Review

Federated Learning: A Survey of Core Challenges, Current Methods, and Opportunities

Madan Baduwal *, Priyanka Paudel , Vini Chaudhary

Posted Date: 5 January 2026

doi: 10.20944/preprints202601.0271.v1

Keywords: federated learning; data heterogeneity; privacy preservation; communication efficiency; client selection; model aggregation

 doi:10.20944/preprints202601.0271.v1

Review

Federated Learning: A Survey of Core Challenges, Current Methods, and Opportunities

Madan Baduwal [1,*], Priyanka Paudel [2] and Vini Chaudhary [1]

[1] Department of CSE, Mississippi State University, Starkville, MS 39762, USA
[2] Department of Computer Science, University of Missouri-St. Louis, St. Louis, MO 63121, USA
* Correspondence: mb4239@msstate.edu; Tel.: +1-(432)-316-1183

Abstract

Federated Learning (FL) has emerged as a transformative distributed learning paradigm that enables collaborative model training without sharing raw data, thereby preserving privacy across large, diverse, and geographically dispersed clients. Despite its rapid adoption in mobile networks, IoT systems, healthcare, finance, and edge intelligence, FL continues to face several persistent and interdependent challenges that hinder its scalability, efficiency, and real-world deployment. In this survey, we present a systematic examination of six core challenges in federated learning: *heterogeneity, computation overhead, communication bottlenecks, client selection, aggregation and optimization, and privacy preservation*. We analyze how these challenges manifest across the full FL pipeline, from local training and client participation to global model aggregation and distribution, and examine their impact on model performance, convergence behavior, fairness, and system reliability. Furthermore, we synthesize representative state-of-the-art approaches proposed to address each challenge and discuss their underlying assumptions, trade-offs, and limitations in practical deployments. Finally, we identify open research problems and outline promising directions for developing more robust, scalable, and efficient federated learning systems. This survey aims to serve as a comprehensive reference for researchers and practitioners seeking a unified understanding of the fundamental challenges shaping modern federated learning.

Keywords: federated learning; data heterogeneity; privacy preservation; communication efficiency; client selection; model aggregation

1. Introduction

Machine learning has become increasingly pervasive across modern digital ecosystems, driven by the rapid proliferation of data-generating devices such as smartphones, wearables, autonomous vehicles (e.g., Tesla, Waymo, XPeng) [1–3], smart appliances, industrial sensors, and emerging AI-enabled hardware including AI smart glasses (e.g., Ray-Ban Meta, Google and HTC AI glasses) [4–6], mixed-reality headsets (e.g., Apple Vision Pro) [7,8], humanoid robots such as Tesla's Optimus and Boston Dynamics' Atlas, agile quadruped robots (e.g., Boston Dynamics Spot, Unitree Go1/G1) [9–11], and autonomous aerial systems including Skydio R1 and automated "drone-in-a-box" platforms [12–14]. These systems continuously sense, interpret, and interact with their physical and digital environments.

Simultaneously, the global Internet of Things (IoT) ecosystem is projected to reach up to 40 billion connected devices by 2030, generating approximately 200 zettabytes of data annually [15]. This unprecedented scale has accelerated the shift toward Edge AI, where data processing and learning are increasingly performed close to data sources, enabling local intelligence while still contributing to global learning objectives. Traditional centralized learning pipelines [16–18], which aggregate raw data on a central server for model development, have achieved state-of-the-art performance across diverse domains [19,20]. However, such pipelines face growing limitations in the modern landscape of strict data governance, distributed computation, and ubiquitous edge intelligence. Privacy regulations such

 doi:10.20944/preprints202601.0271.v1

as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), heightened concerns regarding data ownership, and the communication overhead of transmitting massive raw datasets have rendered centralized learning increasingly impractical or legally constrained in many real-world settings [21–23].

Recent advances in artificial intelligence (AI) have further demonstrated the transformative potential of data-driven learning systems. A prominent milestone is AlphaGo [24], which defeated world-class human Go players and highlighted how large-scale learning, optimization, and reasoning can solve problems once considered intractable for machines. Such breakthroughs reaffirm a central objective of AI: to develop intelligent agents capable of perception, planning, learning, reasoning, and adaptation in complex and dynamic environments [25–31]. As AI systems increasingly permeate real-world applications, they rely on vast volumes of data generated by distributed sources such as mobile devices, sensors, and edge platforms. While centralized machine learning paradigms have been effective in controlled environments [19,20], they are becoming increasingly misaligned with modern deployment realities characterized by privacy constraints, regulatory compliance, communication bottlenecks, and large-scale system heterogeneity [21–23].

Federated Learning (FL) has emerged as a compelling distributed learning paradigm that fundamentally rethinks how collaborative machine learning can be performed without requiring centralized access to raw data [32]. In FL, multiple clients, such as smartphones, IoT devices, or institutional data silos-jointly train a shared global model by performing local updates on private datasets and sharing only aggregated model information with a coordinating server, as illustrated in Figure 1. This decentralized training mechanism preserves data privacy, enhances data security, reduces communication overhead, and supports compliance with data protection regulations such as GDPR and related privacy frameworks [33,34]. Consequently, FL is well suited for large-scale, heterogeneous, and geographically distributed environments, and has catalyzed widespread adoption across application domains including healthcare, finance, IoT networks, mobile and edge computing, and smart-city infrastructures [35]. In healthcare, FL enables collaborative training of diagnostic and predictive models while maintaining strict patient confidentiality [36]; in mobile systems, user devices can collectively improve services such as predictive text and speech recognition without uploading sensitive data to centralized servers [37]; and in smart-city platforms, FL facilitates learning from distributed sensor networks while minimizing bandwidth consumption and preserving data locality [38]. Notably, FL gained significant momentum during the COVID-19 pandemic by enabling privacy-preserving collaboration for medical research, epidemiological modeling, and population-level analytics across institutional boundaries [39,40].

Beyond these established use cases, federated learning is increasingly recognized as a general-purpose collaborative intelligence framework applicable across a broad range of scientific, engineering, and socio-technical disciplines. In the life sciences, FL has been applied to multi-institutional neuroimaging and neuroscience studies [41,42], enabling cross-site learning on distributed MRI and brain imaging data while respecting strict data-sharing and human-subject constraints [43,44]. Closely related domains such as biostatistics, epidemiology, and public health benefit from FL's ability to analyze population-level data across institutions without centralization, a capability that proved particularly valuable during global health crises [45–48]. In chemistry, materials science, and pharmaceutical research, FL enables collaborative training of molecular property prediction and drug discovery models across organizations holding proprietary or sensitive experimental data [49–55]. Industry-scale initiatives such as the MELLODDY project demonstrate that FL can support competitive yet cooperative learning across pharmaceutical partners while preserving intellectual property and data sovereignty [56].

In engineering, IoT, and cyber-physical systems, FL has become a foundational approach for edge intelligence, supporting collaborative learning across heterogeneous sensors, autonomous vehicles, industrial systems, and wireless networks under stringent communication, energy, and privacy constraints [57–59]. Federated optimization enables scalable learning across distributed infrastructures where centralized data aggregation is infeasible or undesirable. Beyond the natural sciences and

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 5 January 2026 doi:10.20944/preprints202601.0271.v1

engineering, FL is increasingly relevant in economics, finance, business, and management, where sensitive transactional data are distributed across institutions and jurisdictions. Federated frameworks enable joint risk modeling, fraud detection, demand forecasting, and financial analytics while preserving confidentiality and regulatory compliance [60–64]. Emerging research also explores FL in social computing, human–computer interaction (HCI), and education, where user-centric data must remain on-device or within institutional boundaries.

