FoT-PDS: A User-Centric Paradigm for Privacy-Preserving IoT

George P. Pinto[1,2], **Cássio V. S. Prazeres**[1]

[1]Institute of Computing – Federal University of Bahia (UFBA)
Salvador – BA – Brazil

[2]Federal Institute of Bahia (IFBA)
Salvador – BA – Brazil

georgepacheco@ifba.edu.br, prazeres@ufba.br

Abstract. *The Internet of Things enables pervasive and ubiquitous data collection, often without users' awareness or consent, reinforcing the so-called privacy paradox, in which technological benefits coexist with associated privacy risks. This thesis introduces FoT-PDS, an original user-centric paradigm that integrates the Fog of Things and Personal Data Stores to empower users with data control through decentralized personal data management, granular data-access control, transparency, and enhanced privacy awareness. The proposal also includes an AI-assisted consent mechanism based on clustering methods to anticipate profiling risks and support informed consent for users. Experimental results show that FoT-PDS significantly improves perceived data control, transparency, and privacy awareness, while the latter mediates the indirect effect of data control on users' perceived trust. Technical evaluation demonstrates the feasibility of the consent mechanism and its potential to mitigate risks associated with profiling. These insights provide empirical evidence supporting the adoption of the FoT-PDS as a viable and effective approach for promoting data control and mitigating privacy risks in the IoT context.*

1. Introduction

Whereas Internet of Things (IoT) technologies have accelerated the digitalization of daily life, from financial transactions to personal communications, they have simultaneously intensified concerns related to personal data privacy. The pervasive nature of connected devices increases individuals' exposure to unauthorized access. This technological scenario presents a trade-off between the convenience of the available technologies and services and user privacy, a phenomenon known as the '*privacy paradox*' [Kokolakis 2017].

In IoT environments, privacy risks are amplified by centralized data management models, where personal data are continuously transferred to third-party platforms with limited transparency and user control. Prior research [Pinto et al. 2024] has identified several privacy threats intrinsic to IoT systems, including identification, localization and tracking, profiling, and linkage. These threats are not merely technical side effects but structural consequences of architectures that decouple data ownership from data usage and decision-making.

In this context, user control, therefore, emerges as a fundamental requirement in addressing privacy risks. However, users typically lack the means to understand

and manage how their personal data are collected, processed, or shared. According to [CISCO 2019], this lack of control is often accompanied by a lack of transparency and awareness. In this manner, enhancing both is essential to increase users' perceived trust and engagement, especially in systems that heavily rely on personal data [Mugariri et al. 2022]. However, current IoT platforms still provide limited practical mechanisms for users to exercise meaningful control over their data, especially in highly distributed and resource-constrained environments.

Personal Data Stores (PDS) [Verborgh 2023] have been proposed as a promising user-centric approach to address this gap by decoupling data storage from service provision and returning control to individuals. In parallel, the Fog of Things (FoT) [Prazeres and Serrano 2016] paradigm has emerged to decentralize IoT processing closer to data sources, reducing latency and dependence on cloud infrastructures. Nevertheless, existing research largely treats PDS and FoT as independent solutions. To date, there is a lack of integrated paradigms that combine decentralized IoT infrastructures with user-controlled data governance and practical consent support mechanisms.

Motivated by this gap, this thesis proposes FoT-PDS, a user-centric paradigm for privacy-preserving IoT that integrates PDS into the FoT architecture. FoT-PDS shifts data governance from service-centric to user-centric models by enabling decentralized personal data management, fine-grained access control, and transparency throughout the IoT data lifecycle. Rather than relying solely on static policies or legal compliance, the paradigm operationalizes privacy through architectural design choices and user-controlled mechanisms.

A central component of FoT-PDS is an AI-assisted consent mechanism designed to support users in making informed data-sharing decisions. The mechanism analyzes users' personal IoT sensor data locally using unsupervised learning techniques to estimate profiling risks. These risks are translated into interpretable indicators that enhance user awareness and autonomy during the consent process, addressing a well-known limitation of traditional consent models that assume high levels of user expertise.

This work goes beyond proposing a technical solution by advancing a conceptual shift in how privacy is addressed in IoT environments. FoT-PDS positions individuals as active agents in the data lifecycle, integrating decentralized storage, semantic data representation, and AI-assisted consent into a coherent paradigm tailored to the constraints and characteristics of IoT systems.

