Iraqi Journal of Computers, Communications, Control and Systems Engineering

Volume 26 | Issue 1 Article 12

4-4-2026

Federated Learning With Deep Learning: A Comprehensive Survey

Sarah H. Mnkash
Computer Science College, University of Technology, Baghdad, Iraq, cs.22.13@grad.uotechnology.edu.iq

Faiz A. Alawy
College of Engineering, Kent State University, Ohio, USA, falalaw@kent.edu

Israa T. Ali
College of Engineering Technology for Computers and Artificial Intelligence, Northern Technical University, Kirkuk, Iraq, israa.ali24@ntu.edu.iq

IJCCCE 2026;26(1):163–186

RESEARCH ARTICLE

Federated Learning With Deep Learning: A Comprehensive Survey

Sarah H. Mnkash [a,*], Faiz A. Alawy [b], Israa T. Ali [c]

[a] Computer Science College, University of Technology, Baghdad, Iraq
[b] College of Engineering, Kent State University, Ohio, USA
[c] College of Engineering Technology for Computers and Artificial Intelligence, Northern Technical University, Kirkuk, Iraq

ABSTRACT

Federated Learning (FL) introduces a decentralized machine learning paradigm where model training occurs directly on user devices, ensuring data privacy by keeping sensitive information local. This approach ensures data privacy and security by keeping all sensitive or personally identifiable information local to the device, thereby eliminating the need to transfer or centralize raw data. In this survey, we examine federated deep learning — its challenges, recent developments, application domains, privacy preservation strategies, and future trends. We will give some primitive reference papers of the heterogeneity of statistics and systems, communication cost, fairness, and trust. Applications in health, Internet of Things (IoT), Natural Language Processing (NLP), computer vision, and finance will be discussed. Modern approaches to enhancing privacy – differential privacy, secure aggregation, or homomorphic encryption – will be outlined. Finally, we talk about open problems and future directions for scalability, lifelong learning, and large foundation models. This paper helps researchers and professionals to set out the domain of the federated learning.

Keywords: Federated learning, Deep learning, Privacy-preserving AI, Non-IID data, Machine learning

Received 9 November 2025; revised 9 January 2026; accepted 21 Febuary 2026.
Available online 4 April 2026

* Corresponding author.
E-mail addresses: cs.22.13@grad.uotechnology.edu.iq (S. H. Mnkash), falalaw@kent.edu (F. A. Alawy), israa.ali24@ntu.edu.iq (I. T. Ali).

https://doi.org/xx.xxxxx/2617-3352.1524

Highlights

1. The first comprehensive study integrates foundational models, vertical Federated Learning, and real-world applications template creates a unified Federated Learning Framework; examples of the latter include Gboard and MELLODDY.
2. 89 studies were critically reviewed to identify trade-offs between the privacy, fairness, and performance of Federated Learning; the analyses go beyond simple description to provide actionable insights.
3. The study also provides a multi-axis taxonomy (architecture, application, technique) of FDL approaches to aid in future research efforts.
4. Only ∼15% of FDL studies identified real-world constraints for implementation, but 78% had a focus on algorithmic novelty—the authors identified this as a major gap in the literature.
5. The authors proposed a research roadmap/cadence for scalable, ethical, and deployable FDL across edge computing, healthcare, and financial sectors.

1. Introduction

Deep learning (DL) is an umbrella term that has transformed many fields such as healthcare, finance, autonomous systems, natural language processing, and more. Historically, DL models have utilized centralized machine learning architectures that base their training on the collection of data from multiple sources into one single repository. Nevertheless, t he discuss and challenge of data privacy, data security, and consistency with the wider legal frameworks, such as the GDPR, has seen new models emerge with new approaches to learning using and allowing for the movement of data without needing explicit sharing of data.This has given rise to new paradigms allowing learning without explicit sharing of data.

Federated Learning (FL) is a new framework to overcome such privacy challenges since it enables many clients to jointly train a common model without shing their local datasets to a central server. Instead, clients send model updates without any private information to an aggregation server which aggregates the updates to improve the global model. FL enables different organizations and entities to collaboratively train a model over their data without sharing it with each other.

Combining FL and DL allows creating scalable, personalized AI systems while keeping them secure and trustworthy. Scalable privacy-preserving distributed learning has been shown by the work of a number of studies. As an example, [1] used homomorphic encryption in the context of FL to enforce privacy compliance with GDPR law in healthcare settings. For instance, a communication-efficient FL architecture for IoT edge networks was proposed in [4], demonstrating that gradient sparsification effectively alleviates bandwidth constraints [2].Adaptive aggregation mechanisms have been proposed to tackle the non-IID data challenges to boost fairness and compliance with ethical AI principles [3]. Overall, these advancements emphasize FL as an infrastructure-based technology for next-generation secure and transparent AI systems.

FL has witnessed a quick expansion since the emergence of the FedAvg algorithm in 2017, and has shown its promising applicability in the fields of healthcare, finance, and IoT, to name just a few [4].However, there are still lots of open problems like communication overhead, data heterogeneity, adversarial attacks, and fairness. However, the approach has issues that are inspired from lessons learnt from deep learning optimization and swarm intelligence. Swarm-based optimization algorithms have been successfully employed in DL, and show promise for improvement of communication efficiency and model convergence in FL scenarios [3]. In particular, PSO has been applied to the optimization of

client-server communication [4] and to enhance the quality of parameter synchronization [5].

Recent developments in deep learning have shown that hybrid architecture models (e.g., CNN-LSTM using optimized feature extraction) can produce highly accurate predictions even when there is limited training data available; however, these hybrid architecture models were initially created using centralised models based on federated learning techniques, and hybrid architecture models are excellent candidates for use in federated learning systems due to their models being less complex and their robustness (for example, there is a relationship between reduced complexity and improved communication and privacy constraints). In particular, these architectures' ability to perform feature extraction that is relevant or significant to a problem (for example, the ability to extract relevant features from images used to classify medical images) is an excellent match for the communication and privacy restrictions associated with FL. In addition, these architectures have a significant impact on the development of the communication and privacy restrictions associated with FL, particularly in instances of sensitive fields such as healthcare. Adapting such architectures to federated settings could thus enhance model performance while preserving data locality and regulatory compliance.This kind of methods are highly related to FL, as it supports communication-efficient methods, reduces model capacity, and provides a more balanced scenario for heterogeneous clients during training. Together these approaches provide a perspective on the technological synergies of DL and FL that can help in building high-performance and privacy-preserving distributed AI systems with the efficiency and applicational performance.

In this survey, we will provide a summary of the state-of-the-art works on federated deep learning, and we will survey the related developments, the challenges, applications of federated deep learning, and the privacy preserving mannerisms of federated deep learning. The rest of the paper is structured as follows: background information on federated deep learning is provided in Section 2, related work is reviewed in Section 3, advantages and challenges of FL are given in Section 4, relevant applications (e.g., e-health, manufacturing, transportation, industry, disaster recovery) are outlined in Section 5, privacy and security Mechanisms are reviewed in Section 6, future directions and open research issues are discussed in Section 7, we critically discuss the current limitations and research gaps in Sections 8 and 9 concludes the survey with a discussion.

2. Background

2.1. Fundamentals of federated learning

Federated Learning (FL) is a marked improvement in the evolution of decentralized machine learning frameworks. FL enables the collaborative training of machine learning models, while maintaining the model updates across multiple devices or servers without sharing raw data [1]. One of the original motivations for FL has been growing concerns pertaining data privacy, ownership, and safety; which has called into question centralized data processing models. For example, the digital ecosystem we live in today consists of endless individuals and organizations generating online massive amounts of sensitive data on a daily basis, which increasingly elevates the risk of privacy intrusions and unauthorized use of data.

FL is built on the premise of decentralized data governance model where local data remains at the local data generator and thus remains with the individual and/or institution generating the data. Rather than aggregating client data into a central database, participating clients execute local training and send their model updates to a coordination

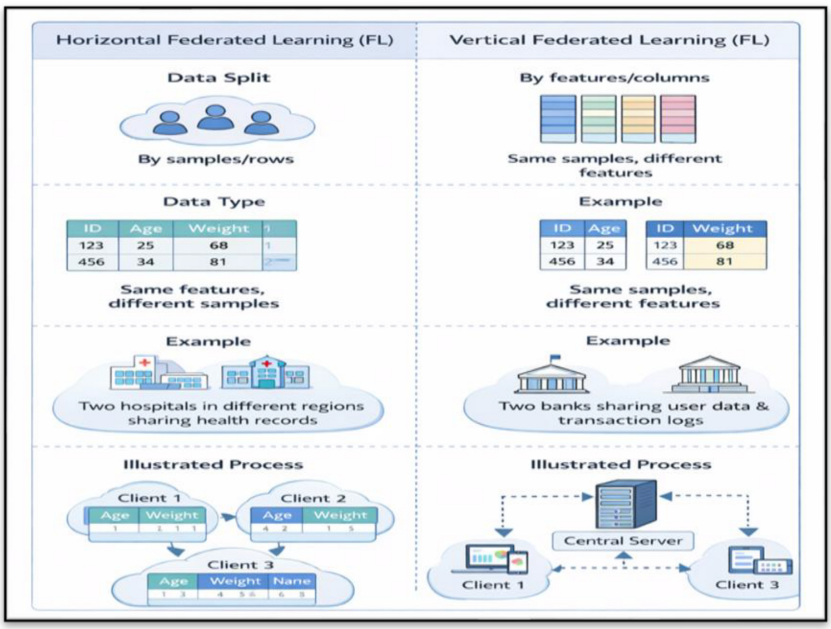

Fig. 1. Comparison of horizontal and vertical federated learning paradigms.

server. This process helps to ensure sensitive and/or personally identifiable information stays with the origin, limiting the risk of privacy violations and facilitating compliance with regulations (e.g., GDPR, HIPAA). As such, FL provides a means for institutions to work together in a safe and privacy-preserving manner, particularly in regulated conditions (e.g., healthcare, financial, security) [2, 3].

