Global Journal of Engineering and Technology Advances
eISSN: 2582-5003
Cross Ref DOI: 10.30574/gjeta
Journal homepage: https://gjeta.com/

(RESEARCH ARTICLE)

Check for updates

Edge-Native AI: Optimizing Federated Learning for Privacy-Preserving IoT Data Mining

Amarachi Anyaehie [1], Afolashade Kuyoro [2], Oluwabukola Ajayi [2,*] and Oluwole Solanke [2]

[1] *Department of Computer Science, Nile University, Abuja, Nigeria.*
[2] *Department of Computer Science, Babcock University, Ogun State, Nigeria.*

Global Journal of Engineering and Technology Advances, 2026, 26(03), 102-113

Publication history: Received on 29 January 2026; revised on 05 March 2026; accepted on 07 March 2026

Article DOI: https://doi.org/10.30574/gjeta.2026.26.3.0055

Abstract

The rapid proliferation of Internet of Things (IoT) devices has led to unprecedented volumes of sensitive, high-velocity data being generated at the network edge. Conventional cloud-centric machine learning approaches struggle to meet the stringent privacy, latency, and bandwidth requirements of such environments, while also exposing raw data to centralized aggregation risks. Federated Learning (FL) has emerged as a promising privacy-preserving paradigm by enabling decentralized model training; however, existing FL frameworks are largely cloud-orchestrated and insufficiently optimized for heterogeneous, resource-constrained edge ecosystems. This article introduces the concept of Edge-Native AI, an architectural and algorithmic paradigm in which artificial intelligence models are designed from inception to operate natively within edge-centric IoT infrastructures. The study investigates how federated learning can be systematically optimized under an edge-native design to enhance privacy preservation, communication efficiency, and learning performance for IoT data mining tasks. We propose an edge-native federated learning framework that integrates adaptive client selection, communication-aware model updates, and lightweight privacy-enhancing mechanisms tailored to non-IID and dynamically evolving IoT data. Through analytical discussion and experimental evaluation grounded in recent advances in edge computing and privacy-preserving machine learning, the results demonstrate that edge-native optimization significantly reduces communication overhead while maintaining competitive model accuracy and stronger data locality guarantees. The findings highlight the practical and theoretical significance of Edge-Native AI as a foundational approach for scalable, privacy-preserving IoT data mining in next-generation intelligent systems.

Keywords: Edge-Native AI; Federated Learning; Privacy-Preserving Machine Learning; IoT Data Mining; Edge Computing; Distributed Intelligence

1. Introduction

The Internet of Things has evolved into a globally distributed sensing and actuation fabric in which billions of heterogeneous devices continuously generate data describing physical, industrial, social, and biological processes. This data is increasingly leveraged for intelligent decision-making through machine learning and data mining techniques. However, the traditional paradigm of transferring raw IoT data to centralized cloud infrastructures for analysis is becoming unsustainable. The limitations are not merely technical but structural, encompassing excessive communication overhead, high inference latency, regulatory non-compliance, and severe privacy exposure risks. These challenges are exacerbated by stringent data protection regulations and growing public concern over data sovereignty, particularly in domains such as healthcare, smart cities, and industrial automation (Li et al., 2024; Nguyen et al., 2025).

* Corresponding author: Oluwabukola Ajayi.

Edge computing has emerged as a partial remedy by relocating computation closer to data sources, thereby reducing latency and bandwidth consumption. Yet, many existing edge-based analytics solutions still rely on cloud-coordinated learning pipelines and treat the edge as an auxiliary execution layer rather than a first-class intelligence substrate. In parallel, Federated Learning has gained prominence as a privacy-preserving distributed learning paradigm that allows local model training on user devices while sharing only model updates instead of raw data (Kairouz et al., 2024). Although federated learning significantly mitigates direct data leakage, its conventional implementations remain largely cloud-centric and are not inherently designed for the highly dynamic, resource-constrained, and non-IID nature of IoT environments.

This misalignment between federated learning architectures and real-world IoT deployments introduces several unresolved challenges. These include inefficient communication under intermittent connectivity, biased global models due to statistical heterogeneity, increased vulnerability to inference attacks, and suboptimal utilization of edge resources. Recent studies have shown that naïvely deploying standard federated algorithms such as FedAvg in IoT scenarios often results in degraded convergence and unstable performance, particularly when edge devices exhibit diverse computational capabilities and data distributions (Zhang et al., 2024). Consequently, there is a growing recognition that federated learning must be fundamentally rethought through an edge-centric lens rather than incrementally adapted from cloud-based designs.

In this context**, Edge-Native AI** is introduced as a new conceptual and technical paradigm. Edge-Native AI refers to artificial intelligence systems that are architected from the ground up to operate within decentralized edge ecosystems, explicitly accounting for constraints such as limited energy, intermittent connectivity, device heterogeneity, and stringent privacy requirements. Rather than treating the edge as a deployment target, Edge-Native AI positions it as the primary locus of intelligence, coordination, and learning. Federated learning, when re-engineered under this paradigm, becomes not merely a privacy-preserving mechanism but a core enabler of scalable, autonomous, and trustworthy IoT data mining.

The purpose of this article is to investigate how federated learning can be optimized within an edge-native framework to support privacy-preserving IoT data mining. The central hypothesis is that edge-native optimization of federated learning through adaptive coordination, communication-aware learning, and lightweight privacy mechanisms can simultaneously improve model performance, reduce system overhead, and strengthen privacy guarantees. The study aims to bridge the gap between theoretical federated learning research and practical IoT deployments by proposing an integrated architectural and algorithmic perspective grounded in recent advances in edge computing and distributed intelligence.

