Privacy-Aware Deep Learning for Real-Time Passenger Counting in Public Transport

ISBN: 978-1-943295-26-5

Carlos Alberto Faria Siva
Edson Gabriel Alves De Jesus
Franklin Silva Rocha
FATEC SBC
(carlos.silva473@fatec.sp.gov.br)
(edson.jesus01@fatec.sp.gov.br)
(franklin.rocha@fatec.sp.gov.br)

William Lopes
Marcelo T Okano
Paulista University and Fatec Sbc
(wilnatelha@gmail.com)
(prof.okano@gmail.com)

This research proposes an AI-based passenger counting system for urban buses using deep learning with YOLOv8 and edge IoT devices. The embedded hardware performs real-time image processing to estimate seated and standing passengers while transmitting only anonymized metadata to a central server. To comply with Brazil's General Data Protection Law (LGPD), the system implements privacy-by-design principles, avoiding any storage of personal images. The project demonstrates how ethical AI deployment can balance technological efficiency with data protection, supporting transparent and responsible smart mobility management.

Keywords: Privacy by Design, General Data Protection Law, Edge Computing, Artificial Intelligence (AI), Intelligent Transport Systems (ITS)

1. Introduction

The digitalization of urban public transport systems has profoundly transformed the way data are collected, processed, and used to support city management. Technologies such as the Internet of Things (IoT), Edge Computing, and Artificial Intelligence (AI) now integrate cameras, sensors, and embedded systems in buses and subways, enabling real-time monitoring of passenger flow, occupancy, and operational efficiency (Xu et al., 2023; Fernández et al., 2024).

However, this technological evolution brings new legal and ethical challenges related to personal data protection. Images captured by cameras, geolocation data, and mobility patterns can all be used to identify individuals, making it essential that their processing complies with privacy regulations such as Brazil's General Data Protection Law (LGPD – Law No. 13,709/2018) (Brasil, 2018).

The LGPD establishes key principles of purpose, necessity, transparency, and security, requiring that any technology handling citizens' data adopt privacy-by-design and data minimization practices from the outset (EDPB, 2020; ISO/IEC 29100:2024). In the context of public transport, these principles are especially critical, since embedded monitoring and ticketing systems process large volumes of sensitive data daily, including facial images, travel routes, and behavioral patterns (Frota de Araújo, 2023).

Implementing LGPD compliance in the transport sector, however, presents practical difficulties. Among them are the lack of clear data governance policies, insufficient technical and legal training among public and private operators, and the inherent challenges of real-time video anonymization (Araújo, 2024; Instituto Federal Catarinense, 2024). These issues call for integrated approaches combining technical, organizational, and regulatory perspectives to ensure that technological innovation advances in harmony with fundamental rights to privacy and data protection.

Accordingly, this study examines the application of the LGPD in urban public transport systems that employ artificial intelligence and IoT technologies, discussing how legal principles of privacy and data protection can be embedded into the design and implementation of intelligent mobility solutions. The research aims to explore how the privacy-by-design paradigm can be effectively applied during the development of AI-based monitoring systems, preventing the inappropriate processing of personal images and ensuring full regulatory compliance.

The paper is organized as follows: Section 2 presents the theoretical framework on LGPD foundations and its implications for intelligent transport systems; Section 3 discusses the main technical and governance challenges related to compliance; Section 4 proposes practical guidelines for implementing ethical and secure AI solutions; and Section 5 concludes with policy implications and directions for future research.

2. Theoretical Framework

The digital transformation and the integration of Artificial Intelligence (AI), Internet of Things (IoT), and embedded systems in public transport redefine the boundaries between efficiency and privacy. Under the requirements of Brazil's General Data

Protection Law (LGPD – Law No. 13.709/2018) (Brasil, 2018), public institutions and service providers must balance data-driven optimization with citizens' fundamental rights.

The theoretical background of this study develops across six interrelated perspectives:
(2.1) LGPD and data-protection foundations;
(2.2) Privacy by Design principles;
(2.3) Secure computing methods;
(2.4) Governance and ethical compliance;
(2.5) Technical architecture for privacy preservation; and
(2.6) an Integrative Perspective on responsible smart mobility.