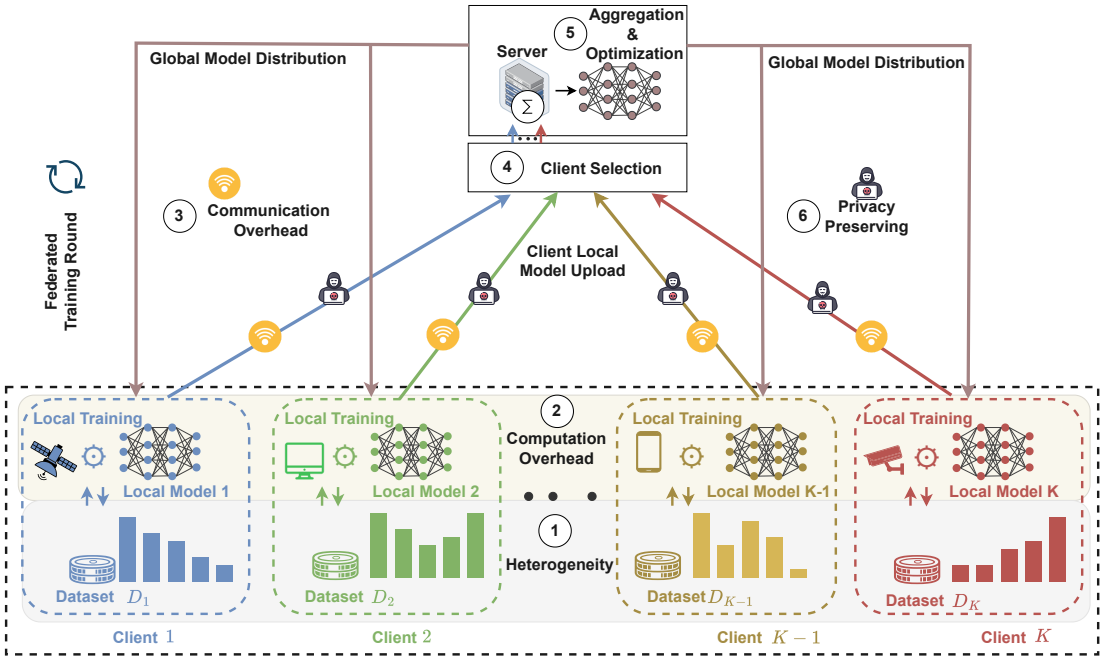

Figure 1. Six core challenges of Federated Learning

While some disciplines, such as philosophy, humanities, law, political science, and the arts, do not yet directly deploy federated learning as a computational tool, FL's core principles of decentralization, data ownership, and privacy-preserving collaboration strongly resonate with these fields' ethical, legal, and societal concerns [34]. As data-driven methodologies increasingly permeate these domains, federated learning offers a foundational framework for responsible, institutionally distributed analytics.

The concept of federated learning was first introduced by Google in 2016 [65], motivated by the need for privacy-aware on-device learning at scale. A canonical real-world deployment is Google Keyboard (Gboard), where millions of Android devices collaboratively learn language models to improve next-word prediction while ensuring that users' private text data never leave their devices. This deployment demonstrated the practical feasibility, scalability, and privacy benefits of FL in real-world systems, catalyzing extensive research and adoption across academia and industry [66]. Beyond mobile applications, federated learning holds particular promise for IoT, edge intelligence, and 5G/6G-enabled wireless networks [67–69], where massive numbers of devices generate data continuously under strict bandwidth, latency, and energy constraints. By enabling local model training and transmitting only compact model updates instead of raw data, FL significantly reduces communication overhead and alleviates network congestion, making it especially attractive for large-scale, resource-constrained, and heterogeneous environments.

Taken together, the promise of federated learning lies in its ability to reconcile three traditionally competing objectives: (i) leveraging distributed data for improved learning performance, (ii) preserving data privacy and ownership, and (iii) enabling communication-efficient and scalable model training. These properties position FL as a foundational paradigm for next-generation distributed intelligence, with successful applications already emerging in mobile computing, healthcare, finance, smart cities, and cyber-physical systems.

 doi:10.20944/preprints202601.0271.v1

Despite its promise, FL faces fundamental challenges that hinder its scalability, efficiency, robustness, and accuracy in real-world deployments. Unlike centralized settings with homogeneous data and controlled infrastructure, FL operates across heterogeneous devices, non-IID and imbalanced datasets, unreliable communication links, and dynamic client participation [70]. These conditions give rise to six interconnected challenges: (i) heterogeneity, (ii) computation overhead, (iii) communication bottlenecks, (iv) client selection and participation management, (v) aggregation and optimization, and(vi) privacy preservation paradigms [71]. Addressing these challenges in isolation is often insufficient, as progress in one dimension may influence or exacerbate others.

While numerous surveys have examined federated learning from perspectives such as privacy, communication efficiency, personalization, system design, and domain-specific applications, most existing works focus on isolated aspects. Prior surveys typically emphasize a single dimension, such as heterogeneity, IoT applications, privacy-preserving mechanisms, aggregation strategies, or communication-efficient optimization, without providing an integrated, cross-layer analysis of how these challenges interact across the FL lifecycle [72]. Moreover, emerging paradigms including meta-learning, multi-task learning, self-supervised learning, contrastive learning, and personalized learning are beginning to reshape FL, yet their connections to core FL challenges remain underexplored (e.g., MAML for personalization [73] or contrastive FL for non-IID mitigation [74]).

This survey addresses these gaps by providing a systematic and holistic examination of six core challenges in federated learning: data heterogeneity, computation overhead, communication efficiency, client selection, aggregation and optimization, privacy preservation, and integration with modern learning paradigms. Our contributions are fourfold:

- We propose a unified and **challenge-centric taxonomy** that systematically organizes federated learning research across the entire FL pipeline, explicitly highlighting the interdependencies and trade-offs among six foundational challenges, rather than treating them in isolation.
- We provide a comprehensive synthesis of **state-of-the-art methods** for each challenge category, critically analyzing their underlying assumptions, algorithmic designs, theoretical guarantees, empirical performance, and practical limitations across diverse deployment settings.
- We conduct an in-depth examination of **emerging learning paradigms**, including meta-learning, personalized federated learning, self-supervised learning, contrastive learning, and continual learning, and elucidate how these paradigms intersect with, extend, and reshape classical federated learning formulations.
- We identify **open research problems and unresolved bottlenecks** at the algorithmic, system, and application levels, and outline promising future research directions toward building scalable, communication-efficient, robust, and trustworthy federated learning systems.