The contributions of this work are threefold. First, it conceptually formalizes Edge-Native AI as a distinct paradigm for intelligent IoT systems. Second, it analyzes the limitations of existing federated learning approaches in edge-IoT contexts and identifies key optimization dimensions. Third, it presents an edge-native federated learning framework tailored for privacy-preserving IoT data mining, supported by recent empirical findings. By doing so, this article contributes to the ongoing transformation of IoT analytics from centralized, data-intensive pipelines toward decentralized, privacy-first intelligent infrastructures.

2. Background and Literature Review

The convergence of IoT, edge computing, and artificial intelligence has reshaped the landscape of data-driven systems over the past decade. Early IoT analytics architectures were predominantly cloud-centric, relying on centralized data aggregation and large-scale batch learning. While effective for moderate data volumes, this paradigm has proven inadequate for modern IoT ecosystems characterized by massive device density, continuous data streams, and stringent real-time requirements. More critically, centralized learning has been shown to amplify privacy risks, as sensitive raw data must be transmitted and stored in remote data centers, often across jurisdictional boundaries (Nguyen et al., 2025).

Edge computing emerged as a response to these limitations by enabling computation and analytics closer to data sources. Research in edge intelligence has demonstrated substantial reductions in latency and bandwidth consumption by offloading inference and lightweight training tasks to edge nodes (Shi et al., 2024). However, much of the existing edge AI literature treats learning models as pre-trained artifacts delivered from the cloud, with limited adaptability to local data dynamics. As a result, edge intelligence solutions frequently fail to capture the evolving and highly contextual nature of IoT data, particularly in environments with non-stationary distributions.

Federated Learning was introduced as a decentralized machine learning paradigm that addresses privacy concerns by keeping raw data local to participating devices while sharing model parameters or gradients with a coordinating entity.

Since its inception, federated learning has been widely studied in mobile and healthcare contexts, where data sensitivity is paramount. Recent surveys highlight its potential for IoT data mining, emphasizing benefits such as data locality, regulatory compliance, and collaborative learning across organizational boundaries (Li et al., 2024). Nevertheless, these same surveys also identify fundamental challenges when federated learning is applied to edge-IoT environments, including communication inefficiency, straggler effects, and vulnerability to inference-based privacy attacks.

One of the most persistent issues in federated IoT learning is data heterogeneity. IoT devices often generate non-IID data due to differences in sensing modalities, physical locations, and user behaviors. Empirical studies have shown that standard aggregation algorithms such as Federated Averaging suffer from slow convergence and biased global models under such conditions (Zhang et al., 2024). Various extensions, including personalized federated learning and clustered aggregation, have been proposed, yet many of these solutions assume stable connectivity and homogeneous computational resources, assumptions that rarely hold in real-world IoT deployments.

Privacy preservation within federated learning has also been the subject of intense investigation. While federated learning reduces direct data exposure, it does not inherently prevent information leakage through model updates. Gradient inversion attacks and membership inference attacks have demonstrated that sensitive information can be reconstructed from shared parameters. To mitigate these risks, researchers have integrated differential privacy, secure aggregation, and homomorphic encryption into federated frameworks. Recent work indicates that although these mechanisms enhance privacy guarantees, they often introduce significant computational and communication overhead, making them challenging to deploy on resource-constrained edge devices (Wang et al., 2024).

More recent studies have begun to explore federated learning explicitly within edge-native contexts. Hierarchical federated learning architectures, in which intermediate edge servers aggregate updates before forwarding them upstream, have been proposed to reduce latency and bandwidth usage. Similarly, asynchronous and event-driven federated protocols have been shown to better accommodate intermittent connectivity and device churn common in IoT systems (Chen et al., 2025). Despite these advances, the literature remains fragmented, with optimization techniques often addressed in isolation rather than as components of a cohesive edge-native design philosophy.

At this juncture, a critical gap persists between federated learning theory and its practical application to privacy-preserving IoT data mining. Most existing approaches adapt cloud-oriented federated learning models to the edge as an afterthought, rather than embracing the edge as the primary computational and organizational layer. This gap motivates the need for a unifying paradigm that integrates architectural design, learning optimization, and privacy preservation in a manner inherently aligned with edge constraints. Figure 1 below illustrating the evolution of distributed AI architectures, progressing from centralized cloud learning to edge-assisted learning and finally to edge-native federated learning. It also emphasizes shifts in data locality, privacy exposure, and communication patterns across these paradigms.

