

FEDERATED MACHINE LEARNING AS A PATHWAY TO PRIVACY-PRESERVING ARTIFICIAL INTELLIGENCE IN SMART DIGITAL ECOSYSTEMS

Kainat Tariq^{*1}, Syed Zeshan Haider², Kohal Deep³, Muhammad Haqan Ali Rai⁴

^{*1}BS Computer Science Virtual University of Pakistan

²BS Computer Science Islamia university Bhawalpur (campus RYK)

³MS Information security NED University of Engineering and Technology, Karachi, Pakistan

⁴BS Software Engineering PMAS-Arid Agriculture University Rawalpindi, Pakistan

^{*1}kainattariq703@gmail.com, ²syedzeeshan123haider@gmail.com, ³kohaldeep93@gmail.com,

⁴haqanali934@gmail.com

DOI: <https://doi.org/10.5281/zenodo.21454681>

Keywords

Federated Machine Learning; Privacy-Preserving Artificial Intelligence; Distributed Learning; Smart Digital Ecosystems; Internet of Things; Edge Computing; Smart Cities; Healthcare Analytics; Industry 5.0; Explainable AI; Blockchain.

Article History

Received: 01 January 2026

Accepted: 17 March 2026

Published: 31 March 2026

Copyright @Author

Corresponding Author: *

Kainat Tariq

Abstract

The ever-increasing trend towards the integration of artificial intelligence, edge computing, and hyper-connectivity has seen the emergence of intelligent digital ecosystems encompassing health care, financial sector, municipal administration, IoT, and industry 5.0. However, such ecosystems have to rely on constant large-scale data collection and, hence, the tradition of data storage on one central server becomes more and more challenging in view of privacy concerns, regulatory compliance, and risk appetite of organizations. In order to address this challenge, federated machine learning (FML) has been developed: it makes it possible to train a global machine learning model jointly by multiple participants without collecting any raw data at one central server. This paper provides a systematic narrative review of FML as a basis for privacy-preserving artificial intelligence based on peer-reviewed literature indexed in IEEE Xplore, SpringerLink, ScienceDirect, ACM Digital Library, Wiley Online Library, Taylor & Francis Online, MDPI, Nature Portfolio, and Google Scholar (articles published from 2020 to 2026). Other than summarizing the architectural and algorithmic advancements made in Federated Machine Learning, there are three key contributions that this review makes, setting it apart from previous literature: first, it clearly highlights why and how the scope and synthesis presented in the current review differ from the previous ones; second, it provides an entirely transparent selection process following the PRISMA guidelines, including details on the search strategy used and inclusion/exclusion criteria for selected articles; third, it presents all significant Federated Optimization algorithms in one comparative table instead of presenting them separately in prose. The review highlights the application of Differential Privacy, Secure Multi-Party Computation, and Homomorphic Encryption to provide model updates privacy protection and the use cases of FML in healthcare, smart cities, IoT, finance, and Industry 5.0. Key challenges associated with FML such as non-IID datasets, communication costs, device diversity, adversarial attacks, and fragmented governance are also discussed, along with emerging trends like Explainable AI,

blockchain-enabled trust layers, edge intelligence, generative and foundation models, as well as next-generation wireless communication technologies.

1. INTRODUCTION

Artificial intelligence technology is no longer just a research-oriented technology but is now a foundational layer under governments, hospitals, banks, and smart cities. The combination of artificial intelligence, machine learning, deep learning, edge computing, 5G wireless technology (Nguyen et al., 2025), and big data analytics has enabled the development of digital ecosystems with capabilities for independent predictions, optimizations, and decisions in a way that would not have been possible ten years ago. These digital ecosystems are constantly powered through sensors, handheld devices, wearables, autonomous vehicles, and industrial machines, which provide them with streams of heterogeneous data (Shen et al., 2022). This data forms the foundation for contemporary AI technologies, allowing them to operate efficiently, automate tasks, and make data-driven decisions.

The functioning of today's AI is closely associated with large datasets of high quality, and traditionally the way of getting access to them is through the process of centralization, when raw data of many users or entities are transferred to one cloud storage for preprocessing, storage, and training purposes. This strategy has proven to be highly effective, delivering good predictive results (McMahan et al., 2018), but at the same time it poses serious risks. The information generated by hospitals, banks, governmental bodies, industrial companies, and end users is inherently confidential, and thus its transfer to a common repository involves significant ethical and legal concerns which cannot be overlooked. Companies are becoming increasingly reluctant to upload their data to central storages due to risks of cyber attacks, insider threats, intellectual property theft, and sanctions (Zhang et al., 2021a).

Digital transformation has only sharpened these concerns. Smart ecosystems are, by construction, distributed across administrative boundaries: smart healthcare networks, intelligent transport systems, industrial IoT deployments (Zhang et al., 2021b), precision agriculture, smart

manufacturing, digital banking, autonomous vehicles, and energy grids all involve multiple independent stakeholders producing and consuming data in real time. The intelligence these ecosystems promise is real, but so is the friction created when centralized architectures require data to cross organizational lines, each crossing adding both technical overhead and exposure (Yang et al., 2019).

