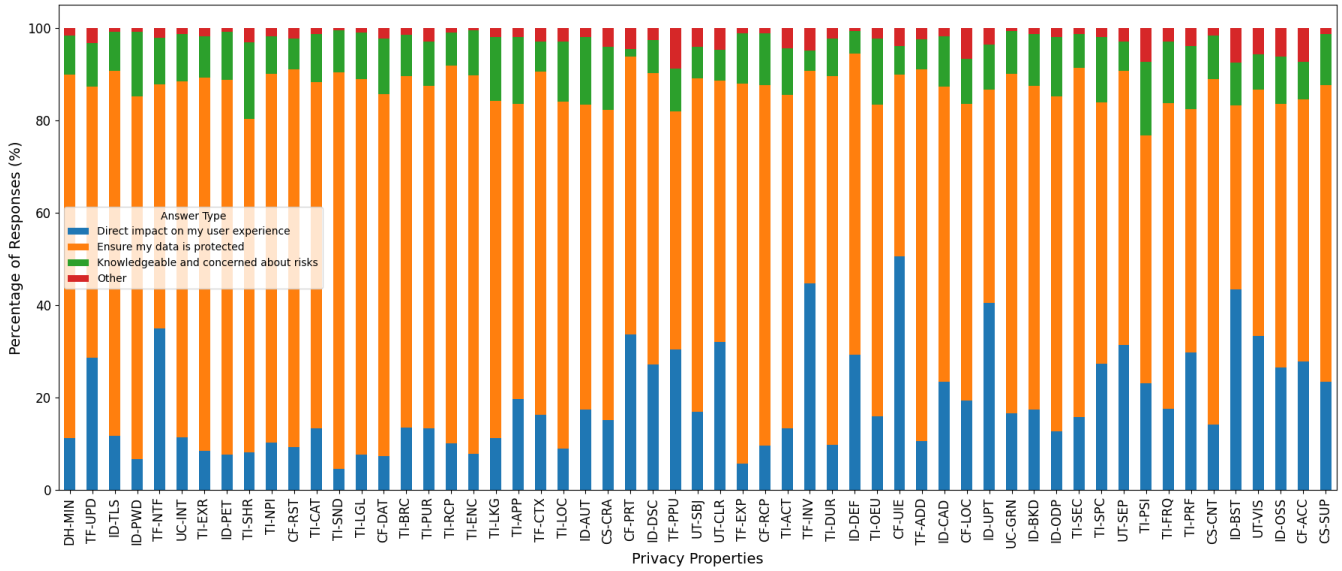

(a) Rationale behind the importance of privacy properties in percentage.

