BILINGUAL PUBLISHING GROUP

Journal of Electronic & Information Systems

https://journals.bilpubgroup.com/index.php/jeis

ARTICLE

Energy-Efficient Federated Learning at the Wireless Edge: A Survey

Christopher Aseer J Albert [1], *Ashli Paul* [1], *Chaitanya V Mahamuni* [2*], *Sophia Chavakula John* [1], *Earnest Ebenezer* [1]

[1] *Department of Computer Engineering, Agnel Charities' Fr. C. Rodrigues Institute of Technology, Navi Mumbai 400703, India*

[2] *Electronics and Telecommunication Engineering, University of Mumbai, Mumbai 400032, India*

ABSTRACT

The recent surge in edge computing and wireless connectivity has accelerated the adoption of Federated Learning (FL), a paradigm that enables privacy-preserving distributed intelligence across resource-constrained devices. However, implementing FL in practical wireless edge networks introduces significant challenges, particularly excessive energy consumption and communication overhead. This survey provides a system-level exploration of energy-efficient FL strategies, examining algorithmic advances and deployment challenges. Core techniques—such as model compression, update sparsification, and adaptive client scheduling—are analyzed with respect to their trade-offs in scalability, convergence, and long-term energy sustainability, especially under non-IID data distributions and heterogeneous device conditions. Practical insights are drawn from case studies in the Internet of Things (IoT), 5G/6G wireless ecosystems, and ultra-low-power device deployments, highlighting both limitations and optimization opportunities for real-world implementations. In addition, the survey explores emerging enablers, including blockchain-based trust frameworks, neuromorphic processors, and reinforcement learning-driven orchestration, which hold potential for achieving robust, sustainable FL in dynamic edge environments. By integrating perspectives from communication theory, distributed systems, and sustainable computing, this work delivers an interdisciplinary roadmap for the realistic deployment of energy-efficient FL in next-generation wireless systems, aiming to guide future research toward scalable, fair, and sustainable federated intelligence at the

*CORRESPONDING AUTHOR:
Chaitanya V Mahamuni, Electronics and Telecommunication Engineering, University of Mumbai, Mumbai 400032, India;
Email: chaitanyamahamuni91@gmail.com

ARTICLE INFO
Received: 5 July 2025 | Revised: 22 August 2025 | Accepted: 1 September 2025 | Published Online: 9 September 2025
DOI: https://doi.org/10.30564/jeis.v7i2.10982

CITATION
Albert, C.A., Paul, A., Mahamuni, C.V., et al., 2025. Energy-Efficient Federated Learning at the Wireless Edge: A Survey. Journal of Electronic & Information Systems. 7(2): 66–86. DOI: https://doi.org/10.30564/jeis.v7i2.10982

wireless edges.

Keywords: Federated Learning; Wireless Edge Networks; Energy Efficiency; Communication Bottlenecks; Client Heterogeneity; IoT; 5G/6G; Privacy-Energy Trade-Offs

1. Introduction

The exponential growth of intelligent connected devices, combined with ubiquitous wireless connectivity, has intensified the demand for sustainable wireless systems. In such systems, energy efficiency, carbon-conscious computation, and device longevity are no longer optional but essential requirements. Traditional machine learning pipelines, which rely on centralized data collection and model training, impose heavy energy and bandwidth costs, making them unsuitable for resource-constrained edge environments [1,2].

Federated Learning (FL) has emerged as a compelling alternative by enabling decentralized model training across edge devices while preserving privacy and reducing raw data transmission [3]. Its sustainability promise lies in supporting collaborative intelligence in energy-limited ecosystems. However, real-world deployment of FL in wireless edge networks remains challenging. Issues such as communication overhead, device heterogeneity, unstable channels, and limited energy reserves directly affect scalability, model accuracy, and system reliability [4,5].

Sustainable wireless edge networks refer to advanced communication infrastructures that are engineered to deliver high performance while simultaneously minimizing environmental impact and operational energy consumption. These systems leverage the capabilities of edge computing, wherein computation is performed closer to data sources, such as IoT devices, embedded sensors, and mobile terminals, to reduce reliance on centralized cloud platforms and to optimize resource utilization [6]. Key characteristics of these networks include distributed processing, dynamic scalability, and low-latency communication. Sustainability, in this context, is typically evaluated using metrics such as energy consumption per operation, carbon footprint, and the operational lifespan of devices and networks [7]. Recent studies, such as Li, L. [8], also highlight the role of green wireless networking in reducing carbon emissions across 6G ecosystems. By relocating computation to the edge, these systems reduce the energy overhead of data transmission and improve responsiveness and fault tolerance [9]. While strategies like energy harvesting, sleep scheduling, and adaptive transmission have been explored to further improve efficiency, their integration into FL remains underdeveloped [10].

This survey critically investigates the current landscape of energy-efficient FL in wireless edge environments. Unlike prior works that primarily catalog methods in isolation, our approach emphasizes the interplay between algorithmic choices and system-level constraints. In contrast to the 2023 JEIS survey on FL, which focused broadly on communication efficiency [11], this work provides a systematic and sustainability-oriented perspective, integrating carbon footprint, hardware heterogeneity, and deployment realities. Furthermore, while several surveys exist in this domain, they often overlook system-level coordination across communication, computation, and orchestration. This paper bridges that gap by examining energy efficiency holistically across layers.

Key contributions of this survey include:

i. A critical assessment of energy and communication bottlenecks that limit the scalability of FL at the wireless edge.
ii. A survey of energy-saving techniques such as model compression, communication sparsification, and asynchronous client scheduling, analyzed under practical constraints such as non-IID data and device dropout.
iii. An exploration of FL deployment challenges in low-power, intermittently connected devices, and energy-harvesting systems.
iv. A comparative synthesis of state-of-the-art approaches, evaluating their trade-offs in energy use, communication cost, and model accuracy across different deployment contexts.
v. A forward-looking analysis of unresolved challenges, including the integration of blockchain-based trust frameworks, green AI models, and neuromorphic hardware into the FL paradigm.
vi. A structured roadmap guiding the reader from background concepts through technical techniques, deployments, trade-offs, applications, benchmarking, re-

search gaps, and future directions.

The remainder of the paper is organized as follows: Section 2 introduces background concepts and related literature. Section 3 outlines communication and energy bottlenecks. Section 4 surveys energy-aware FL strategies. Section 5 examines deployment constraints in edge networks. Section 6 explores system-level optimizations. Section 7 discusses trade-offs among privacy, energy, and model performance. Section 8 presents key applications. Section 9 compares existing methods. Section 10 outlines open challenges, and Section 11 discusses future research avenues. Finally, Section 12 concludes the survey.

2. Federated Learning Overview

Federated Learning (FL) is a decentralized machine learning paradigm in which multiple edge devices participate in training a shared global model using their local data, without exchanging the raw data itself, as shown in **Figure 1**. In a typical FL setup, model parameters are exchanged between clients and a central server for aggregation, preserving data privacy and reducing communication costs.

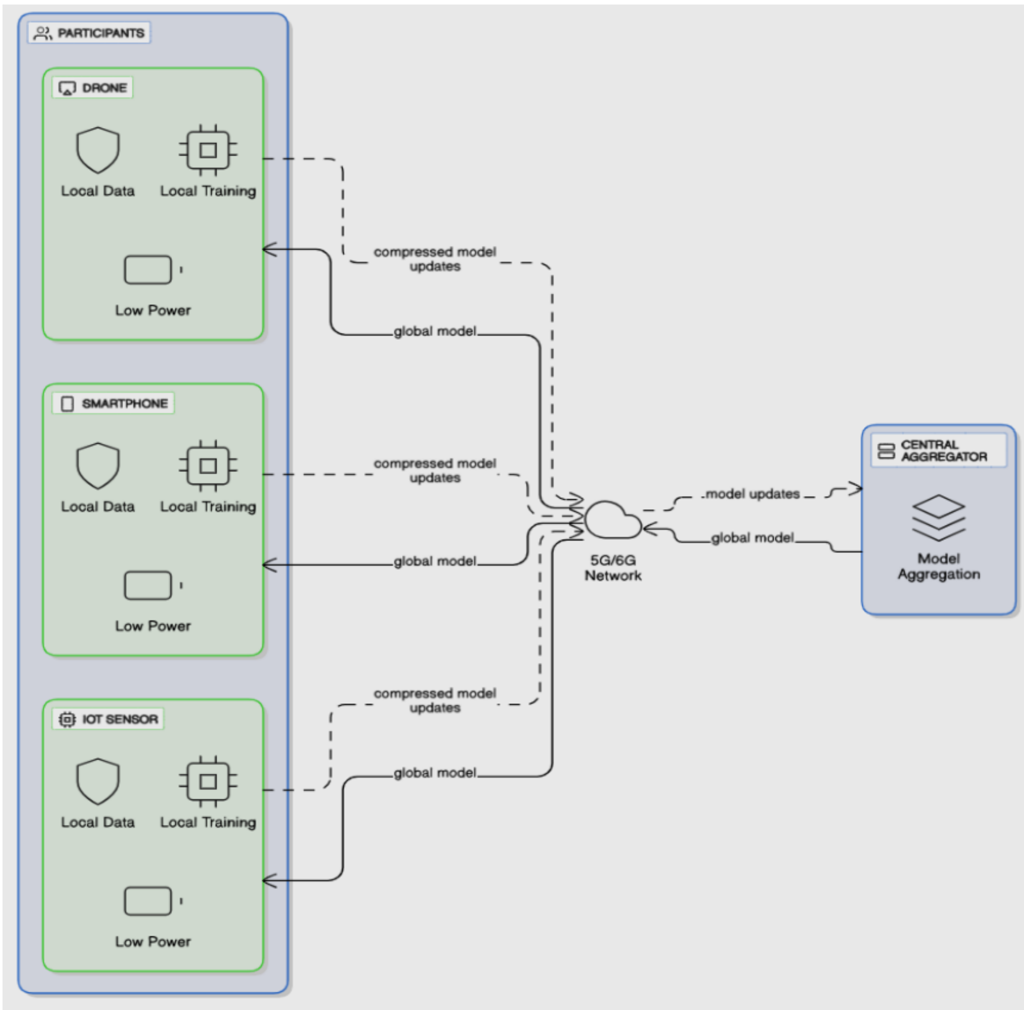

Figure 1. Federated Learning Architecture [12].