Public expectations around privacy have shifted in step with this expansion. People now assume that medical histories, financial transactions, movement patterns, biometric signals, academic records, and personal communications will stay confidential even as they are mined for insight (Abreha et al., 2022). Regulatory regimes have followed the same trajectory: the General Data Protection Regulation (GDPR) in the European Union, the California Consumer Privacy Act (CCPA), and a growing list of national data-protection statutes now constrain how personal data can be gathered, moved, stored, and processed. These frameworks have pushed AI research toward paradigms that can extract useful knowledge without exposing the underlying records, treating privacy preservation as a design requirement rather than a downstream patch (Sun et al., 2021).

Federated Machine Learning (FML) has become the leading response to this requirement. Rather than pooling raw data centrally, FML lets each participating device or organization train a model locally and share only the resulting parameters, gradients, or encrypted updates with a coordinating server, which then aggregates them into a single improved global model without ever seeing the original data (Zhou et al., 2022). This inversion of the usual data flow substantially narrows the attack surface for sensitive information while still letting many participants learn jointly. Contemporary surveys describe federated learning as among the most consequential privacy-aware paradigms precisely because it lets collaborative intelligence emerge

while data stays where it was created (Saha et al., 2024).

In this regard, the current paper seeks to consolidate existing studies on federated machine learning as an enabling tool for privacy-preserving AI in smart digital environments. The review covers the theoretical foundation of the paradigm, its architecture, optimization techniques, privacy solutions, security risks it is susceptible to, communication methods, applications, challenges, and directions of research in this domain. Unlike existing surveys that concentrate on one aspect of federated learning only – algorithms, communication, security – the current review aims at providing a holistic view of federated learning by putting it into a wider perspective of privacy-preserving AI, edge intelligence, IoT, and trustworthy AI governance (Rieke et al., 2020)

1.1 Statement of the Problem

In this regard, the current paper seeks to consolidate existing studies on federated machine learning as an enabling tool for privacy-preserving AI in smart digital environments. The review covers the theoretical foundation of the paradigm, its architecture, optimization techniques, privacy solutions, security risks it is susceptible to, communication methods, applications, challenges, and directions of research in this domain. Unlike existing surveys that concentrate on one aspect of federated learning only – algorithms, communication, security – the current review aims at providing a holistic view of federated learning by putting it into a wider perspective of privacy-preserving AI, edge intelligence, IoT, and trustworthy AI governance (Rieke et al., 2020).

Further structural issue is that many studies of such problems are considered as isolated research directions – privacy mechanisms researched in isolation from the optimization methods, security in isolation from the deployment environment – whereas all of these should be considered in terms of trade-offs between them, taking into consideration the ecosystem they are supposed to operate in. Such fragmentation calls for an integrative review that would cover all of these

aspects and show where the development of Federated Machine Learning has reached and what remains to be done.

1.2 Research Objectives

- To examine the theoretical foundations, architectural variants, and learning mechanisms that underpin Federated Machine Learning as a route to privacy-preserving Artificial Intelligence in smart digital ecosystems.
- To evaluate, comparatively, the privacy-preserving techniques, security mechanisms, and federated optimization algorithms that shape the performance, robustness, and trustworthiness of federated AI systems.
- To map the principal application domains, implementation obstacles, and emerging research directions of Federated Machine Learning across healthcare, IoT, smart cities, finance, and Industry 5.0.

1.3 Research Questions

- How does Federated Machine Learning enable the development of privacy-preserving Artificial Intelligence within smart digital ecosystems?
- How well do existing privacy-preserving techniques, security mechanisms, and federated optimization algorithms perform against the joint demands of accuracy, security, scalability, and trustworthiness?
- What are the leading applications, practical obstacles, and future research opportunities for Federated Machine Learning across smart digital ecosystems?

2. LITERATURE REVIEW

2.1 Evolution of Federated Machine Learning

FML was introduced as a solution to the privacy issues of centralized learning, whereby several parties can create a collaborative model using their raw data locally without transmitting it outside. The early success of FML was due to Google's efforts in developing a predictive model for smartphones' keyboard predictions. Subsequently, the concept quickly became popular in healthcare, finance, smart cities, and industrial Internet of Things. The key difference between FML and

previous distributed computing architectures lies in the fact that in FML, only parameters move across the network, but not the data used to calculate those parameters, thus reducing the cost of transmission. Newer versions of FML, including personal federated learning, hierarchical federated learning, and cross-silo federated learning, were

proposed as solutions to scalability and communication challenges, along with heterogeneous computing environment conditions (Kairouz et al., 2021; Li et al., 2020). The progress from isolated data silos towards contemporary FML is presented in Figure 1 below.




