AD EDUXIAN JOURNAL
A Quarterly Multidisciplinary Blind Peer Reviewed & Refereed Online International Journal
Google Scholar Indexing IF:3.125(IIFS)
Vol (2), Issue (2), May-July 2025 ISSN: 3048-7951

A Comprehensive Review of Federated Learning Frameworks and Their Comparative Performance in Internet of Things Devices

Monika[1], & Prof. Rishi Pal Singh[2]

DOI: https://doi-ds.org/doilink/08.2025-49497881/ADEDJ/V2/I2/MRPS

Review: 08/07/2025 **Acceptance: 20/07/2025** **Publication:14/08/2025**

ABSTRACT

The Internet of Things (IoT) has revolutionized data-driven applications but also introduced critical challenges in terms of privacy, scalability, and computation. Federated Learning (FL), a decentralized machine learning paradigm, has emerged as a promising solution to address these concerns by enabling local model training on devices without transferring raw data. This review provides a comprehensive comparative analysis of existing FL frameworks and techniques in the context of IoT. We evaluate prominent frameworks such as TensorFlow Federated (TFF), Flower, and FATE, analysing their architectural design, communication efficiency, security provisions, and performance across diverse IoT scenarios. We further explore various FL algorithms—such as FedAvg, Fed Prox—and optimization techniques including model compression, differential privacy, and homomorphic encryption. Our findings highlight the trade-offs between accuracy, resource consumption, and scalability, offering insights into framework suitability for applications like smart homes, healthcare, and industrial IoT. Despite FL's potential, challenges remain in communication overhead, device heterogeneity, security vulnerabilities, and system dynamics. We conclude by identifying key research opportunities, such as adaptive personalization, robust privacy-preserving techniques, and scalable FL architectures.

Keywords: Federated learning, Internet of Things, privacy, security, framework comparison, performance evaluation, distributed learning:

1. Introduction

The proliferation of IoT devices has led to an exponential increase in data generation and smart application development. However, centralized machine learning approaches—dependent on aggregating data to central servers—raise significant concerns regarding privacy, bandwidth limitations, and latency. In this context, **Federated Learning (FL)** emerges as a suitable alternative, allowing model training directly on edge devices while preserving user privacy. FL's decentralized nature makes it highly relevant for IoT systems characterized by **resource constraints, heterogeneous data**, and **intermittent connectivity**. This review critically examines how different FL frameworks and techniques perform under these constraints and identifies the trade-offs involved in real-world deployments.

The central research question addressed here is:

> *"How do existing FL frameworks and algorithms compare in their suitability and performance within IoT environments?*

This study is significant because it bridges the gap between theoretical FL models and their practical deployment in IoT. The review is organized as follows: Section 2 discusses FL foundations and benefits, Section 3 presents a comparative analysis of FL frameworks and algorithms in IoT, Section 4 addresses current challenges, Section 5 outlines future research directions, and Section 6 concludes the study.

2. Federated Learning in IoT: Foundations and Benefits

Federated Learning (FL) is a decentralized machine learning technique that enables multiple devices to collaboratively train a shared model without exchanging raw data. In Internet of Things (IoT) environments—where devices such as sensors, wearables, and smart appliances generate vast amounts of sensitive data—FL offers an effective way to protect privacy while leveraging distributed data for model training. Instead of sending data to a central server, FL allows each device to train locally and only share model updates, significantly reducing communication costs and enhancing data

[1] Monika, Research Scholar, Department of Computer Science and Technology, Guru Jambheshwar University, Hisar, Haryana, India, Email-monubhambhuhsr@gmail.com
[2]Professor, Department of Computer Science and technology, Guru Jambheshwar University, Hisar, Haryana, India,

security. This approach aligns well with the constraints of IoT devices, which often face limited bandwidth, processing power, and battery life. Key benefits of FL in IoT include enhanced privacy preservation, as sensitive data remains on the device; lower network congestion through reduced data transmission; and scalability, as it can support a growing number of devices. FL also improves fault tolerance—devices can join or leave without interrupting the process—and allows for model personalization, enhancing relevance for individual devices. However, implementing FL in IoT also introduces challenges, such as non-IID data distribution, device heterogeneity, and intermittent connectivity, which must be addressed to realize its full potential.