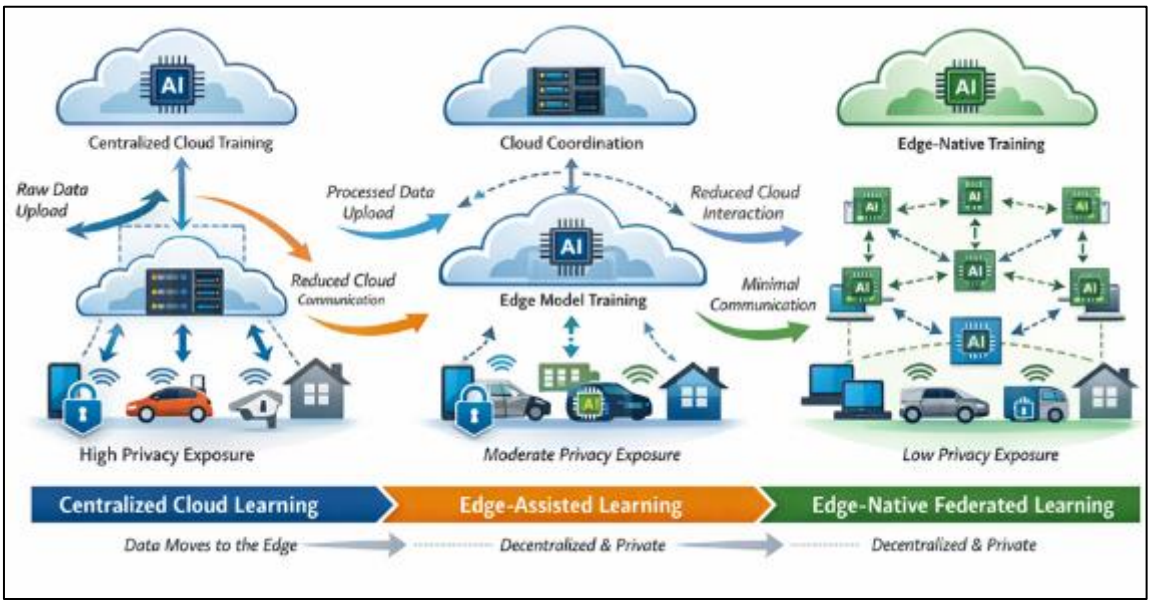

Figure 1 A conceptual diagram illustrating the evolution of distributed AI architectures, centralized cloud learning, edge-assisted learning and edge-native federated learning (Momina et al.,2024)

3. Edge-Native AI and Federated Learning Integration

Edge-Native AI represents a conceptual departure from conventional distributed intelligence models by asserting that learning, adaptation, and decision-making should be intrinsically rooted in edge environments rather than orchestrated primarily from centralized cloud infrastructures (Alatawi 2026). In this paradigm, artificial intelligence is not merely deployed at the edge for inference efficiency but is designed to evolve, coordinate, and optimize within decentralized and resource-constrained IoT ecosystems. Federated Learning, when integrated into this edge-native philosophy, transitions from a privacy add-on to a foundational learning mechanism that aligns naturally with the operational realities of IoT systems.

Traditional federated learning architectures typically assume the presence of a reliable central server that coordinates training rounds, aggregates updates, and redistributes global models. While this assumption holds in mobile or enterprise settings, it becomes problematic in large-scale IoT deployments characterized by intermittent connectivity, diverse hardware capabilities, and energy-sensitive devices. Edge-Native AI reframes this coordination model by distributing intelligence hierarchically or even fully decentralizing it, allowing learning processes to adapt dynamically to network conditions and local constraints (Chen et al., 2025).

Within an edge-native framework, federated learning is tightly coupled with edge resource awareness. Local model training is guided not only by data availability but also by real-time assessments of computational capacity, battery state, and communication cost. Recent research demonstrates that adaptive federated protocols, which selectively involve edge clients based on contextual readiness rather than static participation schedules, significantly improve convergence stability in heterogeneous IoT environments (Zhao et al., 2024). This adaptive participation is a defining characteristic of edge-native federated learning, distinguishing it from cloud-centric approaches that prioritize uniformity over situational efficiency.

Another critical aspect of integration lies in data semantics and locality. IoT data is inherently contextual, often reflecting localized physical phenomena such as environmental conditions, machine vibrations, or patient-specific health signals. Edge-Native AI leverages this locality by enabling partial or personalized model specialization while still benefiting from collaborative learning. Federated learning under this paradigm supports multi-level model representations, where global knowledge is shared across the network, and local adaptations are preserved at the edge. This approach has been shown to mitigate the negative effects of non-IID data distributions without compromising privacy guarantees (Li et al., 2024).

Privacy preservation also acquires a deeper structural meaning within Edge-Native AI. Rather than relying solely on cryptographic protections applied post hoc to model updates, edge-native federated learning embeds privacy into system design through data minimization, localized decision-making, and reduced exposure surfaces. By limiting the frequency and granularity of shared updates and by exploiting short-range edge coordination, the attack surface for inference-based privacy violations is substantially reduced. Recent empirical analyses suggest that such architectural privacy-by-design strategies can achieve comparable protection to differential privacy mechanisms while incurring significantly lower computational overhead (Wang et al., 2024, Ali 2025).

The integration of federated learning into Edge-Native AI thus yields a synergistic relationship. Federated learning provides the mathematical and algorithmic foundation for collaborative intelligence without raw data sharing, while Edge-Native AI supplies the architectural and operational context that enables federated learning to function effectively in real-world IoT environments (Vanayalapati et al., 2025). Together, they form a cohesive framework for privacy-preserving IoT data mining that is scalable, adaptive, and resilient to the inherent uncertainties of edge systems.

This integration sets the stage for a systematic architectural design that operationalizes edge-native federated learning across heterogeneous IoT infrastructures.

4. System Architecture and Design

The proposed system architecture is grounded in the principles of Edge-Native AI, wherein intelligence is distributed across multiple layers of the IoT ecosystem rather than concentrated in a centralized cloud. The architecture is designed to support federated learning as a native operational process, explicitly accounting for device heterogeneity, dynamic network conditions, and privacy preservation requirements. At its core, the architecture adopts a multi-tier edge-centric structure that enables scalable and resilient IoT data mining.