FL architectures are categorized into three primary types:

i. **Horizontal Federated Learning**: Clients share the same feature space but hold different data samples.

ii. **Vertical Federated Learning**: Clients share the same data samples but with different feature sets.

iii. **Federated Transfer Learning**: Applied when both data samples and feature spaces differ across clients [13].

FL's advantages over traditional centralized learning include improved privacy compliance, lower transmission overhead, and support for personalized model development across distributed nodes [14]. However, these benefits are often offset by several critical limitations when FL is deployed in wireless edge environments.

$$\min_{w} \sum_{i=1}^{N} p_i L_i(w) \tag{1}$$

w is the learned model parameters in this global optimization objective, *Li* is the local loss that is experienced on client *i*, and *pi* is the fraction of the training data possession by client *i*. The goal is to optimize the weighted average of local losses of all clients, and it is the foundation of federated learning.

Unreliable communication links, especially in mobile and IoT networks, result in increased latency and synchronization difficulties. Additionally, data across clients is typically non-independent and identically distributed (non-IID), which affects model convergence and generalization. The heterogeneity of edge devices, in terms of computation capacity, energy availability, and memory, further complicates client participation and degrades system-wide performance [15].

Despite a growing body of literature on FL, most methods are designed under idealized assumptions, including stable connectivity, consistent energy availability, and uniform client capabilities. These assumptions rarely hold in practical scenarios, limiting the scalability and energy efficiency of existing FL implementations. A deeper understanding of these constraints is essential to guide the development of sustainable FL techniques, which is the central focus of this survey.

The following part establishes some of the severe energy and communication bottlenecks that constrain scale-up power at the edge of FL.

Sustainability in wireless edge networks is typically evaluated using metrics such as energy consumption per operation, carbon footprint, and device lifespan [7]. To ensure comparability, carbon footprint can be quantified by converting energy usage into CO_2-equivalent emissions using standardized emission-factor models. A simple formulation is given by:

$$CO_2eq\ (g) = Energy\ (kW\text{h}) \times EF\ (gCO_2/kW\text{h})$$

where EF is the emission factor provided by datasets such as the IPCC Guidelines [16] or the ecoinvent database [8]. For example, using the U.S. EPA default emission factor (EF = 233 gCO_2/kWh), an FL training session consuming 0.5 kWh would correspond to approximately 116 gCO_2 emissions. Such quantitative measures allow sustainable FL research to align with climate accountability standards, ensuring that algorithmic efficiency gains are meaningfully tied to environmental impact.

3. Energy and Communication Bottlenecks in Federated Learning

While Federated Learning (FL) presents a decentralized alternative to conventional machine learning with strong privacy guarantees, its deployment in wireless edge networks introduces critical bottlenecks that undermine energy efficiency, scalability, and sustainability. These bottlenecks stem from both communication overhead and device-level limitations, particularly in environments with heterogeneous clients and unstable wireless channels.

3.1. Background and Foundational Concepts

A central challenge in FL is the iterative exchange of model updates between clients and the server, which stresses constrained wireless links and drains device energy. The **communication cost** per round can be expressed as:

$$C = R \times U \times S \tag{2}$$

Here, C represents the total communication cost (in bits). R is the number of communication rounds (dimensionless), U is the number of participating clients (dimensionless), and S is the size of the transmitted model update (bits per round per client). In energy-aware analyses, this communication volume is typically converted into Joules by multiplying by the transmission energy per bit, as specified by the underlying wireless protocol.

Key issues include:

i. Uplink/Downlink Imbalance: FL requires bidirectional communication, yet wireless uplinks are slower and more energy-hungry than downlinks. This leads to latency and uneven energy expenditure, especially on mobile and battery-powered devices [17].

ii. Bandwidth Saturation: As the number of participating devices grows, simultaneous transmissions congest limited spectrum resources. Without congestion-aware scheduling, retransmissions and collisions multiply energy costs [18,19].

iii. Data Heterogeneity: Non-IID client data slows convergence, increasing the required rounds *R*, which inflates communication and energy demand [20].

Although compression and communication reduction methods exist, most do not account for fluctuating wireless conditions such as packet loss, variable topology, or competing traffic.

3.2. Device and Model-Level Constraints

Beyond communication, device-level limitations introduce additional challenges for FL scalability and sustainability:

i. **Battery Depletion and Energy Scarcity**: FL participants—smartphones, IoT sensors, and wearables—operate under strict energy budgets. Extended training rapidly depletes batteries, especially when repeated participation is required. Without energy-aware client selection, resource-constrained devices are often overburdened, leading to dropouts and reduced system availability [21].

ii. **Local Resource Constraints**: Training deep models demands significant memory and compute capacity. On lightweight hardware such as microcontrollers or embedded ARM processors, this competes with critical system tasks, causing instability or a degraded user experience. While model compression and quantization offer relief, many methods ignore hardware profiling and thermal limits, restricting practical deployment [22].

iii. **Operational Trade-offs**: Devices must balance participation in FL with their primary functions (e.g., sensing, monitoring, actuation). The lack of adaptive resource allocation policies in most implementations results in inefficient energy usage and missed application-level deadlines.

Together, these communication and device-level constraints reveal fundamental design gaps in current FL systems. Addressing these requires not just algorithmic refinement but also the development of cross-layer, context-aware FL protocols that integrate network dynamics, hardware heterogeneity, and energy availability into their core scheduling and optimization logic. To combat these limitations, Section 4 examines methods that can be used to enhance the energy consumption of FL.

4. Techniques for Energy-Efficient Federated Learning

Addressing the energy and communication inefficiencies of Federated Learning (FL) in wireless edge environments requires more than isolated algorithmic adjustments—it demands techniques that adapt to heterogeneous devices, unreliable networks, and constrained energy budgets. This section critically examines three broad classes of methods: model compression, communication optimization, and adaptive training paradigms. Each is evaluated in terms of feasibility, trade-offs, and limitations in real-world deployments.

4.1. Device and Model-Level Constraints

Model compression techniques reduce the computational and transmission burden of FL by minimizing the size of models exchanged between clients and the server. Although widely studied in centralized settings, FL introduces unique challenges due to device variability and convergence stability.

$$CR=\frac{||W||_0}{||W||_1} \tag{3}$$

Here, CR denotes the compression rate of a model, $||W||_0$ is the number of non-zero weights after pruning, and $||W||_1$ is the total weight of the original model. A greater *CR* indicates greater compression, which directly improves communication efficiency in edge devices.

Key approaches include:

i. **Weight Pruning**: Eliminates redundant parameters, producing sparse models that lower transmission overhead. However, pruning often leads to non-uniform architectures across clients, complicating aggregation and slowing convergence, especially under non-IID conditions [23].

ii. **Knowledge Distillation**: Transfers knowledge from a large teacher model to a smaller student model, enabling lightweight training and inference on edge devices [24]. Recent FL-specific variants (e.g., federated knowledge distillation and TinyFedKD) address non-IID data and heterogeneous hardware, but they still rely on partial access to representative local data—a constraint in privacy-sensitive environments [25].

iii. **Low-Rank Approximation**: Decomposes weight matrices into lower-dimensional factors, significantly reducing storage and computation [26]. While effective, this method can degrade accuracy and requires computationally heavy preprocessing, making it less suitable for ultra-low-power devices.

Overall, most compression approaches remain validated only in simulations, with limited empirical studies on real low-power edge hardware. This gap raises concerns about deployment feasibility under thermal, latency, and device heterogeneity constraints.

4.2. Communication-Efficient FL

Since communication rounds dominate energy expenditure, several techniques reduce either the frequency or payload of updates exchanged between clients and the server.

i. **Update Sparsification**: Clients transmit only selected gradients (e.g., top-k values), reducing bandwidth usage [27]. While effective, sparsification may hinder convergence under non-IID data or asynchronous updates unless parameters are carefully tuned.

ii. **Event-Triggered Communication**: Updates are transmitted only when changes exceed a threshold [28]. This significantly cuts communication but introduces sensitivity parameters that require runtime tuning, often overlooked in current systems.

iii. **Local SGD (Federated Averaging)**: Clients perform multiple local updates before synchronization [29]. This reduces communication but increases divergence when data is non-overlapping, highlighting the need for regularization mechanisms.

iv. **Over-the-Air (OTA) Computation:** Both analog and digital Over-The-Air (OTA) schemes enable simultaneous transmission and aggregation of model updates over shared channels, drastically reducing communication latency. Over-The-Air (OTA) is particularly promising for large-scale FL at the wireless edge but introduces challenges in error accumulation, synchronization, and noise amplification that remain active areas of research.

Most communication-efficient techniques are tested in controlled settings, with limited evaluation under real-world conditions such as variable bandwidth, packet loss, and client dropouts.

4.3. Adaptive and Asynchronous Federated Learning

To improve resilience and energy efficiency, adaptive and asynchronous FL strategies tailor participation and update mechanisms to each client's local context. While promising in principle, these techniques introduce additional complexity and coordination challenges.

i. **Client Selection Strategies**: Energy-aware and capability-driven client selection algorithms prioritize participants based on device energy levels, compute resources, or data relevance [30]. However, such schemes can introduce systemic biases by over-utilizing high-resource clients, thereby limiting data representativeness and reducing fairness in model training.

ii. **Time and Energy-Aware Scheduling**: These methods dynamically adjust participation windows based on energy forecasts or network quality metrics [31]. Although beneficial in reducing unnecessary client strain, they often assume the availability of real-time telemetry data and prediction models, components that are rarely present in decentralized or privacy-sensitive edge deployments.

iii. **Asynchronous FL**: By removing synchronization barriers, asynchronous FL allows clients to upload updates independently, improving resource utilization and reducing idle time [32]. However, the resulting model staleness and gradient inconsistency remain unresolved challenges that affect convergence speed and final model quality, particularly in large-scale networks with variable client latencies.