2.1 LGPD and Data Protection in Public Transport

The LGPD establishes the principles of purpose, necessity, transparency, and security, creating legal boundaries for personal-data processing and defining the roles of controllers and processors. Its alignment with the European GDPR ensures consistency with international standards while addressing local specificities (Helfer & Rached, 2019).

In intelligent-transport environments, vast sensor networks and camera systems generate high-volume, high-velocity data streams. The central challenge is achieving functional efficiency without violating informational self-determination. As Canedo et al. (2020) note, privacy must be engineered into system design through mechanisms such as data minimization, anonymization, consent management, and auditable access control.

The Mobillens project (Okano, Faria Silva, Alves de Jesus, & Rocha, 2025) demonstrates these principles operationally. Its embedded architecture performs local image processing, discards frames immediately, and transmits only anonymized metadata. This approach exemplifies compliance with the LGPD's Article 6 on necessity and purpose limitation, converting legal mandates into concrete engineering practices.

2.2 Principles of Privacy by Design (PbD)

Privacy by Design (PbD) transforms privacy from a legal requirement into a design philosophy. Antunes and Okano (2025) describe PbD as a systemic framework integrating technical, organizational, and ethical dimensions.

Four PbD principles stand out for transport systems:

- Proactivity and prevention – anticipating risks before they occur.
- Privacy as the default setting – limiting collection to essential data.
- Privacy embedded in architecture – incorporating safeguards at the system level.
- End-to-end lifecycle protection – ensuring security from acquisition to disposal.

Cortina et al. (2019) operationalize these principles through the GDPR Process Assessment Model (PAM), which measures privacy-maturity levels. Such maturity models enable public agencies to translate PbD ideals into measurable governance indicators. PbD thus becomes a bridge between software engineering and regulatory compliance, ensuring technical credibility and institutional legitimacy (EDPB, 2020).

2.3 Privacy by Design and Secure Computing

Recent work links PbD to secure computing approaches that sustain privacy even during computation. Antunes and Okano (2025) highlight homomorphic encryption, which allows operations on encrypted data, ensuring confidentiality without decryption. This technique exemplifies the LGPD's principle of "security throughout processing."

In parallel, Arman et al. (2021) explore automated IoT data ingestion strategies that guarantee anonymization and traceability across distributed devices. When combined with PbD, these methods create a layered defense model, protecting privacy at hardware, software, and network levels.

The Mobillens case (Okano et al., 2025) applies this logic within an edge-AI architecture. By processing inference locally on Raspberry Pi hardware, it ensures that identifiable data never leave the embedded environment. This represents a shift from reactive compliance to proactive privacy assurance, where technical design enforces ethical and legal norms.

2.4 Governance and Ethical Compliance

Long-term compliance requires embedding privacy into organizational culture. Farias et al. (2021) propose the LGPD4BP model, integrating privacy controls into business process modeling to continuously assess risk and compliance maturity.

Ethical oversight is equally vital. Mattiuzzo (2021) warns that algorithmic automation may undermine human dignity if decision-making becomes opaque or dehumanized. Thus, intelligent-transport systems must guarantee explainability, human-in-the-loop validation, and fairness audits.

Governance mechanisms should therefore extend beyond legal conformity, establishing accountability frameworks, independent audits, and training programs that strengthen institutional responsibility. Such initiatives transform compliance into a continuous governance practice, reinforcing public confidence in digital infrastructures.

2.5 Technical Architecture for Privacy Preservation

The convergence of legal mandates and privacy-by-design principles requires translation into technical architecture. Three components are fundamental:

Edge processing, reducing cloud dependency and limiting data exposure.

Secure data pipelines, integrating encryption, hashing, and integrity verification.
Adaptive privacy controls, which dynamically adjust permissions according to operational context.

The Mobillens prototype (Okano et al., 2025) embodies these elements: local inference, encrypted metadata transfer, and automated data deletion. This configuration aligns with Canedo et al. (2020) and Antunes & Okano (2025) by transforming privacy into a measurable system property. Such architectures illustrate how LGPD compliance can coexist with real-time AI analytics in constrained IoT environments.