This survey is intended for researchers, practitioners, system architects, and domain specialists seeking a rigorous and comprehensive understanding of federated learning (FL), encompassing its theoretical foundations, system architectures, deployment challenges, and evolving research landscape. It is particularly relevant to audiences engaged in the development of novel FL algorithms, the large-scale deployment of FL systems, and real-world applications across domains such as healthcare, finance, edge and mobile computing, smart cities, and cyber-physical systems. The survey is organized to guide readers progressively from foundational concepts to advanced challenges. Section 2 introduces the core principles, system architectures, and learning formulations underlying federated learning. Sections 5–10 provide a challenge-centric analysis of the six fundamental challenges that arise across the FL pipeline, synthesizing representative state-of-the-art methods, key design trade-offs, and open limitations. Section 11 surveys major application domains of federated learning, while Section 12 reviews widely used open-source FL frameworks and systems. Finally, Section 13 outlines future research directions and concludes the survey by summarizing the main findings and discussing their broader implications for the design of scalable, efficient, and privacy-preserving federated learning systems.

 doi:10.20944/preprints202601.0271.v1

2. Background & Foundations

Federated Learning (FL) is a decentralized machine learning paradigm that enables multiple clients (e.g., mobile devices, organizations, hospitals, or IoT nodes) to collaboratively train a shared global model under the orchestration of a central server [32,75,76]. In contrast to traditional centralized learning, where raw data from all clients is uploaded to a central repository for training, FL keeps data local to each client and only exchanges model updates (such as gradients or model weights) with the server. By never transmitting personal or sensitive raw data off-device, this approach embodies principles of privacy-by-design and data minimization, thereby reducing the risk of privacy leakage and helping organizations comply with data governance regulations like GDPR and CCPA. Google originally pioneered FL for applications such as Gboard's mobile text prediction, and since then the paradigm has expanded into domains including healthcare, finance, smart transportation, smart cities, and large-scale IoT networks [77–79]. These diverse use cases demonstrate FL's potential to harness distributed data in a privacy-conscious manner across a range of real-world settings.

2.1. Definition of Federated Learning

At its core, federated learning can be understood as a distributed optimization problem that trains a global model without centralizing the data. Formally, suppose there are K clients indexed by $k \in \{1, 2, \ldots, K\}$, where each client k possesses a private dataset $\mathcal{D}_k$ of size n_k (so that the total number of data points across all clients is $n = \sum_{k=1}^{K} n_k$). We denote the model parameters (e.g., the weights of a neural network) by $w \in \mathbb{R}^d$. The goal of federated learning is to minimize a global empirical risk objective that aggregates the losses over all clients' data, without ever directly pooling those data together. This objective can be written as:

$$\min_{w \in \mathbb{R}^d} F(w) = \sum_{k=1}^{K} \frac{n_k}{n} F_k(w), \tag{1}$$

where $F_k(w)$ is the local objective function for client k. The local objective is defined as the empirical risk on client k's dataset,

$$F_k(w) = \frac{1}{n_k} \sum_{(x_i, y_i) \in \mathcal{D}_k} \ell(w; x_i, y_i), \tag{2}$$

with $\ell(w; x_i, y_i)$ representing the loss of model w on a single data sample (x_i, y_i). In simpler terms, $F_k(w)$ measures how well the model w fits the data of client k, and $F(w)$ is a weighted average of these local losses (weighted by the relative size of each client's dataset n_k). By minimizing $F(w)$, one finds model parameters that perform well on the collective data of all clients, all without requiring any client to send its raw data to the server.

A canonical algorithm for solving this federated optimization problem is Federated Averaging (FedAvg), introduced by McMahan et al. [32]. FedAvg is an iterative procedure that proceeds in synchronous communication rounds between the server and a (typically subset of) clients. At the beginning of each round t, the central server holds the current global model w^t. The server first selects a subset S_t of the clients and then sends the current global model parameters w^t to each of those selected clients. Upon receiving the model, each client $k \in S_t$ initializes training from w^t and performs local learning on its own dataset $\mathcal{D}_k$, usually running a few epochs of stochastic gradient descent (SGD), producing an updated model w_k^t. After completing the local training, each client sends its update back to the server, and the server aggregates them as:

$$w^{t+1} = \sum_{k \in S_t} \frac{n_k}{\sum_{j \in S_t} n_j} w_k^t, \tag{3}$$

which gives more weight to clients with larger datasets. This procedure repeats for rounds $t = 0, 1, 2, \ldots$ until convergence or a predefined stopping criterion is reached. FedAvg approximates the effect of SGD on the global dataset while significantly reducing communication overhead.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 5 January 2026 doi:10.20944/preprints202601.0271.v1

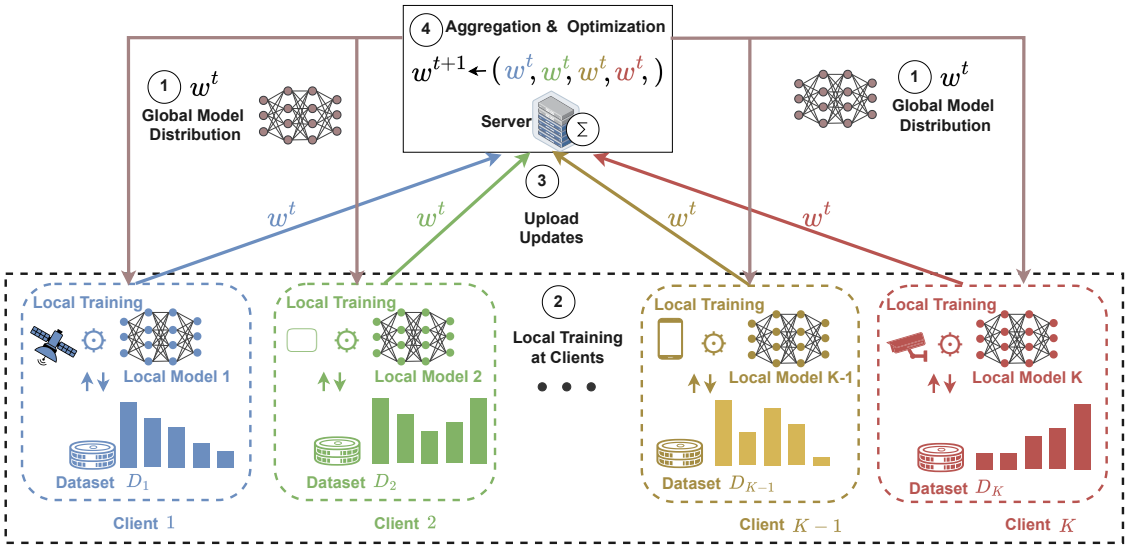

Figure 2. Architecture of federated learning. At communication round t, the server distributes the global model w^t to selected clients. Each client k performs local training on its private dataset $\mathcal{D}_k$ and uploads the resulting model update without sharing raw data. The server aggregates the received updates to produce the next global model w^{t+1}, enabling iterative, collaborative, and privacy-preserving learning across distributed clients.

2.2. Architecture for a Federated Learning System

This subsection outlines the standard client–server workflow of a federated learning system involving K clients, indexed by $k = 1, 2, \ldots, K$, where each client holds a local dataset $\mathcal{D}_k$.