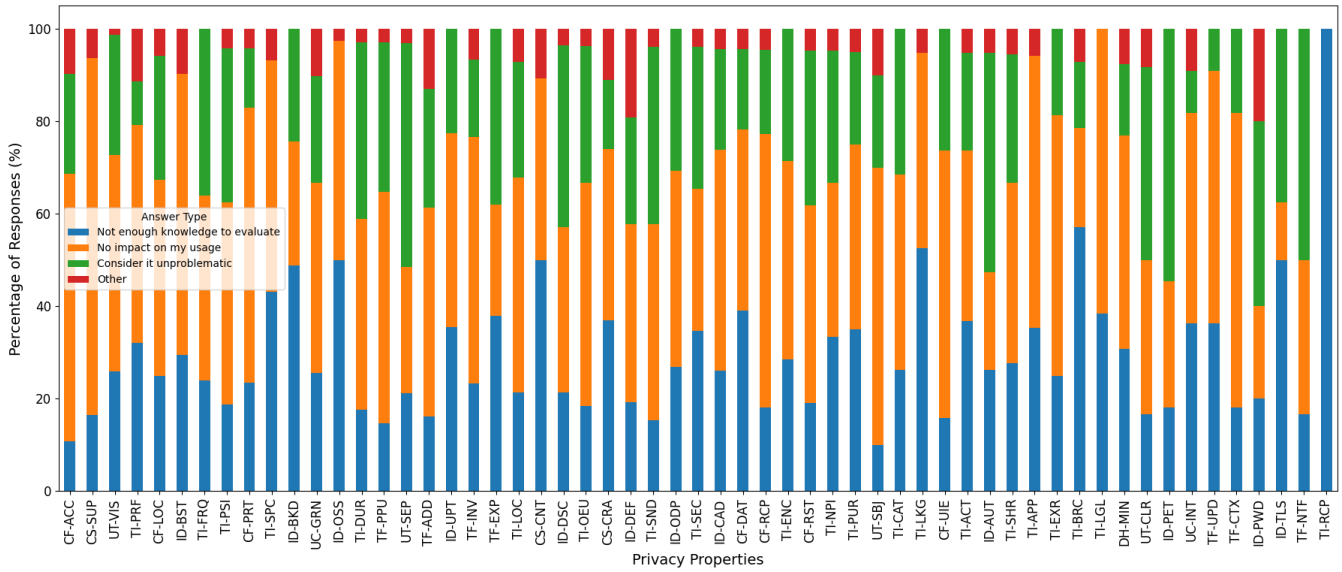

(b) Rationale behind the non-importance of privacy properties in percentage.

Figure 6: The rationale behind privacy properties prioritization in descending order of answer size.

encryption use, and communication frequency. Firmware analysis can reveal authentication mechanisms, encryption primitives, PET usage, and open-source components. Structured analysis of privacy policies [22], app metadata, and UI inspection can, e.g., extract legal grounds, data categories, storage duration, purposes, permissions, defaults, and available controls. This mapping supports implementable audits and comparative tools that populate label-like disclosures with verifiable fields, provides policymakers with a concrete checklist of assessable elements, and helps determine which properties to include and foreground in privacy label designs based on the empirical prioritization.

5.4 Limitations.

Cultural Aspects. Our sample consists solely of German participants, which introduces a cultural constraint. Privacy attitudes and behaviors are known to vary across cultures [43–45], and prior work on privacy labeling has predominantly focused on U.S. samples. By selecting a German cohort, we capture different cultural and regulatory approaches to privacy with differing levels of uncertainty avoidance [42]—A key factor influencing privacy awareness [43].

Willingness to Pay. Our study relies on stated preferences in a hypothetical scenario. We did not introduce economic constraints

(e.g., price premiums for private devices) or brand loyalty factors, which could override privacy concerns in their IoT purchasing decisions. Regardless, the current research is conflicted, as some [6, 13] highlight that users are willing to pay a premium for enhanced security and privacy of IoT devices. In contrast, others report limited willingness to pay for enhanced privacy in other settings [3].

Sample Effect. The high share of properties rated as important may potentially reflect participation effects. Nevertheless, controlled panel recruitment (Norstat) and substantial variance across properties indicate a meaningful hierarchy of preferences rather than just uniformly high ratings.

Properties Tone. The property descriptions were derived from regulatory texts and prior literature. While we simplified terminology where possible, we did not systematically control for the valence of affective wording. As a result, positively framed items (e.g., security or encryption properties such as ID-PWD) may be rated higher than more neutral or technical descriptions. However, we also observe neutrally phrased transparency items (e.g., TI-RCP) ranking above some positively framed security-related items (e.g., TI-ENC), suggesting that the rankings are driven primarily by the underlying content rather than wording alone.

6 Conclusion

As IoT devices grow in complexity, simply increasing the volume of transparency disclosures is no longer a viable approach. Our study demonstrates the value of providing users with a structured hierarchy of privacy and security labels. Across our analyses, we observe that while users rate many privacy properties as important in general, only a small subset, e.g., guest information (CS-SUP), multiple user accounts (CF-ACC), visual flows (UT-VIS), open source (ID-OSS), account requirements (ID-BST), and joint controller contracts (CS-CNT), drives meaningful variance. Many of the remaining properties ranked at the lower end are not necessarily viewed as less important, but rather they are either context-specific or perceived as unverifiable.

Thus, by distinguishing between non-negotiable security baselines and active sovereignty choices, and by recognizing the contextual irrelevance of some standard features, our paper provides insight into the design of the context-aware personalized privacy labels. We therefore propose a shift from 'compliance-driven' labeling, which maximizes data quantity, to 'preference-driven' labeling, which prioritizes the specific risk preferences (such as sovereignty and surveillance) that consumers actually use to differentiate IoT products. In turn, privacy labeling tools can better support user decision-making by directing attention to the features that matter most, as identified in our study.