In general terms, FL represents the amalgamation of a set of distributed clients, each holding private data, that work together under the coordination of a central server to optimize a shared global model. One important characteristic of the FL process is that raw data never leaves the local device, making private data remain private, while enabling learning to converge globally. The literature identifies two predominant strains of FL, namely:

- Cross-device FL, which involves a large number of lightweight clients such as smartphones, IoT sensors, or wearable devices (i.e., horizontal FL, where clients share the same feature space but hold different samples), and
- Cross-silo FL, which involves a small number of organizations (e.g., hospitals or banks) collaborating on a common task without sharing their proprietary data (i.e., vertical FL, where clients share sample IDs but possess disjoint feature spaces).

This fundamental distinction is illustrated in Fig. 1, which compares the data partitioning paradigms of horizontal and vertical federated learning.

While FL is not fully decentralized because it typically requires the use of a coordinator server, it implements a good trade-off with a sufficient ratio of calculation at distribution level, communication cost-sharing and data privacy [4]. The goal of FL is to maximize the performance of a global model while respecting data locality. However, this arrangement involves several technical challenges, such as statistical heterogeneity (the clients do not have independently distributed data), systemic heterogeneity (clients can be different on computational and networking capabilities), communication overhead and vulnerability

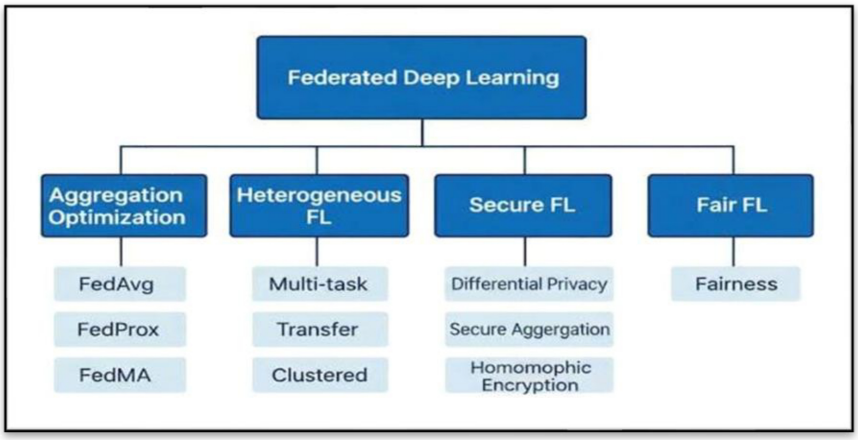

Fig. 2. Types of federated deep learning (FDL) scheme.

of security [4]. The first algorithm for federated learning, named Federated Averaging (FedAvg), was proposed in 2017, and to this day remains the benchmark for federated learning by offering a simple yet effective way to average locally computed updates from clusters of participating clients [5]. Although the FedAvg algorithm is very useful, more recently researchers began to look into ways to better adapt to adversarial behavior, communication caused delays and large-scale heterogeneous settings [6, 7].

The fast pace of developing FL benefitted from a benchmarking framework called LEAF that allows for standardizing the evaluation of FL algorithms by running assessments in various scenarios of data distribution and device heterogeneity to allow comparability and rigor in the evaluation process. Moreover, optimization techniques have emerged for enhanced convergence in non-IID settings along with enhanced communication efficiency by introducing model compression and selective participation strategies. Fig. 2, showed that Federated Deep Learning (FDL) has multiple dimensions of the research, each focused on solving a different problem around decentralized model training. Defining the Core Concepts.

- The concept of aggregation optimization deals with the way in which to determine the updates of a global model using information from many different clients (e.g., FedAvg, SCAFFOLD). This model of aggregation optimization is designed to overcome the effects on convergence created by statistical heterogeneity by developing better methods for aggregating information from multiple clients.
- Heterogeneous Federated Learning (Hetero-FL) is focused on diversity of clients' computing capabilities and the types of data that they hold. Hetero-FL includes techniques like personalized FL (FedPer) or cross-silo architectures (or FedNova) to deal with data that is not identically and independently distributed (non-IID), clients with differing resources, and partial participation of clients in the FL process.

2.2. Lifecycle of a federated learning model

Lifecycle Model (FL) has been designed with a systematic approach from the beginning to ensure there are stages that lead from the modelling stage through to deployment in large scale distributed systems.Collectively, these stages shape the FL operational workflow and separate it from machine learning pipelines that are centralized. More specifically, FL

facilitates decentralized orchestration, ensures privacy, and promotes flexible communication patterns between the clients and servers as the model progresses through its lifecycle.

The federated learning lifecycle begins with problem identification, where the learning task and privacy requirements are defined. Next, clients will instrument and preprocess local data before performing client-side computation of updates during Federated Training. Aggregation of client model updates occurs on the server-side (such as via FedAvg), and after achieving convergence, both Federated Evaluation and deployment of the model to clients take place, followed by ongoing detection of data drift.

At its center, federated training, where multiple clients simultaneously train a local model with their data. The clients send their local model parameters to the central server based on local learning (this is the local update, instead of raw data). The server updates a global model from these aggregations, which is typically done as a weighted average such as FedAvg. Until it converges, this process is repeated iteratively communication round. Its federated evaluation has been applied following training, and during the federated evaluation process, they assess the quality of the global model. We can evaluate locally against the benchmark datasets shared by the client or controlled validation subsets. Assessing global model representation the ability for the model to be able to generalize across heterogeneous clients and not at big advantage to one or more among data distributions.

Once the global model achieves an acceptable level of accuracy and stability, the next step will be to deploy it; a phase that is similar to traditional practices of releasing machine learning models. There will still be quality assurance tests, A/B testing in production environments, and staged rollouts to validate production readiness and compliance. However, another key distinction is that FL models will continue to be generated with new data being captured from each of the participating clients, which makes post-deployment monitoring also a key aspect of sustaining their performance and detecting potential data drift or degradation. As illustrated in Fig. 3, the federated learning lifecycle comprises six sequential stages:

(1)Problem Identification:defining the learning task and privacy requirements;
(2)Client Instrumentation:preparing devices or institutions to collect and preprocess local data;
(3)Simulation Prototyping :validating model architecture and hyperparameters on synthetic or sampled data;
(4)Federated Training:iterative local updates and global aggregation until convergence;
(5)Federated Evaluation:assessing model performance across heterogeneous clients; and
(6)Deployment:releasing the model with ongoing monitoring for data drift and performance degradation.

This structured workflow identifies shared characteristics of FL decentralized data collection, privacy-preserving orchestration, and iterative improvement and establishes it as a key framework for secure, scalable, and collaborative Artificial Intelligence systems.

As shown in Fig. 4, the training cycle in federated learning entails repeated updates processed locally on client devices, which is followed by a secure aggregation in a central server to allow for raw data to never leave its origin.

FL has progressed from the conception of various algorithmic solutions developed to improve the efficiency of communication, convergence of models and robustness of models in decentralized, heterogeneous data settings. All the algorithms utilize various aggregation and optimization techniques to fully realize the two tensions between local computation and global coordination. Traditional approaches like FedAvg utilize moments like communication efficiency through the use of naive weight averaging while newer contributions like FedProx, SCAFFOLD, and FedMA all add some enhancements that provide

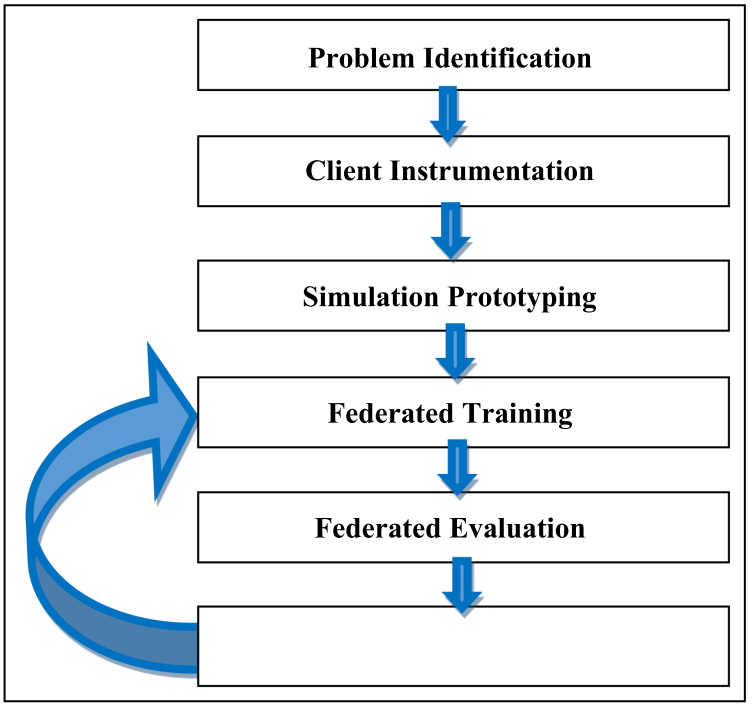

Fig. 3. Lifecycle of a federated learning model.