The contributions of this thesis are validated through both technical evaluation and user-centered empirical analysis. The results demonstrate that FoT-PDS significantly improves users' perceived data control, transparency, and privacy awareness, while trust is indirectly strengthened through increased privacy awareness. Additionally, the technical evaluation confirms the feasibility of the AI-assisted consent mechanism and its potential to mitigate profiling risks in decentralized IoT settings.

By combining architectural innovation, intelligent consent support, and empirical validation, this thesis contributes to the advancement of privacy-preserving IoT systems and provides evidence that user-centric, PDS-based paradigms are viable and effective alternatives to centralized data governance models.

2. Research Questions

This thesis investigates how to mitigate personal data privacy issues in the IoT environments by exploring an approach that integrates PDS into the FoT paradigm. Accordingly, the main research question guiding this study was: ***Can a user-centric PDS-based approach mitigate personal data privacy risks in IoT environments?***

Based on the proposed FoT-PDS paradigm and the experimental evaluations conducted, this question is answered positively. The results demonstrate that a decentralized, user-controlled data management model enhances data control, transparency, privacy awareness, and contributes to increased user trust, while also enabling technical mechanisms to mitigate profiling risks.

To address this central question, the investigation is guided by the following questions. The first three questions adopt a user-centered perspective, focusing on how users perceive and respond to privacy-related dimensions when interacting with the FoT-PDS platform. The last question presents a technical perspective on the AI-assisted consent mechanism proposed in this thesis, focusing on its ability to detect privacy risks through the analysis of personal sensor data.

- **RQ1:** To what extent does the FoT-PDS improve users' perception of data control over collecting, storing, and sharing personal data?
  Answer: The experimental results show that the FoT-PDS significantly improves users' perceived control over their personal data. By decentralizing data storage and enforcing consent-based access through PDSs, users report greater awareness and authority over how their data are collected, stored, and shared.
- **RQ2:** How does the FoT-PDS influence user's perceptions of transparency, privacy awareness, and trust?
  Answer: The results indicate that the FoT-PDS has a positive and direct effect on users' perceptions of transparency and privacy awareness. Additionally, the paradigm indirectly enhances trust in IoT services, demonstrating that user-centric data control mechanisms improve users' confidence in how their data are handled.
- **RQ3:** To what extent do increases in transparency and privacy awareness contribute to users' trust in IoT services?
  Answer: The findings confirm that increases in transparency and privacy awareness significantly contribute to users' trust. Privacy awareness plays a mediating role, indicating that trust is strengthened when users both understand and perceive control over data processing practices.
- **RQ4:** To what extent can the AI-assisted consent mechanism assess the risk of user profiling and support informed consent decisions in IoT environments?
  Answer: The technical evaluation demonstrates that the AI-assisted consent mechanism is capable of assessing profiling risks by analyzing personal sensor data using clustering-based techniques. The generated indicators provide meaningful support for informed consent decisions, enabling users to better understand potential privacy risks before sharing their data.

3. Fog of Things and Personal Data Store paradigm

The FoT-PDS paradigm represents a structural rethinking of how personal data are governed in IoT environments. Traditional IoT architectures are service-centric, with unclear

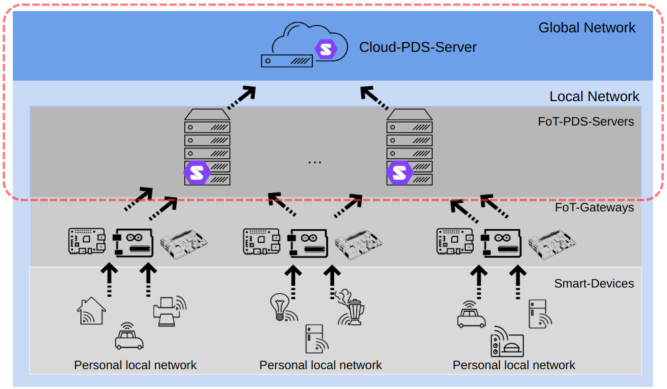

Figure 1. FoT-PDS paradigm.

data flows, centralized storage, and limited or no user agency over their data. In contrast, this paradigm shifts the control to individuals, integrating decentralized data management, semantic enrichment, AI-assisted consent, and risk awareness into a model that redefines how privacy is operationalized.

