

Acta
Universitatis
Danubius



COMMUNICATIO

Privacy Preservation in the Internet of Things: A Scoping Review and Future Research Directions

Tinashe Magara¹, Mampilo Phahlane²

Abstract: Privacy preservation is a critical concern in Internet of Things (IoT) systems due to the rapid expansion of connected devices and sensitive data exchange. Numerous studies have proposed privacy-preserving techniques (PPTs) to counter threats. However, a structured synthesis of these developments has been limited. This scoping review, conducted following PRISMA guidelines, analyzed literature published between 2010 and 2024. Data was collected from six databases, yielding 1,394 studies, of which 260 met predefined eligibility criteria related to privacy mechanisms, IoT layers and application domains. We identified three major categories of PPTs: cryptographic techniques, anonymization and obfuscation methods and decentralized trust models. Results show that most studies focus on e-health and industrial IoT domains, while device identity privacy remains underexplored. Persistent challenges include computational overhead, scalability and data utility trade-offs, particularly in cryptographic and blockchain-based solutions. The review proposes a taxonomy of PPTs and emphasizes the need for lightweight, scalable and trust-free frameworks such as the emerging “Blind Approach.” These findings provide a foundation for advancing efficient and privacy-preserving IoT architectures.

Keywords: Internet of Things (IoT); Privacy-Preserving; Security; Privacy; Taxonomy

¹ Ph.D., Department of Information Systems, College of Science, Engineering and Technology, University of South Africa, Pretoria, South Africa, Corresponding author: tinashenosh@gmail.com.

² Professor, Department of Information Systems, College of Science, Engineering and Technology, University of South Africa, Pretoria, South Africa, E-mail: magart@unisa.ac.za.



Copyright: © 2026 by the authors.

Open access publication under the terms and conditions of the
Creative Commons Attribution-NonCommercial (CC BY NC) license
(<https://creativecommons.org/licenses/by-nc/4.0/>)

1. Introduction

The rapid evolution of the Internet of Things (IoT) from early proprietary telemetry systems of the 1970s into a globally interconnected digital infrastructure has profoundly transformed the technological landscape (Alaba, 2024). This transformation, driven by the convergence of ubiquitous computing, advanced embedded systems and machine learning enabled analytics, has given rise to billions of interconnected devices that continuously generate high-frequency, context-rich and often highly sensitive data (Arif et al., 2024). While this interconnectivity enables efficiency, automation and real-time decision-making across industries, it also exposes users and systems to unprecedented privacy and security risks (Chang et al., 2024). Consequently, privacy preservation has emerged as a critical impediment to the large-scale adoption and societal trust of IoT technologies (Alferidah & Jhanjhi, 2020).

The accelerating urgency of data protection is reflected in the exponential growth of research addressing privacy in IoT systems (Gajjar, 2024; Magara & Zhou, 2024). Despite this growing interest, the literature remains fragmented and domain-specific, with many studies focusing narrowly on particular applications such as e-health, smart transportation and isolated architectural layers. Moreover, the diversity of technical paradigms, ranging from cryptographic primitives and anonymization mechanisms to decentralized trust frameworks, has complicated attempts to integrate findings into a unified body of knowledge. These limitations underscore the need for a comprehensive scoping review that consolidates prior work, maps the heterogeneous landscape of privacy preservation techniques (PPTs) and identifies key research challenges that narrower systematic reviews often overlook.

Existing literature has examined various dimensions of IoT privacy (Demertzi et al., 2023), including privacy goals, enabling technologies and the operational contexts in which privacy-enhancing technologies (PETs) are deployed (Ranaweera et al., 2021; Balboni et al., 2025). Some studies emphasize identity protection through pseudonymization, while others focus on general threats such as data profiling, inference attacks and surveillance across distributed environments (Zuo et al., 2021; Wachter, 2018). However, few have proposed an integrated framework for classifying PPTs that spans multiple layers of IoT architecture and technological domains (Dib et al., 2021). Furthermore, critical and emerging challenges, such as the protection of device identities (Alshboul et al., 2021), mitigation of cross-layer inference attacks and dependence on centralized trust mechanisms in decentralized paradigms like Federated Learning (Liu et al., 2024; Dritsas & Trigka, 2025), remain

underexplored. This review seeks to bridge these knowledge gaps by developing a unified taxonomy of PPTs and analyzing the complex trade-offs influencing their design, deployment and scalability in real-world IoT ecosystems (Chereja et al., 2025).