Acknowledgments

This work was supported by the German Federal Ministry of Research, Technology, and Space (BMFTR) as part of the *Unboxing.IoT.Privacy* project (FKZ: 16KIS1932). We thank our project partners Luigi Lo Iacono, Florian Dehling, and Sebastian Koch from the University of Giessen, Harald Zwingelberg from the Unabhängiges Landeszentrum für Datenschutz (ULD, Independent State Center for Data Protection), and Chris Manivong from OneKey GmbH for their valuable contributions to the review, discussion, and refinement of the privacy properties taxonomy. We also thank Keno Sasse for his assistance with the qualitative coding analysis.

References

[1] Resilience Act. 2024. Regulation (EU) 2024/2847 of the European Parliament and of the Council. *Regulation (eu)* (2024).

[2] Susanne Barth and Menno DT De Jong. 2017. The Privacy Paradox–Investigating Discrepancies Between Expressed Privacy Concerns and Actual Online Behavior–a Systematic Literature Review. *Telematics and Informatics* (2017).

[3] Alastair R Beresford, Dorothea Kübler, and Sören Preibusch. 2012. Unwillingness to Pay for Privacy: A Field Experiment. *Economics letters* 117, 1 (2012).

[4] Carlos Bermejo Fernandez, Lik Hang Lee, Petteri Nurmi, and Pan Hui. 2021. PARA: Privacy Management and Control in Emerging IoT Ecosystems Using Augmented Reality. In *Proceedings of the International Conference on Multimodal Interaction*.

[5] Carlos Bermejo Fernandez, Petteri Nurmi, and Pan Hui. 2021. Seeing Is Believing? Effects of Visualization on Smart Device Privacy Perceptions. In *Proceedings of the ACM International Conference on Multimedia*.

[6] John M Blythe, Shane D Johnson, and Matthew Manning. 2020. What Is Security Worth to Consumers? Investigating Willingness to Pay for Secure Internet of Things Devices. *Crime Science* 9, 1 (2020), 1.

[7] Claire Chen, Hamsini Ravishankar, Dillon Shu, Laxita Jain, Yifeng Zeng, Yuvraj Agarwal, and Lorrie Faith Cranor. 2023. Ask the Consumers: What Should Be On IoT Privacy & Security Labels. In *Poster at the Symposium on Usable Privacy and Security (SOUPS)*.

[8] Chola Chhetri and Vivian Motti. 2022. Designing and Evaluating a Prototype for Data-Related Privacy Controls in a Smart Home. In *International Symposium on Human Aspects of Information Security and Assurance*.

[9] Gordon Corera. 2020. Smart Camera and Baby Monitor Warning Given by Uk's Cyber-Defender. https://www.bbc.com/news/technology-51706631. BBC News.

[10] Lorrie Faith Cranor, Yuvraj Agarwal, and Pardis Emami-Naeini. 2024. Internet of Things Security and Privacy Labels Should Empower Consumers. 67 (2024).

[11] Stefany Cruz, Logan Danek, Shinan Liu, Christopher Kraemer, Zixin Wang, Nick Feamster, Danny Yuxing Huang, Yaxing Yao, and Josiah Hester. 2023. Augmented Reality's Potential for Identifying and Mitigating Home Privacy Leaks. *arXiv preprint arXiv:2301.11998* (2023).

[12] Arianna Di Serio and Fabio Paternò. 2025. Human Control of Privacy and Security Aspects in IoT Settings. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia*.

[13] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 2023. Are Consumers Willing to Pay for Security and Privacy of {IoT} Devices?. In *USENIX Security Symposium*.

[14] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. 2019. Exploring How Privacy and Security Factor Into IoT Device Purchase Behavior. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*.

[15] European Data Protection Board. 2020. *Guidelines 05/2020 on Consent Under Regulation 2016/679.* Guidelines 05/2020. European Data Protection Board.

[16] Manuela Farinosi, Sakari Taipale, et al. 2018. Who Can See My Stuff? Online Self-Disclosure and Gender Differences on Facebook. *Observatorio* 12, 1 (2018).

[17] T Franke, C Attig, and D Wessel. 2018. Affinity for Technology Interaction (ATI) Scale. *Int. J. Human–Computer Interact* 2018 (2018).