The paradigm aims to mitigate privacy challenges in the FoT environment [Prazeres and Serrano 2016] by embracing a user-centric model based on PDS. The FoT-PDS introduces privacy-preserving mechanisms at the edge of the network, close to the data source, and decentralizes data storage to enhance user control. Rather than relying solely on third-party entities, FoT-PDS allows individuals to determine how, when, and by whom their data are accessed, processed, or shared. Furthermore, by enabling data control, the paradigm aims to improve transparency, privacy awareness, and trust. Specifically, FoT-PDS focuses on facing problems directly related to personal data control (or the lack thereof), which are intrinsically related to the misuse of personal data. The paradigm mitigates key IoT privacy threats, including the *identification*, *localization and tracking*, *profiling*, and *linkage* threats, as pointed out by Ziegeldorf et al. [Ziegeldorf et al. 2014].

From the architectural perspective, as illustrated in Figure 1, our paradigm extends the original FoT architecture, focusing on the two major layers: the global network, which includes cloud-based storage solutions (*Cloud-PDS-Server*), and the local network, further divided into three essential components: *Smart-Devices*, *FoT-Gateways*, and *FoT-PDS-Servers*. As highlighted in the figure, the FoT-PDS paradigm operates at both the *Cloud-PDS-Server* and *FoT-PDS-Servers* layers. We adopted a PDS-based data management strategy that empowers users with granular control over how, what, who, and when their data can be accessed, processed, and disclosed, requiring explicit consent.

A key architectural principle of FoT-PDS is that privacy is enforced structurally. User consent is required at every stage where personal data may be accessed or disclosed, and users explicitly determine what data, for which purpose, by whom, and under which conditions it can be processed. By decentralizing storage and enforcing consent close to data sources, the paradigm mitigates critical IoT privacy threats such as identification, profiling, linkage, and tracking.

An integral component of the paradigm is the AI-assisted consent mechanism, which enhances user autonomy by supporting informed decision-making before data dis-

closure. Instead of relying on static consent models, the mechanism locally analyzes personal IoT sensor data using unsupervised learning techniques to estimate profiling risks. These risks are translated into interpretable indicators presented during the consent process, enabling users to better understand potential privacy implications without requiring advanced technical expertise.

On the other hand, from the perspective of the IoT data life cycle, privacy risks do not arise solely at the moment of data generation (DG) but continue and often intensify during subsequent stages, namely Data Storage (DS), Data Processing (DP), and Data Sharing (DSR). The FoT-PDS paradigm may mitigate threats, ensuring data control remains in the hands of users throughout the DS, DP, and DSR stages.

By combining decentralized data management, fine-grained consent enforcement, and intelligent risk awareness within a unified architecture, FoT-PDS operationalizes privacy as a core architectural property of IoT systems rather than as an auxiliary feature. This integration allows privacy protection to be addressed consistently across data generation, storage, processing, and sharing stages, while preserving the performance benefits of fog-based infrastructures.

3.1. AI-assisted Consent

While FoT-PDS defines a way to mitigate privacy issues in the IoT, it positions users as active actors by requiring them to give explicit consent for data access. Performing this task, however, demands a clear understanding of the purposes for which their data will be used, as well as awareness of the associated risks. Relying solely on users to manage consent decisions has proven ineffective in practice, as noted by [Acquisti et al. 2020]. Therefore, we propose an AI-assisted consent mechanism within our paradigm to support users in privacy-related decision-making before data disclosure. The mechanism aims to anticipate potential profiling risks based on users' data, enabling more informed consent decisions.

To achieve this, FoT-PDS applies unsupervised machine learning techniques, specifically clustering, locally and individually to each user's data, preserving decentralization and data isolation. The objective is not to infer or label user profiles, but to evaluate the relevance and consistency of emergent behavioral patterns that may indicate exposure to profiling risks.

Given the absence of reliable ground truth in high-dimensional IoT environments [Hassani and Seidl 2017], the mechanism relies on internal clustering validation metrics to assess the quality and significance of the generated clusters. These validation results are then translated into risk notifications presented to users as part of the consent process, strengthening user awareness and autonomy without requiring them to manually interpret complex privacy implications.