2.1 Core Concepts of Federated Learning (FL)

Federated Learning enables multiple edge devices to collaboratively train a shared global model while keeping data localized. Only model updates (e.g., gradients or weights) are sent to a central aggregator for global update (e.g., via FedAvg) Federated Learning (FL) is a decentralized approach to machine learning that allows multiple devices or clients—such as smartphones, sensors, or IoT endpoints—to collaboratively train a shared global model without transferring raw data to a centralized server. The fundamental goal is to preserve data privacy while still enabling effective model training across distributed data sources.

The process typically consists of the following steps:

- **Local Training**: Each participating device (client) trains the machine learning model on its own local dataset. This allows the device to extract relevant patterns and features from its unique data without exposing that data externally.
- **Model Aggregation**: After local training, each device sends only its trained model parameters (such as weights or gradients) to a central server. The server then aggregates these updates—often using algorithms like *Federated Averaging (FedAvg)*—to update the global model. The updated global model is then redistributed to all devices for the next training round.

This cycle continues iteratively until the model converges to an acceptable level of performance. Importantly, the raw data never leaves the local device, making FL especially attractive in privacy-sensitive environments like healthcare, finance, or smart home systems.

2.2 Key Advantages in IoT

FL presents a number of compelling advantages when applied to IoT ecosystems, where privacy, limited connectivity, and scalability are major concerns:

- **Privacy Preservation**: Since data remains on-device, FL minimizes the risk of sensitive data being intercepted, mishandled, or exposed during transmission. This approach supports compliance with data protection regulations like GDPR and HIPAA.
- **Reduced Communication Overhead**: Instead of transmitting large volumes of raw data, only compact model updates are shared. This is particularly useful in bandwidth-constrained IoT settings, such as rural sensor networks or mobile edge devices.
- **Scalability**: FL supports massive scaling across millions of devices, making it suitable for large IoT deployments. New devices can join or leave the training process dynamically, supporting a wide range of applications, from smart cities to industrial automation.
- **Fault Tolerance**: The federated learning process is inherently robust to the intermittent availability of devices. Devices can drop in or out of the process due to connectivity issues or power constraints without significantly disrupting the overall model training cycle.

2.3 IoT Characteristics Influencing FL: While FL is well-suited to IoT systems, its effectiveness is influenced by the unique characteristics of IoT environments:

- **Data Heterogeneity (Non-IID Data)**: Unlike centralized systems that train on a single, unified dataset, IoT devices often generate data with vastly different distributions. For example, smart thermostats in different geographic locations will encounter distinct temperature patterns. This *non-IID (non-independent and identically distributed)* nature of data makes it harder for a single global model to generalize effectively across all devices.
- **Resource Limitations**: Most IoT devices are lightweight in terms of computing power, memory, and energy. Training even modest-sized models can exhaust battery life or exceed computational capacity, making it necessary to optimize FL algorithms for efficiency.

- **Device Diversity**: IoT ecosystems include a wide variety of devices with different architectures, operating systems, sensors, and firmware. Ensuring compatibility and uniform performance across this diversity is a major challenge for FL deployment.
- **Unreliable Connectivity**: IoT devices often operate in environments with limited or intermittent internet access. For example, remote agricultural sensors or wearable health monitors may lose connection frequently. FL systems must be able to tolerate such conditions, potentially using asynchronous updates or offline training.