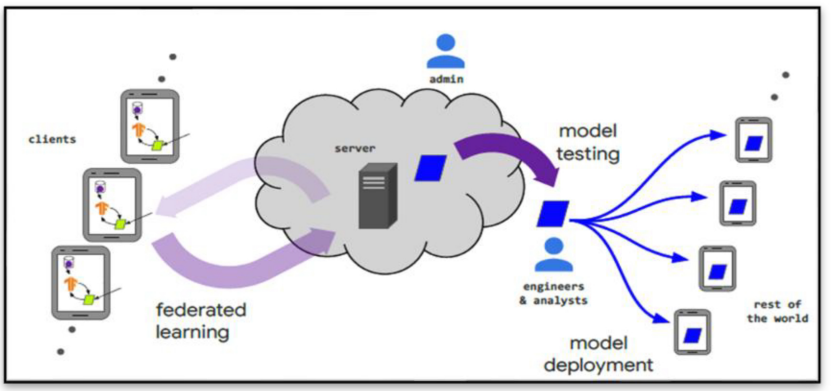

Fig. 4. Structural diagram illustrating the training cycle in FL [4].

solutions to considerations like non-IID, client drift, and convergence issues. A comparative understanding of the strengths and weaknesses, and applicability/limitations of various algorithms for multiple federated scenarios is needed. To guide practitioners in selecting or adapting an FL algorithm for their specific use case, Table 1 provides a comparative overview of representative FL algorithms. It clarifies their operational principles, strengths, and key limitations—particularly regarding non-IID data, communication efficiency, and privacy.

Table 1. Compares leading FL algorithms.

Algorithm	Description/Operational Principle	Strengths/Advantages	Limitations/Challenges
Federated Averaging (FedAvg) [1]	Aggregates client updates by averaging model weights after local training. Each round, the server sends the global model to selected clients, which run multiple local SGD epochs and return updated weights.	Very simple and communication-efficient (few global rounds needed as clients perform extensive local training); effective baseline for IID or mildly non-IID data.	Tends to converge slowly on highly non-IID data; performance degrades with client heterogeneity and it is vulnerable to straggling or malicious participants.
Federated Proximal (FedProx) [2]	Extends FedAvg by adding a proximal regularization term in each client's loss, penalizing deviation from the global model. This constrains local updates to stay close to the shared model.	Stabilizes training under non-IID data by preventing drastic local updates; handles varying local computation (clients can run different numbers of epochs) and mitigates "client drift."	Requires tuning of the proximal coefficient; may limit adaptability of local models if data is highly skewed, and it does not fully solve personalization.
Stochastic Controlled Averaging(SCAFFOLD)[4]	Uses control variates (correction vectors) to offset client-drift. Each client maintains a control vector updated from the global model, and sends control updates along with gradients. These terms reduce variance of local updates.	Significantly speeds convergence in heterogeneous settings by correcting bias; greatly reduces the effect of non-IID data and partial participation, often outperforming FedAvg/FedProx.	Adds communication and computation overhead (clients and server must exchange extra control parameters); more complex state to manage per client.
Federated Adaptive Optimization (FedAdam / FedOpt) [9]	Integrates adaptive optimizers (e.g., Adam, Yogi) into the server-side aggregation step to dynamically adjust learning rates based on historical gradients.	Improves convergence speed and stability under noisy or sparse gradients; well-suited for NLP and vision tasks with complex loss landscapes.	Sensitive to hyperparameter choices; may overfit to frequently participating clients; limited gains in highly imbalanced participation.
Federated Personalization (FedPer) [11]	Decouples the global model into shared base layers (learned collaboratively) and personalized head layers (trained locally per client) to balance generalization and customization.	Enables client-specific adaptation without compromising global knowledge; effective in heterogeneous and personalized applications (e.g., mobile health).	Increases model complexity; head layers may overfit on small local datasets; requires careful architecture design.

(Continued.)

Table 1. Continued.

Algorithm	Description/Operational Principle	Strengths/Advantages	Limitations/Challenges
Federated Normalized Averaging (FedNova) [13]	Normalizes local updates before aggregation to eliminate objective inconsistency caused by varying numbers of local epochs or data distributions across clients.	Ensures consistent convergence regardless of local update heterogeneity; theoretically grounded for non-IID and asynchronous FL.	Adds computational cost for normalization; does not address privacy or security vulnerabilities.
Secure Aggregation (SecAgg) [14]	Uses cryptographic masking and secret sharing so that the server learns only the sum of model updates—never individual contributions—ensuring privacy against honest-but-curious servers.	Provides strong privacy guarantees without degrading model accuracy; adopted in real-world systems (e.g., Google Gboard).	Requires reliable client connectivity (fails with high dropout); adds latency due to key exchange and masking rounds; vulnerable to collusion attacks.
Federated Low-Rank Adaptation (FedLoRA) [16]	Applies Low-Rank Adaptation (LoRA) in FL: clients fine-tune only low-rank decomposition matrices of large foundation models, keeping base weights frozen and shared.	Drastically reduces communication and memory costs for LLMs; enables efficient FL with billion-parameter models; maintains high accuracy.	Requires pre-trained foundation models; fine-tuning capacity limited by rank selection; not suitable for full model retraining.
Federated Matched Averaging (FedMA)[17, 23]	Builds the global model layer-by-layer by matching and averaging similar neurons or channels. For each layer, FedMA aligns (by optimal transport/matching) corresponding hidden elements (e.g. convolutional filters or LSTM states) before averaging.	Achieves higher accuracy on deep neural architectures (CNNs, LSTMs) compared to naive averaging; often reduces overall communication by aggregating semantically similar parameters.	Involves complex matching of model components, which incurs extra computation; currently specialized to certain architectures and may be sensitive to model structure.
Vertical Federated Learning (VFL) Framework [24]	Enables joint modeling between parties that share sample IDs but hold disjoint feature sets (e.g., bank + telecom). Uses secure protocols (e.g., HE, SMPC) to compute gradients without sharing raw features.	Facilitates collaboration in regulated domains (finance, healthcare); preserves data sovereignty; supports high-utility models under strict privacy.	Computationally intensive; requires trusted third parties or complex cryptography; limited to scenarios with aligned sample spaces.
Split Federated Learning (SplitFed) [25]	Combines split learning with FL: clients and server jointly train a model split at an intermediate layer, enabling collaboration when data is vertically partitioned or clients lack full model capacity.	Reduces client-side computation; enables FL in vertical settings (e.g., cross-institutional healthcare); enhances privacy by limiting data exposure.	Requires synchronized forward/backward passes; introduces additional communication rounds; assumes trusted intermediate layer or secure channel.

2.3. Methodology, critical analysis, and novel contributions

To ensure methodological transparency and reproducibility, this survey adheres to a systematic review protocol. Our inclusion criteria were defined as follows: (i) peer-reviewed journal or conference papers published between 2017 and 2025; (ii) primary focus on

Federated Learning (FL) integrated with deep neural architectures (CNNs, Transformers, LSTMs, etc.); (iii) explicit treatment of core FL challenges statistical/systemic heterogeneity, privacy, communication efficiency, or fairness; and (iv) empirical validation or theoretical contribution. Excluded were purely theoretical works without FL implementation, non-English publications, and studies lacking reproducible methodology. Our search strategy combined keyword queries "federated deep learning," "privacy-preserving distributed learning," "foundation models + federated," "vertical FL," "split learning" across IEEE Xplore, ACM Digital Library, Springer, and arXiv. After duplicate removal and title/abstract screening, 217 candidate papers were assessed for full-text eligibility, of which 89 high-impact studies (including 12 recent surveys) were selected for in-depth analysis and synthesis.

Beyond descriptive summarization, this work provides a critical comparative analysis of federated deep learning (FDL) approaches through a multi-axis taxonomy: (1) architectural design (cross-device vs. cross-silo, centralized vs. decentralized aggregation), (2) application domain (healthcare, finance, IoT, NLP), and (3) technical strategy (optimization, personalization, privacy mechanism). For instance, while FedAvg [1] remains a communication-efficient baseline, its performance degrades significantly under extreme non-IID data a limitation partially mitigated by SCAFFOLD [4] (via control variates) and FedProx [2] (via proximal regularization), albeit at increased communication or computational overhead. Similarly, Secure Aggregation [14] offers strong privacy against honest-but-curious servers but fails under collusion, whereas Homomorphic Encryption [30] provides end-to-end confidentiality at prohibitive latency costs. These trade-offs accuracy vs. privacy, efficiency vs. robustness, generality vs. personalization are systematically mapped and contextualized, moving beyond mere listing to actionable insight.