- **Step 1(Global Model Distribution):** At communication round t, the server maintains the current global model w^t and selects a subset of available clients for participation. The server broadcasts w^t along with basic training settings, such as the learning rate and number of local training epochs.
- **Step 2(Local Training at Clients):** Each selected client k updates the received global model using its own local dataset $\mathcal{D}_k$. All clients begin local training from the same model parameters w^t and perform training independently, while all data remain stored and processed locally.
- **Step 3(Model Update Upload):** After completing local training, each participating client sends its updated model parameters (or model changes relative to w^t) back to the server. Only model-related information is communicated; the underlying datasets $\mathcal{D}_k$ are never shared.
- **Step 4(Model Aggregation at the Server):** The server aggregates the updates received from participating clients to form the next global model w^{t+1}. The aggregation reflects the collective contribution of the clients, commonly accounting for differences in local dataset sizes.
- **Step 5(Iterative Model Refinement):** The updated global model is redistributed to clients, and Steps 1–4 are repeated over multiple communication rounds until convergence or a predefined stopping criterion is met. The final outcome is a single global model learned collaboratively across decentralized datasets.

2.3. A Categorization of Federated Learning

Federated learning (FL) deployments are commonly characterized along two complementary axes: *deployment scale* and *data partitioning*. From the deployment perspective, *cross-device FL* involves a very large population of unreliable and resource-constrained clients, such as mobile phones or IoT devices. In this setting, only a randomly selected subset of clients participates in each communication round, and system design prioritizes scalability, fault tolerance, and robustness to client dropout. In contrast, *cross-silo FL* typically involves a small number of powerful and reliable organizations, such as hospitals, banks, or government agencies. In cross-silo settings, most or all clients participate in each round, and the emphasis shifts toward governance, security, compliance, and inter-organizational collaboration. Independent of deployment scale, FL systems are also categorized by how data is

 doi:10.20944/preprints202601.0271.v1

distributed across participating clients. As illustrated in Figure 3, three canonical data-partitioning scenarios arise: *horizontal federated learning*, *vertical federated learning*, and *federated transfer learning*.

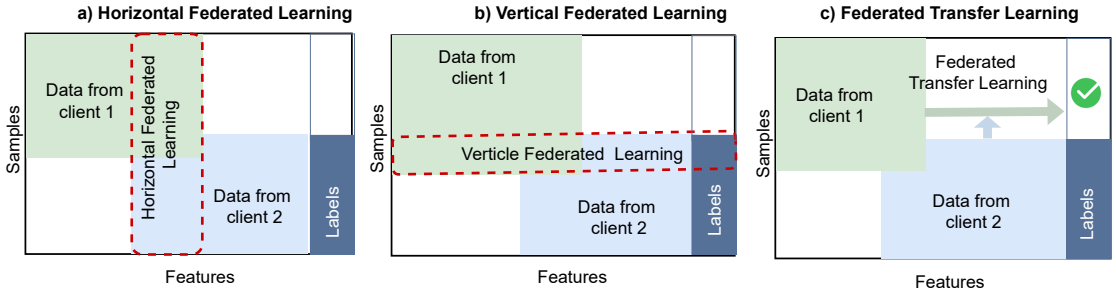

Figure 3. Categorization of federated learning based on data partitioning. (a) Horizontal federated learning, where clients share the same feature and label spaces but hold disjoint subsets of samples. (b) Vertical federated learning, where clients share the same samples but possess different subsets of features. (c) Federated transfer learning, where clients differ in both samples and feature spaces, and knowledge is transferred across domains.

Let $\mathcal{D}_k$ denote the local dataset held by client k, where each dataset can be represented as a matrix whose rows correspond to data samples and whose columns correspond to features. Let $\mathcal{X}_k$ denote the feature space, $\mathcal{Y}_k$ the label space, and $\mathcal{I}_k$ the sample index (or sample ID) space associated with client k. Accordingly, each local dataset can be expressed as

$$\mathcal{D}_k = (\mathcal{I}_k, \mathcal{X}_k, \mathcal{Y}_k),$$

where $\mathcal{I}_k$ identifies the samples owned by client k, $\mathcal{X}_k$ contains the observed features, and $\mathcal{Y}_k$ contains the corresponding labels (if available). The tuple $(\mathcal{I}, \mathcal{X}, \mathcal{Y})$ represents the complete training dataset in a centralized learning setting.

In federated learning, data are distributed across multiple clients such that their feature spaces, label spaces, and sample index spaces may differ. Based on how data are partitioned in the feature space $\mathcal{X}$ and the sample index space $\mathcal{I}$, federated learning can be broadly categorized into horizontal federated learning, vertical federated learning, and federated transfer learning, as illustrated in Figure 3.

2.3.1. Horizontal Federated Learning (HFL)

In horizontal federated learning, all participating clients share a common feature space and a common label space, while each client holds a disjoint subset of samples. Formally, the feature and label spaces satisfy

$$\mathcal{X}_1 = \mathcal{X}_2 = \cdots = \mathcal{X}_K, \qquad \mathcal{Y}_1 = \mathcal{Y}_2 = \cdots = \mathcal{Y}_K,$$

whereas the sample index spaces are mutually disjoint,

$$\mathcal{I}_k \cap \mathcal{I}_j = \emptyset, \qquad \forall k \neq j.$$

Consequently, the global dataset is partitioned *horizontally* across clients along the sample dimension, with each client observing different data instances described by the same set of features and labels (Figure 3(a)). A canonical example of HFL is collaborative language modeling across a population of user devices, where all clients employ the same feature representation (e.g., word or token embeddings) and prediction task, but each device contributes disjoint text samples.

 doi:10.20944/preprints202601.0271.v1

2.3.2. Vertical Federated Learning (VFL)

In vertical federated learning, clients share the same set of samples but possess different subsets of features. Formally, all clients operate over a common sample index space,

$$\mathcal{I}_1 = \mathcal{I}_2 = \cdots = \mathcal{I}_K = \mathcal{I},$$

while their feature spaces are distinct and complementary,

$$\mathcal{X}_1 \neq \mathcal{X}_2 \neq \cdots \neq \mathcal{X}_K.$$

Each client therefore observes a partial feature representation of the same entities. Training typically proceeds via joint optimization over the implicitly concatenated feature space, enabled by secure aggregation, split learning, or cryptographic protocols (Figure 3(b)).

2.3.3. Federated Transfer Learning (FTL)

Federated transfer learning addresses the most general setting, in which clients differ in both their feature spaces and their sample index spaces. Specifically,

$$\mathcal{X}_k \neq \mathcal{X}_j \quad \text{and} \quad \mathcal{I}_k \cap \mathcal{I}_j \approx \emptyset, \qquad \forall k \neq j.$$

FTL exploits partial overlap in label semantics, feature representations, or auxiliary knowledge to enable cross-client knowledge transfer despite minimal data alignment (Figure 3(c)).

2.4. Centralized, Federated, and Decentralized Learning

Figure 4 provides a comparative overview of four learning and federation paradigms that differ in how data, models, and coordination are organized across participants. These paradigms represent a progression from full data centralization to fully decentralized model collaboration.