To guide this investigation, the following research questions (RQs) were formulated:

- RQ1: What are the prevalent Privacy Preservation Techniques (PPTs) utilized in IoT systems and how can they be classified into a comprehensive taxonomy?
- RQ2: In which IoT application domains (e-health, smart cities) and architectural layers (edge, fog, cloud) are specific PPTs most frequently deployed?
- RQ3: What are the critical trade-offs, performance limitations and research gaps constraining the practical deployment and adoption of robust PPTs in IoT ecosystems?

1.1. Theoretical Foundations of Privacy in the Internet of Things

A rigorous theoretical foundation is essential for navigating the complex and multidimensional nature of IoT privacy. Privacy in IoT extends beyond traditional notions of secrecy or physical isolation; it represents the right to informational self-determination, the ability of individuals and organizations to control the collection, use and dissemination of their data and digital identities (Oppl & Stary, 2022). Literature findings show that this multidimensional concept manifests across several critical dimensions:

1. User Privacy: Protecting personal information and behavioral data directly linked to individuals.
2. Device Identity Privacy: Safeguarding unique device identifiers to prevent traceability, spoofing, or unauthorized monetization.
3. Location Privacy: Concealing spatio-temporal information that could reveal user movement or device positioning.
4. Communication Channel Privacy: Ensuring the confidentiality and integrity of transmitted data, mitigating interception and packet-linking attacks.
5. Group Privacy: Preserving the anonymity of collective data concerning multiple users or devices, such as within smart homes or shared IoT environments.

Each of these dimensions is vulnerable to privacy erosion mechanisms, including surveillance, profiling and data aggregation, that demand tailored technical and organizational countermeasures.

A conceptual distinction must be made between the technological enablers of privacy and the methodological systems that implement them:

- Privacy-enhancing technologies (PETs): These form the foundational building blocks, mathematical and computational tools such as cryptographic primitives (such as homomorphic encryption), statistical models (such as differential privacy) and distributed infrastructures (blockchain). PETs are designed to ensure privacy protection at the algorithmic level.
- Privacy preservation techniques (PPTs): These are comprehensive methodologies and protocols that integrate one or more PETs to achieve specific privacy objectives in operational systems for example, pseudonymization schemes, secure multi-party computation, or privacy-preserving federated learning.

The interaction between PETs and PPTs forms the structural basis of privacy preservation in IoT (Lemieux & Werner, 2024). However, challenges such as computational overhead, scalability and cross-layer interoperability continue to limit their effectiveness in large-scale deployments (Adam et al., 2024; Obaidat et al., 2020).

Given these challenges, this review aims to provide a consolidated and systematic understanding of privacy preservation in IoT by:

1. Conducting a comprehensive scoping review of 260 primary studies on PPTs published between 2010 and 2024.
2. Proposing a unified taxonomy classifying PPTs into three primary categories, cryptographic, anonymization and decentralized trust-based models.
3. Mapping the application of PPTs to specific IoT domains and architectural layers.
4. Identifying key research gaps, including high computational costs, overlooked device identity privacy and systemic trust vulnerabilities.

1.2. IoT Layered Architecture and Privacy Threat

IoT systems are typically structured in a layered architecture, often represented as a three-tier model comprising the sensing/edge, network/fog and application/cloud layers (Mrabet et al., 2020; Gupta & Quamara, 2020). Each layer exhibits distinct operational characteristics, data flows and corresponding privacy threats:

- **Sensing/Edge Layer:** Composed of resource-constrained devices such as sensors and actuators. Privacy threats include physical tampering, side-channel attacks and local eavesdropping due to limited computational and cryptographic capabilities.
- **Network/Fog Layer:** Responsible for data aggregation and intermediate processing. This layer faces threats from compromised communication channels, identity spoofing and malicious fog nodes that can infer sensitive information from routing metadata.
- **Application/Cloud Layer:** Encompasses centralized storage and large-scale analytics. Privacy risks arise from data breaches, profiling through inferred data usage and dependency on centralized trusted third parties (TTPs) that may violate data minimization principles.

Understanding these multi-layered threats is fundamental for mapping the corresponding privacy-preserving technologies and their deployment contexts.

The remainder of this paper is structured as follows: Section 2 presents the methodology, including search strategy and selection criteria. Section 3 discusses the research results. Section 4 outlines future research directions and Section 5 concludes the paper.