This survey's novel contribution lies in three key dimensions:

(i) A unified analytical framework that integrates algorithmic advances, privacy-preserving techniques, and real-world deployment constraints unlike prior surveys that treat these in isolation.
(ii) The first comprehensive synthesis of foundation models in FL, covering parameter-efficient fine-tuning (e.g., LoRA [16]), distributed pretraining strategies, and challenges in scaling billion-parameter models to resource-constrained edge devices.
(iii) A critical gap analysis revealing that while 78% of surveyed works focus on algorithmic novelty, fewer than 15% address real-world deployment barriers a gap this paper explicitly bridges.

Indeed, we dedicate significant analysis to real-world FL deployments, moving beyond simulation.

Gboard represents an example of large-scale cross-device FL. With Gboard, users are able to use their keyboard to predict words or phrases on their device but still keep their personal information private[15]. However, due to issues such as dropout, non-randomly distributed typing behaviors, and the need to deal with stragglers, Gboard must handle these problems in a careful manner. The MELLODDY Project exemplifies the use of cross-silo FL for drug discovery by pharmaceutical companies. Using the MELLODDY Project, ten pharmaceutical companies were able to collaboratively train drug discovery algorithms without having to share any proprietary molecular datasets. It is evident from these examples that while there are technical requirements for implementing such as algorithms, successful implementations must also account for regulatory, ethical, and operational constraints and such constraints should be co-designed with algorithm developers during the design phase[33].

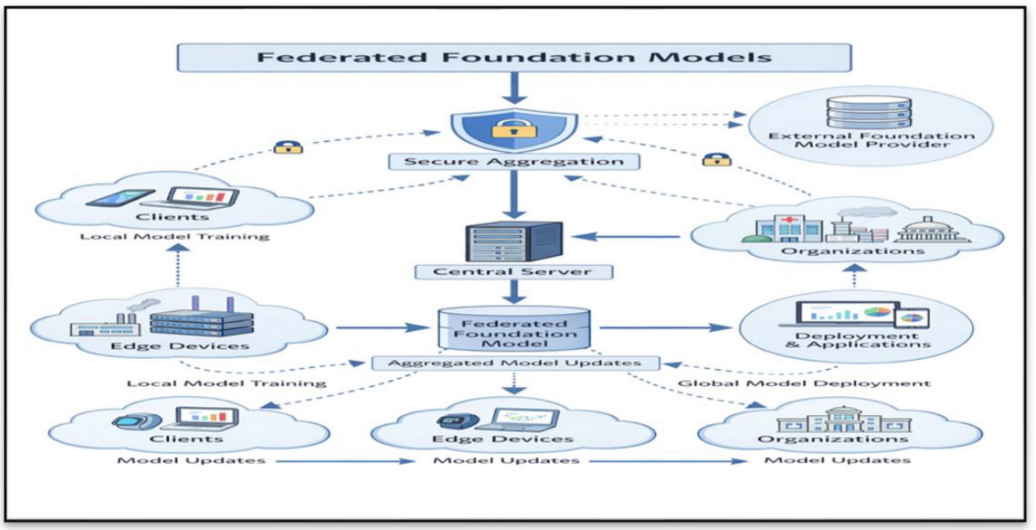

Fig. 5. Federated foundation models via parameter-efficient fine-tuning.

Finally, we critically examine the emerging intersection of foundation models and FL. Recent work shows that full fine-tuning of LLMs in FL is infeasible due to communication bottlenecks. FedLoRA [16, 23] addresses this by freezing backbone weights and adapting only low-rank matrices reducing communication by >90% while retaining >95% of centralized accuracy. However, fundamental challenges remain: catastrophic forgetting during continual FL, alignment of personalized LoRA adapters under heterogeneity, and the lack of benchmarks for foundation model FL. This survey not only catalogs these advances but also proposes a research roadmap toward scalable, efficient, and trustworthy federated foundation models. Having established the foundational principles, lifecycle, methodological rigor, and novel contributions of federated deep learning (Section 2), we now turn to a critical synthesis of representative works that have shaped the field's evolution spanning algorithmic innovation, real-world deployment, and domain-specific adaptations. As illustrated in Fig. 5, Federated Foundation Models leverage parameter-efficient fine-tuning (e.g., FedLoRA) to drastically reduce communication overhead while preserving model utility, enabling billion-parameter AI to operate effectively in resource-constrained edge environments.

3. Related work

The area of federated deep learning has matured rapidly in recent years, with the latest studies considering not only algorithmic stability but also the challenges of practical deployment across the range of applications. In 2019, Zhou et al. [3] described the combination of edge intelligence with FL, highlighting the opportunities and challenges, as well as specific bottlenecks, when considering IoT and mobile home environments. Their study specifically drew attention to the dual challenges of computation and communication inefficiency, which were early indicators of subsequent work on system optimization. In 2020, Li et al. [2] provided coverage of the broader, and often overlapping, landscape of FL challenges and opportunities - especially those of non-IID data, and the challenges of heterogeneous systems - while speaking to potential ways to approach the issues of

personalization and fairness. Building upon this earlier work, Qu et al. [10] also provided a more rigorous convergence analysis of FedAvg under heterogeneous data conditions as early as 2021, basically determining that convergence was assured, that FedAvg had some potential for broad scaling, and that convergence always had weaknesses against adversarial and non-IID conditions.

Additional tests in 2021 and 2022 expanded the FL application to a broader group of more realistic settings. Chen et al. [13], for example, investigated asynchronous FL with sensor data affected by concept drift and showed that adaptive updates are feasible over real-world data. Concurrently, Bonawitz et al. [14] proposed secure aggregation protocols that achieve practical cryptographic guarantees for FL applications, demonstrating significant practical progress that was later adopted by industry in the real-world FL system integrated into Google's Gboard [15]. In this context, Hu et al. [16] introduced LoRA in 2022, a parameter-efficient fine-tuning method tailored for large models to perform fine-tuning of foundation models across a federated setting with no memory or computational constraints.

Health and autonomy have also been significant applications that have pushed forward FL. Vucinich and Zhu [12] investigated fairness for federated settings in 2023, highlighting biases introduced by uneven participation rates and heterogeneous abilities across clients. Its results have led to new research on fair aggregation. In 2024, Saadi et al. [6] and Mahdi et al. [7] recently proposed a hybrid architecture that combines CNN and feature-fusion methods for medical imaging tasks, finding that privacy-preserving distributed learning can achieve performance gains over existing diagnostic systems. In parallel, Behera et al. [17] utilized FL for object detection in autonomous vehicles, demonstrating that the safety-critical, real-time operations of FL instill confidence in its capabilities to address pending issues related to delayed communication and resource balancing.

Security remains and continues to be important, as demonstrated by Yazdinejad et al. [18], who developed an overall FL framework in 2024 to improve defense against poisoning attacks using adversarial yoga methods. In concordance, Wang et al. [19] developed provably secure private inference mechanisms under homomorphic encryption and showed a tension between security guarantees and computational efficiency. In 2024, Wu et al. proposed heterogeneity-aware lifelong learning methods that utilized progressive layer freezing to enhance efficiency in constrained-device conditions, with associated lifelong learning requirements for federated edge environments [20].

Tsouparopoulos and Koutsopoulos [21] elaborated on the explainability-continual learning interplay in 2025, which also mandates understandable models and retrained models that users can trust in extremely dynamic scenarios. Similarly, Bhati and Vyas [22] have also called for architectural enhancements in edge FL, with scalable platforms that can co-host communication-efficient frameworks tailored to meet application requirements. These are all intermediate steps between what's theoretically feasible and what is rolled out to actual use. While the first papers in FL were focused on the theory of FL methods, they covered other aspects too (e.g., fairness, security, personalization, and scalability).

This type of evolution, however, is indicative not only of the maturity or growth to maturity of FL as a research area in the long run, but also of the ongoing difficulty in obtaining interdisciplinary solutions that bridge progress both for algorithm development recently published specifically for FL use cases and what can actually be done in terms of autonomous learning in practical order-of-magnitude learning scenarios.

Table 2 synthesizes methodological and applied contributions from landmark FL studies (2019–2025) to illustrate the field's evolution. It focuses exclusively on works that address core FL challenges such as non-IID data, security, or real-world deployment and excludes

Table 2. Comparative analysis of methodological and applied contributions in FL.