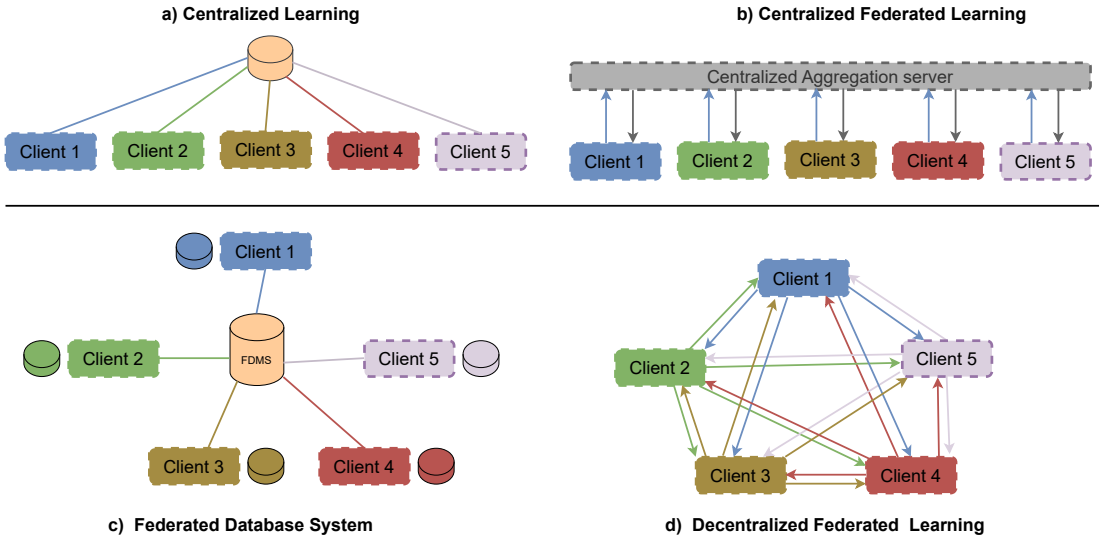

Figure 4. Overview of learning and federation paradigms: (a) centralized learning with data aggregation at a single server; (b) centralized federated learning with server-based model aggregation; (c) federated database systems enabling coordinated access to distributed databases; and (d) decentralized federated learning with peer-to-peer model collaboration.

2.4.1. Centralized Learning

In centralized learning, all data generated by distributed clients are collected and stored at a single central server, where model training is performed. Clients act primarily as data sources, and learning

 doi:10.20944/preprints202601.0271.v1

occurs exclusively on centrally aggregated datasets. As illustrated in Figure 4(a), this paradigm offers a simple training workflow and strong optimization control, but it introduces significant drawbacks related to data privacy, regulatory compliance, communication cost, and scalability. In modern large-scale and privacy-sensitive applications, centralized learning is often impractical or legally restricted.

2.4.2. Centralized Federated Learning

Centralized federated learning retains a central coordination server but fundamentally differs from centralized learning in that raw data remain local to each client. As shown in Figure 4(b), clients perform local model training on private data and transmit only model updates (e.g., weights or gradients) to a centralized aggregation server. The server coordinates training rounds, aggregates updates, and redistributes the global model. This paradigm balances scalability and privacy while maintaining global orchestration and has become the dominant architecture for practical federated learning deployments.

2.4.3. Federated Database Systems

Federated database systems represent an earlier and conceptually distinct form of federation. Rather than collaboratively training models, these systems focus on coordinated access, querying, and management of distributed databases. As illustrated in Figure 4(c), a federation manager mediates queries across multiple autonomous data sources, enabling integrated data views without physically centralizing storage. While federated database systems support data governance and autonomy, they do not inherently address collaborative machine learning or iterative model optimization.

2.4.4. Decentralized Federated Learning

Decentralized federated learning removes the central aggregation server entirely and replaces it with peer-to-peer model of collaboration among clients. As depicted in Figure 4(d), clients exchange model updates directly with neighboring peers according to a communication topology. Model consensus emerges through repeated local aggregation and information propagation rather than centralized coordination. This paradigm improves robustness to single points of failure and enhances fault tolerance, but introduces new challenges related to convergence guarantees, communication overhead, and coordination complexity.

Together, these paradigms highlight fundamental trade-offs between centralization, privacy, scalability, and robustness. Centralized learning prioritizes simplicity but sacrifices privacy, while centralized federated learning offers a practical balance between coordination and data locality. Federated database systems emphasize data access rather than learning, and decentralized federated learning pushes collaboration to its most distributed form, enabling resilience at the cost of increased system complexity. This spectrum of paradigms provides essential context for understanding the design choices and challenges underlying modern federated learning systems.

2.5. Federated Learning Versus Edge Computing

Federated Learning (FL) and Edge Computing are often discussed together due to their shared emphasis on decentralized data processing and reduced reliance on cloud-centric architectures. However, they represent fundamentally different concepts: edge computing is a *system and infrastructure paradigm*, whereas federated learning is a *machine learning paradigm*. Understanding their distinction and interaction is essential for designing scalable, privacy-aware intelligent systems.

2.5.1. Edge Computing

Edge computing moves computation, storage, and analytics closer to data sources, such as mobile devices, IoT sensors, access points, or edge servers. Its primary goal is to reduce end-to-end latency, bandwidth consumption, and dependency on centralized cloud infrastructure. Edge computing supports a wide range of workloads, including data filtering, real-time inference, stream processing,

 doi:10.20944/preprints202601.0271.v1

and control tasks. Importantly, edge computing by itself does not prescribe *how* learning is performed; it merely provides a distributed execution environment for computation near the data origin [80,81].

2.5.2. Federated Learning

Federated learning, in contrast, explicitly defines a collaborative learning protocol in which multiple clients jointly train a shared model while keeping raw data local. FL specifies how local training, model update exchange, and aggregation are orchestrated across distributed participants. From a systems perspective, FL can be deployed on top of cloud, edge, or hybrid cloud–edge infrastructures. Thus, FL addresses *learning coordination and privacy*, while edge computing addresses *where computation takes place*.

2.5.3. Conceptual Relationship

The relationship between FL and edge computing can be summarized as follows: edge computing provides the *execution substrate*, while federated learning provides the *learning algorithmic framework*. In practice, many FL systems operate in edge environments, where clients or edge nodes perform local training and exchange model updates. However, edge computing does not require FL, and FL does not strictly require edge computing; for example, FL can be deployed across geographically distributed data centers or institutional silos without edge devices.

2.5.4. Learning and Communication Perspective

From an optimization viewpoint, federated learning minimizes a global objective of the form

$$\min_{w} \; F(w) = \sum_{k=1}^{K} p_k F_k(w),$$

where the coordination of local objectives F_k is governed by a learning protocol. Edge computing imposes no such objective and may instead support inference-only pipelines, centralized training with edge inference, or task-specific analytics. Communication in edge computing typically involves raw or partially processed data, whereas FL exchanges model parameters or compressed updates, leading to distinct communication patterns and system trade-offs.

2.5.5. Complementarity and Integration

In modern deployments, FL and edge computing are often complementary. Edge resources enable efficient local training and inference under latency, energy, and bandwidth constraints, while FL enables privacy-preserving collaboration across those distributed edge nodes. Hierarchical architectures further combine both paradigms, where edge servers aggregate updates from nearby devices before forwarding them to a cloud-level coordinator [82,83]. Such integration is particularly relevant in large-scale IoT networks, smart cities, and next-generation wireless systems.