[18] Thomas Groß. 2021. Validity and Reliability of the Scale Internet Users' Information Privacy Concerns (IUIPC). *Proceedings on Privacy Enhancing Technologies* (2021).

[19] Mariea Grubbs Hoy and George Milne. 2010. Gender Differences in Privacy-Related Measures for Young Adult Facebook Users. *Journal of Interactive Advertising* 10, 2 (2010).

[20] Anna Ida Hudig, Chris Norval, and Jatinder Singh. 2023. Transparency in the Consumer Internet of Things: Data Flows and Data Rights. *London: Information Commissioner's Office* (2023).

[21] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W Reeder. 2009. A "Nutrition Label" for Privacy. In *Proceedings of the Symposium on Usable Privacy and Security*.

[22] Sebastian Koch. 2025. Machine Learning Ansätze zur Untersuchung von Privatheitseigenschaften von IoT-Produkten: Eine Literaturstudie. In *SPRING Graduate Workshop*.

[23] Alex Koohang, Carol Springer Sargent, David Fuller, and Tara Underwood. 2023. Internet of Things (Iot): Users' Concerns About Privacy and Security. *Issues in Information Systems* 24, 2 (2023).

[24] Tingya Kuo and Hung-Lian Tang. 2015. Gender Differences in Facebook's Privacy Settings. *Issues in Information Systems* 16, 1 (2015).

[25] Dave Lee. 2018. Amazon Alexa Heard and Sent Private Chat. https://www.bbc.com/news/technology-44248122. BBC News. Accessed: 2026-01-15.

[26] Dave Lee. 2019. Google Admits Error Over Hidden Microphone. https://www.bbc.com/news/technology-47303077. BBC News. Accessed: 2026-01-15.

[27] Tianshi Li, Lorrie Faith Cranor, Yuvraj Agarwal, and Jason I Hong. 2024. Matcha: An Ide Plugin for Creating Accurate Privacy Nutrition Labels. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 1 (2024).

[28] Yanzi Lin, Jaideep Juneja, Eleanor Birrell, and Lorrie Faith Cranor. 2023. Data Safety vs. App Privacy: Comparing the Usability of Android and iOS Privacy Labels. *arXiv preprint arXiv:2312.03918* (2023).

[29] Naresh K Malhotra, Sung S Kim, and James Agarwal. 2004. Internet Users' Information Privacy Concerns (IUIPC): The Construct, the Scale, and a Causal Model. *Information Systems Research* 15, 4 (2004).

[30] Chris Norval and Jatinder Singh. 2023. A Room With an Overview: Toward Meaningful Transparency for the Consumer Internet of Things. *IEEE Internet of Things Journal* 11, 5 (2023).

[31] Yong Jin Park. 2013. Digital Literacy and Privacy Behavior Online. *Communication Research* 40, 2 (2013).

[32] Parks Associates. 2022. Data Privacy and Security in the Connected Home – Quantified Consumer. https://www.parksassociates.com/products/support-services/data-privacy. Accessed: 2025-11-25.

[33] David Phelan. 2019. Amazon Admits Listening To Alexa Conversations: Why It Matters. https://www.forbes.com/sites/davidphelan/2019/04/12/amazon-confirms-staff-listen-to-alexa-conversations-heres-all-you-need-to-know/. Forbes. Accessed: 2026-01-15.

[34] Alexandr Railean and Delphine Reinhardt. 2018. Let There Be Lite: Design and Evaluation of a Label for Iot Transparency Enhancement. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*.

[35] Alexandr Railean and Delphine Reinhardt. 2020. OnLITE: On-Line Label for Iot Transparency Enhancement. In *Proceedings of the Nordic Conference on Secure IT Systems*.

[36] Alexandr Railean and Delphine Reinhardt. 2021. Improving the Transparency of Privacy Terms Updates: Opinion Paper. In *Annual Privacy Forum*.

[37] Hirak Ray, Flynn Wolf, Ravi Kuber, and Adam J Aviv. 2021. "Warn Them" or "Just Block Them"?: Investigating Privacy Concerns Among Older and Working Age Adults. *Proceedings on Privacy Enhancing Technologies* (2021).

[38] Protection Regulation. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council. *Regulation (eu)* 679, 2016 (2016), 10–13.