2. Methodology

This study adopted a systematic scoping review methodology to provide a comprehensive overview of privacy preservation research in the internet of things domain. Scoping reviews are particularly well-suited for multidisciplinary and rapidly evolving fields characterized by methodological heterogeneity and conceptual fragmentation. This approach enables the identification of dominant research trends, theoretical frameworks and knowledge gaps without imposing the restrictive parameters typical of systematic reviews.

The methodological protocol followed the preferred reporting items for systematic reviews and meta-analyses (PRISMA) 2020 guidelines (Wirani et al., 2024) and incorporated adaptations from the Joanna Briggs Institute (JBI) framework for scoping reviews (Buglio et al., 2024). These methodological standards ensure transparency, reproducibility and robustness across all stages of the review process. The study period was delimited to 2010–2024, corresponding to the decade in which privacy-preserving technologies (PPTs) and privacy-enhancing technologies (PETs) gained significant traction within IoT systems.

2.1. Search Strategy

A comprehensive search strategy was employed to capture all relevant and peer-reviewed literature on IoT privacy preservation. Six major academic databases were selected for their breadth of coverage across engineering, computer science and multidisciplinary research: IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, Scopus and Web of Science.

Search strings were iteratively refined to maximize the retrieval of relevant studies. The final Boolean search query was as follows:

("Internet of Things" OR "IoT") AND ("privacy" OR "confidentiality" OR "data protection") AND ("technique" OR "model" OR "framework" OR "protocol") AND ("privacy-preserving" OR "privacy-enhancing" OR "PPT" OR "PET").

The search included peer-reviewed journal articles and conference papers published in English between 2010 and 2024. Gray literature, book chapters, preprints and white papers were excluded to maintain methodological consistency and ensure the inclusion of rigorously vetted sources.

2.2. Inclusion and Exclusion Criteria

To ensure the quality and relevance of the included studies, a two-phase screening process was applied. Articles were first assessed by title and abstract, followed by a full-text evaluation. The inclusion and exclusion criteria were designed to capture studies that directly addressed privacy-preserving methodologies within IoT contexts. Table 1, summarizes the inclusion and exclusion criteria that was applied in this paper.

Table 1. Inclusion and Exclusion Criteria

Inclusion Criteria	Exclusion Criteria
Peer-reviewed journal or conference papers.	Non-peer-reviewed materials
Published between 2010 and 2024.	Studies not related to IoT systems or architectures.
Explicit discussion or implementation of PPTs/PETs.	Generic cybersecurity or encryption works lacking privacy focus.
Empirical, experimental, or comparative studies evaluating PPTs.	Conceptual essays without methodological evaluation.
Research addressing IoT application domains or architectural layers (edge, fog, cloud).	Duplicate or republished content.
After applying these criteria, 260 primary studies were selected for synthesis.	

2.3. Screening and Selection Procedure

The systematic search yielded 1,394 records. Duplicate entries were first removed using Mendeley reference management software, followed by manual validation. Titles and abstracts were screened to assess relevance and remaining studies underwent full-text review by the authors. Disagreements were resolved through discussion. Table 2 shows the summary of the flow of study selection.

Table 2. PRISMA Flow Summary of Study Selection

Step	Initial Records	Included	Excluded (Reason)
Identification (Search)	1,394	–	–
Screening (Title & Abstract)	–	412	982 (Non-relevant, duplicates)
Eligibility (Full-Text Review)	–	260	152 (Conceptual or lacking novel PPTs)
Included in Final Synthesis	–	260	–

The process was conducted in accordance with PRISMA 2020 standards to ensure reproducibility and methodological transparency.

2.4. Data Extraction and Coding Framework

A structured data extraction protocol was developed to standardize data collection across all selected studies. Each article was coded across five analytical dimensions that collectively characterize privacy-preserving research within IoT:

1. Implemented Technique Type – encryption-based, anonymization, perturbation, decentralized trust mechanisms.
2. IoT Architectural Layer – Edge, Fog, or Cloud.
3. Application Domain –e-health, industrial IoT, smart home, transportation, agriculture.
4. Targeted Privacy Dimension – user privacy, device identity privacy, location privacy, communication privacy, or group privacy.
5. Performance and Trade-offs – computational overhead, latency, scalability, energy consumption, or privacy–utility balance.