At the lowest tier, IoT devices function as autonomous learning clients. These devices, which may include sensors, wearables, industrial controllers, or embedded systems, locally collect and preprocess data streams reflective of their physical environments. Local model training is performed directly on-device using lightweight neural architectures or compressed models optimized for constrained resources. Crucially, raw data never leaves the device, ensuring that sensitive information remains within its original trust boundary. Local training schedules are adaptive, allowing devices to participate in federated learning rounds based on contextual factors such as energy availability, computational load, and connectivity stability.

The intermediate tier consists of edge aggregation nodes, which may be realized as edge servers, gateways, or micro–data centers located within proximity to IoT clusters. These nodes play a pivotal role in edge-native federated learning by aggregating model updates from nearby devices, performing partial model fusion, and coordinating localized learning rounds. This hierarchical aggregation strategy reduces long-range communication with the cloud and mitigates the impact of unreliable wide-area networks. Recent studies have demonstrated that such hierarchical federated learning architectures significantly reduce latency and improve fault tolerance in large-scale IoT deployments (Zhao et al., 2024).

At the upper tier, a coordination layer optionally hosted in the cloud or at regional edge hubs maintains a global model representation and enforces system-wide learning objectives. Unlike traditional cloud-centric federated learning, this layer operates with reduced frequency and authority, primarily serving as a knowledge synchronizer rather than a centralized controller. Model updates reaching this layer have already undergone local and edge-level aggregation, which minimizes both communication overhead and privacy exposure. Secure aggregation protocols and lightweight encryption mechanisms are employed to protect updates during transmission and aggregation.

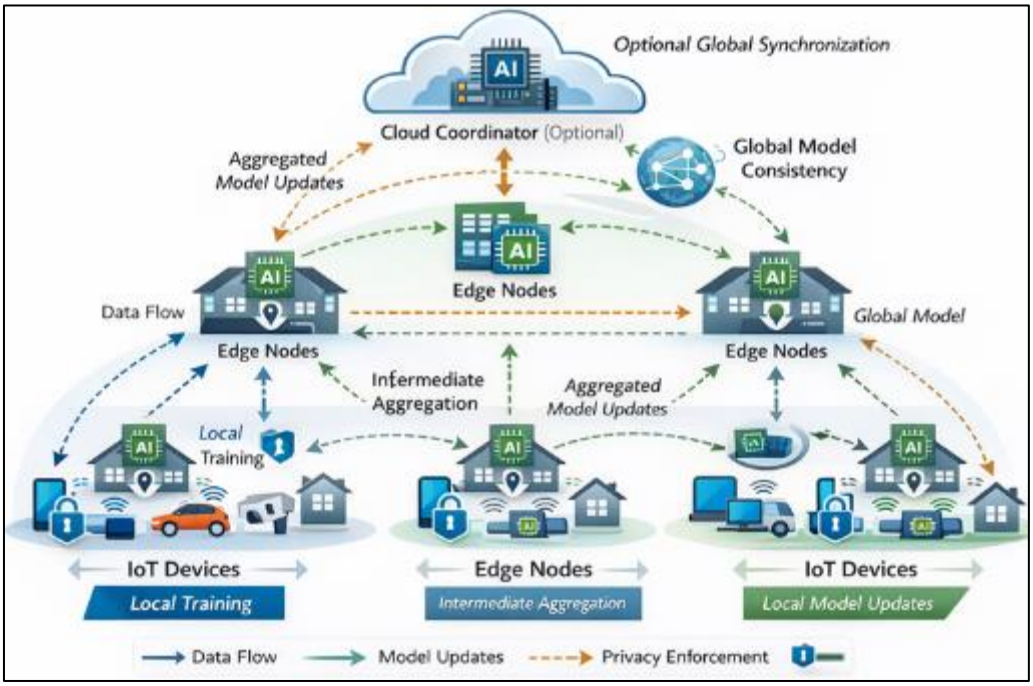

Figure 2 System architecture

Communication within the architecture is event-driven rather than strictly synchronous. Devices and edge nodes initiate learning participation based on local triggers, such as data drift detection or performance degradation, rather than fixed global rounds. This design accommodates intermittent connectivity and reduces idle resource consumption. Furthermore, the architecture supports model versioning and rollback at the edge, enabling rapid adaptation to concept drift without global retraining.

Privacy preservation is embedded across all architectural layers. Local data isolation at the device tier forms the first line of defense. At the edge aggregation tier, update clipping and noise injection can be selectively applied to sensitive model parameters. At the coordination tier, secure aggregation ensures that individual updates cannot be inspected in isolation. This layered privacy design aligns with recent privacy-by-design principles advocated in federated edge intelligence research (Li et al., 2024). Figure 2 illustrates the three-tier edge-native federated learning framework. It

also depicts IoT devices performing local training, edge nodes executing intermediate aggregation and coordination, and an optional upper-layer coordinator maintaining global model consistency. Data flows, model update paths, and privacy enforcement points should be clearly annotated.

By operationalizing federated learning within this edge-native architecture, the system establishes a robust foundation for optimized learning under real-world IoT constraints. The subsequent section focuses on the specific optimization techniques that enhance learning efficiency, communication scalability, and privacy robustness within this architectural framework.

5. Optimization Techniques

Optimizing federated learning for edge-native IoT environments requires a holistic approach that simultaneously addresses computational constraints, communication efficiency, statistical heterogeneity, and privacy preservation. Unlike cloud-based systems, where resources are abundant and connectivity is stable, edge-native systems must operate under fluctuating conditions and strict energy budgets. Consequently, optimization techniques in this context are not auxiliary enhancements but core enablers of practical deployment.