Figure 1. Maturation Timeline of Federated Machine Learning

2.2 Privacy-Preserving Artificial Intelligence

Privacy-preserving AI has moved from being a nice-to-have to being a baseline expectation, as organizations lean more heavily on AI systems to process data that is personal, medical, financial, or otherwise sensitive. Centralized AI architectures, by requiring a single point of data storage, create a single point of failure for breaches, regulatory violations, and reputational damage. Federated Machine Learning addresses this by decoupling model improvement from data centralization, and by layering in dedicated privacy-enhancing technologies – Differential Privacy, Secure Multi-Party Computation, Homomorphic Encryption, and Secure Aggregation – that harden the exchanged information itself. These mechanisms are not free; each buys stronger confidentiality at some cost in computation or communication, but collectively they are what makes privacy-preserving AI a credible foundation for trustworthy,

responsible AI systems rather than an aspirational label (Yang et al., 2019; Mothukuri et al., 2021).

2.3 Federated Learning in Smart Digital Ecosystems

Smart digital ecosystems – consist of IoT devices, edge nodes, cloud infrastructures, self-driving cars, hospitals, banks, and urban infrastructures produce massive amounts of data continuously in ways that would be difficult for any one organization to centrally store and analyze. Federated learning offers a solution to enable collaborative intelligence between these parties without having to centralize the data, allowing each party to control its own data and bear the liability for it. The use of federated learning along with edge computing further strengthens its advantages because processing gets nearer to the source of the data, reducing latency, bandwidth

utilization, and computation costs (Nguyen et al., 2022; Lim et al., 2020).

2.4 Challenges and Emerging Trends in Federated Machine Learning

Federated Machine Learning's However, benefits and coexist with a series of persistently occurring technical and operational obstacles that hinder the widespread application of this technology. The data of clients are usually non-independent and non-identically distributed, which reduces the speed of convergence and prediction quality. Communication complexity, heterogeneous capacity of devices, limited edge computing resources, model poisoning, gradient exposure, and the uneven engagement of clients complicate coordination within a federated framework. At present, the main focus of academic study is on combining federated learning with blockchain, Explainable AI, reinforcement learning, digital twins, and large foundation models in an attempt to receive transparency, security, scalability, and personalization in one solution.

2.5 Research Gap

Federated Machine Learning has received considerable attention in the literature, but most reviews are confined to examining just one layer of the federated machine learning process - be it the optimization layer, the communication layer or the security layer - without recognizing the need for a solution where privacy, accuracy, scalability, robustness and explainability constraints should be balanced. Comparisons between different application areas including medicine, smart cities, financial services, and Industry 5.0 are generally not well explored, and the combination of federated learning with more recent technologies like edge intelligence, blockchain, Explainable AI, and large language models are only briefly discussed.

This review is different from previous ones on each of three specific grounds. First, instead of describing algorithms or privacy technologies in isolation, it provides an integrated story where the layers of architecture, privacy technology, optimization algorithm, and deployment domain are all interconnected such that trade-offs in any

of them can be understood in terms of implications in another one. Second, unlike most other narrative reviews, which seldom explain how their corpus of the source literature was collected and selected, this review presents an explicit strategy for searching literature and specifies the selection and inclusion/exclusion criteria along with a modified PRISMA selection flow chart (Section 3). Finally, this review brings six main federated learning algorithms – FedAvg, FedProx, FedNova, FedOpt, SCAFFOLD, and FedDyn, into a single comparative table by communication cost, robustness to non-iid data, and common deployment domains (Section 4.6). Such a comparative table is provided partially and, as a rule, implicitly in other surveys.

3. METHODOLOGY

3.1 Research Design

This paper is structured as a systematic narrative review, which is used due to the fact that it is possible to synthesize theory, empirical results, architecture, and practical examples from literature that is heterogeneous enough in methodology to conduct a meta-analysis. While this paper is not claiming any rigorous statistics of a systematic review, the method of choosing the studies used in the analysis is based on the PRISMA principles of explicitness in terms of the methodology – searching strategy, inclusion and exclusion criteria, and flow of the selection process to prove its reproducibility.

3.2 Search Strategy and Data Sources

Literature was retrieved from nine internationally recognized scientific databases: IEEE Xplore, SpringerLink, ScienceDirect (Elsevier), the ACM Digital Library, Wiley Online Library, Taylor & Francis Online, MDPI, Nature Portfolio, and Google Scholar. These sources were chosen for their concentration of peer-reviewed work in artificial intelligence, distributed computing, cybersecurity, edge computing, and privacy-preserving technologies.

Search strings combined federated-learning terminology with privacy and domain-specific keywords using Boolean operators, following the general pattern:

("federated learning" OR "federated machine learning") AND ("privacy-preserving" OR "differential privacy" OR "secure aggregation" OR "homomorphic encryption") AND ("smart city" OR "Internet of Things" OR "healthcare" OR "finance" OR "Industry 5.0")

Searches were restricted to work published between 2020 and 2026 to keep the review current with the field's rapid pace, and titles, abstracts, and keyword indices were the primary fields queried, supplemented by citation-chasing through the reference lists of highly cited surveys (Xu et al., 2021).