2.6 Integrative Perspective: Toward Responsible Smart Mobility

Synthesizing these dimensions reveals that legal norms, technical innovation, and ethical governance must operate synergistically. The LGPD (Brasil, 2018) provides the normative foundation; Privacy by Design offers the methodological pathway; and secure computing with edge AI delivers the technological feasibility. Together, they enable the construction of responsible mobility ecosystems where innovation and individual rights coexist.

As Antunes and Okano (2025) argue, privacy protection is most effective when embedded in organizational structures and reinforced through technical safeguards. Following Mattiuzzo (2021), ethical compliance must remain centered on human oversight and transparency.

Therefore, the integration of LGPD principles, PbD, and secure computing establishes a foundation for trustworthy AI in transportation, balancing efficiency with dignity, automation with accountability, and data innovation with the preservation of civil rights.

3. Methodology

This research adopts a Design Science Research (DSR) approach, which is particularly suited for the creation and validation of artifacts that solve real-world problems at the intersection of technology, organization, and regulation. The study aims to design a Privacy-by-Design (PbD)-based framework that enables the implementation of the Brazilian General Data Protection Law (LGPD – Law No. 13.709/2018) within intelligent transport systems employing embedded and edge computing technologies.

The methodological design integrates three analytical dimensions, legal-normative, technical-computational, and organizational-ethical, ensuring that compliance with the LGPD becomes a structural component of AI system design rather than an external constraint. This alignment between regulation and innovation underpins the concept of ethical-by-design engineering (Antunes & Okano, 2025).

3.1 **Phase 1:** Problem Identification and Contextualization

The research begins with a contextual analysis of the privacy challenges associated with data-driven public transport systems. These systems typically employ IoT devices, computer vision, and AI algorithms to collect operational data (e.g., occupancy, vehicle status, or passenger flow). However, such systems often process images that may contain identifiable personal data, which falls under the protection scope of the LGPD.

The Problem Definition Guiding this Study can be Stated as

"How can LGPD principles be embedded into the design and operation of AI-based transport systems to prevent personal data exposure and ensure ethical compliance?"

A bibliographic review and legal analysis were conducted using sources from the Revista de Direito Civil Contemporâneo (Helfer & Rached, 2019), Entropy (Canedo et al., 2020), and Communications in Computer and Information Science (Cortina et al., 2019). The review identified recurring gaps between legal compliance frameworks and their practical application in AI system development.

This stage established the theoretical foundation for developing a PbD-based model capable of transforming abstract privacy principles (purpose limitation, data minimization, and accountability) into measurable design criteria.

3.2 **Phase 2:** Conceptual Design and Artifact Construction

The second phase focused on designing the artifact, a Privacy-by-Design Framework for Intelligent Transport Systems (PbD-ITS), built upon the Design Science Research (DSR) cycle of construction, demonstration, and evaluation.

The framework is composed of three integrated layers

1. Legal and Normative Layer

Defines compliance requirements based on the LGPD (Brasil, 2018) and GDPR principles, emphasizing purpose limitation, consent management, and data minimization. The PAM model (Cortina et al., 2019) and LGPD4BP (Farias et al., 2021) served as methodological references for compliance maturity assessment.

2. Technical Layer

Implements privacy-preserving technologies at the system architecture level. Drawing from Antunes and Okano (2025), the design integrates homomorphic encryption and edge inference to maintain data confidentiality during processing. Arman et al. (2021) contribute methods for automated IoT data ingestion that preserve anonymity.

3. Organizational and Ethical Layer

Embeds privacy governance and human oversight mechanisms into institutional routines. Following Mattiuzzo (2021), the artifact incorporates explainability and human-in-the-loop decision structures to prevent the dehumanization of algorithmic systems.

The framework translates legal obligations into operational system properties, ensuring that each development stage, from data acquisition to model deployment, adheres to PbD and LGPD principles.

3.3 Phase 3: Validation through Case Study – The Mobillens Prototype

Validation was carried out through an empirical application using the Mobillens project (Okano, Faria Silva, Alves de Jesus, & Rocha, 2025) as a real-world case study. The project implements embedded edge computing on Raspberry Pi hardware to perform local AI inference for passenger detection and counting while transmitting only non-identifiable metadata.