[39] Yun Shen and Pierre-Antoine Vervier. 2019. IoT Security and Privacy Labels. In *Annual Privacy Forum*.

[40] Statista. 2025. Smart Home Worldwide. https://www.statista.com/outlook/cmo/smart-home/worldwide?currency=USD. Accessed: 2025-10-15.

[41] Philipp Thalhammer, David Müller, Alexander Schmidt, Michael Huber, Albrecht Schmidt, and Sebastian Feger. 2023. ConnectivityControl: A Model Ecosystem for Advanced Smart Home Privacy. In *Proceedings of the International Conference on Mobile and Ubiquitous Multimedia*.

[42] The Culture Factor Group. Hofstede Insights. 2025. Country Comparison Tool. https://www.theculturefactor.com/country-comparison-tool. Accessed: 2026-01-15.

[43] Sabine Trepte, Leonard Reinecke, and Nicole B. et al. Ellison. 2017. A Cross-Cultural Perspective on the Privacy Calculus. *Social Media+ Society* (2017).

[44] Anran Xu, Zhongyi Zhou, Kakeru Miyazaki, Ryo Yoshikawa, Simo Hosio, and Koji Yatani. 2023. DIPA: An Image Dataset with Cross-cultural Privacy Concern Annotations. In *Proceedings of the International Conference on Intelligent User Interfaces (IUI Companion)*.

[45] Anran Xu, Zhongyi Zhou, Kakeru Miyazaki, Ryo Yoshikawa, Simo Hosio, and Koji Yatani. 2024. DIPA2: An Image Dataset with Cross-cultural Privacy Perception Annotations. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* (2024).

[46] Shikun Zhang, Yuanyuan Feng, Yaxing Yao, Lorrie Faith Cranor, and Norman Sadeh. 2022. How Usable Are Ios App Privacy Labels? *Proceedings on Privacy Enhancing Technologies* (2022).

[47] Shikun Zhang and Norman Sadeh. 2023. Do Privacy Labels Answer Users' Privacy Questions?. In *Workshop on Usable Security and Privacy*.

A User Study Details: Privacy and Security in IoT Products

A.1 Study Information and Data Processing

Dear Ladies and Gentlemen,

The aim of this study is to identify the most important privacy and security properties relevant to users of Internet of Things (IoT) products. To this end, we are conducting a survey to capture user preferences and concerns and to better understand which aspects are particularly important to them. The ultimate goal of our study is to develop a platform that allows users to check which privacy and security features IoT products offer. This is intended to facilitate informed purchasing decisions and create greater transparency in this area.

The estimated duration of the study is approximately 15 to 20 minutes. The survey should only be completed by adult participants with full legal capacity.

Participation in this study is strictly voluntary, and you may cancel your participation at any time without giving reasons. Your data will then be deleted immediately.

Participation requires your consent and confirmation that you are of legal age and have full legal capacity.

Data Processing Information

Controller of Processing
Georg-August-Universität Göttingen, Wilhelmsplatz 1, 37073 Göttingen

Scientific Contacts
Lindrit Kqiku: kqiku@cs.uni-goettingen.de
Sebastian Schillinger: s.schillinger@stud.uni-goettingen.de

Purpose of Processing
The survey serves as a **scientific study** with potential publications on the preferences of IoT device users.

Collected Personal Data
For the scientific study: Weightings of transparency, control, and security aspects of IoT devices, gender, age group, and types of IoT devices used. This data is collected pseudonymously and anonymized upon completion of the survey.
Separately: The operator of the survey server `https://tinyurl.com/gwdg-dptn` records the access date, IP address, browser type, and a random number (session cookie) for operation and standard usage statistics.

Rights Regarding Consent
You have the right to refuse consent and to revoke it at any time. For this purpose, every survey page offers the link "*Exit survey and clear responses*".

Storage Location in Europe
All personal data is processed in accordance with the European General Data Protection Regulation (GDPR).

No Automated Decisions
The processing of your data does not involve any automated decision-making.

Further Information
Detailed privacy policy information can be accessed via the Article 13 PDF document linked in the survey interface.