3. Comparative Analysis of FL Frameworks and Techniques in IoT:

3.1 Overview of Frameworks

Federated Learning (FL) frameworks are software tools designed to facilitate distributed model training across multiple clients while maintaining data privacy. Selecting the right FL framework is critical for effective deployment, especially in resource-constrained and privacy-sensitive environments like IoT. Below is a detailed comparison of three widely adopted FL frameworks—**TFF (TensorFlow Federated)**, **Flower**, and **FATE (Federated AI Technology Enabler)**—evaluated across architecture, machine learning support, IoT suitability, and key features.

TensorFlow Federated (TFF)

- **Architecture**: TFF uses a centralized aggregator model. All updates from clients are sent to a central server for aggregation.
- **ML Support**: Natively supports TensorFlow-based models.
- **Suitability for IoT**: Moderate; while it is a powerful research tool, its computational overhead and tight coupling with TensorFlow make it less ideal for lightweight IoT devices.
- **Key Features**: Excellent for prototyping and experimentation; supports custom aggregation logic and federated simulations.

Flower

- **Architecture**: Flower is designed with a flexible and modular topology. It can operate in centralized or decentralized environments.
- **ML Support**: Compatible with multiple ML libraries, including PyTorch and TensorFlow.
- **Suitability for IoT**: High; Flower is well-suited for real-world FL deployment, including heterogeneous and resource-constrained IoT devices.
- **Key Features**: Lightweight, scalable, supports real-device simulations, easy-to-extend APIs, and built-in support for cross-device communication.

FATE (Federated AI Technology Enabler)

- **Architecture**: Designed for enterprise-scale deployments, FATE supports a modular, privacy-centric architecture.
- **ML Support**: Compatible with TensorFlow, XGBoost, and other advanced ML models.
- **Suitability for IoT**: Medium; FATE is more suitable for industrial use cases with strong security demands, such as finance or healthcare.
- **Key Features**: Built-in support for advanced privacy-preserving techniques such as Homomorphic Encryption (HE), Secure Multiparty Computation (SMPC), and Differential Privacy (DP). Supports both horizontal and vertical federated learning.

3.2 FL IoT systems present unique challenges for FL due to data distribution, device limitations, and privacy concerns. This has led to the development of various algorithms and supporting techniques designed to improve model performance, security, and communication efficiency.

1. Core Federated Learning Algorithms

a. FedAvg (Federated Averaging):

- **Description:** The foundational FL algorithm proposed by Google, where local models are trained independently and the server aggregates them by averaging the model weights**.**
- Advantages: Simple, efficient, and effective in homogeneous data settings.
- Limitations: Performs poorly with non-IID (non-identically distributed) data, which is common in IoT environments where each device collects different types of data**.**

b. FedProx (Federated Proximal)

- **Description:** Extends FedAvg by adding a proximal term during local training to limit deviation from the global model.
- **Advantages:** Improves convergence and stability in the presence of heterogeneous data.
- **Use Case in IoT:** Suitable for devices with skewed data distributions, such as smart home sensors or wearable devices.

c. FedMA (Federated Matched Averaging) and FedDyn (Federated Dynamics)

- **Description:** These are personalized FL algorithms designed to adapt to device-specific data.
  - FedMA matches and averages hidden units of neural networks.
  - FedDyn dynamically adjusts optimization to minimize divergence from global training.

  Advantages: High personalization, better performance on non-IID data.

  Limitations: Increased complexity and computational cost, which may challenge resource-limited IoT devices.

2. Communication Optimization Techniques: IoT devices often operate under bandwidth and energy constraints. The following techniques are used to reduce communication overhead:

a. Model Pruning:

- Removes unnecessary model weights or neurons to reduce model size.
- Reduces update size and computation on IoT devices.

b. Quantization:

- Converts high-precision weights to lower-bit formats (e.g., float32 → int8).
- Decreases data transfer size with minimal loss in accuracy.

c. Sparsification

- Sends only the most significant gradients or weights (e.g., top-k updates).
- Improves bandwidth efficiency, especially important in unstable network conditions.