3.3 Inclusion and Exclusion Criteria

Inclusion criteria:

- Peer-reviewed journal articles, conference proceedings, systematic reviews, survey papers, and book chapters.
- Publications written in English and dated between 2020 and 2026.
- Work centered on federated learning, privacy-preserving AI, smart digital ecosystems, edge computing, IoT, blockchain integration, or related security mechanisms.
- Studies reporting sufficient methodological detail to be evaluated for relevance and quality.

Exclusion criteria:

- Publications unrelated to federated learning or its privacy-preserving applications.
- Duplicate records retrieved across multiple databases.
- Non-peer-reviewed sources, editorials, opinion pieces, and preprints without subsequent peer review.
- Studies lacking sufficient methodological detail to assess rigor or relevance.

3.4 Study Selection Process

A total of 486 records were obtained from the nine databases in the first search. Exclusion of 97 duplicates and 34 records that were not written in English or were not peer-reviewed left 355 records, of which 168 were excluded from title and abstract screening due to being off-topic, editorial in nature, or purely opinionated. Of the 187 records that remained after full-text assessment using the criteria discussed above, 92 were excluded mainly because of lack of enough information on methodology or relevance to the topic of federated learning in a smart-ecosystem environment. A total of 95 studies qualified for thematic synthesis in the rest of this paper. Figure 2 provides a flowchart.




Figure 2. Study Selection Flow Diagram (PRISMA-Adapted)

3.5 Data Extraction and Thematic Analysis

Each of the 95 articles reviewed has been classified based on a common extraction template covering: publication year, database used, methodology used, domain of application, privacy preserving technologies, and optimization algorithms. Based on this classification of codes, the articles have been organized under the thematic categories of this review: evolution of the paradigm, architecture framework, privacy preserving

technologies, optimization algorithms, domains of applications, implementation issues, and future research direction. Cross tabulation of these categories along with the domains of application and the algorithms helped pinpoint the convergence, divergence, and gaps within the literature (for instance, near unanimous consensus on differential privacy and lack of cross-domain governance frameworks), thereby informing the comparative tables presented in sections 4 and 5.

4. FEDERATED LEARNING ARCHITECTURE AND TECHNIQUES

4.1 Federated Learning Framework Architecture

The Federated Learning approach relies on training that is carried out using a decentralized process instead of having a centralized collection system. Data is no longer pooled from multiple places into a single repository; a coordinating server sends the initial global model to various client devices, who train the global model using their respective local data without sending the data outside the device or organization. The parameters of the updated model are sent back to the server. This process is illustrated by Figure 3 in a form of a wheel, where the server acts as the center and

sends the global model to hospitals, banks, IoT hubs, and smart cities.

In this case, it would have to aggregate these updates through Federated Averaging (FedAvg) to generate a superior global model that will once again be distributed out to start a new round of local training. This cycle goes on until the global model becomes sufficiently good. It is the fact that each client fully owns its data that makes it possible for such architecture to scale well for a wide variety of participants – from consumer devices like personal phones and wearable devices through hospitals and banks, all the way to self-driving cars and industrial IoT devices, edge servers, and even smart cities, all contributing on equal terms to build the global model.

Figure 3. Privacy-Preserving Framework Workflow




Figure 3. Radial Architecture of the Federated Learning Cycle

This foundational concept is not often used by itself in today's implementations. These concepts include Differential Privacy, Secure Multi-Party Computation, Homomorphic Encryption, Secure Aggregation, Blockchain-based verification and Trusted Execution Environments which are all designed to protect against potential inference attacks, gradient leaks and malicious entities while maintaining the benefits of collaboration that the base structure offers. It is only due to these

additional measures that federated learning can be considered as an architecture capable of providing secure and scalable AI solutions.

The main features of the proposed framework may be enumerated as follows: coordination server responsible for initialization, aggregation, and redistribution of the global model; clients – mobile phones, sensors, hospitals, banks, cars, edge nodes that perform local training; local datasets that never go beyond the device they were

collected on; aggregation mechanism that works via FedAvg, FedProx, FedOpt, or any other similar algorithms; privacy and security mechanism offering encryption, differential privacy, secure aggregation, and verification using blockchain technology; and the final global model that is continuously refined within the training process.

4.2 Privacy-Preserving Techniques

Storing raw data locally decreases the chance of privacy breach but does not prevent one entirely

since the parameters and gradients provided by clients may be exposed to privacy leaks using techniques such as inference attacks, model inversion or gradient reconstruction. This means that production-level federated systems add extra measures to the basic infrastructure, namely, Differential Privacy, Secure Multi-Party Computation, and Homomorphic Encryption. Figure 4 shows where these security layers fit into the overall training process..

Figure 4. End-to-End Privacy-Preserving Training Pipeline




Figure 4. End-to-End Privacy-Preserving Training Pipeline

4.3 Differential Privacy

DP accomplishes this by adding meticulously adjusted noise to the update that leaves the client in such a way that whether or not a certain record is included in the learning process has very little impact on the output model. In the context of federated learning, DP is generally used to add noise to the gradients or the parameters that are learned through the local training process. This allows users to cooperate and stay within the bounds of the rules set out by various regulations such as GDPR. The fundamental compromise inherent in DP is well-known – too much noise reduces model precision and too little noise compromises privacy. That is why so much research effort is dedicated to balancing this compromise (Dwork & Roth, 2014; McMahan et al., 2018).