A.2 Part 1: Usage of Smart Devices Questions

Q1. Do you own or use one or more of the following smart devices? *[Type: Matrix / Yes-No]*

- Smart Speaker (e.g., Alexa, Google Home)
- Smart Air Quality Monitor
- Smart Lighting
- Smart Thermostat
- Smart Plug
- Smart Door Lock
- Baby Monitor
- Surveillance Camera
- Smart Fridge
- Smart Washing Machine
- Smart Coffee Machine
- Robot Vacuum
- Smartwatch (e.g., Apple Watch)
- Fitness Tracker (e.g., Oura Ring)
- Smart Glasses
- Smart Dashcam
- GPS Tracker for Vehicles
- Smart Insulin Pump
- Smart Blood Pressure Monitor
- Smart Sleep Tracker
- Smart Scale
- Smart Toothbrush

Q2 - Q23. How often do you use [Device Name]? *[Type: Single Choice]*

Condition: Displayed for each device marked "Yes" in Q1.

- Daily
- Several times a week
- Once a week
- Monthly
- Annually
- Never

Q24. Do you use smart devices privately or professionally? *[Type: Single Choice]*

- Private
- Professional
- Both private and professional
- Neither private nor professional

Q25. Where do you use this/these smart device(s)? *[Type: Single Choice]*

Condition: Displayed if Q24 is not "Neither".

- At home
- At work
- Both at home and at work
- Other [Free Text]

A.3 Part 2: Feature Evaluation (Core Study)

Note: The participants were presented with a randomized selection of privacy and security features listed in Tab. 1. For each feature, the following question logic was applied:

Question Logic Pattern. *Scenario: Suppose you want to buy a new smart device. How important is the following function or property of the new device to you?*

- **[Feature Description inserted here (Tab. 1)]**

[Type: Likert Scale (5-point)]

- Not at all important
- Not important
- Neutral
- Important
- Very important

Follow-up A (Positive): Please explain why the previously asked property is important to you.

Condition: Triggered if Main Question = "Important" OR "Very important"

- I am well versed in this area and have concerns about the risks.
- I want to ensure that my data is protected.
- This property has a direct impact on my user experience.
- Other [Free Text]

Follow-up B (Negative): Please explain why the previously asked property is not important to you.

Condition: Triggered if Main Question = "Not at all important" OR "Of little importance"

- I view this property as unproblematic.
- I do not know enough about this property to evaluate it.
- This property has no influence on my usage.
- Other [Free Text]

A.4 Part 3: Attention Checks

Interspersed within the feature evaluation section to ensure data quality.

Check 1: Please select 'Very important' to show that you are paying attention. *[Type: Likert Scale]*

Check 2: Please select 'Not important' to show that you are paying attention. *[Type: Likert Scale]*

A.5 Part 4: Standardized Scales

A.5.1 Affinity for Technology Interaction (ATI) Scale. **Intro:** The following is about your interaction with technical systems... *[Type: Likert Scale (6-point: Strongly Disagree to Strongly Agree)]*

- I like to occupy myself in greater detail with technical systems.
- I like to test the functions of new technical systems.
- I predominantly deal with technical systems because I have to.
- When I have a new technical system in front of me, I try it out intensively.
- I enjoy spending time becoming acquainted with a new technical system.
- It is enough for me that a technical system works; I don't care how or why.
- I try to understand how a technical system works in detail.
- It is enough for me to know the basic functions of a technical system.
- I try to make full use of the capabilities of a technical system.

A.5.2 Internet Users' Information Privacy Concerns (IUIPC-8). **Intro:** The following is about your concerns regarding the protection of your personal data on the Internet... *[Type: Likert Scale (7-point: Strongly Disagree to Strongly Agree)]*

- Consumer online privacy is really a matter of consumers' right to exercise control and autonomy over decisions about how their information is collected, used, and shared.
- Consumer control of personal information lies at the heart of consumer privacy.
- Companies seeking information online should disclose the way the data are collected, processed, and used.
- A good consumer online privacy policy should have a clear and conspicuous disclosure.
- It usually bothers me when online companies ask me for personal information.
- When online companies ask me for personal information, I sometimes think twice before providing it.
- It bothers me to give personal information to so many online companies.
- I am concerned that online companies are collecting too much personal information about me.