In summary, edge computing and federated learning address different layers of distributed intelligence. Edge computing focuses on *system placement and execution*, whereas federated learning focuses on *collaborative model training under data locality constraints*. Their combination enables scalable, low-latency, and privacy-aware learning pipelines, but they should not be conflated as equivalent paradigms.

 doi:10.20944/preprints202601.0271.v1

Table 1. Summary of Notations

Notation	Description
K	Total number of clients participating in FL
k	Client index, $k \in \{1, \ldots, K\}$
$\mathcal{D}_k$	Local dataset stored at client k
n_k	Number of samples at client k
n	Total number of samples, $n = \sum_{k=1}^{K} n_k$
$x_{k,i}$	i-th data sample (feature vector) at client k
$y_{k,i}$	Corresponding label of $x_{k,i}$
w	Global model parameters
w^t	Global model at communication round t
w_k^t	Local model of client k at round t
d	Dimensionality of model parameters, $w \in \mathbb{R}^d$
$F(w)$	Global objective function
$F_k(w)$	Local objective function at client k
$\ell(w; x, y)$	Sample-wise loss function
p_k	Aggregation weight of client k, $p_k = \frac{n_k}{n}$
η	Learning rate
E	Number of local training epochs per round
t	Communication round index
$\mathcal{S}^t$	Set of clients selected at round t

Table 2. List of Acronyms

Acronym	Meaning
FL	Federated Learning
HFL	Horizontal Federated Learning
VFL	Vertical Federated Learning
FTL	Federated Transfer Learning
PFL	Personalized Federated Learning
DFL	Decentralized Federated Learning
FedAvg	Federated Averaging
IID	Independent and Identically Distributed
Non-IID	Non-Identically Distributed Data
SGD	Stochastic Gradient Descent
DP	Differential Privacy
SMPC	Secure Multi-Party Computation
HE	Homomorphic Encryption
TEE	Trusted Execution Environment
IoT	Internet of Things
P2P	Peer-to-Peer
QoS	Quality of Service
NAS	Neural Architecture Search
GNN	Graph Neural Network

3. Related Surveys

Federated learning has attracted substantial attention in recent years, and a large body of survey literature has emerged to systematize its foundations, applications, and challenges. Early comprehensive works, such as Yang et al. [84], introduced the basic concepts, architectures, and application scenarios of FL, with an emphasis on secure FL frameworks among a relatively small number of institutional participants. Li et al. [85] focused on FL in massive networks of mobile and edge devices, highlighting practical challenges around efficiency, heterogeneity, and privacy in cross-device environments. Kairouz et al. [86] provided a landmark survey summarizing advances and open problems in FL from multiple research angles, including optimization, privacy, robustness, and systems, and

 doi:10.20944/preprints202601.0271.v1

has become a canonical reference in the field. Other general surveys [87–93] further consolidate core FL concepts, system components, and application domains, offering broad overviews of the field's evolution.

Beyond these general treatments, many surveys concentrate on particular aspects or deployment settings of FL. Several works focus on platforms, protocols, and engineering aspects [94–98], reviewing enabling hardware/software infrastructures, communication architectures, and practical system design considerations. Others study FL in specific domains such as IoT and edge computing [99–106], wireless networks and 6G systems [99,101,104], or application areas including healthcare and disease prediction [107–109]. These works provide valuable domain-oriented perspectives but typically adopt a scenario- or application-centric taxonomy rather than organizing the literature around core technical challenges across the FL pipeline.

A second cluster of surveys targets specific technical dimensions of FL. Communication efficiency has been studied extensively, with surveys reviewing compression, sparsification, structured updates, and resource-aware protocols [97,106,110–112]. Privacy and security have motivated another large body of work, including surveys on threat models, inference and poisoning attacks, and defense mechanisms such as differential privacy, secure aggregation, homomorphic encryption, trusted hardware, and blockchain-based designs [101,113–121]. Complementary efforts focus on trustworthy and robust FL, covering interpretability, fairness, robustness, accountability, and attack-resilient aggregation [112,116,122]. These surveys offer in-depth views of individual problem dimensions but largely treat them in isolation from other system-level and algorithmic challenges.

Heterogeneity and personalization have also been recognized as central issues in FL. Dedicated surveys on heterogeneous FL [112,123] categorize methods for handling statistical, model, communication, and device heterogeneity, while works on personalized FL [124,125] focus on strategies for tailoring global models to client-specific data distributions. Additional surveys study FL in conjunction with complementary paradigms such as neural architecture search [126], blockchain-based coordination [119,120], or multimodal and graph learning [127]. Although these contributions illuminate important subfields of FL, they typically emphasize one primary axis, e.g., heterogeneity, personalization, communication, or security, rather than providing a unified analysis of how these aspects interact across the full FL lifecycle.

In summary, existing surveys can be broadly categorized into: (i) general FL overviews that cover concepts, architectures, and applications at a high level [84–93]; (ii) domain- or system-specific surveys focusing on IoT, edge computing, wireless networks, healthcare, or industrial applications [99–101,103–108]; and (iii) dimension-specific surveys centered on communication efficiency [97,110–112], heterogeneity and personalization [123–125], or privacy, security, and trust [113–122]. While these works have significantly advanced our understanding of federated learning, they often either (i) provide broad but coarse-grained overviews, or (ii) deliver deep but narrow analyses focused on a single challenge, domain, or technique.

Positioning of our survey: In contrast, this survey adopts a *challenge-centric* viewpoint that systematically organizes the FL literature around Six tightly coupled core challenges: data heterogeneity, computation overhead, communication bottlenecks, client selection, aggregation and optimization, and privacy preservation. Rather than treating these aspects separately, we explicitly analyze their interdependencies across the FL pipeline and examine how progress in one dimension (e.g., communication or privacy) propagates to others (e.g., optimization dynamics, system scalability, or personalization quality). Table 3 summarizes representative existing FL surveys and contrasts their scope, focus, and taxonomies with the holistic perspective proposed in this work.

 doi:10.20944/preprints202601.0271.v1

Table 3. Comparison between our survey and representative existing FL surveys.

Survey	Year	Scope / Domain	Main Focus / Taxonomy	Difference from Our Survey
Yang et al. [84]	2019	General FL; data distribution types	Divides FL into three categories according to data distribution characteristics.	Overview of FL but lacks detailed classification and summary of existing methods.
Li et al. [85]	2020	General FL; efficiency, heterogeneity, privacy	Challenges of FL from efficiency, heterogeneity, and privacy perspectives; several future research directions.	Our survey provides a more comprehensive and integrated challenge-centric taxonomy, including finer-grained treatment of heterogeneity.
Lim et al. [100]	2020	Mobile edge networks	Survey of FL in mobile edge networks and edge-computing scenarios.	Scenario-specific; our survey is cross-domain and challenge-centric.
Niknam et al. [128]	2020	Wireless communication networks	Applications and challenges of FL in wireless communication environments.	Domain-centric; our survey is broader and integrates multiple challenges across the FL pipeline.
Kulkarni et al. [129]	2020	Statistical heterogeneity; personalization	Shows how statistical heterogeneity can hinder FL and highlights the need for personalized FL.	Heterogeneity-focused; our survey treats heterogeneity as one of multiple coupled core challenges.
Wu et al. [124]	2020	Personalized FL; cloud–edge IoT	Personalized FL framework in a cloud–edge architecture for intelligent IoT applications.	Personalization-centric; our survey covers broader FL schemes and cross-challenge interactions.
Aledhari et al. [94]	2020	Enabling technologies, protocols, applications	Reviews FL-enabling platforms, protocols, use-cases, and key challenges.	Enabling-tech focus; our survey provides a broader pipeline-wide challenge-centric taxonomy.
Li et al. [107]	2020	FL applications	Reviews major FL applications in industrial engineering and computer science, outlining key research fronts.	Application-focused; our survey emphasizes challenge-centric analysis beyond application categorization.
Nguyen et al. [99]	2021	IoT, smart services	FL applications in IoT (smart healthcare, transport, UAVs, smart cities); FL-enabled IoT services (caching, offloading, attack detection).	IoT-only; our survey analyzes cross-domain and cross-challenge interactions across the FL pipeline.
Yin et al. [130]	2021	Privacy-preserving FL	5W taxonomy; privacy leakage risks; privacy-preservation mechanisms.	Privacy-focused; our survey situates privacy within a broader set of interconnected challenges.
Li et al. [87]	2021	FL systems	Categorization by data distribution, privacy mechanism, communication architecture, federation scale.	Systems-centric; our survey provides a unified challenge-centric view spanning systems + algorithms + applications.