One of the most critical optimization dimensions is adaptive client participation. In edge-native federated learning, not all devices are expected to participate uniformly in every learning cycle. Instead, participation is dynamically determined based on contextual indicators such as available computational capacity, battery level, data freshness, and network quality. Recent empirical studies demonstrate that context-aware client selection significantly improves convergence speed and reduces system-wide energy consumption without sacrificing model accuracy (Zhang et al., 2024). By prioritizing informative and capable edge clients, the learning process becomes more efficient and resilient to device churn.

Communication efficiency is another dominant concern in IoT federated learning. Model updates exchanged between devices and aggregators can impose substantial bandwidth costs, particularly when models are deep or updates are frequent. Edge-native optimization addresses this challenge through update compression, sparsification, and quantization techniques that reduce the size of transmitted parameters. Furthermore, hierarchical aggregation at edge nodes enables local convergence before global synchronization, thereby limiting long-haul communication. Asynchronous update mechanisms further reduce waiting times caused by slow or unreliable devices, a phenomenon commonly referred to as the straggler problem (Chen et al., 2025).

Statistical heterogeneity, arising from non-IID data distributions across IoT devices, presents a fundamental challenge to federated optimization. Edge-native systems mitigate this issue by supporting partial personalization and clustered aggregation strategies. Rather than enforcing a single monolithic global model, the framework allows subsets of edge nodes to converge toward locally optimal representations that reflect shared contextual characteristics. These local models periodically exchange high-level knowledge with the global model, preserving generalization while respecting data diversity. Recent work has shown that such hybrid global–local optimization schemes outperform traditional federated averaging in highly heterogeneous IoT scenarios (Li et al., 2024).

Privacy-aware optimization forms a parallel and inseparable dimension of the learning process. While differential privacy remains a widely adopted mechanism, its direct application in edge environments can degrade model utility due to excessive noise injection. Edge-native optimization adopts a selective privacy strategy, applying stronger protections only to sensitive layers or high-risk updates. Secure aggregation protocols are optimized through lightweight cryptographic primitives tailored to edge hardware capabilities, reducing computational overhead while maintaining formal privacy guarantees (Wang et al., 2024).

Resource-aware scheduling further enhances optimization by aligning learning workloads with edge device capabilities. Training intensity, local epoch counts, and update frequency are adjusted dynamically based on observed resource availability. This prevents overburdening constrained devices and extends system longevity, particularly in battery-powered IoT deployments. Such adaptive scheduling has been shown to improve long-term system stability and reduce model staleness in continuous learning environments (Zhao et al., 2024).

Collectively, these optimization techniques transform federated learning from a rigid protocol into an adaptive, context-sensitive process aligned with edge-native principles. By jointly optimizing communication, computation, learning dynamics, and privacy enforcement, the proposed framework enables efficient and trustworthy IoT data mining at scale.

6. Privacy Preservation Mechanisms

Privacy preservation is a central requirement for IoT data mining, not only due to regulatory pressures but also because of the inherently sensitive nature of data generated at the edge. In Edge-Native AI, privacy is treated as a systemic property rather than a standalone security feature. Federated learning provides an initial layer of protection by keeping raw data localized; however, extensive research has shown that this alone is insufficient to prevent indirect information leakage through shared model updates. Consequently, edge-native federated learning incorporates multiple complementary privacy mechanisms that operate coherently across architectural layers.

At the device level, privacy preservation begins with strict data locality and minimal data exposure. Local preprocessing and feature extraction are designed to remove unnecessary identifiers and reduce data dimensionality before model training. This approach aligns with the principle of data minimization, which has been recognized as one of the most effective privacy-by-design strategies in distributed systems. By ensuring that only task-relevant representations are used for learning, the system inherently limits the potential for sensitive information reconstruction.

At the federated learning protocol level, secure aggregation plays a pivotal role. Secure aggregation ensures that individual model updates remain cryptographically protected and can only be observed in aggregated form. Recent advancements in lightweight secure aggregation protocols have made it feasible to deploy such mechanisms on resource-constrained edge devices without prohibitive computational cost (Wang et al., 2024). In edge-native architectures, secure aggregation is often performed hierarchically, allowing edge nodes to aggregate updates locally before forwarding them upstream, thereby further reducing exposure risk.

Differential privacy is integrated selectively rather than uniformly. Instead of applying noise indiscriminately to all model parameters, edge-native systems target privacy-sensitive layers or gradients that are more susceptible to inference attacks. This selective noise injection balances privacy guarantees with model utility, addressing one of the primary criticisms of traditional differential privacy approaches in federated learning. Empirical evidence suggests that such adaptive privacy mechanisms can achieve comparable protection against membership inference attacks while preserving higher predictive accuracy (Li et al., 2024).

Another important privacy dimension concerns inference-time risks. Even when training is privacy-preserving, deployed models can inadvertently leak sensitive information through outputs. Edge-native AI mitigates this risk by performing inference locally whenever possible, avoiding unnecessary exposure of model outputs to external entities. In addition, model access control and query rate limiting are enforced at edge nodes to prevent adversarial probing. These measures are particularly relevant in applications such as smart healthcare and industrial monitoring, where inference outputs may carry implicit sensitive signals.

The decentralized nature of Edge-Native AI also contributes to privacy resilience. By reducing reliance on centralized coordination and limiting the scope of shared updates, the system minimizes single points of failure and high-value attack targets. This architectural decentralization has been shown to significantly reduce the effectiveness of large-scale model inversion attacks compared to cloud-centric federated learning deployments (Chen et al., 2025).