Adaptive techniques offer a path toward realistic deployment but demand cross-layer integration and runtime intelligence that is still nascent in current FL frameworks. Their success hinges on the development of lightweight, decentralized schedulers that can balance global learning objectives with per-device constraints.

Taken together, these energy-aware techniques represent an evolving toolkit for sustainable federated learning. Yet, the lack of rigorous testing in real-world conditions, such as on mobile, intermittently connected, and energy-harvesting devices, underscores the need for system-level co-design and deployment-focused evaluation. The next sections examine how these strategies translate to practical implementation in wireless edge systems.

5. FL Deployment in Sustainable Wireless Edge Networks

The deployment of Federated Learning (FL) in real-world wireless edge environments extends far beyond algorithm design. It depends on the readiness of the communication infrastructure, the heterogeneity of hardware ecosystems, and the ability to align learning processes with stringent energy and resource constraints. While the emergence of 5G/6G networks and the proliferation of IoT devices provide new opportunities for scalable distributed intelligence (**Figure 2**), they also expose practical limitations that are often overlooked in academic literature.

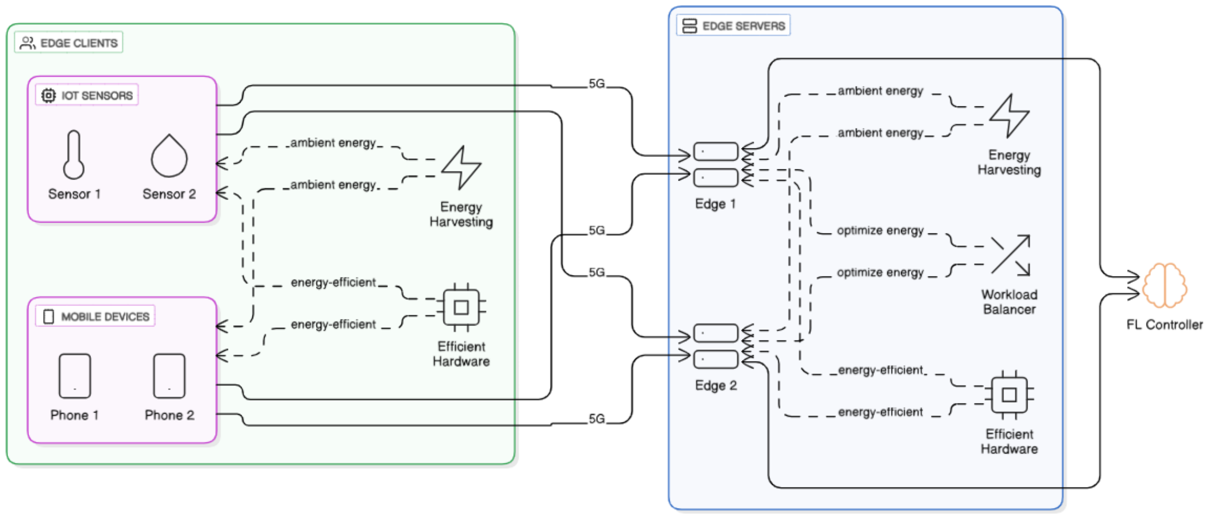

Figure 2. FL at the Wireless Edge with Energy Optimization.

5.1. 5G/6G Integration for Green FL

Next-generation mobile communication systems are frequently cited as enablers of scalable, low-latency FL deployments. However, the effectiveness of these networks in supporting energy-aware FL remains contingent on several underexplored factors.

i. **Network Slicing and Edge Intelligence**: 5G and 6G promise network slicing and Multi-access Edge Computing (MEC), which, in theory, allow FL tasks to be isolated and prioritized at the edge [33,34]. Yet, runtime adaptability is limited, and cross-layer orchestration of FL workloads within network slices remains an open challenge, particularly when competing Quality of Service (QoS) demands exist.

ii. **Energy Harvesting and Protocol Compatibility**: Cellular standards now support low-power communication and even energy harvesting features. However, these are rarely integrated with FL scheduling [35]. For example, while solar-assisted scheduling can reduce energy costs by up to **28%** in mobile sensor networks, such co-optimization has not been systematically studied in FL. Emerging 6G concepts such as ambient backscatter and AI-optimized protocols remain largely experimental, with minimal evaluation in federated contexts [36].

iii. **Deployment Gap**: Despite theoretical compatibility, real-world FL deployments over 5G/6G remain scarce. Most proposed frameworks assume ideal connectivity and predictable device behavior. In practice, mobile networks exhibit temporal variability, service fragmentation, and competing traffic loads. Without adaptive orchestration at the protocol level, the energy-saving potential of 5G/6G features such as slicing or ambient intelligence is severely constrained.

In summary, while next-generation wireless networks provide necessary building blocks for FL, leveraging them for sustainable and robust deployment will require significant innovation in orchestration, adaptive slicing, and energy-feedback integration—none of which are fully addressed by current research.

5.2. IoT and Low-Power Devices

The IoT ecosystem represents a natural setting for FL due to its scale, privacy sensitivity, and latency-critical applications. However, IoT devices introduce constraints that complicate scalability and fairness:

i. **Computational and Power Constraints**: Many IoT nodes, such as sensors, wearables, and controllers, operate with minimal processing capacity, volatile power supply, and limited memory [37]. FL's resource demands often exceed these safe operating limits. While compression and event-driven updates help, they are rarely adapted to the heterogeneous capabilities of diverse device classes.

ii. **Deployment Asymmetry and Participation Instability**: Practical deployments show significant disparities in participation due to inconsistent energy availability, intermittent connectivity, and fluctuating workloads. As a result, FL models face unstable convergence, biased updates, and stale aggregations. Dropout-resilient aggregation methods are rarely integrated into current FL frameworks [38].

iii. **Case Studies and Limitations**: Pilot deployments in smart grids, digital agriculture, and health monitoring highlight both promise and limitations. For instance, in digital agriculture, intermittent device connectivity caused up to 15% slower convergence compared to controlled testbeds. Long-tail participation, where a few devices dominate training, remains a recurring issue [39]. Real-time inference needs in applications such as patient monitoring are difficult to reconcile with FL's iterative updates.

The integration of FL into the IoT domain cannot rely on existing learning architectures alone—it demands co-design across hardware, firmware, networking, and learning protocols. Without such holistic optimization, the use of FL in resource-constrained environments risks becoming impractical or unsustainable at scale.

While 5G/6G infrastructure and IoT proliferation are often presented as natural allies of FL, their current form exposes critical gaps in coordination, resilience, and contextual adaptability. Without robust aggregation methods, energy-aware dropout handling, and adaptive orchestration, FL deployments risk instability and bias during convergence. Section 6 explores system-level optimizations designed to address these constraints.

6. Optimization Strategies for Energy-Aware Federated Learning

In wireless edge environments characterized by device heterogeneity, unstable connectivity, and strict energy constraints, optimizing Federated Learning (FL) requires more than localized algorithmic tweaks; it calls for holistic, system-level coordination. Effective FL deployment depends on intelligent orchestration across both clients and the aggregation server, aiming to minimize energy expenditure without sacrificing model convergence, fairness, or scalability. This section critically examines two primary optimization domains: energy-aware client participation and adaptive model aggregation, while interrogating the assumptions and limitations of current solutions.

6.1. Resource-Aware Client Participation

The assumption that all clients are equally capable of contributing has been invalidated by real-world deployments. FL frameworks must adapt to device-specific variability in energy, compute, and connectivity.

$$U_i = A_i - E_i \tag{4}$$

Here, U_i is the participation utility, A_i is the expected accuracy contribution of client i, E_i is its estimated energy cost, and α, β are tunable weights representing the trade-off between accuracy and energy.

i. **Dynamic Client Scheduling**: Context-aware scheduling prioritizes clients based on battery life, compute load, and link quality [40]. While this improves energy efficiency, it risks biasing model updates toward powerful clients, limiting fairness in domains like healthcare and disaster response.

ii. **Incentive-Based Participation**: Credit- or token-based systems encourage participation [41]. However, blockchain-backed incentive schemes add cryptographic overhead; for instance, lightweight blockchain consensus can consume 10–15% additional energy per round if not optimized, highlighting the trade-off between trust and sustainability[16].

Despite their promise, current strategies lack runtime telemetry feedback and fairness-aware control loops. Frequent prioritization of high-resource clients risks systemic bias and limits generalization.

6.2. Energy-Constrained Model Aggregation

While much of the optimization focus remains client-centric, server-side aggregation also plays a pivotal role in shaping the energy profile and convergence efficiency of FL systems. The challenge lies in balancing aggregation frequency and responsiveness with communication overhead and energy expenditure.

i. **Aggregation Frequency Control**: Federated Averaging reduces communication by allowing multiple local updates before synchronization[42]. Yet, static intervals often cause divergence under non-IID data. Few approaches adapt frequency dynamically based on gradient variance or learning curves.
ii. **Adaptive Aggregation Mechanisms**: Real-time adjustment of aggregation intervals based on client energy levels or loss dynamics shows promise[43]. However, such methods require telemetry reporting and lightweight controllers—capabilities not widely available in current deployments.

Moreover, server-side operations—secure aggregation, differential privacy, and validation—incur energy costs rarely quantified in the literature. Without accounting for this footprint, system-level sustainability evaluations remain incomplete.