2.5. Data Synthesis and Analysis

A mixed-methods synthesis approach was employed, integrating thematic and quantitative analyses to capture both conceptual breadth and empirical depth.

- Thematic analysis was used to identify and group recurring privacy-preserving paradigms, classify PPTs into functional clusters and reveal conceptual patterns across IoT layers and domains.
- Descriptive quantitative mapping was applied to visualize the distribution of PPTs across architectural layers and application domains, identifying areas of research concentration and underrepresentation.
- Critical comparative analysis was conducted to evaluate reported limitations and performance trade-offs, particularly focusing on computational efficiency, latency constraints and vulnerabilities in decentralized trust architectures.

This triangulated synthesis enabled a holistic interpretation of the IoT privacy preservation landscape, forming the foundation for the taxonomy and research gap analysis presented in subsequent sections.

2.6. Methodological Limitations

Although the study adhered to PRISMA and JBI standards, several methodological limitations must be acknowledged.

First, despite the use of six major databases, the scope of coverage may have omitted relevant studies published in less-indexed venues or in non-English languages, leading to potential selection and language bias. Second, while inter-reviewer reliability was high, subjectivity in thematic interpretation may have introduced minimal bias in categorizing hybrid techniques that spanned multiple architectural layers or privacy objectives. Third, the time-bound nature (2010–2024) of the review excludes the most recent unpublished or preprint studies, which could slightly limit the representation of emerging topics such as post-quantum privacy mechanisms and privacy-preserving machine learning.

Furthermore, variations in reporting standards and evaluation metrics across studies constrained the ability to perform quantitative meta-analysis or compare technique performance directly. Finally, the exclusion of gray literature may have omitted some practical or industry-driven innovations not yet formalized in academic publications.

Despite these limitations, the review’s methodological rigor, characterized by its systematic search, dual-reviewer validation and triangulated synthesis, ensures that the results accurately represent the breadth and depth of contemporary IoT privacy preservation research. The study thus provides a robust empirical foundation for subsequent analysis, taxonomy development and future research directions.

3. Results

This section presents the findings of the scoping review and introduces the proposed taxonomy of privacy preservation techniques (PPTs) in the internet of things (IoT). The taxonomy organizes the identified techniques into three principal methodological pillars, based on their underlying privacy-enhancing technologies (PETs):

- a) cryptographic-based approaches,
- b) anonymization and perturbation-based approaches and
- c) Decentralized and trust-based models.

Each category reflects a distinct conceptual foundation and addresses specific privacy challenges across IoT architectural layers and application domains.

Cryptographic-Based Techniques

Cryptographic approaches remain the most mature and widely implemented category of PPTs in IoT systems. These techniques employ mathematical mechanisms to ensure data confidentiality, integrity and authenticity, particularly for information in storage or transit.

Lightweight Cryptography and Homomorphic Encryption

Due to the inherent heterogeneity of IoT architectures, cryptographic solutions diverge according to computational capability. At the Edge layer, resource-constrained devices rely on lightweight cryptography (LWC), which optimizes ciphers such as ASCON v1.2, PRINCE and SPECK for minimal power consumption and limited memory availability. These algorithms favor energy efficiency and implementation simplicity over maximal cryptographic strength.

Conversely, homomorphic encryption (HE), which allows computation directly on encrypted data, offers strong confidentiality guarantees but incurs significant computational and latency costs. As a result, HE is predominantly adopted within Cloud or Fog layers where greater processing capacity is available. Although additively homomorphic encryption (AHE) has been experimentally implemented on lightweight platforms such as Raspberry Pi, the practical deployment of fully homomorphic encryption (FHE) remains constrained by its high overhead, bootstrapping latency and ciphertext expansion.

Attribute-Based and Searchable Encryption

Attribute-based encryption (ABE) and searchable encryption (SE) represent advanced cryptographic control mechanisms that enable fine-grained access control and secure data retrieval. ABE facilitates access decisions based on user attributes or contextual rules, making it particularly valuable in multi-stakeholder environments such as the internet of medical things (IoMT). Similarly, SE allows encrypted data to be queried without revealing the underlying plaintext or search criteria, an essential feature for cloud-based IoT platforms where sensitive data must remain confidential.