A.6 Part 5: Demographics

D1 Please select the age group you belong to. *[Type: Single Choice]*

- 18–24 ... 70+ (in 5-year increments)

D2 How would you describe your gender? *[Type: Single Choice]*

- Male
- Female
- Diverse
- Prefer not to say

B Privacy Properties

B.1 Category Definitions

Tab. 4 provides the definitions of the eight taxonomy categories used to classify the privacy properties in Tab. 1.

C Additional Results

C.1 The Rationale Behind Privacy Properties Perceived Importance By Number of Responses

Fig. 7 illustrates the rationale underlying the importance of the privacy properties based on the number of responses rather than percentages, as shown in Fig. 6, thereby making the impact of individual properties more visible, particularly in cases where percentage-based representations may obscure properties with fewer and more responses.

C.2 Device Ownership and Property Prioritization: Detailed Analysis

This appendix complements the effect of device ownership in property prioritisation with additional details. We tested whether ownership of specific IoT device categories is associated with the perceived importance of individual privacy and security properties, without separately eliciting ratings for each device type.

For each device category d and property p, we compared owners and non-owners using a binary outcome (rated *important* or *highly important* vs. otherwise; neutral excluded). Because participants rated only a subset of properties, each test uses the subset

Table 4: Definitions of the eight taxonomy categories used in Tab. 1.

Category	Description
Control Function	Functions that allow users to control privacy (e.g., control of sent data, recipients).
Usability of Controls	Usability characteristics of privacy controls (e.g., granularity of privacy settings in the UI).
Transparency Function	Functions that enable or enhance transparency (e.g., contextual data disclosure).
Transparency Information	Information disclosed by the manufacturer/service about data practices and actors (e.g., recipient of information, app Permissions Required).
Usability of Transparency	Usability characteristics of transparency information/functions (e.g., separation of data protection information from Terms & Conditions).
Data Handling	Properties describing how data is processed (e.g., minimization).
Compliance/Support	Compliance with regulations (e.g., CRA).
Implementation Detail	Technical properties impacting privacy/security (e.g., backend service demands).

Table 5: Device–property pairs significant under per-device BH FDR correction ($p_{\text{dFDR}} < 0.05$). No pair survived global Bonferroni or global BH FDR correction. V = Cramér's V; OR = odds ratio (Haldane–Anscombe corrected); $\%_{\text{own}}/\%_{\text{non}}$ = percentage rating the property as important among owners / non-owners; n = total participants with valid binary rating for that pair.

Device	Property Label	n	$\%_{\text{own}}$	$\%_{\text{non}}$	$\Delta\%$	V	OR	p_{dFDR}
GPS Tracker	CF-ACC	199	83.3	42.6	40.7	.278	6.24	.005
Smart Door Lock	CF-ACC	199	83.3	44.0	39.3	.241	5.79	.032
Smart Washing Machine	ID-BST	178	90.0	54.7	35.3	.255	6.51	.038

of participants who both reported ownership for d and provided a non-neutral rating for p. We applied a χ^2 test of independence and used Fisher's exact test when expected cell counts were small. We report effect sizes using Cramér's V and odds ratios (OR). OR was computed with a Haldane–Anscombe correction to avoid undefined values in sparse 2×2 tables. Three device-property pairs were significant under per-device BH FDR correction with small effect sizes, as detailed in Tab 5.