6.3. Emerging System-Level Enablers

Beyond client and server optimizations, emerging technologies offer additional levers for sustainable FL:

i. Blockchain-Integrated FL: Blockchain provides tamper-proof logging and trust in decentralized FL, but consensus mechanisms (e.g., Proof-of-Work) are energy-intensive. Recent work on lightweight or hierarchical blockchains reduces this cost by up to 40%, making blockchain-based FL more practical for edge deployments[11].
ii. Neuromorphic Hardware: Event-driven spiking neural networks (SNNs) and neuromorphic chips such as Intel Loihi perform sparse, asynchronous updates instead of dense matrix multiplications. This reduces inference energy by 60–80% compared to GPUs in certain workloads, offering a promising substrate for energy-constrained FL training[16].
iii. 6G Slicing and Orchestration: Dynamic 6G slicing can reserve dedicated spectrum and compute resources for FL tasks. Simulation studies suggest that adaptive slicing policies can cut FL training latency by 20–25% under mixed-traffic conditions, making large-scale deployments more sustainable[16].

While client scheduling and adaptive aggregation remain central to energy-aware FL, sustainable system-level optimization demands cross-layer design. Blockchain can enhance trust, neuromorphic hardware can reduce energy costs, and 6G slicing can allocate resources intelligently. Yet, current implementations remain at the conceptual or simulation stage, with limited validation in real-world deployments. Bridging this gap requires predictive telemetry, fairness-aware scheduling, and runtime coordination across devices, protocols, and hardware platforms.

7. Security, Privacy, and Sustainability Trade-Offs

In Federated Learning (FL), security and privacy preservation are not ancillary features but foundational requirements—particularly in wireless edge environments that routinely handle sensitive, personally identifiable, or regulated data. However, the implementation of these protective mechanisms imposes significant computational, memory, and communication overheads, which can directly undermine the energy efficiency and scalability goals of sustainable FL. Despite an expanding body of work on privacy-preserving FL, relatively few studies address the inherent trade-offs between strong security guarantees and constrained device capabilities. This section critically examines these tensions and outlines the design limitations of current solutions.

7.1. Privacy-Preserving FL Mechanisms

FL systems typically rely on one or more privacy-preserving techniques to limit information leakage during local training and model exchange. Among the most commonly employed are differential privacy (DP) and secure aggregation protocols, both of which introduce additional energy and performance costs that are rarely quantified in system-level evaluations.

i. Differential Privacy (DP). By injecting calibrated statistical noise into model updates, DP prevents the reverse engineering of client data [44]. However, its implementation introduces non-negligible computational overhead, particularly when noise is sampled per-parameter or per-layer in deep models. More importantly, DP can impair model accuracy, especially in data-scarce or highly non-IID environments, leading to prolonged training and increased communication rounds.
ii. Secure Aggregation. This cryptographic technique enables the central server to aggregate encrypted model updates without accessing individual client contributions, thereby enhancing confidentiality and robustness against inference attacks [45]. While secure aggregation strengthens system-level trust, its execution incurs additional encryption, decryption, and key management costs, which can strain low-power devices such as wearables, sensors, or microcontrollers.

Although these mechanisms are essential for user trust and regulatory compliance, most implementations assume devices have sufficient compute and energy resources to support them—an assumption often unrealistic in edge deployments.

7.2. Energy Implications of Security Protocols

The sustainability of FL deployments is fundamentally shaped by the overheads introduced by privacy-preserving protocols. These costs often go unaccounted for in theoretical models and are inadequately addressed in system-level design.

i. Encryption and Authentication Overheads. Advanced encryption schemes—such as homomorphic encryption and secure multi-party computation—are computationally intensive and increase transmission latency, particularly with high-dimensional model updates [46]. Even lightweight authentication requires additional CPU cycles and memory bandwidth, which can translate into higher battery drain and thermal stress.
ii. Sustainability–Trust Trade-Offs. The trade-off between rigorous security and energy efficiency is increasingly recognized as a core challenge in FL. Emerging approaches advocate the use of lightweight cryptographic primitives (e.g., elliptic curve cryptography, hardware-accelerated AES) [47], selective privacy enforcement (where devices adjust privacy levels based on data sensitivity), and context-aware privacy protocols that adapt to device energy state or network conditions.

Dynamic Trade-Off Strategies. Beyond static mechanisms, there is growing interest in sensitivity-based privacy adjustment, where privacy budgets and noise levels are tuned dynamically based on the sensitivity of local data, task criticality, or residual device energy. For instance, health-monitoring devices may apply stronger privacy when handling patient-identifiable ECG data, but relax privacy (and thus save energy) when transmitting aggregated activity summaries. Similarly, adaptive privacy frameworks can lower cryptographic intensity for non-critical updates while preserving high protection for sensitive or high-risk transactions. Despite their promise, such dynamic adjustment strategies remain underexplored and lack standardized benchmarks for evaluating both energy and privacy guarantees.

7.3. Summary and Outlook

Achieving privacy-preserving yet energy-efficient FL requires integrated approaches that align cryptographic safeguards with device-level realities and application-specific risk tolerances. The path forward lies in developing lightweight, context-aware, and dynamically adaptive privacy mechanisms, validated on resource-constrained hardware and benchmarked not only by privacy strength but also by energy and latency efficiency (**Table 1**).

Table 1. Comparison between security and privacy-preserving mechanisms in FL based on the aspect of energy impact and edge deployment feasibility.

Protocol/Model	Energy Cost	Privacy Strength	Suitability for Edge Devices	Trade-Offs
Differential Privacy (DP)	High	Strong	Medium	May degrade accuracy and increase training time
Secure Aggregation	Medium–High	Strong	Low–Medium	Adds cryptographic overhead; burdens low-end devices
Lightweight ECC/AES	Low	Moderate	High	Well-suited for IoT; efficient and fast
Selective Privacy Enforcement	Low–Medium	Adaptive	High	Adjusts protection level based on device or data sensitivity
Context-Aware Privacy	Low–Medium	Adaptive	High	Dynamic adjustment using energy state or network condition
Sensitivity-Based Adjustment	Low–Medium	Adaptive–Strong	High	Dynamically tunes privacy budget vs. energy

8. Applications and Case Studies

The adoption of energy-efficient Federated Learning (FL) is expanding across several application domains where privacy, sustainability, and localized intelligence are paramount. While theoretical frameworks for FL are widely developed, their real-world implementation reveals a complex interplay between deployment constraints, system heterogeneity, and domain-specific performance requirements. This section critically examines how FL is being adapted across sectors such as smart cities, agriculture, healthcare, and disaster management, highlighting both operational benefits and persistent limitations.

8.1. Smart Cities and Infrastructure Monitoring

The vision of smart cities relies on massive deployments of interconnected sensors and edge devices for continuous monitoring of traffic, energy, pollution, and infrastructure health. In this context, FL enables decentralized training, reducing reliance on centralized cloud servers and preserving privacy. Energy-efficient FL is particularly valuable in extending the lifespan of distributed sensor nodes, many of which are deployed in hard-to-reach or power-limited environments[48].

However, scalability is limited by device diversity, unstable urban wireless networks, and non-stationary data. For example, traffic light control systems trained with FL must adapt to variable congestion patterns influenced by events or weather, while structural health monitoring systems face challenges in aggregating asynchronous updates from heterogeneous sensors[14]. Addressing these challenges will require robust scheduling, hierarchical aggregation, and fault-tolerant communication protocols.

8.2. Smart Agriculture and Precision Farming

Agriculture is well-suited for FL due to localized data generation and the reliance on solar-powered or battery-limited devices. Edge devices such as drones and soil sensors can collaboratively train yield predictors or irrigation models without transmitting raw data[49].

Energy-aware FL reduces data transfer costs and extends device lifetimes, with case studies in irrigation scheduling showing efficiency gains[50]. However, current deployments often assume uniform data availability, ignoring microclimate variability and soil diversity. To achieve sustainable agricultural FL, future systems must integrate personalized federated models and energy-adaptive participation strategies.

8.3. Healthcare IoT and Remote Monitoring

Healthcare is a high-stakes domain where FL ensures data privacy while enabling predictive health models acr0oss wearables, home sensors, and diagnostic tools. Applications include early arrhythmia detection, elderly fall prediction, and COVID-19 symptom modeling[51,52].

Yet, FL in healthcare faces acute trade-offs between privacy, accuracy, and energy consumption. Devices such as wearables must balance sensing, inference, and training within strict thermal and energy budgets. Pilot deployments show potential but lack longitudinal evaluation under patient variability, irregular sensor usage, and cross-hospital data heterogeneity. Future efforts must integrate energy-aware pruning, selective updates, and context-driven privacy mechanisms.

8.4. Disaster Management and Environmental Sensing

In disaster-prone areas, FL enables decentralized learning on edge devices such as seismic monitors, drones, and flood sensors, which can continue operating even when central infrastructure is unavailable [53].

Energy efficiency is crucial due to harsh environments and unpredictable power. For instance, wildfire-monitoring drones equipped with FL can adaptively update models based on energy levels and detection confidence, while federated flood prediction models can improve local alerts with reduced communication [54]. The main barriers are coordinated aggregation under unreliable connectivity and the scarcity of labeled training data, which highlights the importance of transfer learning and unsupervised FL approaches.

9. Comparative Analysis

To advance the design of energy-efficient Federated Learning (FL) systems, it is essential to systematically evaluate and compare existing approaches along key operational dimensions: energy consumption, communication cost, model accuracy, and deployment feasibility. As demonstrated in **Figure 3**, the performance of FL strategies varies significantly across energy and accuracy dimensions, with optimization methods like FedOpt and hybrid models offering improved energy-efficiency without compromising accuracy. This section provides a comparative analysis of prevalent techniques and critically examines the benchmarks and evaluation practices that currently shape this research space. While progress has been made across multiple fronts, a closer inspection reveals significant gaps in evaluation consistency, deployment realism, and methodological rigor.