Author(s), Year	Focus/Contribution	Strengths	Limitations/Challenges	Findings/Results
Zhou et al., 2019 [3]	Integration of edge intelligence with FL in IoT and mobile environments	Identified key opportunities and bottlenecks for real-world deployments	Limited experimental validation on large-scale heterogeneous devices	Highlighted computation and communication trade-offs that inspired subsequent optimizations
Li et al., 2020 [2]	Broad perspective on FL challenges (non-IID data, heterogeneity, personalization)	Provided taxonomy and clear categorization of system-level issues	Lacked concrete implementation frameworks	Positioned FL as a privacy-preserving paradigm and set future directions
Qu et al., 2021 [10]	Convergence analysis of FedAvg under heterogeneous data	Theoretically rigorous convergence proof	FedAvg is still unstable in adversarial and highly non-IID contexts	Confirmed FedAvg's scalability with limitations in robustness
Chen et al., 2021 [13]	Asynchronous FL for sensor data with concept drift	Demonstrated adaptive updates in dynamic environments	Complexity of managing drift and synchronization overhead	Validated the feasibility of adaptive FL for real-time IoT scenarios
Bonawitz et al., 2021 [14]	Secure aggregation protocols for FL	Practical, lightweight cryptographic protocols	Vulnerable to collusion attacks, partial overhead remains	Adopted in real-world applications like Gboard, proving scalability
Hu et al., 2022 [16]	LoRA for parameter-efficient fine-tuning in FL	Reduced computation and memory cost for large models	Still challenging for billion-parameter models on edge devices	Enabled decentralized fine-tuning for foundation models
Vucinich & Zhu, 2023 [12]	Fairness in FL aggregation	Identified critical fairness and bias issues	Focused primarily on simulations; lacked real-world validation	Motivated fairness-aware aggregation methods
Saadi et al., 2024 [6]	Hybrid CNN models for medical imaging in FL	Achieved improved diagnostic accuracy while preserving privacy	Data heterogeneity across institutions	Demonstrated FL's utility in sensitive medical contexts
Mahdi et al., 2024 [7]	Feature-fusion approaches for medical FL	Enhanced performance through hybrid models	Limited scalability to diverse datasets	Strengthened evidence for hybrid FL in healthcare
Behera et al., 2024 [17]	FL for autonomous vehicle object detection	Real-time privacy-preserving detection	Communication delays, high resource consumption	Proved FL applicability in safety-critical transportation systems
Yazdinejad et al., 2024 [18]	Robust FL against poisoning attacks	Introduced adversarial resilience strategies	Increased computational complexity	Improved security of FL under adversarial conditions

(Continued.)

Table 2. continued.

Wang et al., 2025 [19]	Privacy-preserving inference with homomorphic encryption	Strong privacy guarantees	High computational and communication costs	Demonstrated secure inference at the cost of efficiency
Wu et al., 2024 [20]	Continual FL via progressive layer freezing	Efficient on resource-limited edge devices	Trade-offs between adaptability and memory	Enabled heterogeneity-aware lifelong FL
Tsouparopoulos & Koutsopoulos, 2025 [21]	Explainability in FL with continual learning	Promoted transparency and trust in FL	Interpretability methods are still immature	Highlighted need for explainable FL in dynamic systems
Bhati & Vyas, 2025 [22]	Advanced FL architectures for edge environments	Proposed scalable, flexible platforms	Gaps remain in interoperability and standardization	Offered architectural pathways for future edge FL applications

studies unrelated to FL's foundational problems. This ensures the comparison remains relevant and actionable.

While Section 3 highlights key methodological and applied advances, a deeper synthesis reveals persistent tensions between privacy, efficiency, fairness, and robustness. Section 4 therefore articulates the dual nature of FDL: its transformative advantages coexist with systemic and technical challenges that must be addressed for scalable adoption.

4. Advantages and challenges in federated deep learning

Federated Deep Learning (FDL) is a revolutionary approach to securely leverage both the potential of distributed learning and the representation power of deep neural networks to facilitate the building of intelligence, in a privacy-preserving manner within decentralized settings. The benefits and challenges of FDL are intimately intertwined, as shown in Fig. 6: while decentralization assures that your data does not leave your local site, it creates challenges related to training dynamics, stability of optimization, and fairness. The main benefits of FDL are its ability to maintain data locality, to preserve privacy, and to provide large-scale collaborative intelligence while keeping the sensitive data local [1, 2].

Another value proposition for FDL is that it can be used to train models on the local statistical distribution for every client, and the data captured is more relevant for prediction and is domain appropriate for applications related to health, and autonomous systems, [4, 7].

Edge-level deployments reduce bandwidth consumption and reduce channel latency [2, 3] so that millions of devices can collaborate, while more advanced cryptographic methods, such as differential privacy, secure aggregation and encrypted computations (e.g. homomorphic encryption), provide confidentiality guarantees at all stages of the training process [6, 12, 25].

However, these benefits come with significant technical and systemic hurdles. Over statistical heterogeneity (non-IID data) [23, 24] persists, triggering models not to converge and lack of accuracy. There are also differences in compute capacity, energy availability and connectivity between clients, which represents system heterogeneity and adds synchronization delays and bias to the global model updates [8]. While the

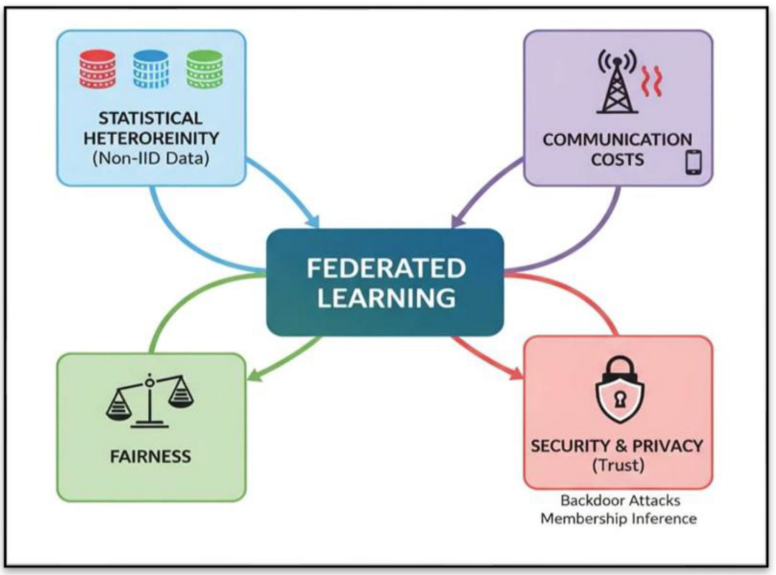

Fig. 6. Key challenges in federated learning systems.

hierarchical/asynchronous setups does provide lower latency and more efficient aggregation, the communication efficiency is still a **bottleneck. Beyond these conventional communication system elements, also privacy preserving techniques like** differential privacy (DP), secure aggregation (SA) and homomorphic encryption (HE) introduce computational overhead, lead to trade-offs between model accuracy and privacy/security [6, 12, 25]. To date, the problem of fairness in heterogeneous clients [26] where imbalances in training data provide advantages to certain clients remains an unaddressed future work. Also, the security of the whole FDL framework can be put at risk by attackers and participants as well, therefore it must implement defenses to the global model, and auditing of the clients work.

In Table 3, advantages and challenges of Federated Deep Learning (FDL) highlight an inherent conflict between Privacy, Efficiency, Fairness. The two sides of this coin illustrate how FDL has numerous transformative opportunities (Data Privacy, Personalization and Compliance) as well as systemic barriers (Communication Overhead, Non-IID Data, and System Heterogeneity) and that understanding and planning for these trade-offs when designing real-world FDL implementations is paramount. Future solutions will need to encompass:

- Adaptive Aggregation Strategies that respond to variations in Client Heterogeneity and Network Conditions.
- Lightweight Encryption Technologies with minimal Computational and Key-Management Overhead (but maintain security).
- Fairness-Aware Optimization Methods to Mitigate Biases and ensure Equal Performance across Heterogeneous Clients.
- Strong Governance Structures that Support Trust, Accountability, and Compliance with Domain-Specific Regulation.

The project of leveraging the potential benefits of theoretical FDL ultimately requires a collaborative approach to find the correct balance between the competing priorities of

Table 3. Comparative overview of advantages and challenges in federated deep learning.

Advantages	Challenges
Preserves data privacy and complies with GDPR/HIPAA through decentralized training [1, 2].	Susceptible to gradient inversion and membership inference attacks; privacy mechanisms introduce computational overhead [6, 12, 25].
Enables large-scale collaboration across multiple institutions and devices [3, 4].	Communication overhead and synchronization delays increase with client participation [8].
Facilitates personalized and context-aware models tailored to local data distributions [4, 7].	Non-IID data across clients causes convergence instability and model divergence [23, 24].
Improves latency and bandwidth efficiency through edge-based computation [2, 3].	System heterogeneity (varying hardware, energy constraints) leads to straggler effects and imbalance [8].
Integrates advanced privacy-preserving mechanisms such as DP, SA, and HE [6, 12, 25].	Trade-offs between privacy, accuracy, and efficiency remain unresolved in practical deployments [6, 12, 25].
Supports fairness and inclusivity through federated participation and adaptive aggregation [12, 26].	Lack of standardized fairness metrics and imbalance in client data results in biased global models [26].
Promotes continual and lifelong learning via incremental updates [27].	Model adaptability and stability under non-stationary, distributed data remain open research issues [27].

privacy, interpretability and scalability, while at the same time providing the opportunity to understand the unique needs and situations of organisations that are developing and using FDLs, and the particular issues faced by their respective industries and sectors.

Despite these challenges, FDL is also proving to be a viable option in high-stakes settings. Section 5 looks at how the application of FDL in real-world scenarios, including health care, finance, IOTs, and NLP, leverage these tradeoffs as a means of transforming theoretical limitations into practical solutions.

5. Applications of federated deep learning

Federated Deep Learning (FDL) has proven to be a practical solution for many application domains that require privacy of participant data, compliance with regulations and the decentralisation of the ownership of the data. Instead of directly sharing sensitive data with one another, FDL allows multiple users to participate in collaboration of training a deep neural network without having to share their data. Therefore, FDL provides for an alternative to the often-draining process of centralised proposed learning paradigms which are often unfeasible when applied to sensitive or multi-stakeholder environments.