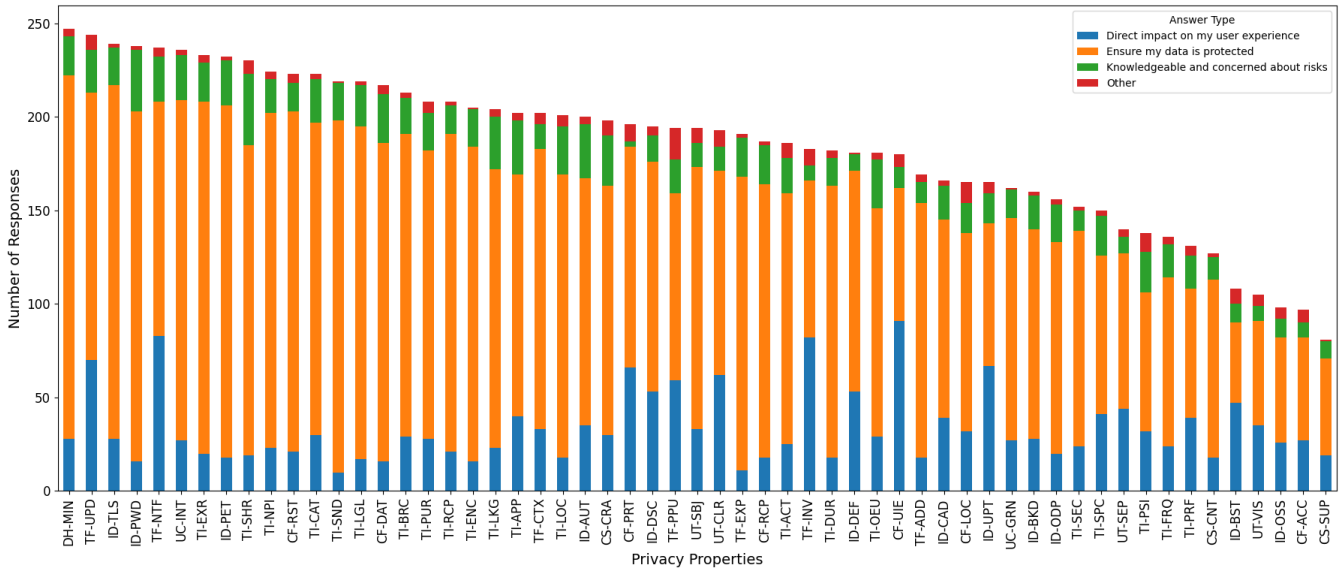

(a) Rationale behind the importance of privacy properties.

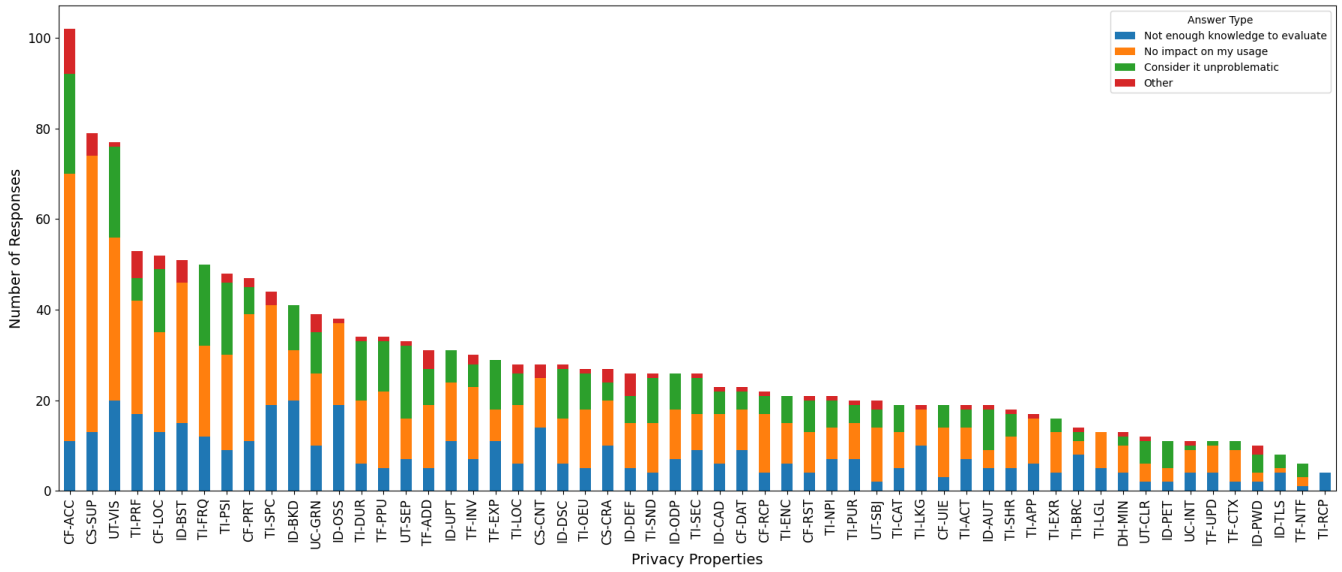

(b) Rationale behind the non-importance of privacy properties.

Figure 7: The rationale behind privacy properties prioritization in descending order of answer size.