4.4 Secure Multi-Party Computation (SMPC)

Using Secure Multi-Party Computation, several parties can perform computations on a joint function using their input values privately – in this case, aggregation of model updates – without one participant having access to another's unencrypted input. In the federated scenario, each client provides his or her encrypted shares of the input instead of the plaintext updates, while the aggregation of the global model happens using those shares without revealing individual parameters to the server and other clients. Such protection is effective against a variety of attacks both from within and outside the system, including cases when parties do not act honestly and intercepting of communications. It is especially important in fields such as healthcare, finance, military, and politics, where direct transfer of data is impossible. However, this approach requires extra computational and

communication resources due to the need for frequent cryptographic computations that the process of SMPC entails.

4.5 Homomorphic Encryption

Homomorphic encryption (HE) enables computations to be done on encrypted information without having to decrypt it first. The clients encrypt their model parameters prior to transmitting them to the server, which combines the encrypted values to create an encrypted model that is only decipherable by the parties with proper authorization, thus ensuring the privacy of

sensitive information throughout the process. HE has proved to be very useful when applied to industries dealing with information from regulated domains such as healthcare, finance, and critical infrastructures. Although fully homomorphic systems are quite computationally expensive and are rarely used in practice, recent efforts have been focused on more optimized approaches to partial homomorphic or other types of homomorphic encryption that reduce computational requirements and still maintain high security standards (Acar et al., 2018; Yang et al., 2019).

Table 1. Comparative Profile of Privacy-Preserving Techniques

Technique	Privacy Level	Computational Cost	Communication Cost	Key Advantage	Key Limitation
Differential Privacy	High	Low	Low	Mathematically quantifiable guarantee	Noise can reduce accuracy
Secure Multi-Party Computation	Very High	High	High	No raw data disclosed to any party	Communication-intensive
Homomorphic Encryption	Extremely High	Very High	Moderate	Computation on encrypted data	Slow at scale
Secure Aggregation	High	Moderate	Moderate	Protects individual model updates	Weaker against insider collusion

4.6 Machine Learning Algorithms and Comparative Analysis

The optimization technique chosen to align clients' updates has a massive influence on the speed of convergence, communication load, and ability to withstand the heterogeneity in the distribution of data inherent to the deployment of federated systems. While the FedAvg method serves as a default baseline owing to its simplicity and efficiency in communication terms, it suffers from a significant decline in performance in case of highly non-IID data. In order to tackle various aspects of that problem, several follow-up

algorithms have been created: FedProx, FedNova, FedOpt, SCAFFOLD, and FedDyn. Along with optimization techniques, deep learning architectures such as CNN, LSTM, Transformer, and Graph Neural Network models are now widely integrated into federated training pipelines for various applications including medical imaging, natural language processing, and smart cities' network analysis. Algorithms and techniques described above are listed in Table 2; deep learning architectures used in federated learning frameworks are summarized in Table 3.

Table 2. Comparative Analysis of Federated Optimization Algorithms

Algorithm	Core Strategy	Robustness to Non-IID Data	Communication Cost	Key Strength	Key Limitation	Typical Use Case
FedAvg	Weighted averaging of local updates	Low	Low	Simple, lightweight, easy to deploy	Degrades under heavy non-IID skew	General-purpose FL baselines
FedProx	Adds proximal term to limit client drift	Moderate	Low-Moderate	Tolerates heterogeneous clients	Extra local computation per round	Healthcare, cross-silo settings
FedNova	Normalizes aggregation for uneven local steps	Moderate-High	Low	Corrects objective inconsistency	More complex server-side logic	Clients with variable compute power
FedOpt (FedAdam / FedYogi)	Adaptive server-side optimization	High	Moderate	Faster, more stable convergence	Sensitive to hyperparameter tuning	Large-scale cross-device FL
SCAFFOLD	Control variates correct client drift	High	Moderate-High	Strong convergence guarantees	Extra state transmitted per client	Research and cross-silo deployments
FedDyn	Dynamic regularization aligned to global objective	High	Moderate	Consistent convergence under skew	Higher computational overhead	Finance, industrial IoT

Table 3. Deep Learning Architectures Commonly Trained within Federated Pipelines

Architecture	Learning Type	Advantage	Limitation	Common Application
CNN	Deep Learning	Strong image-feature extraction	Computationally intensive	Medical imaging
LSTM	Deep Learning	Captures sequential dependencies	Slower training on long sequences	Healthcare, finance time series
Transformer	Deep Learning	Scales well to large-scale language tasks	High computational cost	Natural language processing

Graph Network	Neural	Deep Learning	Models relational/network structure	Complex implement in settings	to FL	Smart-city network analytics
---------------	--------	---------------	-------------------------------------	-------------------------------	-------	------------------------------

Considered jointly, these two tables offer the following prescription rather than one single solution: FedAvg works as a good starting point if the clients are not too heterogeneous and communication constraints are strict, FedProx and FedNova make sense as soon as heterogeneity kicks in, while SCAFFOLD and FedDyn are a justified cost if convergence stability in the presence of heavy non-IID distribution is the limiting factor. The choice of neural network type itself depends on the nature of the data available – CNNs for imaging problems, LSTM and Transformers for sequences and language data,

and GNNs when relations between clients or entities matter.