While the following categorization (**Table 2**) offers an entry point into method comparison, it abstracts away nuanced differences in hardware performance, environmental variability, and multi-objective trade-offs—all of which are critical in real-world settings.

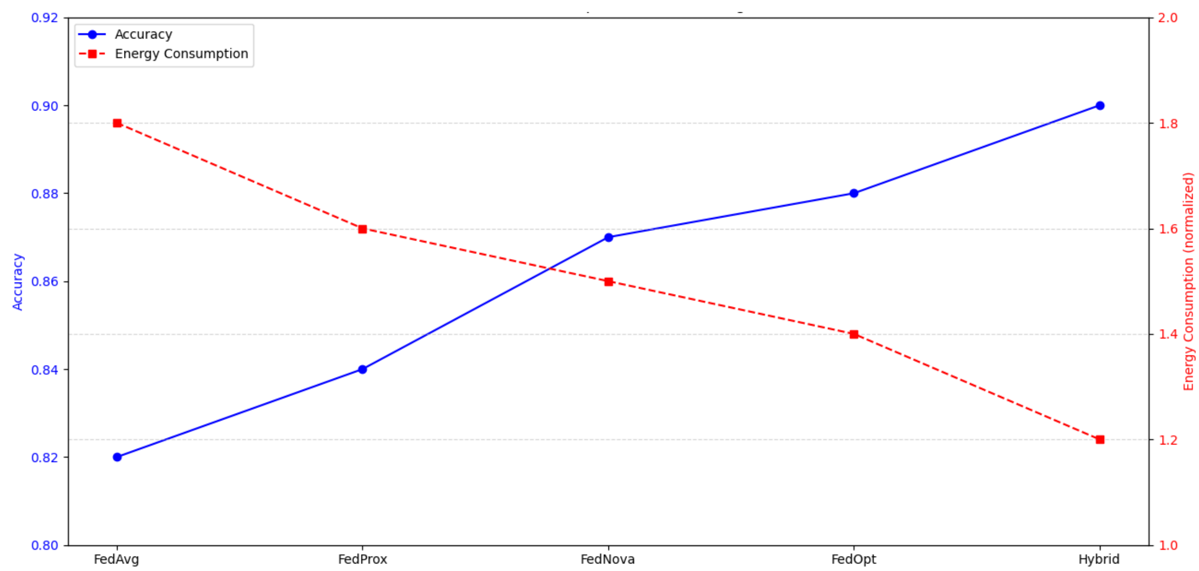

Figure 3. Accuracy and energy consumption comparison of federated learning strategies (FedAvg, FedProx, FedNova, FedOpt, Hybrid) on LEAF benchmark dataset with 100 clients.

Table 2. Summary of Energy-Efficient FL Techniques.

Technique	Energy Use	Communication Cost	Model Accuracy	Remarks
Model Pruning and Quantization	Low	Low	Moderate–High	Effective on constrained devices [55].
Update Sparsification	Medium	Very Low	Moderate	May degrade convergence on non-IID data [56].
LocalSGD (FedAvg)	Medium	Medium	High	Standard FL baseline; may diverge under data heterogeneity [57].
Adaptive Client Selection	Low–Medium	Low	High	Balances accuracy and resource allocation [58].
Event-Triggered Communication	Low	Very Low	Moderate	Best suited for non-critical updates [59].
Asynchronous FL	Medium	Low	Moderate–High	Reduces straggler effects; vulnerable to staleness [60].

9.1. Benchmarks and Evaluation Metrics

Empirical evaluation in energy-efficient FL remains largely fragmented. While many studies employ standardized datasets such as CIFAR-10, MNIST, and FEMNIST for image classification tasks, these benchmarks offer limited insight into deployment-scale variability or device heterogeneity. The LEAF benchmark [57] has emerged as a more representative alternative, simulating FL-specific challenges such as statistical and system heterogeneity. Likewise, time-series and healthcare datasets like HAR and MIMIC-III are increasingly adopted for realistic use cases.

Despite these efforts, inconsistencies persist in the choice and reporting of evaluation metrics. Commonly used performance indicators include:

i. Energy per training round (in Joules or mWh), often estimated via simulation or limited hardware profiling.
ii. Communication volume, measured in megabytes or message count, without accounting for retransmissions or network failures.
iii. Convergence time, typically tracked as epochs or communication rounds, yet often omitting variance across devices.
iv. Test accuracy or task-specific performance, usually benchmarked on IID validation data.
v. System-wide latency, device dropout rates, and availability metrics, which are crucial but rarely standardized.

Critically, these metrics are rarely integrated into a unified evaluation framework, which makes cross-study comparison unreliable and undermines reproducibility.

9.2. Limitations of Existing Works

Despite growing interest in energy-aware FL, several foundational limitations persist across the current literature:

i. **Absence of Deployment-Scale Validation**: A significant proportion of studies rely exclusively on simulations or emulators. As noted by Mohammadi et al. [61], very few methods are tested on real edge hardware with constrained power profiles, and even fewer account for environmental variability such as wireless interference or thermal throttling.
ii. **Inconsistent Baselines and Evaluation Setups**: Differences in hyperparameter tuning, model architectures, and dataset preprocessing across studies introduce inconsistencies that hinder fair performance comparisons [62]. The lack of open-source reproducibility compounds this issue.
iii. **Neglect of Cross-Layer Optimization**: Many approaches target improvements at the algorithmic level without addressing the communication-computation-learning interdependencies inherent in wireless edge systems. This siloed optimization neglects critical bottlenecks such as MAC-layer contention or energy drain from wireless transmission [63].
iv. **Oversimplification of Data Distributions**: Most works assume IID or near-IID data across clients—an assumption that rarely holds in practical applications. Techniques that appear energy-efficient under synthetic or controlled conditions often falter when faced with

real-world, skewed, or sparse client data [64].

v. **Limited Integration of Privacy and Energy Efficiency**: Security and privacy concerns, while acknowledged in isolation, are seldom addressed in tandem with energy efficiency. As Ma et al. [65] highlights, only a small fraction of FL frameworks co-optimizes for both privacy protection (e.g., secure aggregation, differential privacy) and energy constraints, leading to either compromised performance or excessive energy overheads.

Collectively, these limitations underscore the urgent need for multi-dimensional evaluation frameworks, grounded in realistic testbeds and backed by standardized benchmarks. Future work must move beyond isolated optimizations and pursue co-design approaches that account for algorithmic robustness, device heterogeneity, communication topology, and energy availability, all within an integrated FL lifecycle. Moreover, the energy used by the server-side process of secure aggregation or model validation is not measured in the majority of evaluations, or, at best, only partially. Without this, the field risks stagnating in theoretical progress without meaningful practical translation. Nonetheless, a number of open research questions continue to exist, which are discussed below.

10. Open Research Challenges

Despite notable progress in energy-efficient Federated Learning (FL), the pathway to realizing scalable, robust, and ecologically sustainable FL deployments remains obstructed by several unresolved challenges. These gaps span infrastructure limitations, algorithmic complexity, systemic disintegration across computing layers, and emerging integration hurdles with complementary technologies. This section critically outlines these open research challenges and identifies promising directions for future work.

10.1. Lack of Standardized Testbeds and Representative Datasets

One of the most critical bottlenecks in advancing the field is the absence of standardized, reproducible test environments for evaluating energy-efficient FL under realistic conditions. Most existing studies rely on simulation frameworks or synthetic datasets that fail to replicate key constraints found in actual wireless edge deployments, such as device heterogeneity, intermittent connectivity, and unpredictable energy availability [66].

The current dependence on idealized hardware settings leads to overestimation of performance and underreporting of energy overheads. Moreover, there is no consensus on how to benchmark trade-offs between accuracy, energy consumption, latency, and client fairness. This limits cross-paper comparability and impedes methodological rigor.

There is an urgent need for open-source, modular FL testbeds equipped with energy metering, wireless emulation (e.g., Wi-Fi, LoRaWAN, LTE), and diverse hardware profiles. Such testbeds must support heterogeneous data distributions and incorporate realistic dropout and failure patterns to enable system-level validation and fair benchmarking [67].

10.2. FL Under Extreme Resource Constraints

Many promising FL use cases—such as environmental sensing, rural healthcare, and disaster response—necessitate operation on ultra-constrained hardware platforms with limited memory, compute power, and power supply. Existing FL algorithms, even those optimized for edge AI, often carry overheads in communication, cryptography, and model synchronization that render them infeasible for such deployments.

Several technical challenges remain unresolved:

i. Designing ultra-lightweight model architectures capable of training and inference within kilobyte-level memory budgets.

ii. Implementing incremental or partial model updates to reduce communication and computation loads.

iii. Developing resilient aggregation strategies that handle frequent client dropout, delayed updates, and intermittent energy supply [68].

Furthermore, these scenarios are characterized by volatile environmental factors and a lack of network redundancy, further complicating the deployment of synchronous or high-frequency update schemes. Current FL frameworks rarely incorporate energy-aware fallback strategies or runtime adaptability to address these edge cases [69].

10.3. Cross-Layer Optimization Deficiencies

Energy-efficient FL remains largely segmented by research discipline, with separate communities addressing algorithmic design, network protocol engineering, and embedded systems. This siloed approach results in disjointed optimizations that fail to deliver system-wide efficiency.

Cross-layer co-design remains underexplored and presents a high-impact research frontier. Specifically, future work must address:

i. Joint optimization across communication, computation, and learning layers, considering real-time device state and network conditions.
ii. Runtime adaptation of FL workflows, where model complexity, training frequency, and transmission protocols adjust dynamically based on energy and latency feedback.
iii. Integration of FL logic with networking advancements, such as 5G/6G network slicing, cooperative scheduling, and mesh-based edge routing [70].