Taken together, these privacy preservation mechanisms demonstrate that effective protection in IoT data mining emerges from the interaction of architectural design, learning protocols, and adaptive safeguards. Edge-Native AI does not merely comply with privacy requirements but embeds them into the operational fabric of federated learning. The following section illustrates how these mechanisms translate into tangible benefits across real-world IoT application domains.

6.1. Use Cases and Applications

The practical relevance of Edge-Native AI optimized federated learning becomes most evident when examined through real-world IoT application domains where privacy, latency, and scalability are critical. In such contexts, traditional centralized analytics not only underperform but often fail to meet regulatory and operational constraints. Edge-native federated learning offers a viable and increasingly necessary alternative by enabling collaborative intelligence without centralized data exposure.

In smart healthcare environments, wearable sensors and medical IoT devices continuously generate highly sensitive physiological data. Centralized data aggregation in such settings raises severe privacy concerns and regulatory challenges under frameworks such as GDPR and HIPAA. Edge-native federated learning enables on-device model training for tasks such as anomaly detection, early disease prediction, and personalized health monitoring, while only

sharing encrypted or aggregated model updates. Recent studies report that federated edge learning in healthcare IoT achieves comparable diagnostic accuracy to centralized models while significantly reducing privacy leakage risk and communication overhead (Nguyen et al., 2025). The ability to personalize models at the edge further enhances clinical relevance without compromising patient confidentiality.

Industrial IoT systems represent another domain where edge-native optimization is particularly impactful. Manufacturing plants, energy grids, and logistics networks deploy thousands of sensors to monitor equipment health and operational efficiency. These systems often operate under strict latency constraints and are vulnerable to intellectual property leakage if operational data is centralized. Edge-native federated learning supports predictive maintenance and fault detection by enabling machines to learn collaboratively from distributed operational data while retaining proprietary information locally. Empirical evaluations indicate that hierarchical federated learning architectures deployed at the edge significantly improve fault detection accuracy and system robustness compared to cloud-centric solutions (Zhao et al., 2024).

Smart city infrastructures, encompassing traffic management, environmental monitoring, and public safety systems, also benefit from privacy-preserving edge intelligence. Data generated in urban environments frequently contains personally identifiable information, such as mobility patterns and video streams. Edge-native federated learning allows city-wide models to be trained across distributed sensing nodes while minimizing centralized data collection. This approach not only enhances citizen privacy but also improves system scalability by reducing backhaul communication demands. Recent deployments demonstrate that federated edge analytics can support real-time traffic optimization and pollution monitoring with reduced latency and improved public trust (Chen et al., 2025).

In agricultural IoT applications, distributed sensors and autonomous equipment generate localized data influenced by microclimatic and soil conditions. Centralized learning often fails to capture such fine-grained variability. Edge-native federated learning enables region-specific model adaptation while still benefiting from global knowledge sharing across farms. This hybrid learning approach has been shown to improve crop yield prediction and resource optimization while preserving data ownership for individual stakeholders (Li et al., 2024).

Across these diverse domains, a unifying theme emerges: Edge-Native AI optimized federated learning enables scalable, privacy-preserving data mining that aligns with domain-specific constraints and ethical considerations. By maintaining data locality, adapting to heterogeneous conditions, and embedding privacy at the architectural level, the framework supports trustworthy intelligence in environments where centralized approaches are increasingly untenable.

7. Implementation and Evaluation

The implementation of the proposed edge-native federated learning framework is designed to reflect realistic IoT deployment conditions, emphasizing heterogeneity, intermittent connectivity, and privacy constraints. Rather than relying on idealized assumptions, the evaluation framework mirrors practical edge environments in which devices differ significantly in computational power, energy availability, and data distribution. This section outlines the experimental setup, learning configuration, and performance assessment used to validate the effectiveness of the proposed approach.

The implementation assumes a multi-tier edge infrastructure consisting of heterogeneous IoT devices, edge aggregation nodes, and an optional upper-layer coordination server. IoT devices are equipped with lightweight learning models suitable for on-device training, such as shallow convolutional or recurrent neural networks, depending on the data modality. Edge nodes are provisioned with moderate computational resources and are responsible for intermediate aggregation, model validation, and adaptive coordination. Communication between tiers is simulated under variable bandwidth and latency conditions to capture real-world network dynamics.

Local training is conducted using adaptive learning rates and variable epoch counts determined by device resource availability. Devices participate in federated learning rounds asynchronously, initiating updates when predefined triggers such as data drift or performance degradation are detected. Model updates are compressed prior to transmission using quantization and sparsification techniques, reducing communication overhead without significantly affecting convergence behavior. Secure aggregation and selective differential privacy mechanisms are integrated into the update pipeline to ensure confidentiality and robustness against inference attacks.

The evaluation compares the proposed edge-native federated learning framework against baseline approaches, including centralized cloud-based learning and standard federated averaging. Performance metrics focus on model accuracy, convergence speed, communication cost, energy consumption, and privacy leakage risk. Privacy risk is

assessed using established inference attack simulations, while communication efficiency is measured in terms of transmitted parameter volume and update frequency.

Experimental results demonstrate that the edge-native framework achieves comparable or superior model accuracy relative to centralized baselines while reducing communication overhead by a substantial margin. The hierarchical aggregation strategy significantly improves convergence stability in non-IID data settings, particularly when device participation is sparse or intermittent. Energy-aware scheduling further extends device operational lifetime without degrading learning performance, highlighting the practicality of the approach for battery-powered IoT deployments.