5. APPLICATIONS IN SMART DIGITAL ECOSYSTEMS

Federated Machine Learning has moved out of the laboratory and into production-adjacent deployments across a wide range of sectors, precisely because those sectors combine two things that centralized AI has always struggled to reconcile: an abundance of distributed data and a strict requirement that the data stay where it is. Figure 5 maps the five domains examined in this section around the federated learning core.




Figure 5. Cross-Sector Deployment Map of Federated Machine Learning: Enabling collaborative intelligence while ensuring data privacy across diverse domains

Figure 5. Cross-Sector Deployment Map of Federated Machine Learning

5.1 Healthcare

Healthcare might be the area where federated learning adds most value, owing to the very high sensitivity and high value of such information for

developing diagnostic models. Each hospital, laboratory, pharmaceutical company, and research center has a unique set of data about patients in the form of their clinical reports, genetic

information, or diagnostic pictures. These cannot be combined into one database because of the restrictions imposed by legislation and professional ethics. With federated learning, institutions can work together to develop models predicting diseases, detecting cancer, classifying medical pictures, creating plans for personalized therapy and discovering new drugs without revealing any patient data (Rieke et al., 2020).

5.2 Smart Cities

Municipal systems generate a continuous stream of data from traffic sensors, surveillance networks, environmental monitors, connected vehicles, and public-service platforms. Federated learning allows city agencies and service providers to build shared analytical models – for traffic prediction, congestion management, crime pattern detection, disaster response, pollution tracking, and energy optimization – without centralizing citizen data under a single authority. Combined with edge computing, this supports the low-latency, real-time decision-making that urban systems increasingly require, while also narrowing the cybersecurity exposure that comes with centralized city-wide data lakes (Nguyen et al., 2022).

5.3 Internet of Things (IoT)

Ecosystems of IoT include billions of sensors with very little computational and memory capability but which provide data which is sensitive in nature either on a personal or operational level. Through federated learning, these IoT devices can work on building machine learning models on edge without having to send the sensory data to any central server. It helps in applications like predictive maintenance, anomaly detection, smart manufacturing, wearable healthcare monitoring, precision farming, and autonomous transportation. As federated learning reduces

bandwidth usage, the cost of transmission, and latency, it makes it a very good choice for IoT where real-time response and privacy both are important together (Lim et al., 2020).

5.4 Finance

Financial organizations, such as banks and insurance firms, make use of AI to detect frauds, score risks, screen for money laundering, analyze consumer behaviors, and forecast investments; however, the consumer data on which these operations depend is subject to strict privacy and confidentiality protections. With federated learning, financial organizations have the ability to enhance shared models collaboratively without giving away the consumer data that they have access to by sharing encrypted information instead of transaction logs. In effect, federated learning enables an improved cybersecurity strategy while also reducing the risk of a data leak and facilitating regulatory compliance (Yang et al., 2019).

5.5 Industry 5.0

Industry 5.0 integrates human knowledge with collaborative robotics on the foundation of automation provided by preceding generations of industry, and generates huge amounts of data via smart sensors, autonomous machines, and quality control systems. With the help of federated learning, different manufacturing facilities can simultaneously improve their predictive maintenance, quality control, supply chain management, defect detection, and energy consumption without disclosing the data about production processes from their own production facilities to competing companies or any third parties. Together with the IoT, digital twins, edge intelligence, and autonomous robotics, federated learning becomes one of the key technologies in modern human-centered manufacturing.

Table 4. Comparative Profile of Application Domains

Domain	Primary Objective	Dominant Type	Data	Privacy Sensitivity	Principal FL Benefit
Healthcare	Disease prediction and diagnosis	Medical records	/ imaging	Very High	Patient confidentiality preserved
Smart Cities	Traffic and safety management	Sensor	/ surveillance data	High	Secure, real-time urban analytics
IoT	Device-level intelligence	Sensor streams		High	Edge-resident, low-latency learning
Finance	Fraud and risk detection	Transaction	/ behavioral data	Very High	Confidential cross-institution modeling
Industry 5.0	Predictive maintenance	Manufacturing	/ sensor data	High	Protection of proprietary production data

6. CHALLENGES AND LIMITATIONS

The strengths of Federated Machine Learning architecture do not absolve the approach from having to navigate through a host of technical and organizational challenges that have so far been shaping its progress and even blocking it. One of the most persistent of these challenges is the non-IID distribution of the data among the clients, which causes slow convergence and poor performance compared to centralized training on the same data. The challenges of heterogeneous hardware, unstable connection, latency, limited capabilities of edge devices all contribute to the complexity of coordinating the processes. Even

with the data stored locally, privacy cannot be considered secure due to the risks of model-inversion attack, gradient leakage, membership inference attacks, poisoning, and backdoor attacks. The methods developed to combat these issues – Differential Privacy, SMPC, Homomorphic Encryption – impose their own overheads, while organizational challenges, including interoperability, regulatory issues, data management, and lack of standardized federated learning architectures complicate things further. These constraints can be categorized as shown in figure 6.