Anonymization and Perturbation-Based Techniques

These techniques aim to obscure identifiable information within data streams, thereby preventing the re-identification of individuals or devices while preserving analytical utility.

k-Anonymity and l-Diversity Models

Classical anonymization models such as k -Anonymity and l -Diversity form the foundation of data masking strategies. In IoT, they are adapted to ensure that each record remains indistinguishable from at least k others, or that quasi-identifiers maintain at least l distinct sensitive values. These approaches are often used to protect static device or user attributes; however, they struggle to cope with the high dimensionality, temporal variability and continuous nature of IoT data streams. Maintaining anonymity over dynamically evolving datasets remains a key challenge.

Differential Privacy

Differential Privacy (DP) has emerged as a dominant PET for guaranteeing anonymity and unlinkability through the addition of mathematically calibrated random noise to data inputs or model updates. This technique is particularly valuable in federated learning (FL) environments, where local model updates must remain private while contributing to global learning.

Despite its strong theoretical guarantees, DP introduces a fundamental privacy–utility trade-off. Excessive noise degrades analytical accuracy and can hinder model convergence, studies have reported accuracy reductions from 99.2% to 66.3% under high privacy budgets. Recent research suggests applying noise at the “split layer”, where data transitions from device to server, as the most effective compromise between privacy and utility.

Decentralized and Trust-Based Models

Decentralized approaches address the core limitation of traditional IoT architectures, dependence on centralized trusted entities, by distributing trust and control across the network.

Blockchain for Auditing and Data Provenance

Blockchain technology offers immutable, transparent and verifiable data management for IoT systems. It ensures auditability, enforces access accountability and establishes trustworthy provenance of sensor data. In Industrial IoT (IIoT) environments, Blockchain is often integrated with federated learning (FL) to secure

model updates and enhance resilience against tampering. This combination strengthens decentralized trust and improves the reliability of data exchange among heterogeneous nodes.

Trust-Free and Blind Protocols

Emerging trust-free architectures seek to mathematically remove reliance on trusted third parties (TTPs). Blind protocols exemplify this approach, enabling necessary services, such as data query or aggregation, to operate without revealing underlying content or identity information to the controlling server. These designs embed privacy guarantees directly into the system architecture, providing strong resistance to insider and collusion threats while supporting essential IoT functionalities.

The analysis of the 260 reviewed studies reveals a clear evolution in IoT privacy preservation, from resource-constrained cryptographic optimization toward distributed, privacy-aware intelligence. Early research primarily addressed lightweight encryption for embedded devices, while recent efforts increasingly focus on combining federated learning (FL) with differential privacy (DP) to enable privacy-preserving analytics at scale. This progression reflects the growing need to process massive, real-time IoT data streams without incurring the communication and security burdens of centralized cloud processing.

The research activity is concentrated in domains handling sensitive or mission-critical information, particularly smart homes, medical IoT (IoMT/E-Health) and industrial IoT (IIoT):

- **IoMT/E-Health:** Prioritizes ultra-lightweight cryptographic solutions, such as ASCON optimized through hardware-oriented scheduling algorithms, to ensure confidentiality and device availability in battery-limited medical systems.
- **IIoT:** Emphasizes decentralized architectures combining blockchain and FL to guarantee data provenance, integrity and trust across complex industrial supply chains.
- **Smart cities and transportation:** Continue to face persistent challenges related to location privacy and secure communication, driven by extensive mobility and heterogeneous device interactions.

Across all studies, a recurring challenge involves balancing privacy protection, computational performance and data utility. Table 3 summarizes the comparative trade-offs among key PET categories. Techniques offering higher confidentiality

(FHE) tend to incur prohibitive overheads, while those optimizing efficiency (LWC) provide weaker resistance to sophisticated attacks. Similarly, differential privacy, despite ensuring statistical anonymity, must negotiate the privacy–utility boundary, as excessive noise directly reduces model accuracy and limits usability.

Table 3. Performance and Trade-Off Comparison

PET Category	Example Technology	Primary Goal(s)	Layer Suitability	Key Limitation / Trade-Off
Lightweight Cryptography	ASCON v1.2, SPECK, PRINCE	Confidentiality, Integrity, Availability	Edge Sensing	Requires continual optimization to balance / energy efficiency and cryptographic strength; moderate speed trade-offs.
Homomorphic Encryption (HE)	FHE, AHE schemes	Confidentiality, Computation on Encrypted Data	Cloud / Fog on (AHE feasible at Edge)	High computational overhead, latency and ciphertext expansion; unsuitable for real-time use.
Statistical Models	Differential Privacy (DP)	Anonymity, Unlinkability	Edge / Fog Aggregation	Privacy-Utility Trade-off; excessive noise degrades model accuracy and convergence.
Decentralized Infrastructure	Blockchain, Federated Learning (FL)	Data Provenance, Trust Management	Fog / Cloud	Communication bottlenecks, scalability challenges and heterogeneous device constraints.