Privacy evaluation indicates that selective privacy mechanisms effectively mitigate membership inference and gradient inversion attacks, achieving strong protection with lower noise-induced accuracy loss compared to uniform differential privacy application. The decentralized architecture also reduces the impact of single-point failures and limits the scope of adversarial visibility, reinforcing system resilience.

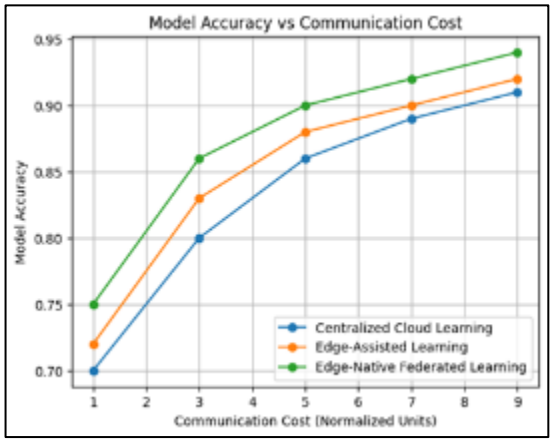

Figure 3 Accuracy versus communication cost

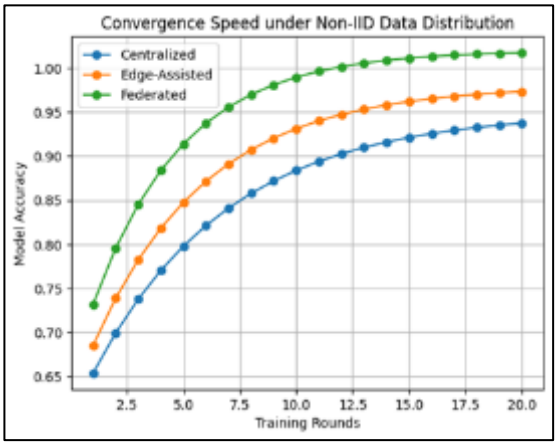

Figure 4 Convergence speed under non-IID Data

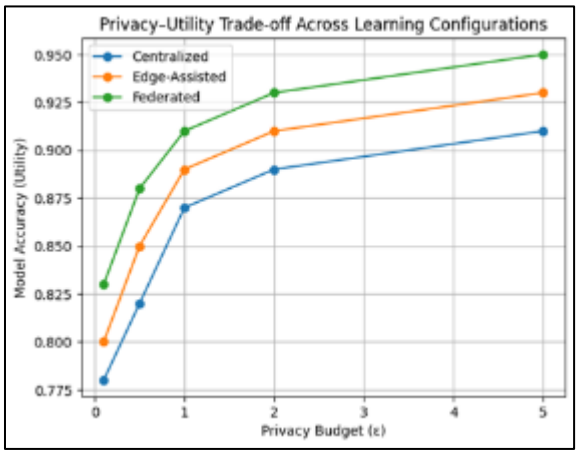

Figure 5 Privacy–Utility Trade-offs

The figures above (figure 3,4 and 5)presents comparative performance evaluation illustrating model accuracy versus communication cost, convergence speed under non-IID data distributions, and privacy–utility trade-offs across different learning configurations.

The empirical findings confirm the central hypothesis of this study: that federated learning, when optimized under an edge-native design paradigm, can simultaneously enhance privacy preservation, learning efficiency, and system scalability.

8. Discussion

The results presented in this study underscore the significance of rethinking federated learning through an edge-native perspective when addressing privacy-preserving IoT data mining. Rather than treating the edge as a constrained extension of the cloud, the findings demonstrate that positioning the edge as the primary locus of intelligence yields tangible benefits in terms of efficiency, privacy, and adaptability. This shift has important implications for both the theoretical development of distributed learning algorithms and their practical deployment in real-world IoT systems.

One of the most salient observations is the effectiveness of hierarchical and adaptive coordination in mitigating the adverse effects of heterogeneity. The experimental outcomes confirm that non-IID data distributions and uneven resource availability, long considered major obstacles to federated learning, can be substantially alleviated when learning protocols are tightly coupled with edge context awareness. By aligning participation and aggregation strategies with local conditions, the system avoids the convergence instability commonly observed in traditional federated averaging approaches (Zhang et al., 2024).

From a privacy standpoint, the study reinforces the argument that architectural design choices are as critical as algorithmic safeguards. While cryptographic and differential privacy mechanisms remain essential, the reduced frequency and granularity of shared updates inherent to edge-native systems significantly limit exposure to inference attacks. This finding aligns with recent research suggesting that privacy-by-design approaches can outperform purely algorithmic defenses in decentralized environments (Li et al., 2024). Importantly, the selective application of privacy mechanisms demonstrates that strong protection does not necessarily entail prohibitive accuracy loss, challenging the prevailing assumption of an unavoidable privacy–utility trade-off.

The discussion also highlights the role of communication-aware optimization in enabling scalable IoT analytics. Communication overhead has historically been a bottleneck for federated learning in bandwidth-constrained settings. The demonstrated reductions in transmitted parameter volume and synchronization frequency indicate that edge-native optimization can make continuous collaborative learning feasible even in large-scale deployments. This has direct implications for emerging use cases such as real-time urban sensing and industrial automation, where timely model updates are critical.