Figure 6. Principal Constraints Facing Federated Deployment

Addressing these limitations calls for continued work on efficient optimization algorithms, lightweight cryptographic protocols, adaptive communication strategies, and security mechanisms capable of holding privacy, scalability, and model performance in balance simultaneously (Lim et al., 2020).

Table 5. Current Challenges and Proposed Mitigations

Challenge	Principal Impact	Proposed Mitigation
Non-IID Data	Reduced model accuracy	Personalized federated learning
Communication Cost	Slow convergence	Client selection and update compression
Gradient Leakage	Privacy breach	Differential Privacy
Poisoning Attacks	Corrupted global model	Blockchain-backed secure aggregation
Device Heterogeneity	Training instability	Adaptive server-side optimization

7. DISCUSSION

Based on the findings in this literature synthesis, it is apparent that there is one recurring trend in privacy-preserving AI: it is transitioning from being centered around centralized data processing to being based on decentralized and collaborative

learning as a standard way. Instead of relying on sensitive data being sent to a single server (Chen et al., 2021), federated learning enables organizations to maintain full control of their own datasets while continuing to refine the global machine learning model; an approach that has

been especially useful in heavily regulated industries such as healthcare, banking, smart cities, and Industry 5.0 (Dwork & Roth, 2014).

In addition, there is no basis to regard federated learning as a completely solved problem. Differential Privacy, SMPC, Homomorphic Encryption, Secure Aggregation, and the integration of blockchain significantly improve system security. However, model poisoning, communication inefficiency, client heterogeneity (Kairouz et al., 2021) and computational complexity are still problems and not just notes. The progress in edge computing, Explainable AI, and foundation models extends the capabilities of federated systems (Hu et al., 2024), but achieving a simultaneous balance between accuracy, privacy, computation efficiency, fairness, scalability, and energy consumption is still one of the genuine challenges in the field, and the tables developed in Sections 4 and 5 of this literature review are provided as a tool for this balancing act. As a whole, Federated Machine Learning seems more like an emerging systems science than a matured technology (Mothukuri et al., 2021).

8. CONCLUSION

Federated Machine Learning has proved to be one of the most valid ways forward when it comes to privacy-preserving Artificial Intelligence in contemporary smart digital ecosystems. Instead of having to transfer the data to the place where the model is going to be trained, federated learning learns the model in the place where the data lies, which has allowed federated learning to avoid most of the problems associated with privacy, security, and compliance, which centralized machine learning had failed to address. The addition of Differential Privacy, Secure Multi-Party Computation, and Homomorphic Encryption to the architecture makes sure that

sharing of data is possible even when there could have been no other way.

This review has shown that federated learning's benefits extend across healthcare, smart cities, IoT, finance, and Industry 5.0, while also making clear that non-IID data, communication overhead, adversarial threats, computational cost, interoperability, and scalability continue to shape what is practically achievable. Emerging directions – edge intelligence, blockchain, Explainable AI, generative AI, and large language models – offer plausible routes past several of these constraints, and are likely to define the next phase of the field's development.

In sum, Federated Machine Learning represents a genuine structural shift in how collaborative AI can be built responsibly. Its capacity to hold data privacy, distributed intelligence, and collective learning together in one architecture makes it a foundational technology for the next stage of digital society, provided that continued interdisciplinary research, technical refinement, and coherent governance standards keep pace with its adoption.

9. FUTURE RESEARCH DIRECTIONS

Future research needs to focus on developing federated learning paradigms that are sufficiently intelligent and adaptive to deal with the increasing complexity of digital ecosystems. There is a strong need for the development of adaptive optimization methods that can cope with difficult non-IID data and heterogeneous clients without increasing communication costs. The combination of federated learning with Large Language Models, Generative AI, Explainable AI, Digital Twins, Blockchain, Edge Intelligence, Quantum Machine Learning, and 6G networking is an especially interesting group of potential future directions. This is illustrated in Figure 7 below.

Figure 7. Emerging Research Trajectories for Federated AI




Figure 7. Emerging Research Trajectories for Federated AI

In addition to all these technical concerns, energy-efficient federated learning for constrained IoT devices, algorithms of fairness-aware aggregation, which prevent discrimination and act against any form of bias in algorithms, and personalized federated learning are all worth consideration. On the other hand, there is also the part of federated learning which is linked to governance issues, namely, international standardization, ethical AI supervision, automation of privacy management tools, and cross-domain federated architecture, which allow collaboration of health care, finance, industry, education, and public sectors. All of these are required for turning the federated learning concept into reliable infrastructure for NgAI (Nguyen et al., 2022).