Continued on next page

 doi:10.20944/preprints202601.0271.v1

Table 3. Comparison between our survey and representative existing FL surveys (continued).

Survey	Year	Scope / Domain	Main Focus / Taxonomy	Difference from Our Survey
Kairouz et al. [86]	2021	General FL; foundations and open problems	Recent advances in FL: comprehensive survey of open problems and challenges.	Broad overview; lacks fine-grained method classification under a unified challenge framework.
Wahab et al. [101]	2021	General FL; challenges and approaches	Fine-grained classification scheme of existing FL challenges and approaches.	Different organizing principle; our survey emphasizes six tightly coupled core challenges and their interdependencies.
Khan et al. [131]	2021	IoT applications	Advances in FL for IoT applications and a taxonomy using various parameters (e.g., robustness, privacy, communication cost).	IoT-centric; our survey is cross-domain and pipeline-wide challenge-centric.
Zhu et al. [126]	2021	FL + NAS	Surveys FL, NAS methods, and emerging federated NAS approaches with a taxonomy of online/offline and single/multi-objective variants.	Focuses on FL–NAS intersection; our survey provides broader FL challenge coverage beyond architecture search.
Blanco-Justicia et al. [113]	2021	Security & privacy in FL	Surveys privacy and security attacks in FL and mitigation strategies, highlighting challenges in achieving both simultaneously.	Security/privacy-focused; our survey integrates these aspects within a broader, multi-challenge FL taxonomy.
Lo et al. [95]	2021	FL from a software engineering perspective	Systematic review of FL system development lifecycle: requirements, architecture, implementation, and evaluation.	SE-focused lifecycle view; our survey provides a broader, challenge-centric taxonomy across the full FL pipeline.
Liu et al. [88]	2022	General FL systems	From distributed ML to FL; system architecture; parallelism; aggregation; communication; security; taxonomy of FL systems.	System-architecture oriented; our survey is challenge-centric and integrates computation, communication, heterogeneity, privacy, and optimization.
Gao et al. [132]	2022	Heterogeneous FL (data, system, model)	Investigates heterogeneous FL in terms of data-space, statistical, system, and model heterogeneity.	This work classifies existing methods based on problem settings and learning objectives, while our survey classifies methods based on specific techniques.
Tan et al. [125]	2022	Personalized FL; taxonomy	Explores the field of personalized FL and conducts a taxonomic survey of existing methods.	This work briefly explains statistical heterogeneity, but lacks a comprehensive taxonomy and analysis of the challenges in FL.

Continued on next page

 doi:10.20944/preprints202601.0271.v1

Table 3. Comparison between our survey and representative existing FL surveys (continued).

Survey	Year	Scope / Domain	Main Focus / Taxonomy	Difference from Our Survey
Pouriyeh et al. [97]	2022	Communication efficiency in FL	Reviews communication constraints, efficiency challenges, and secure communication strategies in FL.	Communication-focused; our survey integrates communication with other key FL challenges in a unified framework.
Mahlool et al. [98]	2022	General FL: concepts and applications	Covers FL components, challenges, and applications with emphasis on medical use-cases.	Application-oriented; our survey offers a broader, structured challenge-centric taxonomy beyond specific domains.
Zhang et al. [114]	2022	Security & privacy threats in FL	Classifies FL attacks by adversary type, reviews major threat models and mitigation techniques, including DGL, GAN-based attacks, and TEE/blockchain defenses.	Threat-focused; our survey integrates security/privacy with broader FL challenges across the entire pipeline.
Bharati et al. [108]	2022	General FL; applications & challenges	Reviews FL frameworks, architectures, applications (especially healthcare), and key pri-vacy/security/heterogeneity challenges.	Application-heavy; our survey provides a broader, structured challenge-centric classification beyond domain-specific analyses.
Abreha et al. [103]	2022	FL in edge computing	Systematic survey of FL implementation in edge environments, covering architectures, protocols, hardware, applications, and challenges.	Edge-computing–focused; our survey provides a broader, cross-environment challenge-centric taxonomy.
Gupta et al. [96]	2022	FL in distributed environments	Reviews centralized, decentralized, and heterogeneous FL frameworks, focusing on privacy, DP techniques, and distributed optimization.	Distributed-environment focus; our survey provides a broader, unified challenge-centric taxonomy across all FL settings.
Wen et al. [90]	2023	General FL; challenges and applications	Surveys FL basics, privacy/security mechanisms, communication issues, heterogeneity, and practical applications.	Covers core challenges and applications broadly; our survey offers a more structured, challenge-centric taxonomy across all FL dimensions.
Moshawrab et al. [110]	2023	Aggregation algorithms in FL	Reviews FL aggregation strategies and algorithms, their implementations, limitations, and future directions.	Aggregation-focused; our survey covers aggregation as one component within a broader, multi-challenge FL taxonomy.

Continued on next page

 doi:10.20944/preprints202601.0271.v1

Table 3. Comparison between our survey and representative existing FL surveys (continued).

Survey	Year	Scope / Domain	Main Focus / Taxonomy	Difference from Our Survey
Beltrán et al. [111]	2023	Decentralized FL (DFL)	Examines DFL fundamentals, architectures, communication mechanisms, frameworks, and application scenarios.	DFL-specific focus; our survey provides a broader, unified view across both centralized and decentralized FL challenges.
Ye et al. [123]	2023	Heterogeneous FL (HFL)	Surveys challenges and solutions in statistical, model, communication, and device heterogeneity, with a taxonomy of HFL methods.	Focused solely on heterogeneity, our survey treats heterogeneity as one challenge within a broader, integrated FL taxonomy.
Neto et al. [118]	2023	Secure FL; attacks and defenses	Systematic review of FL security vulnerabilities, attack types, mitigation strategies, and secure FL applications.	Security-focused, our survey integrates security alongside other core FL challenges in a unified framework.
Almanifi et al. [112]	2023	Communication + computation efficiency in FL	Surveys communication- and computation-efficiency techniques, challenges, and optimization strategies in FL.	Efficiency-focused, our survey integrates efficiency with broader FL challenges across the full pipeline.
Gupta et al. [117]	2023	Game-theoretic FL	Reviews game-theory–based FL models for incentives, authentication, privacy, trust, and threat detection, with bibliometric analysis.	GT-focused; our survey provides a broader, multi-challenge perspective beyond incentive mechanisms.
Moshawrab et al. [109]	2023	FL for disease prediction	Reviews FL concepts, aggregation approaches, and medical applications, highlighting limitations and future directions.	Healthcare-focused, our survey provides a broader, cross-domain challenge-centric taxonomy beyond specific medical applications.
Asad et al. [133]	2023	Communication-efficient FL	Surveys communication-reduction techniques, including compression, structured updates, resource management, and client selection.	Communication-specific; our survey integrates communication with broader FL challenges in a unified taxonomy.
Che et al. [127]	2023	Multimodal FL	Surveys multimodal FL methods, categorizing congruent vs. incongruent MFL, with benchmarks, applications, and future directions.	Modality-focused, our survey provides a broader challenge-centric taxonomy beyond multimodal considerations.
Sirohi et al. [104]	2023	FL for 6G secure communication systems	Analyzes vulnerabilities, threats, and defenses in FL across 6G application domains.	Domain-specific security focus; our survey provides a broader, unified challenge-centric taxonomy across all FL settings.