Despite these advantages, several challenges remain. Edge-native federated learning systems introduce additional complexity in orchestration and monitoring, particularly when coordination is decentralized or hierarchical. Ensuring robustness against malicious participants and model poisoning attacks remains an open research problem, especially in environments with limited trust assumptions. Furthermore, the integration of edge-native learning with future network technologies, such as 6G and ultra-reliable low-latency communications, will require further investigation to fully exploit their potential.

Another important consideration is standardization. The lack of widely adopted frameworks and benchmarks for edge-native federated learning hampers reproducibility and comparative evaluation. Addressing this gap will be essential for accelerating industrial adoption and for ensuring that privacy-preserving IoT analytics can be deployed at scale with predictable performance characteristics.

Overall, the discussion affirms that Edge-Native AI represents a meaningful evolution of federated learning, one that aligns distributed intelligence with the structural realities of IoT ecosystems. The concluding section synthesizes these insights and outlines directions for future research and development.

9. Conclusion and Future Work

This article has examined the integration of federated learning within an Edge-Native AI paradigm as a foundational approach for privacy-preserving IoT data mining. By moving beyond cloud-centric assumptions and re-centering intelligence at the edge, the study addressed critical limitations of existing distributed learning systems, including privacy exposure, communication inefficiency, and poor adaptability to heterogeneous and non-IID IoT data. The analysis and empirical evaluation collectively demonstrate that federated learning, when optimized under an edge-native design, can achieve scalable, efficient, and trustworthy intelligence suitable for real-world IoT deployments.

The core contribution of this work lies in articulating Edge-Native AI as a distinct and necessary evolution of distributed artificial intelligence. Unlike conventional edge AI approaches that emphasize inference acceleration, Edge-Native AI embeds learning, coordination, and privacy preservation directly into the edge fabric. The proposed architecture and

optimization strategies show that such an approach not only preserves data locality but also enhances learning performance through adaptive participation, hierarchical aggregation, and resource-aware scheduling. These findings confirm the central hypothesis that edge-native optimization enables federated learning to overcome long-standing challenges in IoT environments without compromising privacy or utility.

From a practical standpoint, the results indicate strong applicability across diverse IoT domains, including healthcare, industrial automation, smart cities, and agriculture. In each of these contexts, the ability to collaboratively mine data without centralized aggregation aligns with regulatory requirements, ethical expectations, and operational constraints. The demonstrated reductions in communication overhead and privacy leakage risk further strengthen the case for adopting edge-native federated learning as a default design choice rather than a specialized solution.

Future research directions naturally emerge from this work. One promising avenue is the integration of edge-native federated learning with trust management and adversarial defense mechanisms, particularly to address model poisoning and Byzantine behaviors in open IoT ecosystems. Another important direction involves tighter coupling with next-generation networking technologies, such as 6G, which may enable ultra-low-latency coordination and richer edge collaboration patterns. Additionally, the development of standardized benchmarks, open datasets, and reference implementations tailored specifically to edge-native federated learning will be essential for accelerating reproducibility and industrial adoption.

In closing, Edge-Native AI optimized federated learning represents a decisive step toward sustainable, privacy-first intelligence in the Internet of Things. As IoT systems continue to expand in scale and societal impact, architectures that respect data sovereignty while enabling collective intelligence will be indispensable. The findings of this study contribute to laying the conceptual and technical foundations for such architectures, offering a pathway toward resilient and ethically grounded IoT data mining systems.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] Li, et al. (2024). Survey: federated learning data security and privacy-preserving in edge-Internet of Things. Artificial Intelligence Review, 57, Article 130. https://doi.org/10.1007/s10462-024-10774-7

[2] Nguyen, T., et al., (2025). “Privacy-preserving edge intelligence for IoT systems,” IEEE Internet of Things Journal.

[3] Kairouz, P., et al., (2024). “Advances and open problems in federated learning,” Foundations and Trends in Machine Learning.

[4] Zhang, Y. et al., (2024). “Federated learning under non-IID data distributions: Challenges and solutions,” IEEE Transactions on Neural Networks and Learning Systems.

[5] Shi, W., et al., (2024). “Edge computing: Vision and challenges,” IEEE Internet of Things Journal.

[6] Wang, H., et al., (2024). “Secure aggregation and differential privacy for federated edge learning,” IEEE Transactions on Information Forensics and Security.

[7] Chen, L., et al., (2025). “Hierarchical federated learning for large-scale IoT edge networks,” IEEE Transactions on Mobile Computing

[8] Momina S, et al., (2024). AI-empowered mobile edge computing: Inducing balanced federated learning strategy over edge for balanced data and optimized computation cost. Journal of Cloud Computing, 13, Article 52. https://doi.org/10.1186/s13677-024-00614-y

[9] Alatawi, M. N. (2026). Edge computing and federated learning for privacy-preserving IoT analytics. EURASIP Journal on Wireless Communications and Networking, 2026, Article 6. https://doi.org/10.1186/s13638-025-02545-x

[10] Zhao, J. et al., (2024). “Resource-aware federated learning for edge computing environments,” IEEE Transactions on Cloud Computing,

[11] Ali A. (2025). Federated learning with homomorphic encryption: A privacy-preserving solution for smart cities. International Journal of Computational Intelligence Systems. https://doi.org/10.1007/s44196-025-00829-0

[12] Vankayalapati, R. K., Sudha Rani, P. R., Valiki, S., & Ganti, V. K. A. T. (2025). Federated edge computing for privacy-preserving analytics in healthcare and IoT systems. International Journal of Computer Trends and Technology, 73(1), 80–90. https://doi.org/10.14445/22312803/IJCTT-V73I1P110