Without such coordination, FL deployments will continue to suffer from inefficiencies that nullify energy savings achieved at any individual layer.

10.4. Integration with Blockchain and Green AI: Trade-Offs and Barriers

The convergence of FL with blockchain technologies and green AI design principles is widely regarded as a promising route toward trustworthy and eco-conscious distributed learning. However, the integration of these technologies introduces significant systemic and energy trade-offs.

i. Blockchain for FL coordination and auditability offer transparency and trust, particularly in untrusted or multi-stakeholder environments. Yet, traditional consensus mechanisms such as Proof-of-Work (PoW) are computationally intensive and energy-prohibitive for edge deployment.
ii. There is a need for lightweight, low-latency consensus mechanisms (e.g., Proof-of-Authority, Directed Acyclic Graph (DAGs), or federated Byzantine agreement) that align with FL's energy and bandwidth constraints.
iii. Additionally, leveraging blockchain for incentive systems or secure logging introduces communication and storage overhead, which must be offset by architectural efficiency [71].
iv. Parallel efforts in green AI focus on designing inherently efficient models using pruning, quantization, and neuromorphic computing. However, their application in FL settings remains under-evaluated, particularly in scenarios with non-IID data and device variance [72].

Without rigorous co-evaluation of these integration efforts on constrained hardware, the benefits of blockchain and green AI in FL remain largely speculative. These open challenges highlight the pressing need for interdisciplinary collaboration that bridges systems engineering, hardware design, wireless networking, cryptography, and applied machine learning. Only through such integrated efforts can FL transition from controlled research environments to sustainable, large-scale deployments in real-world wireless edge networks. These challenges should be addressed creatively by highly adaptive, cross-disciplinary innovations in the future.

11. Future Research Directions

Ensuring the long-term viability of energy-efficient Federated Learning (FL) in sustainable wireless edge networks requires a shift from isolated optimizations toward intelligent, adaptive, and hardware-conscious system designs. This section outlines emerging research directions that hold the potential to transform FL into a truly self-optimizing and energy-aware framework—provided that rigorous validation, cross-disciplinary integration, and deployment realism are pursued.

11.1. Integration with Emerging 6G Communication Technologies

The future of FL at the edge is closely tied to advances in 6G communication infrastructure. Capabilities such as network slicing, intelligent reflecting surfaces, and ambient backscatter communication provide opportunities to reduce energy consumption and enhance adaptability.

In particular, network slicing allows dynamic allocation of FL tasks through slice-specific resource reservation mechanisms, where resources are assigned according to device energy budgets, latency requirements, or application priorities (e.g., healthcare vs. entertainment). Such mechanisms enable context-aware scheduling, latency-sensitive aggregation, and channel-aware model transmission, making FL adjustable in real time to network dynamics and device constraints. Future research should explore how slicing reservation policies can be co-optimized with federated schedulers to jointly minimize energy consumption and training delays constraints.

11.2. Hybridization with Adaptive Learning Paradigms

Hybridizing FL with adaptive learning paradigms such as Reinforcement Learning (RL) and Lifelong Learning offers the potential for self-organizing systems that respond to resource variability, client churn, and evolving data distributions.

i. Federated Reinforcement Learning (FRL) models FL coordination as a sequential decision-making process for client selection, communication frequency, and local training configuration [73]. However, existing approaches often assume stable reward signals and simple state spaces—assumptions that break down in dynamic wireless networks. Moreover, RL agent overhead must be carefully optimized to avoid offsetting FL's energy gains.
ii. Lifelong FL supports continuous learning across evolving data streams, mitigating catastrophic forgetting and client churn [74]. Yet, current implementations often depend on costly regularization or rehearsal strategies unsuited to energy-limited devices.

Future work should develop lightweight, communication-aware adaptive algorithms that incorporate real-time edge feedback (e.g., device battery level, connectivity status). At the same time, lightweight blockchain frameworks (e.g., Proof-of-Authority, Directed Acyclic Graph (DAG) consensus) can support incentive and trust mechanisms with reduced energy costs, enabling secure, auditable interactions in multi-stakeholder domains such as healthcare and finance services.

11.3. Hardware-Centric Innovations and Alternative Energy Sources

Achieving sustainable FL also requires innovations at the hardware and power management levels.

i. Energy-Harvesting Devices. Leveraging renewable energy sources (e.g., solar, RF, kinetic) enables devices to schedule participation opportunistically [75]. However, naive energy-aware scheduling risks model staleness. Future research should integrate predictive energy harvesting models (e.g., solar forecasting) with model compression so that compressed updates are prioritized when energy availability is low, achieving joint optimization of accuracy and sustainability [76].
ii. Neuromorphic Computing. Brain-inspired processors, such as spiking neural networks (SNNs), provide ultra-low-power, event-driven computation [12]. Empirical studies suggest up to 10× energy savings compared to GPU-based training, but their integration into FL is still immature. Future work must quantify these gains in FL- specific settings, and develop co-design frameworks that enable SNN-based clients to interoperate with conventional aggregators. Toolchains and spike-compatible aggregation processes will be crucial to unlock neuromorphic FL at scale [77].

Overall, hardware-conscious FL requires multi-objective co-design of algorithms, devices, and communication layers. Building real-world testbeds is critical, incorporating heterogeneous IoT devices, variable network conditions, and power-profiling tools to benchmark performance realistically. Frameworks such as LEAF [8] and OARF [66] provide standardized federated settings that can be extended into multi-layer, energy-sensitive benchmarking environments. By leveraging these open-source benchmarks, our survey ensures transparent and comparable evaluation across studies without requiring the release of a standalone codebase.

12. Conclusions

This survey has critically examined the emerging field of energy-efficient Federated Learning (FL) within sustainable wireless edge networks. While FL presents a promising framework for decentralized intelligence and data privacy, its deployment in energy- and resource-constrained environments introduces significant challenges that demand systemic, cross-layer solutions.

We began by analyzing the structural limitations of current wireless systems in supporting FL, identifying communication overheads, device heterogeneity, and energy scarcity as primary barriers to scalability. The survey then evaluated core strategies—such as model compression, sparsification, adaptive scheduling, and asynchronous updates—not only for their theoretical efficiency but also for their practical viability on edge hardware. Our review emphasized the trade-offs between accuracy, communication cost, and energy use, noting that benchmark studies on the LEAF dataset reveal energy reductions of up to 75% under event-triggered updates, though often at the cost of reduced accuracy and fairness.

The analysis extended to deployment contexts involving IoT, 5G/6G infrastructures, and low-power devices, where sustainability goals intersect with practical constraints. Case studies in smart cities, agriculture, healthcare, and disaster management demonstrated the applicability of FL, while exposing persistent gaps in synchronization, robustness, and fairness. Comparative synthesis of existing methods underscored the absence of unified evaluation metrics, reproducible testbeds, and cross-domain benchmarks in the literature.

Looking forward, sustainable FL requires a shift from algorithm-centric optimization toward co-designed, multi-objective systems that jointly consider hardware capabilities, renewable energy sources, network dynamics, and trust mechanisms. Promising directions include:

i. 6G-enabled network slicing to dynamically allocate resources for FL training,
ii. Neuromorphic processors spiking neural networks (SNNs) that can deliver up to 10× energy efficiency compared to GPUs,
iii. Blockchain-based trust frameworks with lightweight consensus to ensure integrity, and
iv. Green AI models that integrate model compression with energy harvesting predictions for joint optimization.

Ultimately, achieving this vision will demand interdisciplinary collaboration across embedded systems, wireless networking, machine learning, and privacy engineering. Without such integration, FL risks remaining a theoretically compelling but practically constrained solution. The future of energy-efficient FL lies in its ability to evolve into a fully adaptive, privacy-preserving, and energy-conscious ecosystem—one that is not only scalable and robust, but also aligned with the global imperative of sustainable intelligent infrastructure.

Author Contributions

Conceptualization, C.A.J.A., A.P., C.V.M., S.C.J. and E.E.; methodology, C.A.J.A., A.P., C.V.M., S.C.J. and E.E.; software, C.A.J.A., A.P. and C.V.M.; validation, C.A.J.A., A.P., C.V.M. and S.C.J.; formal analysis, C.A.J.A., A.P., C.V.M. and S.C.J.; investigation, A.P., C.A.J.A., S.C.J. and E.E.; resources, C.V.M.; data curation, A.P. and S.C.J.; writing—original draft preparation, A.P., C.A.J.A. and E.E.; writing—review and editing, C.V.M.; visualization, E.E. and A.P.; supervision, C.V.M.; project administration, C.V.M. All authors have read and agreed to the published version of the manuscript.

Funding

This research received no external funding.

Institutional Review Board Statement

Not applicable. This study did not involve humans or animals.

Informed Consent Statement

Not applicable. This research did not involve human participants.

Data Availability Statement

No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest

The authors declare no conflict of interest.

References

[1] Lorincz, J., Capone, A., Wu, J., et al., 2019. Greener, Energy-Efficient and Sustainable Networks: State-of-the-Art and New Trends. Sensors. 19(22), 4864. DOI: https://doi.org/10.3390/s19224864
[2] Murshed, M.G.S., Murphy, C., Hou, D., et al., 2021. Machine Learning at the Network Edge A Survey.