3.1. Research Gaps

Despite significant advances, several systemic challenges persist across the IoT privacy landscape as presented in Table 4:

- **Technical and Performance Constraints:** Fully homomorphic encryption remains computationally prohibitive for latency-sensitive applications and sustainability metrics such as residual energy or network scalability receive limited attention in prototype evaluations.
- **Architectural Trust Deficits:** Many privacy-preserving machine learning (PPML) frameworks continue to depend on centralized servers for

aggregation or key management, creating vulnerability to collusion and undermining the decentralization ideal.

- **Neglected Privacy Dimensions:** Device identity privacy remains insufficiently addressed, even as IoT security frameworks increasingly rely on persistent identifiers under zero-trust architectures. This tension between identifiability (for security) and non-identifiability (for privacy) introduces a critical design paradox.

Table 4. Key Research Gaps and Deployment Challenges

Challenge Category	Representative Gaps and Barriers
Technical & Performance Barriers	• Computational overhead of FHE and complex encryption schemes. • Limited attention to sustainability metrics (energy, scalability).
Architectural Trust Deficits	• Continued dependence on centralized TTPs and aggregation servers. • Susceptibility to collusion or insider compromise.
Neglected Privacy Dimensions	• Inadequate protection of device identity attributes. • Conflict between identity persistence and anonymity requirements.

The literature indicates a clear progression from lightweight device-level encryption toward holistic, decentralized and privacy-aware architectures. However, achieving a sustainable balance between privacy strength, computational feasibility and data utility remains an open research frontier. Addressing these interconnected challenges is critical to realizing scalable, trustworthy and privacy-preserving IoT ecosystems.

4. Future Research

The findings of this work reveal that while significant progress has been made in developing privacy preservation techniques (PPTs) for the internet of things (IoT), several critical frontiers remain underexplored. Future research must not only address the persistent trade-offs between privacy strength, computational efficiency and data utility but also pursue holistic, interoperable and explainable privacy-preserving ecosystems. Three overarching directions are outlined below.

The next phase of IoT privacy research should focus on the design of hybrid privacy-enhancing technologies (PETs) that integrate the complementary strengths of cryptographic, anonymization and decentralized models while mitigating their respective weaknesses. For example, the fusion of federated learning (FL) and

differential privacy (DP) with lightweight encryption (LWC) can create scalable, edge-optimized frameworks that preserve data locality while safeguarding confidentiality and anonymity.

To achieve this, researchers must explore adaptive cryptographic schemes capable of dynamically adjusting security parameters according to device capability and network context. Additionally, the development of resource-aware PET orchestration frameworks, that automatically balance privacy and latency constraints, will be essential to realize privacy-preserving intelligence at the Edge layer.

Another crucial direction involves optimizing PETs for high-dimensional big data environments (HDBDE). Emerging IoT ecosystems, such as autonomous transportation and industrial monitoring, generate large, multi-modal datasets requiring efficient privacy-preserving analytics. Integrating advanced optimization algorithms such as the sine cosine swarm optimization (SCSO) for feature selection with learning ensembles like the temporal convolutional networks (TCN), mean absolute error (MAE) models and XGBoost may help maintain analytical efficacy while enforcing privacy constraints. Such hybrid PETs will enable real-time processing of complex data streams without compromising performance or security.

As artificial intelligence (AI) increasingly underpins privacy management systems, through mechanisms like anomaly detection, noise injection, or adversarial defense, the opacity of these models introduces a new dimension of trust and accountability challenges. The application of explainable AI (XAI) to privacy-preserving systems represents a pivotal future direction.

Integrating XAI principles into privacy architectures can provide interpretable and verifiable insights into how decisions about data access, anonymization, or encryption are made. Such explainability can demystify the operation of complex AI-driven PETs and foster user and stakeholder confidence, especially in safety-critical domains such as healthcare and industrial automation. Moreover, explainable privacy mechanisms will be crucial for ensuring regulatory transparency, aligning technical outcomes with emerging AI ethics and data governance frameworks.