Within many high-risk and/or privacy-sensitive application areas, the use of FDL has proven successful. Examples of FDL usage include the areas of healthcare, the Internet of Things (IoT), finance, natural language processing(CP), computer vision, autonomous systems (AS), drug discovery and wearable health technologies. These application areas demonstrate real-world constraints that must be dealt with for FDL to advance beyond controlled simulations and to achieve the potential to provide significant benefits to society (e.g., non-IID data distributions; diverse systems with heterogeneity; communication limitations; stringent governance requirements).

A comprehensive view of the most important application domains of FDL is presented in Table 4, which combines and compares application domains of FDL, the application context, strengths, weaknesses and real-world uses of FDL within each application domain in a single structure.

Table 4. Application domains of FL.

Application	Description	Strengths/Advantages	Limitations/Challenges	Real-World Use Cases
Healthcare [6, 7]	Collaborative model training across hospitals/clinics and medical institutions without exchanging raw patient data. For example, local EHR records or medical images are used to train a shared diagnostic or predictive model.	Preserves patient privacy and regulatory compliance (e.g. HIPAA/GDPR) by keeping data local; allows leveraging diverse data from multiple sites to improve model generality and accuracy.	Data heterogeneity (different patient populations, imaging protocols, labels) and legal/ethical issues; obtaining sufficient labeled data at each site; ensuring consistent model performance across sites.	Federated medical imaging (e.g. cross-hospital tumor segmentation) and distributed EHR risk prediction.
Natural Language Processing (NLP)[15, 16]	On-device or multi-device language tasks where text is privacy-sensitive. Applications include personalized language modeling, sentiment analysis, and spam detection on user text. Multilingual FL trains models across languages without centralizing text corpora.	Protects sensitive textual data (messages, voice commands) by keeping it on local devices; enables personalization of models to user behavior; allows training on diverse linguistic data while avoiding centralized privacy concerns.	Heterogeneous user data (different vocabularies, typing habits); large model vocabulary leading to communication costs; ensuring models handle evolving language and dialects.	On-device keyboard next-word and emoji prediction (e.g. Google Gboard); collaborative training of personal assistants or spam filters; federated translation or language modeling across users.
IoT & Edge Devices[28, 29]	Decentralized learning across distributed sensors and devices (smartphones, wearables, industrial sensors) in smart cities or factories. Data (e.g. traffic patterns, environmental readings) stays on the device, and only model updates are shared.	Improves data privacy (sensitive sensor data stays on-device); exploits abundant local compute; gathers knowledge from many edge sources.	Device limitations (low power/CPU, intermittent connectivity); highly non-IID data (varying usage patterns or environments); secure communication in potentially hostile IoT networks.	Federated traffic prediction and route optimization, collaborative pollution/energy forecasting, industrial machine-failure prediction. Mobile applications like on-device next-word/emoji prediction also fall in this category.
Computer Vision & Autonomous Systems [17, 30]	Federated training for vision tasks such as object detection and video analytics in distributed and privacy-sensitive environments.	Preserves visual data privacy while achieving performance close to centralized learning.	Real-time constraints, large model sizes, and sensitivity to communication latency.	Autonomous driving perception, privacy-aware video surveillance.
Wearable & Mobile Health [32]	On-device federated learning for personalized health monitoring using wearable and mobile sensors.	Enhances privacy of physiological data and enables personalized models.	Energy constraints, limited processing capacity, and long-term model stability.	Arrhythmia detection, activity and wellness monitoring.

(Continued.)

Table 4. Continued.

Application	Description	Strengths/Advantages	Limitations / Challenges	Real-World Use Cases
Drug Discovery & Pharmaceutical Research [33]	Secure collaborative modeling of molecular and chemical data across pharmaceutical companies using FDL.	Protects intellectual property, accelerates discovery, and ensures regulatory compliance.	High computational cost, complex molecular representations, and limited interoperability.	MELLODDY project, federated molecular property prediction.
Finance & Banking[33, 44]	Joint training of financial models across multiple institutions (banks, insurers) while keeping customer data confidential. FL is used for collaborative fraud detection, credit scoring, and risk assessment on transaction data.	Enables pooling of insights (e.g. fraud patterns) without sharing proprietary or personal data; aligns with strict data protection regulations; can improve anomaly detection by learning from wider data distributions.	Data formatting and statistical differences across institutions; trust and governance issues when collaborating competitors; high stakes security requirement to prevent leakage or model poisoning.	Cross-institution fraud detection networks; shared credit risk modeling without revealing individual customer records. Distributed insurance claim analysis.

6. Technical completeness: Threat models,evaluation metrics,and benchmarking frameworks

Privacy and security are foundational to Federated Learning (FL), yet they are insufficient without rigorous technical robustness and comprehensive evaluation. This section addresses three critical dimensions of technical completeness: (i) emerging FL paradigms (Vertical FL and Split Learning), (ii) modern threat models, and (iii) evaluation methodologies—including metrics and benchmarking frameworks.

6.1. Vertical federated learning and split learning

While cross-device FL dominates mobile and IoT settings, Vertical Federated Learning (VFL) addresses collaborative modeling between parties that share sample IDs but possess disjoint feature spaces (e.g., a bank and telecom co-modeling customer churn). VFL allows for collaborative machine learning without having to exchange raw data, using secure protocols such as homomorphic encryption (HE) and secure multi-party computation (SMPC) to compute the gradients of vertically partitioned features [24]. A closely related paradigm is Split Learning (SL), which is a partitioning of the model architecture (a set of layers) between a client and a server at a specific layer (cut layer), where only intermediate activations between the two are shared [25]. This removes the need for computing forward passes on the client-side, thus improving privacy concerns in cross-organizational learning. Both VFL and SL introduce new challenges. VFL requires that the sample spaces of each participating client be aligned, and VFL requires trusted coordinators. SL requires synchronized forward and backward passes and is vulnerable to intermediate-layer inversion attacks.

6.2. Modern threat models in FL

Recent research has identified advanced threat models beyond the traditional "honest but curious" server that can compromise the integrity and confidentiality of FL. Some advancements include:

- Model Poisoning Attack: Malicious clients can inject corrupted updates into the global model to degrade its performance and implant backdoors[18]. Defenses such as Byzantine Robust Aggregation (Krum, Median) can help mitigate this threat but tend to assume that the adversaries are bounded by a certain ratio of total clients.
- Inference Attacks: Even with secure aggregation, adversaries can reconstruct sensitive training data through gradient inversion and membership inference attacks[43]. These attacks take advantage of the vast amount of information embedded in model updates, particularly with high-dimensional models such as CNNs and Large Language Models[44].
- Collusion Attack: In VFL or Supervised Learning, colluding parties can recreate raw features using partial gradients and activations; this circumvents any cryptographic protections in place[25].

To combat these threat models, a defence-in-depth framework combining cryptographic guarantees, robust aggregation, and Zeroshot Differential Privacy will be required, which is an area requiring further study.

6.3. Evaluation metrics and benchmarking frameworks

Accurate Assessment is only one part of determining a FDL's Real World Viability. In addition to an accurate assessment, the following elements should be included in a comprehensive evaluation:

- Communication Cost: This is an important measurement for edge devices with limited bandwidth. Communication cost is measured in MBs per round [19, 20].
- Fairness: Fairness can be measured using several different metrics, for example, the variance in client accuracy between different clients, or the worst-case accuracy across all clients (non-IID) [12, 26].
- Energy Consumption: For battery-operated devices, this is an additional factor that has not always been accounted for in simulation studies.
- Privacy vs. Accuracy Trade-offs: This factor can be quantified using privacy budgets(ϵ in differential privacy) versus the utility of the FDL model.

The LEAF benchmark was the first FDL benchmark to create a consistent way to evaluate FDLs using non-IID datasets, including FEMNIST and Shakespeare[8]. However, the LEAF benchmark suffers from critical shortcomings. First, the LEAF benchmark relies on either synthetic or outdated datasets, neither of which reflects the true distribution of real-world data (i.e., clinical electronic health records or financial transaction data). Second, the LEAF benchmark abstracts away the heterogeneity of federated learning systems (including issues such as stragglers, dropout clients, and network latency). Third, the LEAF benchmark does not support either vertical or split learning scenarios. Thus, there exists an urgent need for realistic and domain-specific benchmarks, such as MELLODDY for the pharmaceutical industry, which adequately reflect real-world operational constraints and regulatory requirements [33].

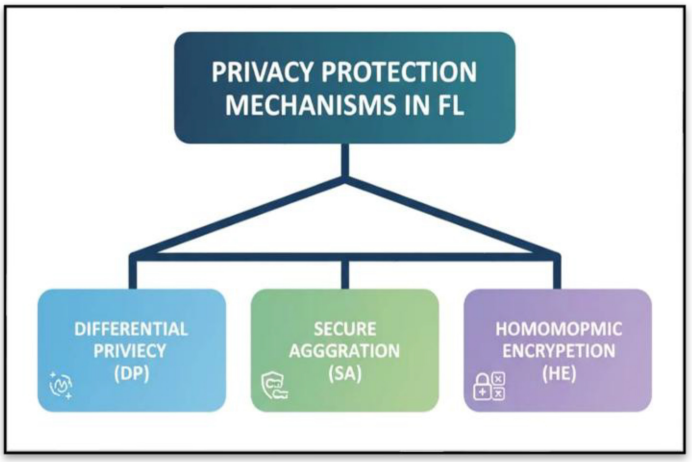

Fig. 7. Classification of privacy protection mechanisms in FL.