ACM Computing Surveys. 54(8), 1–37. DOI: https://doi.org/10.1145/3469029

[3] Beltrán, E.T.M., Pérez, M.Q., Sánchez, P.M.S., et al., 2023. Decentralized Federated Learning: Fundamentals, State of the Art, Frameworks, Trends, and Challenges. IEEE Communications Surveys & Tutorials. 25(4), 2983–3013. DOI: https://doi.org/10.1109/COMST.2023.3315746

[4] Cheikh, I., Roy, S., Sabir, E., et al., 2025. Energy, Scalability, Data and Security in Massive IoT: Current Landscape and Future Directions. arXiv preprint. arXiv:2505.03036. DOI: https://doi.org/10.48550/arXiv.2505.03036

[5] Yang, Z., Chen, M., Wong, K.-K., 2022. Federated Learning for 6G: Applications, Challenges, and Opportunities. Engineering. 8, 33–41. DOI: https://doi.org/10.1016/j.eng.2021.12.002

[6] Rahmani, H., Shetty, D., Wagih, M., et al., 2023. Next-Generation IoT Devices: Sustainable Eco-Friendly Manufacturing, Energy Harvesting, and Wireless Connectivity. IEEE Journal of Microwaves. 3(1), 237–255. DOI: https://doi.org/10.1109/JMW.2022.3228683

[7] López, O.L.A., Rosabal, O.M., Ruiz-Guirola, D.E, et al., 2023. Energy-Sustainable IoT Connectivity: Vision, Technological Enablers, Challenges, and Future Directions. IEEE Open Journal of the Communications Society. 4, 2609–2666. DOI: https://doi.org/10.1109/OJCOMS.2023.3323832

[8] Li, L., 2025. Research on Future 6G Green Wireless Networks. Green Technologies and Sustainability. 3(5), 100156. DOI: https://doi.org/10.1016/j.grets.2024.100156

[9] Chen, Y., Huang, S., Gan, W., et al., 2023. Federated Learning for Metaverse: A Survey. In Proceedings of the WWW '23 Companion: Companion Proceedings of the ACM Web Conference 2023, Austin, TX, USA, 30 April–4 May 2023. pp. 1151–1160. DOI: https://doi.org/10.1145/3543873.3587584

[10] Trigka, M., Dritsas, E., 2025. Edge and Cloud Computing in Smart Cities. Future Internet. 17(3), 118. DOI: https://doi.org/10.3390/fi17030118

[11] Lu, M., Fu, G., Osman, N.B., et al., 2021. Green Energy Harvesting Strategies on Edge-Based Urban Computing in Sustainable Internet of Things. Sustainable Cities and Society. 75, 103349. DOI: https://doi.org/10.1016/j.scs.2021.103349

[12] Sharma, A., Sharma, P., 2021. Energy Harvesting Technology for IoT Edge Applications. In: Kheng, T.Y. (Ed.). Smart Manufacturing-When Artificial Intelligence Meets the Internet of Things. IntechOpen: London, UK. DOI: https://doi.org/10.5772/intechopen.92565

[13] Nguyen, D.C., Ding, M., Pathirana, P.N., et al., 2021. Federated Learning for Internet of Things: A Comprehensive Survey. IEEE Communications Surveys & Tutorials. 23(3), 1622–1658. DOI: https://doi.org/10.1109/COMST.2021.3075439

[14] Banabilah, S., Aloqaily, M., Alsayed, E., et al., 2022. Federated Learning Survey: Fundamentals, Enabling Technologies, and Future Applications. Information Processing & Management. 59(6), 103061. DOI: https://doi.org/10.1016/j.ipm.2022.103061

[15] Mestoukirdi, M., 2023. Reliable and Communication-Efficient Federated Learning for Future Intelligent Edge Networks [PhD thesis]. Sorbonne Université: Paris, France. Available from: https://theses.hal.science/tel-04470748v1

[16] Zhu, J., Cao, J., Saxena, D., et al., 2023. Blockchain-Empowered Federated Learning: Challenges, Solutions, and Future Directions. ACM Computing Surveys. 55(11), 1–31. DOI: https://doi.org/10.1145/3570953

[17] Salehi, M., Hossain, E., 2021. Federated Learning in Unreliable and Resource-Constrained Cellular Wireless Networks. IEEE Transactions on Communications. 69(8), 5136–5151. DOI: https://doi.org/10.1109/TCOMM.2021.3081746

[18] Evgenidis, N.G., Mitsiou, N.A., Koutsioumpa, V.I., et al., 2024. Multiple Access in the Era of Distributed Computing and Edge Intelligence. Proceedings of the IEEE. 112(9), 1497–1526. DOI: https://doi.org/10.1109/JPROC.2024.3417528

[19] Wang, Y., Xu, Y., Shi, Q., et al., 2021. Quantized Federated Learning under Transmission Delay and Outage Constraints. IEEE Journal on Selected Areas in Communications. 40(1), 323–341. DOI: https://doi.org/10.1109/JSAC.2021.3126081

[20] Fekri, M.N., Grolinger, K., Mir, S., 2023. Asynchronous Adaptive Federated Learning for Distributed Load Forecasting with Smart Meter Data. International Journal of Electrical Power & Energy Systems. 153, 109285. DOI: https://doi.org/10.1016/j.ijepes.2023.109285

[21] Khajehali, N., Yan, S., Chow, Y.-W., et al., 2023. A Comprehensive Overview of IoT-Based Federated Learning: Focusing on Client Selection Methods. Sensors. 23(16), 7235. DOI: https://doi.org/10.3390/s23167235

[22] Jouini, O., Sethom, K., Namoun, A., et al., 2024. A Survey of Machine Learning in Edge Computing: Techniques, Frameworks, Applications, Issues, and Research Directions. Technologies. 12(6), 81. DOI: https://doi.org/10.3390/technologies12060081

[23] Hoefler, T., Alistarh, D., Ben-Nun, T., et al., 2021. Sparsity in Deep Learning: Pruning and Growth for Efficient Inference and Training in Neural Networks. Journal of Machine Learning Research. 23, 1–124.

[24] Paparounas, F., Christofas, V., Amanatidis, P., et al., 2024. Evaluating Knowledge Distillation and Compression Techniques for Edge Class Devices. In Proceedings of the 28th Pan-Hellenic Conference on Progress

in Computing and Informatics, Athens, Greece, 13–15 December 2024; pp. 444–449.
[25] Caldas, S., Duddu, S.M.K., Wu, P., 2019. LEAF: A Benchmark for Federated Settings. Available from: https://github.com/TalwalkarLab/leaf (28th June 2025)
[26] Huang, S., 2020. Low Rank Approximations for Efficient DNN Training [Master's thesis]. The George Washington University: Washington, DC, USA.
[27] Singh, S., Kumar, S., 2024. Efficient Distributed Training through Gradient Compression with Sparsification and Quantization Techniques. arXiv prepreint. arXiv:2502.07634. DOI: https://doi.org/10.48550/arXiv.2502.07634
[28] Peng, C., Li, F., 2018. A Survey on Recent Advances in Event-Triggered Communication and Control. Information Sciences. 457–458, 113–125. DOI: https://doi.org/10.1016/j.ins.2018.04.055
[29] Nadiradze, G., Sabour, A., Alistarh, D., et al., 2019. Asynchronous Decentralized SGD with Quantized and Local Updates. arxiv preprint. arXiv:1910.12308. DOI: https://doi.org/10.48550/arXiv.1910.12308
[30] Fu, L., Zhang, H., Gao, G., et al., 2023. Client Selection in Federated Learning: Principles, Challenges, and Opportunities. IEEE Internet of Things Journal. 10(24), 21811–21819. DOI: https://doi.org/10.48550/arXiv.2211.01549
[31] Ertem, M., 2024. Renewable Energy-Aware Machine Scheduling under Intermittent Energy Supply. IEEE Access. 12, 23613–23625. DOI: https://doi.org/10.1109/ACCESS.2024.3365074
[32] Liu, J., Xu, H., Wang L., et al., 2021. Adaptive Asynchronous Federated Learning in Resource-Constrained Edge Computing. IEEE Transactions on Mobile Computing. 22(2), 674–690. DOI: https://doi.org/10.1109/TMC.2021.3096846
[33] Sun, H., Liu, Y., Al-Tahmeesschi, A., et al., 2025. Advancing 6G: Survey for Explainable AI on Communications and Network Slicing. IEEE Open Journal of the Communications Society. 6, 1372–1412. DOI: https://doi.org/10.1109/OJCOMS.2025.3534626
[34] Pham, Q.-V., Fang, F., Ha, V.N., et al., 2020. A Survey of Multi-Access Edge Computing in 5G and Beyond: Fundamentals, Technology Integration, and State-of-the-Art. IEEE Access. 8, 116974–117017. DOI: https://doi.org/10.1109/ACCESS.2020.3001277
[35] Mushtaq, M.U., Venter, H., Singh, A., et al., 2025. Advances in Energy Harvesting for Sustainable Wireless Sensor Networks: Challenges and Opportunities. Hardware. 3(1), 1. DOI: https://doi.org/10.3390/hardware3010001
[36] Xu, F., Hussain, T., Ahmed M., et al., 2023. The State of AI-Empowered Backscatter Communications: A Comprehensive Survey. IEEE Internet of Things Journal. 10(24), 21763–21786. DOI: https://doi.org/10.1109/JIOT.2023.3299210
[37] Hameed, R.T., Mohamad, O.A., 2023. Federated Learning in IoT: A Survey on Distributed Decision Making. Babylonian Journal of Internet of Things. 2023, 1–7. DOI: https://doi.org/10.58496/BJIoT/2023/001
[38] Bambagini, M., Marinoni, M., Aydin, H., et al., 2016. Energy-Aware Scheduling for Real-Time Systems: A Survey. ACM Transactions on Embedded Computing Systems. 15(1), 1–34. DOI: https://doi.org/10.1145/2808231
[39] Berkani, M.R.A., Chouchane, A., Himeur, Y., et al., 2025. Advances in Federated Learning: Applications and Challenges in Smart Building Environments and Beyond. Computers. 14(4), 124.
[40] Kanellopoulos, D.N., Sharma, V.K., 2022. Dynamic Load Balancing Techniques in the IoT: A SurveyReview. Symmetry. 14(12), 2554. DOI: https://doi.org/10.3390/sym14122554
[41] Ahmed, A., Choi, B.J., 2023. FRIMFL: A Fair and Reliable Incentive Mechanism in Federated Learning. Electronics. 12(15), 3259.
[42] Nanayakkara, S.I., Pokhrel, S.R., Li, G., 2024. Understanding Global Aggregation and Optimization of Federated Learning. Future Generation Computer Systems. 159, 114–133. DOI: https://doi.org/10.1016/j.future.2024.05.009
[43] Bayram, F., Aupke, P., Ahmed, B.S., et al., 2023. DA-LSTM: A Dynamic Drift-Adaptive Learning Framework for Interval Load Forecasting with LSTM Networks. Engineering Applications of Artificial Intelligence. 123, 106480. DOI: https://doi.org/10.1016/j.engappai.2023.106480
[44] Demelius, L., Kern, R., Trügler, A., 2025. Recent Advances of Differential Privacy in Centralized Deep Learning: A Systematic Survey. ACM Computing Surveys. 57(6), 1–28. DOI: https://doi.org/10.1145/3712000
[45] Zhang, Y., Behnia, R., Yavuz, A.A., et al., 2024. Uncovering Attacks and Defenses in Secure Aggregation for Federated Deep Learning. arXiv preprint. arXiv:2410.09676. DOI: https://doi.org/10.48550/arXiv.2410.09676
[46] Kokila, M., Reddy, S.K., 2025. Authentication, Access Control and Scalability Models in Internet of Things Security–A Review. Cyber Security and Applications. 3, 100057. DOI: https://doi.org/10.1016/j.csa.2024.100057
[47] Aljrees, T., Kumar, A., Singh, K.U., et al., 2023. Enhancing IoT Security through a Green and Sustainable Federated Learning Platform: Leveraging Efficient Encryption and the Quondam Signature Algorithm. Sensors. 23(19), 8090.
[48] Jiang, J.C., Kantarci, B., Oktug, S., et al., 2020. Federated Learning in Smart City Sensing: Challenges and Opportunities. Sensors. 20(21), 6230.