The convergence of XAI and PETs thus presents a dual opportunity: enhancing both the technical robustness and perceived legitimacy of privacy-preserving IoT systems.

Beyond technological innovation, the future of IoT privacy depends on the establishment of unified regulatory and design standards. The current fragmentation of privacy protocols across hardware manufacturers and service providers undermines interoperability and trust, creating systemic vulnerabilities. Future policy and standardization efforts must act as technical accelerators, mandating comprehensive privacy-by-design (PbD) frameworks that embed privacy preservation into every stage of IoT system development.

A key step in this process is the global harmonization of privacy assessment methodologies. Future frameworks should be both domain-agnostic and storage-agnostic, capable of evaluating privacy exposure based on data behavior and contextual risk rather than rigid platform specifications. The development of metrics such as a privacy exposure index (PEI), which quantifies system-level privacy risk, could serve as a universal benchmark for assessing compliance and guiding policy enforcement.

By grounding privacy assessment in measurable, technology-neutral criteria, these standardized frameworks would provide legal certainty, technical interoperability and consumer trust necessary for the safe deployment of next-generation IoT ecosystems.

Finally, future research must recognize sustainability and trust as foundational dimensions of privacy preservation. Lightweight, energy-efficient cryptographic schemes and scalable decentralized models should be evaluated not only for security resilience but also for their long-term environmental and operational sustainability. Integrating trust management systems with sustainability-aware metrics, such as residual node energy, communication efficiency and adaptive trust scoring, can lead to more resilient and socially responsible IoT infrastructures.

Developing socio-technical frameworks that balance privacy, transparency and sustainability will be vital for advancing global digital trust. Ultimately, privacy preservation must evolve from a technical add-on to a core design philosophy, underpinning the ethical, secure and sustainable growth of the IoT ecosystem.

5. Conclusion

This work provides a comprehensive synthesis of privacy preservation research in the internet of things (IoT), systematically analyzing 260 primary studies published between 2010 and 2024. The study introduces a unified taxonomy of privacy preservation techniques (PPTs) that classifies existing approaches into three core

methodological pillars: cryptographic-based, anonymization and perturbation-based and decentralized trust-based models. This taxonomy offers a consolidated framework for understanding the diversity of privacy-enhancing technologies (PETs) deployed across IoT architectures and domains.

The findings reveal an evolving research trajectory, from early emphasis on lightweight cryptography for constrained devices to a growing focus on distributed intelligence and hybrid privacy mechanisms, particularly those integrating federated learning (FL) and differential privacy (DP). Despite these advances, persistent challenges remain, including computational inefficiency, architectural trust dependencies and the underrepresentation of device identity privacy within IoT privacy models. These gaps highlight the need for more holistic, sustainable and explainable approaches.

The review further identifies three key imperatives for future research and practice:

- Advancing hybrid and adaptive PET architectures that balance privacy robustness with system efficiency.
- Embedding explainable AI (XAI) to enhance transparency, interpretability and trust in AI-driven privacy systems.
- Promoting global standardization and privacy-by-design (PbD) frameworks to harmonize technical and regulatory efforts across industries.

In conclusion, privacy preservation in IoT must evolve beyond isolated technical safeguards toward integrated, interoperable and trust-centric ecosystems. Achieving this vision will require sustained collaboration between researchers, engineers, policymakers and industry practitioners to ensure that IoT technologies remain both innovative and ethically responsible in their protection of user and device privacy.

References

- Adam, M., Hammoudeh, M., Alrawashdeh, R., & Alsulaimy, B. (2024). A survey on security, privacy, trust, and architectural challenges in IoT systems. *IEEE Access*, 12, 57128-57149.
- Alaba, F. A. (2024). In Internet of Things: A Case Study in Africa. In *The Evolution of the IoT* (pp. 1-18). Springer Nature Switzerland.
- Alferidah, D. K., & Jhanjhi, N. Z. (2020). A review on security and privacy issues and challenges in internet of things. *International Journal of Computer Science and Network Security IJCSNS*, 20(4), 263-286.