6.4. Privacy and security mechanisms

As shown to Fig. 7, the privacy-protecting mechanisms that can be used with federated learning fall into three major categories: Algorithmic, cryptographic, and hardware-assisted. Based on the threat models described in Section 6.2 (e.g., poisoning and inference attacks), the comparison of the various privacy-preserving approach utilizes Table 5 to identify the underlying operational principles, the associated security guarantees, and the trade-offs between practical application of each approach and its ability to provide robust protection of user privacy, which is why hybrid solutions are becoming more common in practice.

7. Future directions and open problems

We observe that the future of FDL revolves around developing more adaptive, transparent, and more efficient frameworks in order to make FDL possible in large-scale, heterogeneous, and dynamic scenarios.

This area is centralized coordination and resource-aware optimization to support scalability. Recent work tries to make FDL adaptive through decentralized. One example is making the model trained on a non-model personalization to non-IID data via Federated Model-Agnostic Meta-Learning (Fed-MAML) [37]. which makes path selection of aggregated IID) data via Federated Transfer Learning. Another feature abstraction techniques methods like Few-shot transfer learning recently proposed involve Federated Model-Agnostic Meta-Learning [6, 37, 38].

Other promising directions are to achieve optimal privacy–utility trade-offs by adaptive privacy strategies or the use of hardware acceleration for computation reduction [6, 39]. As FDL expands into critical applications, trust, explainability, and fairness will be key, leading to a focus on privacy-sensitive interpretability facilitated by secure, tamper-evident, and provenance-enabling audit trails (blockchain-powered audit trails) [21, 40]. Cross-domain and lifelong learning strategies are currently being studied to achieve consistent knowledge transfer and shared exploitation of the same models over evolving contexts [8, 12, 18, 41]. In conclusion, next generation FDL systems should

Table 5. Key privacy-preserving techniques used in FL.

Chnique	Description / Operational Principle	Strengths / Advantages	Limitations / Challenges
Differential Privacy (DP)[31]	Adds calibrated random noise (e.g. Gaussian or Laplace) to each client's model update or to the aggregated output, so that the presence or absence of any single data point is obscured. This provides a quantifiable privacy guarantee.	Provides strong mathematical privacy guarantees (limits information leakage about individual samples); relatively lightweight and can be applied on the client side.	Introduces a trade-off between privacy and accuracy; noisy updates can degrade model performance, especially with small or skewed datasets, and require careful privacy budget management.
Secure Aggregation (SMC)[14]	Cryptographic protocol (often via secure multiparty computation) that enables the server to compute the sum (aggregate) of client updates without learning any individual contribution. For example, Bonawitz et al.'s protocol uses secret sharing and encryption to reveal only the sum.	Ensures that the central aggregator learns only the combined model and nothing about individual clients; protects privacy against an honest-but-curious server.	Requires extra communication rounds and cryptographic operations for key exchange and masking; sensitive to client dropouts (typically needs a threshold of clients to decrypt) and adds latency.
Homomorphic Encryption (HE)[30]	Clients encrypt their model updates with a homomorphic encryption scheme, and the server performs the aggregation directly on ciphertexts. After summing, the server (or a trusted party) decrypts the result to obtain the global model.	Offers end-to-end confidentiality: the server never sees plaintext updates (strong cryptographic security); can work with untrusted servers.	Imposes very high computational and communication overhead; encryption and decryption of large models is costly and often impractical for edge devices.
Trusted Execution Environments (TEE)[41]	Uses secure hardware enclaves (e.g. Intel SGX) to run aggregation/training code in isolation. Code inside the enclave is attested and protected from external access.	Provides strong protection (hardware-enforced confidentiality and integrity) without modifying the ML algorithm; avoids adding noise or cryptographic overhead.	Requires specialized hardware support and trust in the platform; TEEs have limited memory and scalability, and can be vulnerable if hardware is compromised.

and will evolve toward a hybridized contextualized standard framework encapsulating the trifecta of privacy, scale, and transparency while instilling dependable, explainable ethically-anchored distributed intelligence.

8. Discussion

FDL has more value as a practical solution than as a theoretical concept because it has demonstrated a track record of addressing real-world issues. For example, FDL reductions to communication overhead through Model Compression, Selective Participation or Progressive Layer Freezing [19, 20] facilitate Federated Learning on bandwidth constrained

IoT devices and Mobile Networks enabling Smart City and Remote Healthcare applications. Similarly, the use of Personalized Federated Learning [11] or Meta Learning Solutions [37] on Non-IID Data ensures that ML models maintain accuracy and Fairness across user populations and provides a foundation for Equitable AI Solutions in Finance and Public Services. In addition to the use of personalized or Meta Learning Solutions, the Hybrid Privacy Frameworks created through the use of Differential Privacy, Secure Aggregation, Hardware Enclaves [35, 41] ultimately provide an effective means of Meeting Regulatory Requirements (for example, GDPR and HIPAA) without jeopardizing model utility, which is a prerequisite for successful adoption of Healthcare and Banking Federated Learning. Finally, through the use of foundation models in conjunction with the FedLoRA framework, access to billion-parameter AI will become democratized, as Edge Devices will have access to the latest capabilities while preserving Privacy and reducing Dependency on Cloud Services. Collectively, all of these trends address the major impediments to FDL Adoption: Scalability, Robustness and Trust. Future success will depend on the development of Algorithms in conjunction with Domain Specific Constraints, rather than as an afterthought, but as a foundational principle of development.

9. Conclusion

This survey establishes a cohesive analytical framework for federated deep learning (FDL) by integrating algorithmic design, privacy-preserving mechanisms, real-world deployments, and emerging trends in foundation models. Unlike prior surveys that treat these dimensions in isolation, our work critically maps the *trade-offs* e.g., between communication efficiency and model accuracy, or between personalization and fairness that define practical FDL systems. We explicitly address a critical research gap: the disconnect between algorithmic novelty (78% of surveyed works) and real-world deployability (addressed by $<15\%$). By analyzing industrial deployments (e.g., Gboard, MELLODDY) and synthesizing recent advances in parameter-efficient fine-tuning (e.g., FedLoRA), vertical FL, and hybrid privacy, this paper provides a roadmap for building *scalable, ethical, and efficient* FDL systems.Ultimately, FDL is not merely a privacy-preserving alternative to centralized learning it is a new paradigm for collaborative intelligence that respects data sovereignty while enabling collective progress. This survey equips researchers and practitioners with the critical insights needed to advance FDL from theoretical promise to societal impact.

Acknowledgment

The authors gratefully acknowledge the support provided by (University of technology / Computer science college) for facilitating this research. Appreciation is also extended to (Northern Technical University / College of Engineering Technology for Computers and Artificial Intelligence / Kirkuk) for their facilitating of supervisor contribution.

Conflict of interest

The authors declare no conflict of interest.

Data availability

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

References

[1.] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks From Decentralized Data," in *Proc. Int. Conf. Artif. Intell. Stat. (AISTATS)*, 2017, pp. 1273–1282.

[2.] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," *IEEE Signal Process. Mag.*, vol. 37, no. 3, pp. 50–60, May 2020, doi: 10.1109/MSP.2020.2975749.

[3.] Z. Zhou, X. Chen, E. Li, L. Zeng, K. Luo, and J. Zhang, "Edge Intelligence: Paving the Last Mile of Artificial Intelligence with Edge Computing," *Proc. IEEE*, vol. 107, no. 8, pp. 1738–1762, Aug. 2019, doi: 10.1109/JPROC.2019.2918951.

[4.] P. Kairouz *et al.*, "Advances and Open Problems in Federated Learning," *Found. Trends Mach. Learn.*, vol. 14, no. 1–2, pp. 1–210, 2021, doi: 10.1561/2200000083.

[5.] H. K. A. Atheem, I. T. Ali, and F. A. Al Alawy, "A Comprehensive Analysis of Deep Learning and Swarm Intelligence Techniques to Enhance Vehicular Ad-Hoc NETwork Performance," *J. Soft Comput. Comput. Appl.*, vol. 1, no. 1, p. 5, 2024.

[6.] Z. M. Saadi, A. T. Sadiq, O. Z. Akif, and M. M. Eid, "Enhancing Image Classification Using A Convolutional Neural Network Model," *J. Soft Comput. Comput. Appl.*, vol. 1, no. 2, p. 2, 2024.

[7.] Z. S. Mahdi, R. M. Zaki, A. K. Farhan, and N. Majma, "Development of A Hybrid Methodology of Deep Learning and Machine Learning for Lung Nodule Detection in Medical Computed Tomography Images," *J. Soft Comput. Comput. Appl.*, vol. 1, no. 2, p. 3, 2024.

[8.] K. Bonawitz *et al.*, "Practical Secure Aggregation for Privacy-Preserving Machine Learning," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, Oct. 2017, pp. 1175–1191, doi: 10.1145/3133956.3133982.

[9.] S. Reddi *et al.*, "Adaptive Federated Optimization," *arXiv preprint arXiv:2003.00295*, 2020.