3.3 Security and Privacy Enhancements in Federated Learning: While Federated Learning (FL) offers inherent privacy advantages by ensuring that raw data remains on local devices, it is not immune to security threats. Even without direct data sharing, model updates can leak sensitive information through various attacks, such as:

- Model Inversion Attacks: Attempt to reconstruct original data from shared gradients.
- Membership Inference Attacks: Determine whether specific data was used to train the model.
- Poisoning Attacks: Malicious devices send manipulated updates to corrupt the global model.

To defend against these threats, several advanced privacy-preserving techniques are used in FL:

a. Differential Privacy (DP): What it does: Differential Privacy introduces random noise into the model updates (e.g., gradients or weights) before they are sent to the server. This noise ensures that individual data points cannot be accurately reconstructed or identified.

- Prevents an attacker from determining whether a particular record was used in training.
- Provides strong, quantifiable privacy guarantees.
- Widely adopted in privacy-critical sectors like healthcare, finance, and surveillance.

DP is lightweight and can be implemented on constrained IoT devices (e.g., wearables, smart home hubs). However, the added noise may slightly reduce model accuracy, requiring a balance between privacy and utility.

b. Homomorphic Encryption (HE): HE allows mathematical operations to be performed directly on encrypted data, without requiring decryption. For FL, this means devices can encrypt their model updates, and the server can still aggregate them correctly—without ever seeing the raw parameters.

- Raw model updates are never exposed in plaintext, even during aggregation.
- Strong cryptographic protection against eavesdropping or man-in-the-middle attacks.

While HE offers high privacy, it is computationally intensive. Most low-power IoT devices may struggle to perform homomorphic encryption due to processing and memory constraints. It is better suited to powerful edge servers or gateway devices in the IoT architecture.

c. Secure Multiparty Computation (SMPC): SMPC splits data or computations into encrypted fragments that are distributed across multiple parties. No single party has access to the full dataset or model update, but they can jointly compute the desired result.

- Eliminates the need for a central trusted aggregator.
- Ensures strong privacy in collaborative FL across multiple data owners or organizations.
- Resistant to insider attacks.

SMPC is ideal for multi-organization or cross-vendor IoT networks, such as federated medical systems or smart traffic networks where data ownership is distributed. However, the communication overhead and complexity make it more appropriate for server-level coordination rather than device-level training.

4. Handling Heterogeneity in IoT Devices

IoT devices vary widely in data types, computation power, and network stability. To address this heterogeneity:

a. Clustered Federated Learning

- Groups devices with similar data distributions or capabilities.
- Each cluster trains its own specialized model.
- Reduces the negative impact of non-IID data on global model convergence.

b. Personalized Federated Learning

- Each device gets a customized local model fine-tuned from the global model.
- Techniques include meta-learning, transfer learning, and regularized local objectives.
- Balances global knowledge with local adaptation, improving prediction accuracy per device.

4. Challenges and Limitations of Federated Learning in IoT: Despite the promise of Federated Learning (FL) in preserving privacy and enabling distributed intelligence, its deployment in Internet of Things (IoT) environments presents several significant challenges. These limitations stem from the unique properties of IoT systems, including device diversity, connectivity issues, limited resources, and dynamic environments.

- **Communication Overhead and Latency:** FL involves frequent exchanges of model updates between edge devices and a central server or aggregator. In IoT systems:
- **Limited bandwidth** and **intermittent connectivity** (e.g., rural sensors or battery-powered wearables) often lead to delays or failures in communication.
- The repetitive transfer of model weights, even in compressed form, can strain network resources, especially in large-scale deployments.
- **Impact**: Slower training, increased energy consumption, and dropped updates during transmission.
- **Potential Solutions**: Compression techniques (e.g., sparsification, quantization), asynchronous communication protocols, or edge-level aggregation.
- **Non-IID Data Distributions:** IoT devices often generate non-IID (non-independent and identically distributed) data due to their diverse functions, locations, and users. For example:
- A smart thermostat in a cold region records very different temperature patterns than one in a tropical region.
- Slower convergence of the global model.
- Reduced generalizability and accuracy, as the global model may not fit well across all devices.
- **Potential Solutions**: Clustered FL, personalized FL, or weighting updates based on data quality and contribution.
- **Resource Constraints of IoT Devices:** Most IoT devices are **resource-constrained**, with:
- Limited **CPU**, **RAM**, and **battery** power.
- Inability to train large or complex models locally.
- Devices may drop out of training cycles.
- Limits the complexity of models that can be trained locally, impacting performance.
- **Potential Solutions**: Model compression, lightweight architectures (e.g., MobileNet), offloading partial computation to edge servers.
- **Security Vulnerabilities:** While FL improves data privacy by keeping data local, it introduces new **security threats**:

- **Model Inversion Attacks**: Adversaries reconstruct sensitive data from shared model gradients.
- **Membership Inference**: Identifying if a specific record was part of the training dataset.
- Breach of user privacy, regulatory non-compliance.
- **Potential Solutions**: Incorporate Differential Privacy (DP), Homomorphic Encryption (HE), or secure aggregation protocols.
- **System Dynamics and Scalability:** IoT environments are **highly dynamic**:
- Devices may **join or leave** the network unpredictably due to power cycles, disconnections, or mobility (e.g., vehicles, drones).
- Synchronization and model consistency become difficult when participants frequently change.
- Unstable training cycles, leading to poor or delayed convergence.
- Redundant or outdated model updates.
- **Potential Solutions**: Use **asynchronous FL protocols**, enable local buffering of updates, adopt decentralized learning architectures.
- **Data Poisoning Attacks:** Malicious devices in the FL network may upload **manipulated updates** to corrupt the global model intentionally
- Degraded model accuracy or biased decision-making.
- Potential real-world consequences in critical IoT applications (e.g., healthcare, autonomous driving).
- **Defense Mechanisms**:
- **Anomaly detection**: Monitor updates for unusual patterns.
  Robust aggregation: Techniques like Krum, Trimmed Mean, or Median eliminate or downweight suspicious contributions.

5. Future Directions and Research Opportunities:

- **Efficient communication**: Explore adaptive update frequency, edge aggregation.
- **Personalization**: Develop methods balancing global performance with local relevance.
- **Advanced privacy**: Improve HE and SMPC scalability for IoT-scale deployment.
- **Scalable FL**: Design asynchronous, decentralized protocols for dynamic environments.
- **Application-specific models**: Tailor FL methods to domains like:
  - **Healthcare IoT**: Sensitive data, need for privacy.
  - **Smart cities**: Real-time responsiveness, edge-cloud hybrid models.
  - **Smart agriculture**: Intermittent data, remote deployment.

6. Conclusion: This review presents a critical comparison of FL frameworks and techniques tailored for IoT environments. Federated Learning shows considerable promise in overcoming the privacy, communication, and resource-related challenges endemic to IoT. Among frameworks, Flower stands out for flexibility and resource efficiency, while FATE excels in security-centric deployments. However, key challenges—including data heterogeneity, communication inefficiency, and dynamic system behavior—necessitate ongoing innovation. Future work should focus on lightweight, secure, and personalized FL mechanisms that accommodate the unique characteristics of IoT systems across domains.

7. References:

- Kairouz, P., et al. (2021). "Advances and Open Problems in Federated Learning." *Foundations and Trends® in Machine Learning*.
- Li, T., et al. (2020). "Federated Optimization in Heterogeneous Networks." *Proceedings of MLSys*.
- Bonawitz, K., et al. (2019). "Towards Federated Learning at Scale: System Design." *SysML*.
- Liu, Y., et al. (2022). "A Comprehensive Survey on Federated Learning Systems." *ACM Computing Surveys*.
- Zhang, C., et al. (2021). "Federated Learning for the Internet of Things: A Comprehensive Survey." *IEEE Communications Surveys & Tutorials*.