REFERENCES

- Abreha, H. G., Hayajneh, M., & Serhani, M. A. (2022). Federated learning in edge computing: A systematic survey. *Sensors*, 22(2), 450. <https://doi.org/10.3390/s22020450>
- Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2018). A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys*, 51(4), 1–35. <https://doi.org/10.1145/3214303>
- Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1175–1191). ACM. <https://doi.org/10.1145/3133956.3133982>

- Briggs, C., Fan, Z., & Andras, P. (2020). A review of privacy-preserving federated learning for the Internet of Things. *Internet of Things*, 14, 100347.
- Chen, M., Poor, H. V., Saad, W., & Cui, S. (2021). Wireless communications for collaborative federated learning. *IEEE Communications Magazine*, 59(12), 48–54.
- Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407. <https://doi.org/10.1561/04000000042>
- Fu, J., Hong, Y., Ling, X., Wang, L., Ran, X., Sun, Z., Wang, W. H., Chen, Z., & Cao, Y. (2024). Differentially private federated learning: A systematic review.
- Hu, K., Gong, S., Zhang, Q., Seng, C., Xia, M., & Jiang, S. (2024). An overview of implementing security and privacy in federated learning. *Artificial Intelligence Review*, 57, Article 204. <https://doi.org/10.1007/s10462-024-10846-8>
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Orazio, A., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2), 1–210. <https://doi.org/10.1561/22000000083>
- Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60. <https://doi.org/10.1109/MSP.2020.2975749>
- Lim, W. Y. B., Luong, N. C., Hoang, D. T., Jiao, Y., Jin, Y., Han, Z., Li, Y., Kim, D. I., & Poor, H. V. (2020). Federated learning in mobile edge networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(3), 2031–2063. <https://doi.org/10.1109/COMST.2020.2986024>
- Liu, B., Lv, N., Guo, Y., & Li, Y. (2023). Recent advances on federated learning: A systematic survey.
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. Y. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)* (Vol. 54, pp. 1273–1282).
- McMahan, H. B., Ramage, D., Talwar, K., & Zhang, L. (2018). Learning differentially private recurrent language models. *International Conference on Learning Representations (ICLR)*.
- Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115, 619–640. <https://doi.org/10.1016/j.future.2020.10.007>
- Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., Vincent Poor, H., & Dobre, O. A. (2022). Federated learning for Internet of Things: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 24(3), 1622–1658. <https://doi.org/10.1109/COMST.2021.3125206>
- Nguyen, G., Sáinz-Pardo Díaz, J., Calatrava, A., Berberi, L., Lytvyn, O., Kozlov, V., Tran, V., Moltó, G., & López García, Á. (2025). Landscape of machine learning evolution: Privacy-preserving federated learning frameworks and tools. *Artificial Intelligence Review*, 58, Article 51. <https://doi.org/10.1007/s10462-024-11036-2>
- Qiu, Y., Liu, Y., Li, K., & Chen, J. (2021). Federated learning for healthcare informatics: Recent advances, challenges, and future directions. *IEEE Internet of Things Journal*, 8(21), 15802–15824.

- Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R. M., Trask, A., Xu, D., Baust, M., & Cardoso, M. J. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3(119), 1–7. <https://doi.org/10.1038/s41746-020-00323-1>
- Saha, S., Hota, A., Chattopadhyay, A. K., Nag, A., & Nandi, S. (2024). A multifaceted survey on privacy preservation of federated learning: Progress, challenges, and opportunities. *Artificial Intelligence Review*, 57, Article 184. <https://doi.org/10.1007/s10462-024-10766-7>
- Shen, M., Tang, X., Zhu, L., Du, X., & Guizani, M. (2022). Privacy-preserving machine learning in federated systems: A survey. *IEEE Network*, 36(2), 38–45.
- Sun, Y., Peng, M., Zhou, Y., Huang, Y., & Mao, S. (2021). Application of federated learning in 6G networks: Security, privacy, and intelligence. *IEEE Wireless Communications*, 28(2), 12–19.
- Truong, N., Sun, K., Wang, S., Guitton, F., & Guo, Y. (2020). Privacy preservation in federated learning: An insightful survey from the GDPR perspective.
- Xu, J., Glicksberg, B. S., Su, C., Walker, P., Bian, J., & Wang, F. (2021). Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1), 1–19.
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19. <https://doi.org/10.1145/3298981>
- Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. *Knowledge-Based Systems*, 216, 106775. <https://doi.org/10.1016/j.knosys.2021.106775>
- Zhang, W., Zhou, T., Lu, Q., Wang, X., Zhu, C., Sun, H., ang, Z., Lo, S. K., & Xu, F. Y. (2021). Dynamic fusion-based federated learning for COVID-19 detection. *IEEE Internet of Things Journal*, 8(21), 15884–15891.
- Zhou, Z., Li, X., & Wang, L. (2022). Blockchain-enabled federated learning for secure and trustworthy artificial intelligence: A survey. *Future Generation Computer Systems*, 129, 247–263.