- Alshboul, Y., Bsoul, A. A. R., Al Zamil, M., & Samarah, S. (2021). Cybersecurity of smart home systems: Sensor identity protection. *Journal of Network and Systems Management*, 29(3), 22.
- Arif, M., Maya, J. A., Anandan, N., Pérez, D. A., Tonello, A. M., Zangl, H., & Rinner, B. (2024). Resource-efficient ubiquitous sensor networks for smart agriculture: A survey. *IEEE Access*, 12, 193332-193364.
- Balboni, P., Bella, G., Capparelli, F., & Barata, M. T. (2025). Privacy-enhancing technologies. In *Computer and Information Security Handbook* (pp. 891-905). Morgan Kaufmann.
- Bughio, K. S., Cook, D. M., & Shah, A. (2024). Investigating the intersections of vulnerability detection and Internet of Medical Things (IoMT) in healthcare, a scoping review protocol for Remote Patient Monitoring. *Research, Society and Development*, 13(6), e11313646080-e11313646080.
- Chang, K. C., Niu, H., Kim, B., & Barber, S. (2024). Iot privacy risks revealed. *Entropy*, 26(7), 561.
- Chereja, I., Erdei, R., Delinschi, D., Pasca, E., Avram, A., & Matei, O. (2025). Privacy-conductive data ecosystem architecture: by-design vulnerability assessment using privacy risk expansion factor and privacy exposure index. *Sensors*, 25(11), 3554.
- Demertzi, V., Demertzis, S., & Demertzis, K. (2023). An overview of privacy dimensions on the industrial internet of things (IIoT). *Algorithms*, 16(8), 378.
- Dib, M., Torabi, S., Bou-Harb, E., & Assi, C. (2021). A multi-dimensional deep learning framework for IoT malware classification and family attribution. *IEEE Transactions on Network and Service Management*, 18(2).
- Dritsas, E., & Trigka, M. (2025). Federated learning for IoT: A survey of techniques, challenges, and applications. *Journal of Sensor and Actuator Networks*, 14(1), 9.
- Gajjar, N. T. (2024). Data privacy and protection in the digital age: Emerging trends and technologies. *International Journal of Engineering Applied Science and Management*, 5(4), 1-4.
- Gupta, B. B., & Quamara, M. (2020). An overview of Internet of Things (IoT): Architectural aspects, challenges, and protocols. *Concurrency and computation: Practice and experience*, 32(21), e4946.
- Lemieux, V. L., & Werner, J. (2024). Protecting privacy in digital records: The potential of privacy-enhancing technologies. *ACM Journal on Computing and Cultural Heritage*, 16(4), 1-18.
- Liu, J., Chen, C., Li, Y., Sun, L., Song, Y., Zhou, J., ... & Dou, D. (2024). Enhancing trust and privacy in distributed networks: a comprehensive survey on blockchain-based federated learning. *Knowledge and Information Systems*, 66(8), 4377-4403.
- Magara, T., & Zhou, Y. (2024). Internet of things (IoT) of smart homes: privacy and security. *Journal of Electrical and Computer Engineering*, 2024(1), 7716956.
- Mrabet, H., Belguith, S., Alhomoud, A., & Jemai, A. (2020). A survey of IoT security based on a layered architecture of sensing and data analysis. *Sensors*, 20(13), 3625.
- Obaidat, M. A., Obeidat, S., Holst, J., Al Hayajneh, A., & Brown, J. (2020). A comprehensive and systematic survey on the internet of things: Security and privacy challenges, security frameworks, enabling technologies, threats, vulnerabilities and countermeasures. *Computers*, 9(2), 44.

- Oppl, S., & Sary, C. (2022). Motivating users to manage privacy concerns in cyber-physical settings—A design science approach considering self-determination theory. *Sustainability*, 14(2), 900.
- Ranaweera, P., Jurcut, A. D., & Liyanage, M. (2021). Survey on multi-access edge computing security and privacy. *IEEE Communications Surveys & Tutorials*, 23(2), 1078-1124.
- Wachter, S. (2018). Normative challenges of identification in the Internet of Things: Privacy, profiling, discrimination, and the GDPR. *Computer law & security review*, 34(3), 436-449.
- Wirani, Y., Eitiveni, I., & Sucahyo, Y. G. (2024). Framework of smart and integrated household waste management system: a systematic literature review using PRISMA. *Sustainability*, 16(12), 4898.
- Zuo, Z., Watson, M., Budgen, D., Hall, R., Kennelly, C., & Al Moubayed, N. (2021). Data anonymization for pervasive health care: systematic literature mapping study. *JMIR medical informatics*, 9(10), e29871.