The validation employed three evaluation dimensions

- Legal compliance, verified by the absence of any personal image storage or transmission, fulfilling the LGPD's principles of necessity and data minimization (Brasil, 2018).
- Technical efficiency, measured through the system's ability to perform real-time inference on embedded devices without external cloud processing, following the PbD approach (Antunes & Okano, 2025).
- Governance maturity, assessed using the PAM and LGPD4BP indicators (Cortina et al., 2019; Farias et al., 2021) to evaluate the traceability and auditability of anonymization processes.

The results demonstrated that privacy can be integrated as a design property, enabling compliance not only through policies but through technological infrastructure itself.

3.4 Phase 4: Evaluation and Knowledge Consolidation

The evaluation phase synthesized empirical findings and theoretical insights, confirming that Privacy by Design can be systematically implemented within intelligent transport systems. The study contributes both scientific knowledge (advancing PbD theory under the LGPD context) and practical knowledge (providing guidelines for engineers, policymakers, and municipal authorities).

This phase consolidated the final artifact as a Privacy-Aware Edge Framework (PAEF), characterized by:

- Local data processing via embedded AI (Raspberry Pi);
- Real-time anonymization and encrypted data transfer;
- Integrated compliance indicators aligned with PbD and LGPD principles; and
- Continuous feedback loops for auditing and governance maturity improvement.

The study confirmed that privacy and operational performance are not mutually exclusive, when embedded in architecture, legal compliance becomes a catalyst for trustworthy and ethical technological innovation.

4. Results and Discussions

4.1 Overview of Empirical Findings

The implementation of the Privacy-by-Design Framework for Intelligent Transport Systems (PbD-ITS) within the Mobillens project confirmed that LGPD compliance can be operationalized through embedded system design. The Raspberry Pi–based architecture successfully processed video streams locally using deep learning inference while ensuring that no identifiable personal data were stored or transmitted.

The system's output consisted solely of numerical and anonymized metadata, such as total seated and standing passengers, ensuring compliance with the LGPD's principles of necessity, purpose limitation, and security (Brasil, 2018). This architecture not only maintained privacy integrity but also achieved high processing performance, demonstrating that privacy preservation and efficiency are not mutually exclusive.

The implementation of homomorphic encryption (Antunes & Okano, 2025) and secure metadata transmission reinforced confidentiality across the entire data lifecycle. Furthermore, the local edge configuration reduced dependency on cloud computing, minimizing risks associated with data transfer and third-party access , a central concern discussed by Helfer and Rached (2019) in their analysis of Brazil's data protection framework.

4.2 Comparative Analysis with Theoretical Framework

The empirical results validate the conceptual pillars presented in Section 2. Specifically, the framework demonstrates that the combination of Privacy by Design (PbD), secure computing, and ethical governance provides a viable path toward implementing the LGPD in complex socio-technical systems such as public transportation.

1. Alignment with LGPD and GDPR Principles

The results confirmed that PbD principles can translate directly into legal compliance mechanisms. Following Cortina et al. (2019) and Farias et al. (2021), the implementation of audit trails, consent minimization, and anonymization procedures exemplifies the concept of embedded compliance, transforming privacy from an administrative checklist into a technical design feature.

2. Validation of PbD as a Governance Mechanism

The project validated Antunes and Okano (2025), who proposed that Privacy by Design functions as an operational governance model. In Mobillens, privacy controls were integrated into every phase of data processing, including collection, inference, transmission, and deletion.

3. Ethical Oversight and Human Dignity

The project maintained human oversight in line with Mattiuzzo (2021), ensuring that algorithmic decisions did not undermine human dignity or autonomy. The design incorporated explainable AI modules and human validation checkpoints to promote accountability and transparency.

4. Organizational Learning and Maturity

Using PAM (Cortina et al., 2019) and LGPD4BP (Farias et al., 2021) models, the assessment revealed a maturity level equivalent to proactive privacy governance (level 4). This finding demonstrates that organizations adopting PbD evolve from reactive compliance to continuous learning and adaptation, essential for sustainable governance.

5. Technical Performance and Data Integrity

Consistent with Arman et al. (2021), the automated IoT data ingestion process maintained performance levels compatible with real-time applications. The combination of edge AI and encryption-based anonymization ensured high accuracy while minimizing exposure risk, supporting Canedo et al. (2020) in emphasizing privacy as a measurable engineering property.