The mechanism quantifies the profiling risks through our Profile Metric (PM), derived primarily from the Silhouette index. This metric estimates how clearly a user's data can be organized into distinct behavioral patterns, which directly relates to the potential for profiling. The Silhouette index was selected due to its interpretability and suitability for localized, per-user analysis, aligning with the decentralized nature of FoT-PDS, where clustering is performed individually within each user's Pod. Although additional clustering indices (Davies-Bouldin and Calinski–Harabasz) are used during algorithm selection

to complement performance assessment, they provide only global evaluations and lack the local interpretability required for user-centric risk communication. For this reason, the Silhouette-based PM is adopted as the core indicator presented to users during the consent process.

The AI-assisted consent architecture comprises three components: a REST API, an ML Client, and the User's PDS. The REST API mediates external data access requests and enforces authentication and authorization policies. The ML Client operates locally over data stored in the PDS, performing clustering, computing the PM, and persisting the resulting metrics back into the user's storage. This design ensures that profiling risk assessment remains decentralized, privacy-preserving, and fully under user control.

4. Main Contributions

In our thesis, we investigated how to mitigate personal data privacy risks in IoT environments, particularly in the FoT paradigm, by adopting a user-centric approach based on PDS. The main contribution of this thesis is the proposal and validation of FoT-PDS, an original user-centric paradigm for privacy-preserving Internet of Things (IoT). The work introduces a novel integration of Personal Data Stores (PDS) into Fog of Things (FoT) architectures, shifting IoT data governance from service-centric models to a paradigm in which users retain control over their personal data while promoting transparency, privacy awareness, and users' trust.

A second original contribution is the design of an AI-assisted consent mechanism embedded in the FoT-PDS paradigm. Unlike traditional static consent or policy-based approaches, the proposed mechanism employs unsupervised learning (clustering) to anticipate profiling risks based on personal IoT data, providing users with interpretable indicators that support informed and context-aware consent decisions. This represents an innovative use of AI not for data exploitation, but for privacy risk awareness and mitigation.

From a technical standpoint, the thesis presents a fully specified and implemented architecture, detailing the interaction between Smart Devices, FoT Gateways, FoT PDS Servers, and Cloud PDS Servers. The proposal demonstrates how decentralized storage, semantic data annotation, selective synchronization, and fine-grained access control can be combined to enforce user consent across distributed infrastructures.

Additionally, the evaluation encompasses both the technical validation of the consent mechanism and a quantitative user study, supported by statistical modeling, to assess perceived data control, transparency, privacy awareness, and trust. This combination ensures strong internal validity and reproducibility, meeting high scientific standards.

4.1. Publications

Table 1 summarizes journal articles directly derived from the core contributions of this thesis, covering the proposed paradigm, the AI-assisted consent mechanism, and their empirical evaluation.

Table 2 presents additional publications that, while not directly derived from the thesis contributions, are closely related to the application scenarios, experimental contexts, and broader research agenda in which this doctoral work is embedded.

Table 1. Publications directly derived from the core contributions of this thesis.

Title	Year	Venue	Qualis	Contribution Focus	Citations*
My Data, My Rules: An Experimental Study on a User-Centric Approach to Data Privacy in the Internet of Things[Pinto et al. 2026]	2026	Computing	A1	FoT-PDS empirical validation and experimental user study	–
A User-Centric IoT Platform for Privacy with AI-Assisted Consent [Pinto and Prazeres 2025b]	2025	IEEE Open Journal of the Computer Society	A1	Practical implementation of the FoT-PDS platform	1
Enhancing IoT Data Privacy: AI-Assisted Consent Mechanism in a PDS-Based Solution [Pinto et al. 2025b]	2025	Internet of Things	A1	AI-assisted consent mechanism and profiling risk analysis	2
Data Privacy in the Internet of Things: A Perspective of Personal Data Store-Based Approaches [Pinto and Prazeres 2025a]	2025	Journal of Cybersecurity and Privacy	A1	Conceptual foundations of PDS-based privacy	6
Towards Data Privacy in a Fog of Things [Pinto and Prazeres 2024]	2024	Internet Technology Letters	A4	FoT-PDS foundational concepts and architecture	8
A Systematic Review on Privacy-Aware IoT Personal Data Stores [Pinto et al. 2024]	2024	Sensors	A1	Systematic analysis of PDS-based IoT privacy solutions	43

*Citation counts are based on Google Scholar.