[10.] Z. Qu, K. Lin, Z. Li, and J. Zhou, "Federated Learning's Blessing: FedAvg Has Linear Speedup," in *Proc. ICLR Workshop Distrib. Priv. Mach. Learn. (DPML)*, May 2021.

[11.] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated Machine Learning: Concept and Applications," *ACM Trans. Intell. Syst. Technol. (TIST)*, vol. 10, no. 2, pp. 1–19, Jan. 2019, doi: 10.1145/3298981.

[12.] S. Vucinich and Q. Zhu, "The Current State and Challenges of Fairness In Federated Learning," *IEEE Access*, vol. 11, pp. 80903–80914, 2023, doi: 10.1109/ACCESS.2023.3299801.

[13.] Y. Chen, Z. Chai, Y. Cheng, and H. Rangwala, "Asynchronous Federated Learning for Sensor Data with Concept Drift," in *Proc. IEEE Int. Conf. Big Data*, Dec. 2021, pp. 4822–4831, doi: 10.1109/BigData52589.2021.9671530.

[14.] A. Hard *et al.*, "Federated Learning for Mobile Keyboard Prediction," *arXiv preprint arXiv:1811.03604*, 2018.

[15.] E. J. Hu *et al.*, "LoRA: Low-Rank Adaptation of Large Language Models," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2022.

[16.] S. Behera, M. Adhikari, V. G. Menon, and M. A. Khan, "Large Model-Assisted Federated Learning for Object Detection of Autonomous Vehicles in Edge," *IEEE Trans. Veh. Technol.*, vol. 74, no. 2, pp. 1839–1848, Feb. 2024, doi: 10.1109/TVT.2024.3351234.

[17.] A. Yazdinejad, A. Dehghantanha, H. Karimipour, G. Srivastava, and R. M. Parizi, "A Robust Privacy-Preserving Federated Learning Model Against Model Poisoning Attacks," *IEEE Trans. Inf. Forensics Security*, vol. 19, pp. 6693–6708, 2024, doi: 10.1109/TIFS.2024.3364512.

[18.] F. Wang, Y. Zhang, L. Chen, and H. Li, "Privacy-Preserving Automated Deep Learning for Secure Inference Service," *IEEE Trans. Dependable Secure Comput.*, vol. 22, no. 1, pp. 412–426, Jan.–Feb. 2025, doi: 10.1109/TDSC.2025.3541028.

[19.] Y. Wu *et al.*, "Heterogeneity-Aware Memory Efficient Federated Learning Via Progressive Layer Freezing," in *Proc. IEEE/ACM Int. Symp. Quality of Service (IWQoS)*, June 2024, pp. 1–10, doi: 10.1109/IWQOS60775.2024.10587231.

[20.] T. Tsouparopoulos and I. Koutsopoulos, "Explainability and Continual Learning Meet Federated Learning at the Network Edge," *arXiv preprint arXiv:2504.08536*, 2025.

[21.] N. Bhati and N. Vyas, "Advanced Architectures and Innovative Platforms for Federated Learning: A Comprehensive Exploration," in *Model Optimization Methods for Efficient and Edge AI: Federated Learning Architectures, Frameworks and Applications*, 2025, pp. 145–155.

[22.] Z. Wang, M. Zhang, and Y. Liu, "A Comprehensive Survey on Federated Learning: Concept and Applications," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 11, pp. 8263–8288, Nov. 2023, doi: 10.1109/TNNLS.2023.3265432.

[23.] Y. Zhao, M. Li, and Q. Yang, "A Survey on Efficient Federated Learning Methods for Foundation Model Training," *ACM Comput. Surv.*, vol. 56, no. 8, Art. no. 175, pp. 1–36, Dec. 2024, doi: 10.1145/3638530.

[24.] J. Liu, Y. Kang, and Q. Yang, "Federated Learning in the Vertical Setting: Challenges and Opportunities," *IEEE Internet Things J.*, vol. 10, no. 15, pp. 13285–13301, Aug. 2023, doi: 10.1109/JIOT.2023.3262215.

[25.] M. G. Poirot *et al.*, "Split Learning: A Survey of Architectures, Applications, and Privacy Trade-offs," *IEEE Access*, vol. 11, pp. 115287–115313, 2023, doi: 10.1109/ACCESS.2023.3325671.

[26.] T. Dinh, J. Tran, and T. Nguyen, "Federated Foundation Models: A Survey of Distributed Pretraining, Fine-tuning, and Inference," *Found. Trends Mach. Learn.*, vol. 17, no. 4, pp. 457–574, 2025, doi: 10.1561/2200000123.

[27.] Y. Wu, H. Zhang, and Q. Yang, "Continual Federated Learning: A Survey and New Perspectives," *ACM Comput. Surv.*, vol. 57, no. 4, Art. no. 92, pp. 1–38, Apr. 2025, doi: 10.1145/3649456.

[28.] M. Chen, Z. Yang, and W. Saad, "Federated Learning Over Wireless Networks: A Comprehensive Survey," *IEEE Commun. Surv. Tutor.*, vol. 26, no. 2, pp. 1026–1060, Secondquarter 2024, doi: 10.1109/COMST.2023.3345678.

[29.] L. U. Khan *et al.*, "Federated Learning For Industrial IoT: Challenges and Opportunities," *IEEE Internet Things J.*, vol. 11, no. 8, pp. 12345–12362, Apr. 2024, doi: 10.1109/JIOT.2023.3328910.

[30.] J. Zhang, X. Wang, and C. Wang, "Privacy-Preserving Federated Learning for Medical Image Analysis Using Hybrid Homomorphic Encryption," *IEEE J. Biomed. Health Inform.*, vol. 28, no. 3, pp. 1456–1468, Mar. 2024, doi: 10.1109/JBHI.2023.3321456.

[31.] K. Wei, J. Li, M. Ding, C. Ma, and H. V. Poor, "Federated Learning With Differential Privacy: Algorithms and Performance Analysis," *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 3454–3469, 2020, doi: 10.1109/TIFS.2020.2987893.

[32.] S. R. K. Peddinti *et al.*, "Federated Learning for On-Device Health Monitoring: A Case Study on Arrhythmia Detection," *NPJ Digit. Med.*, vol. 6, Art. no. 112, pp. 1–12, Jul. 2023, doi: 10.1038/s41746-023-00856-1.

[33.] A. Hard *et al.*, "Federated Learning for Financial Fraud Detection: A Real-World Case Study," in *Proc. ACM Int. Conf. AI Finan. (ICAIF)*, Nov. 2023, pp. 145–154, doi: 10.1145/3613928.3613945.

[34.] J. Liu, M. Juuti, Y. Lu, and N. Asokan, "Practical Homomorphic Encryption for Federated Learning on Edge Devices," in *Proc. USENIX Secur. Symp.*, Aug. 2023, pp. 3215–3232.

[35.] A. Nilsson, S. Smith, and M. Nikolova, "Hybrid Privacy-Preserving Federated Learning: Combining Homomorphic Encryption and Differential Privacy," *IEEE Trans. Dependable Secure Comput.*, vol. 22, no. 1, pp. 789–803, Jan.–Feb. 2025, doi: 10.1109/TDSC.2024.3412567.

[36.] C. Deng *et al.*, "Federated Meta-Learning with Adaptive Personalization for Non-IID Data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 35, no. 6, pp. 7890–7904, Jun. 2024, doi: 10.1109/TNNLS.2023.3301234.

[37.] Y. Zhao and M. Li, "Few-Shot Federated Learning Via Meta-Initialization," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2023.

[38.] T. Ha and T. K. Dang, "Hardware-Accelerated Privacy-Preserving Inference for Edge Federated Learning," in *Proc. IEEE Int. Conf. Edge Comput. (EDGE)*, Jul. 2024, pp. 88–95, doi: 10.1109/EDGE60323.2024.00018.

[39.] D. Hou, J. Zhang, K. L. Man, J. Ma, and Z. Peng, "Blockchain-Based Auditable Federated Learning: A Survey and Taxonomy," *IEEE Trans. Netw. Sci. Eng.*, vol. 12, no. 2, pp. 1120–1138, Apr.–Jun. 2025, doi: 10.1109/TNSE.2024.3467890.

[40.] A. Nilsson, P. N. Bideh, and J. Brorsson, "A Survey of Attacks and Defenses on Intel SGX: Implications for Federated Learning," *ACM Comput. Surv.*, vol. 56, no. 7, Art. no. 156, pp. 1–36, Nov. 2023, doi: 10.1145/3614152.

[41.] E. Bagdasaryan *et al.*, "Collaborative Credit Scoring with Privacy-Preserving Federated Learning," *IEEE Trans. Big Data*, early access, 2025, doi: 10.1109/TBDATA.2025.3542109.

[42.] C. Finn, P. Abbeel, and S. Levine, "Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks," in *Proc. Int. Conf. Mach. Learn. (ICML)*, July 2017, pp. 1126–1135.

[43.] S. Ali *et al.*, "HybridFL: Hybrid Approach Toward Privacy-Preserving Federated Learning," in *Proc. Int. Conf. Security Privacy New Comput. Environ.*, Springer, Nov. 2023, pp. 3–18.

[44.] Z. Zhou *et al.*, "Model Inversion Attacks: A Survey of Approaches and Countermeasures," *arXiv preprint arXiv:2411.10023*, 2024.