[49] Wen, J., Zhang, Z., Lan, Y., et al., 2023. A Survey on Federated Learning: Challenges and Applications. International Journal of Machine Learning and Cybernetics. 14, 513–535. DOI: https://doi.org/10.1007/s13042-022-01647-y

[50] Žalik, K.R., Žalik, M., 2023. A Review of Federated Learning in Agriculture. Sensors. 23(23), 9566. DOI: https://doi.org/10.3390/s23239566

[51] Dembani, R., Karvelas, I., Akbar, N.A., et al., 2025. Agricultural Data Privacy and Federated Learning: A Survey of Challenges and Opportunities. Computers and Electronics in Agriculture. 232, 110048. DOI: https://doi.org/10.1016/j.compag.2025.110048

[52] Wang, T., Du, Y., Gong, Y., et al., 2023. Applications of Federated Learning in Mobile Health: Scoping Survey. Journal of Medical Internet Research. 25, e43006. DOI: https://doi.org/10.2196/43006

[53] Rani, S., Kataria, A., Kumar, S., et al., 2023. Federated Learning for Secure IoMT-Applications in Smart Healthcare Systems: A Comprehensive RSeviewurvey. Knowledge-Based Systems. 274, 110658. DOI: https://doi.org/10.1016/j.knosys.2023.110658

[54] Gambito, M.A., Carnevale, L., Jabbarpour, M.R., et al., 2024. Hierarchical Federated Learning for Natural Disaster Management. In Proceedings of the IEEE/ACM 17th International Conference on Utility and Cloud Computing (UCC), Sharjah, United Arab Emirates, 16–19 December 2024; pp. 282–289.

[55] El Hoffy, A., Kwon, S.S.-C., Yeh, H.-G., 2023. Federated/Deep Learning in UAV Networks for Wildfire Surveillance. In Proceedings of the 2023 Wireless Telecommunications Symposium (WTS), Boston, MA, USA, 19–21 April 2023; pp. 1–9.

[56] Li, A., Markovic, M., Edwards, P., et al., 2024. Model Pruning Enables Localized and Efficient Federated Learning for Yield Forecasting and Data Sharing. Expert Systems with Applications. 242, 122847. DOI: https://doi.org/10.1016/j.eswa.2023.122847

[57] Li, J., Zhang, Y., Li, Y., et al., 2024. FedSparse: A Communication-Efficient Federated Learning Framework Based on Sparse Updates. Electronics. 13(24), 5042. DOI: https://doi.org/10.3390/electronics13245042

[58] Glasgow, M.R., Yuan, H., Ma, T., 2022. Sharp Bounds for Federated Averaging (Local SGD) and Continuous Perspective. In Proceedings of the 25th International Conference on Artificial Intelligence and Statistics, online, 28–30 March 2022; pp. 9050–9090.

[59] Marfo, W., Tosh, D.K., Moore, S.V., 2025. Adaptive Client Selection in Federated Learning: A Network Anomaly Detection Use Case. arXiv preprint. arXiv:2501.15038. DOI: https://doi.org/10.48550/arXiv.2501.15038

[60] Zehtabi, S., Hosseinalipour, S., Brinton, C.G., 2022. Event-Triggered Decentralized Federated Learning over Resource-Constrained Edge Devices. arXiv preprint. arXiv:2211.12640. DOI: https://doi.org/10.48550/arXiv.2211.12640

[61] Mohammadi, S., Symeonidis, I., Balador, A., et al., 2025. Empirical Analysis of Asynchronous Federated Learning on Heterogeneous Devices: Efficiency, Fairness, and Privacy Trade-offs. arXiv preprint. arXiv:2505.07041. DOI: https://doi.org/10.48550/arXiv.2505.07041

[62] Bi, Y., 2024. Empirical Validation of Federated Learning with YOLO v7-Tiny for Road Sign Detection: A Simulation-Based Comparative Study. Applied and Computational Engineering. 49, 92–101. DOI: https://doi.org/10.54254/2755-2721/49/20241067

[63] Zhao, H., Peng, P., Chen, C., et al., 2025. FedRS-Bench: Realistic Federated Learning Datasets and Benchmarks in Remote Sensing. arXiv preprint. arXiv:2505.08325. DOI: https://doi.org/10.48550/arXiv.2505.08325

[64] Sani, M., 2022. Cross-Layer Optimization Techniques in Federated Learning for Improving Privacy and Performance. DOI: https://doi.org/10.13140/RG.2.2.26907.68648

[65] Ma, X., Zhu, J., Lin, Z., et al., 2022. A State-of-the-Art Survey on Solving Non-IID Data in Federated Learning. Future Generation Computer Systems. 135, 244–258. DOI: https://doi.org/10.1016/j.future.2022.05.003

[66] Lee, J., Seif, M., Cho, J., et al., 2024. Exploring the Privacy-Energy Consumption Tradeoff for Split Federated Learning. IEEE Network. 38(6), 388–395. DOI: https://doi.org/10.1109/MNET.2024.3395904

[67] Hu, S., Li, Y., Liu, X., et al., 2022. The OARF Benchmark Suite: Characterization and Implications for Federated Learning Systems. ACM Transactions on Intelligent Systems and Technology. 13(4), 1–32. DOI: https://doi.org/10.1145/3510540

[68] Agiollo, A., Bellavista, P., Mendula M., et al., 2024. EneA-FL: Energy-Aware Orchestration for Serverless Federated Learning. Future Generation Computer Systems. 154, 219–234. DOI: https://doi.org/10.1016/j.future.2024.01.007

[69] Millar, J., Haddadi, H., Madhavapeddy, A., 2025. Energy-Aware Deep Learning on Resource-Constrained Hardware. arXiv preprint. arXiv:2505.12523. DOI: https://doi.org/10.48550/arXiv.2505.12523

[70] Kumar, P., Debele, S.E., Sahani, J., et al., 2021. An Overview of Monitoring Methods for Assessing the Performance of Nature-Based Solutions Against Natural Hazards. Earth-Science Surveys. 217, 103603. DOI: https://doi.org/10.1016/j.earscirev.2021.103603

[71] Nikolaidis, F., Symeonides, M., Trihinas, D., 2023. Towards Efficient Resource Allocation for Federated Learning in Virtualized Managed Environments. Future Internet. 15(8), 261. DOI: https://doi.org/10.3390/fi15080261

[72] Chauhan, M., Sahoo, D.R., 2024. Towards a Greener

Tomorrow: Exploring the Potential of AI, Blockchain, and IoT in Sustainable Development. Nature Environment and Pollution Technology. 23(2), 1105–1113. DOI: https://doi.org/10.46488/NEPT.2024.v23i02.044

[73] Alghamedy, F.H., El-Hagga, N., Alsumayt, A., et al., 2024. Unlocking a Promising Future: Integrating Blockchain Technology and FL-IoT in the Journey to 6G. IEEE Access. 12, 115411–115447. DOI: https://doi.org/10.1109/ACCESS.2024.3435968

[74] Qi, J., Zhou, Q., Lei, L., et al., 2021. Federated Reinforcement Learning: Techniques, Applications, and Open Challenges. arXiv preprint. arXiv:2108.11887. DOI: https://doi.org/10.48550/arXiv.2108.11887

[75] Wang, Z., Wu, F., Yu, F., et al., 2024. Federated Continual Learning for Edge-AI: A Comprehensive Survey. arXiv preprint. arXiv:2411.13740. DOI: https://doi.org/10.48550/arXiv.2411.13740

[76] Zhang, K., Cao, X., 2024. Federated Learning With Energy Harvesting Devices: An MDP Framework. arXiv preprint. arXiv:2405.10513. DOI: https://doi.org/10.48550/arXiv.2405.10513

[77] Kamble, G.U., Patil, C.S., Alman, V.V., et al., 2024. Neuromorphic Computing: Cutting-Edge Advances and Future Directions. In: Bai, K.J., Yi, Y. (Eds.). Recent Advances in Neuromorphic Computing. DOI: https://doi.org/10.5772/intechopen.1006712