5. Advances Over the State of the Art

Existing research on privacy preservation in IoT environments remains limited in scope, as most approaches focus on isolated mechanisms, such as access control, encryption, anonymization, or regulatory compliance, while continuing to rely on centralized data management models. As a result, the se solutions provide limited support for meaningful user control, transparency, privacy awareness, users' trust, and informed consent across the IoT data lifecycle, reinforcing limited user autonomy, profiling risks, and static, coarse-grained consent models that assume users can assess complex privacy implications [Pinto et al. 2024].

The FoT-PDS paradigm advances the state of the art by addressing these gaps in a unified manner. By integrating PDS directly into the FoT architecture, FoT-PDS bridges two previously disconnected research lines, enabling decentralized data processing while preserving continuous user control over personal data. This integration shifts privacy from an add-on feature to a structural property of the IoT architecture, enforced across data storage, processing, and sharing stages.

Moreover, FoT-PDS introduces an AI-assisted consent mechanism that overcomes

Table 2. Additional publications contributing to complementary validation scenarios and domain applications related to this thesis.

Title	Year	Venue	Qualis	Relation to Thesis	Citations*
Internet of Things Devices Management for Smart Cities [Sousa et al. 2025]	2025	International Conference on Internet of Things, Big Data and Security (IoTBDS)	A4	Smart city application scenario	–
Bridging the Cost Gap: A Comprehensive Analysis of CAPEX and OPEX for Smart Home Transition from a Provider's Perspective [Seixas et al. 2025]	2025	International Conference on Internet of Things, Big Data and Security (IoTBDS)	A4	Smart home deployment and evaluation context	2
Model and Service for Privacy in Decentralized Online Social Networks [Pinto et al. 2025a]	2025	Journal of Electronic Science and Technology	A3	Deccentralized privacy models and data governance	–
A Case Study of Smart Home Development [Martins et al. 2025]	2024	IEEE Software	A3	Empirical smart home case study	3
Designing, Implementing, and Testing AI-Oriented Smart Home Applications: Challenges and Best Practices [Campos et al. 2024]	2024	European Conference on Software Architecture	A3	AI-based IoT application design context	4

*Citation counts are based on Google Scholar.

the limitations of static and policy-based consent models. The mechanism locally analyzes users' personal IoT sensor data using unsupervised learning techniques to estimate profiling risks. These risks are translated into interpretable indicators that support context-aware and risk-informed consent decisions, reducing the cognitive burden on users while preserving decentralization and data sovereignty.

Finally, unlike most related work, which evaluates privacy solutions primarily from a technical perspective, this thesis combines architectural innovation with empirical user-centered evaluation. By assessing technical feasibility and users' perceptions of data control, transparency, privacy awareness, and trust, FoT-PDS provides evidence that privacy-preserving IoT architectures must be evaluated as socio-technical systems, rather than purely technical artifacts.

In summary, the key advance of this thesis lies not in proposing another privacy mechanism but in delivering an empirically validated paradigm that unifies decentralized IoT architectures, user-controlled data governance, and intelligent consent support. This integration addresses fundamental limitations identified in the literature and represents a positive step forward in the design of privacy-preserving IoT systems.

6. Conclusion

This thesis investigated how personal data privacy risks in IoT environments can be mitigated through a user-centric approach based on Personal Data Stores (PDS) within the

Fog of Things (FoT) paradigm. To address this challenge, the FoT-PDS paradigm was proposed, combining decentralized data management, user-controlled storage, and an AI-assisted consent mechanism to enhance data control, transparency, privacy awareness, and trust.

The empirical evaluation demonstrated that FoT-PDS significantly improves users' perceived control over their personal data. Increased data control positively influenced transparency and privacy awareness, while trust was indirectly affected through privacy awareness, highlighting its central role in shaping users' trust in privacy-preserving IoT systems. In addition, the technical evaluation confirmed that the AI-assisted consent mechanism can effectively estimate profiling risks using clustering-based analysis of personal sensor data, generating interpretable indicators that support informed consent decisions without requiring advanced technical knowledge.

Overall, the results validate that a PDS-based, user-centric approach can mitigate IoT privacy risks by combining architectural decentralization with intelligent consent support. The work contributes both conceptually and practically to the advancement of privacy-preserving IoT systems, reinforcing data control as a key enabler of user trust.