4.3 Discussion: Integration of Legal, Technical, and Ethical Dimensions

The integration of legal, technical, and ethical dimensions in this research shows that compliance frameworks like the LGPD can be effectively operationalized through systemic design. The Mobillens prototype embodies a privacy-first architecture that translates law into code, bridging the gap between normative intent and computational reality.

Moreover, the study supports the claim by Antunes and Okano (2025) that Privacy by Design is a multidimensional framework adaptable to both cloud and edge computing contexts. The real-world implementation proved that privacy protection can enhance, rather than hinder, innovation, security, and governance transparency.

From a broader perspective, this convergence strengthens trust in AI-enabled public systems, establishing a foundation for ethical smart mobility that aligns with Brazil’s LGPD and international data protection norms.

5. Conclusions

This research demonstrates that LGPD compliance in AI-based transport systems is not only a legal necessity but a strategic design choice that reinforces public trust and ethical accountability. By integrating Privacy by Design (PbD) principles into the architectural core of embedded systems, the study transformed privacy from a reactive concern into a proactive system property.

The Mobillens case study validated the practical feasibility of implementing the Privacy-by-Design Framework for Intelligent Transport Systems (PbD-ITS) under real operating conditions. Local edge processing, anonymized inference, and encrypted metadata transfer ensured full compliance with the LGPD’s requirements of necessity, transparency, and security, while maintaining system performance and scalability.

The study contributes to both academic and practical domains by

- Extending PbD and DSR methodologies to the context of edge computing and smart mobility;
- Demonstrating that privacy-aware AI systems can balance efficiency, governance, and human rights; and
- Providing a reference model for public agencies seeking to modernize transportation infrastructure under privacy-centric frameworks.

Future work should focus on expanding the model toward federated learning, differential privacy, and cross-platform interoperability, enhancing the capacity for privacy-preserving analytics across broader urban mobility networks.

Ultimately, this research reinforces the idea that privacy is not a limitation but an enabler of sustainable, trustworthy, and human-centered digital transformation.

6. References

1. Antunes, S. N., & Okano, M. T. (2025, August). Enhancing collaborative cloud computing security: A privacy by design approach with homomorphic encryption. In IFIP International Conference on Advances in Production Management Systems (pp. 447–461). Cham: Springer Nature Switzerland.
2. Arman, N., López-Iturri, P., Celaya-Echarri, M., Azpilicueta, L., & Falcone, F. (2021). Automating IoT data ingestion enabling visual representation. Sensors, 21(18), 6173. https://doi.org/10.3390/s21186173
3. Brasil. (2018). Lei nº 13.709, de 14 de agosto de 2018. Lei Geral de Proteção de Dados Pessoais (LGPD). Diário Official da União, Brasília, DF.

4. Canedo, E. D., Pinheiro, D. M., de Almeida, R. P., & Castro, M. S. (2020). Perceptions of ICT practitioners regarding software privacy. Entropy, 22(12), 1336. https://doi.org/10.3390/e22121336
5. Cortina, S., da Silva, M. M., de Oliveira, M. F., & Rocha, A. R. (2019). Designing a data protection process assessment model based on the GDPR. Communications in Computer and Information Science, 1101, 160–176. Springer.
6. Farias, G. B., Araújo, E., & Santoro, F. (2021). Are my business process models compliant with LGPD? The LGPD4BP method to evaluate and to model LGPD-aware business processes. ACM International Conference Proceeding Series, 199–207. https://doi.org/10.1145/3477314.3477326
7. Helfer, R., & Rached, D. (2019). Privacy in Brazil: Analysis on the new law on data protection. Revista de Direito Civil Contemporâneo, 22(3), 79–102.
8. Mattiuzzo, M. (2021). Let the algorithm decide: Is human dignity at stake? Revista Brasileira de Políticas Públicas, 11(2), 115–133. https://doi.org/10.5102/rbpp.v11i2.7324
9. Okano, M. T., Faria Silva, C. A., Alves de Jesus, E. G., & Rocha, F. S. (2025). Mobillens: Arquitetura de hardware e implementação de um sistema embarcado para aquisição de dados em transporte coletivo (Trabalho de Conclusão de Curso). FATEC São Bernardo do Campo “Adib Moisés Dib”.