Despite these contributions, the evaluation was conducted in controlled environments, and the consent mechanism was assessed with a limited set of algorithms and features. Future work includes deploying FoT-PDS in real-world IoT scenarios, establishing baseline consent strategies for comparative evaluation, exploring additional clustering techniques and richer contextual features, and integrating Federated Learning to strengthen data minimization further.

References

Acquisti, A., Brandimarte, L., and Loewenstein, G. (2020). Secrets and likes: The drive for privacy and the difficulty of achieving it in the digital age. *Journal of Consumer Psychology*, 30(4):736–758.

Campos, D., Martins, L., Mota, J., et al. (2024). Designing, implementing, and testing ai-oriented smart home applications: Challenges and best practices. In *Software Architecture. ECSA 2024 Tracks and Workshops*, pages 83–99, Cham. Springer Nature Switzerland.

CISCO (2019). Consumer privacy survey: The growing imperative of getting data privacy right.

Hassani, M. and Seidl, T. (2017). Using internal evaluation measures to validate the quality of diverse stream clustering algorithms. *Vietnam J. of Computer Science*, 4(3):171–183.

Kokolakis, S. (2017). Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64:122–134.

Martins, L., Campos, D., Júnior, J. M., et al. (2025). A case study of smart home development. *IEEE Software*, 42(6):64–73.

Mugariri, P., Abdullah, H., García-Torres, M., Parameshchari, B. D., and Abdul Sattar, K. N. (2022). Promoting information privacy protection awareness for internet of things (iot). *Mobile Information Systems*, 2022(1):4247651.

Pinto, G. P., Donta, P. K., Dustdar, S., and Prazeres, C. (2024). A systematic review on privacy-aware iot personal data stores. *Sensors*, 24:2197.

Pinto, G. P., Leles, J. R., da Costa Souza, C., de Souza, P. R., Durão, F. A., and Prazeres, C. (2025a). Model and service for privacy in decentralized online social networks. *Journal of Electronic Science and Technology*, 23(1):100302.

Pinto, G. P. and Prazeres, C. (2024). Towards data privacy in a fog of things. *Internet Technology Letters*.

Pinto, G. P. and Prazeres, C. (2025a). Data privacy in the internet of things: A perspective of personal data store-based approaches. *Journal of Cybersecurity and Privacy*, 5(2).

Pinto, G. P. and Prazeres, C. (2025b). A user-centric iot platform for privacy with ai-assisted consent. *IEEE Open Journal of the Computer Society*, 6:1834–1846.

Pinto, G. P., Sousa, N. R., Da Silva, C. N., Peixoto, M. L., Figueiredo, G. B., and Prazeres, C. V. (2025b). Enhancing iot data privacy: Ai-assisted consent mechanism in a pds-based solution. *Internet of Things*, 34:101807.

Pinto, G. P., Sousa, N. R., and Prazeres, C. V. (2026). My data, my rules: an experimental study on a user-centric approach to data privacy in the internet of things. *Computing*, 108(3):33.

Prazeres, C. and Serrano, M. (2016). SOFT-IoT: Self-Organizing FOG of Things. In *2016 30th International Conference on Advanced Information Networking and Applications Workshops*, pages 803–808.

Seixas, N. F. S., Maia, A. H. O., Pinto, G. P., et al. (2025). Bridging the cost gap: A comprehensive analysis of capex and opex for smart home transition from a provider's perspective. In *Proceedings of the 10th International Conference on Internet of Things, Big Data and Security - IoTBDS*, pages 27–38. INSTICC, SciTePress.

Sousa, N. R., Pinto, G. P., and Prazeres, C. V. S. (2025). Internet of things devices management for smart cities. In *Proceedings of the 10th International Conference on Internet of Things, Big Data and Security - IoTBDS*, pages 15–26. INSTICC, SciTePress.

Verborgh, R. (2023). Re-decentralizing the Web, for good this time. In Seneviratne, O. and Hendler, J., editors, *Linking the World's Information: Essays on Tim Berners-Lee's Invention of the World Wide Web*, pages 215–230. ACM.

Ziegeldorf, J. H., Morchon, O. G., and Wehrle, K. (2014). Privacy in the internet of things: threats and challenges. *Security and Communication Networks*, 7:2728–2742.