Continued on next page

 doi:10.20944/preprints202601.0271.v1

Table 3. Comparison between our survey and representative existing FL surveys (continued).

Survey	Year	Scope / Domain	Main Focus / Taxonomy	Difference from Our Survey
Qammar et al. [119]	2023	Blockchain-based FL	Systematic review of integrating blockchain with FL to enhance security, privacy, accountability, and robustness.	Blockchain-specific focus; our survey provides a broader, multi-challenge FL taxonomy beyond decentralized ledger integration.
Zhu et al. [120]	2023	Blockchain-empowered FL	Surveys how blockchain addresses coordination, trust, incentives, and security issues in FL, with a taxonomy of BlockFed system models.	Blockchain-focused; our survey provides a broader challenge-centric analysis beyond ledger-integrated FL architectures.
Liu et al. [89]	2024	General FL; recent advances	Systematic review of recent FL methods, applications, taxonomy, and frameworks.	Broad recent-advances survey; our work provides a more integrated, challenge-centric analysis.
Yurdem et al. [91]	2024	General FL; overview and strategies	Comprehensive overview of FL principles, strategies, applications, tools, and future directions.	Broad introductory overview; our survey provides deeper, challenge-focused analysis across the full FL pipeline.
Alotaibi et al. [134]	2024	Non-IID + communication challenges in FL	Systematic mapping of techniques for handling non-IID data and improving communication efficiency in FL.	Focuses on two specific challenges; our survey provides a broader, integrated challenge taxonomy.
Tariq et al. [122]	2024	Trustworthy FL (interpretability, fairness, robustness)	Reviews trustworthiness foundations in FL, proposing a taxonomy covering interpretability, transparency, fairness, privacy/robustness, and accountability.	Trust-focused; our survey integrates trustworthiness alongside broader technical FL challenges within a unified framework.
Saha et al. [115]	2024	Privacy-preserving FL	Surveys privacy risks, attacks, and defenses in FL.	Privacy-focused; our survey situates privacy within a broader, multi-challenge FL taxonomy.
Hu et al. [116]	2024	Security & privacy in FL	Analyzes FL threat models, vulnerabilities, and defense strategies.	Security/privacy-focused; our survey integrates these aspects with other key FL challenges in a unified perspective.
Xie et al. [135]	2024	HE-based privacy-preserving FL	Surveys efficiency optimization strategies for HE-based FL.	HE-specific efficiency focus; our survey situates HE within a broader, multi-challenge FL landscape.
Kaur et al. [136]	2024	General FL; recent advances & applications	Reviews FL framework, categories, benefits, and diverse applications, highlighting recent advances and open concerns.	Broad application-oriented review; our survey provides a more detailed, challenge-centric taxonomy across the full FL pipeline.

Continued on next page

 doi:10.20944/preprints202601.0271.v1

Table 3. Comparison between our survey and representative existing FL surveys (continued).

Survey	Year	Scope / Domain	Main Focus / Taxonomy	Difference from Our Survey
Albshaier et al. [105]	2025	FL for cloud & edge security	Systematically reviews FL applications for cloud/edge security.	Domain-specific; our survey provides a broader, cross-domain challenge-centric taxonomy.
Jia et al. [106]	2025	Communication-efficient FL (mobile edge)	Surveys methods for reducing communication overhead in FL in mobile edge settings.	Communication-centric and edge-focused; our survey integrates communication with broader FL challenges across the full pipeline.
Chaudhary et al. [92]	2025	General FL systems	Provides a detailed overview of FL systems, architectures, frameworks, applications, and prospects.	Systems-focused; our survey offers a broader, challenge-centric taxonomy across all FL dimensions.
Our Survey	**2026**	General FL; cross-domain	Systematic survey of six core challenges: heterogeneity, computation, communication, client selection, aggregation/optimization, privacy, and integration.	Holistic challenge-centric viewpoint, covering cross-layer interactions, emerging FL paradigms, and multi-domain applications.

4. Survey Protocol and Taxonomy

This section details the methodology used to conduct our literature survey and the taxonomy that structures our analysis. We first describe how we formulated the guiding research questions and collected relevant literature, and then we present the six core challenge taxonomy that emerged from this process.

4.0.1. Research Methodology and Research Questions

We adopted a structured and systematic review protocol inspired by established best practices in large-scale literature surveys. The primary objective of this survey is to present a *challenge-centric* perspective on Federated Learning (FL) that is both comprehensive in scope and explicitly focused on the algorithmic, system-level, and deployment constraints that limit real-world adoption. Rather than organizing prior work solely by application domain or algorithmic family, our methodology emphasizes the identification, categorization, and interrelation of fundamental challenges that arise throughout the FL lifecycle.

Guided by this objective, we formulated five research questions (RQs) that collectively frame the scope and direction of our investigation.

- **RQ1:** What are the major research directions, system architectures, and application domains of federated learning across academia and industry?
- **RQ2:** What fundamental challenges arise when deploying federated learning in realistic, large-scale, and heterogeneous environments?
- **RQ3:** What algorithmic techniques, system designs, and optimization strategies have been proposed to address these challenges?
- **RQ4:** How do these challenges interact across the federated learning pipeline, and what trade-offs emerge among communication efficiency, optimization performance, privacy guarantees, fairness, and robustness?
- **RQ5:** Which challenges remain insufficiently addressed, and what open problems and research opportunities emerge from current limitations?

 doi:10.20944/preprints202601.0271.v1

To answer these questions, we systematically collected, filtered, and analyzed relevant literature using a multi-stage search and selection process. The resulting body of work was then synthesized and organized into a coherent taxonomy consisting of six core challenges. This taxonomy serves as the structural backbone of the survey and enables a unified analysis of federated learning methods across algorithmic, system, and application layers.

4.0.2. Search Strategy

We employed a multi-stage search strategy to capture both foundational and recent contributions in federated learning (FL) research. Specifically, we queried the official websites of major machine learning conferences and journals, scholarly databases, and digital libraries, including the top ML conferences and journals, IEEE Xplore, the ACM Digital Library, SpringerLink, ScienceDirect, and arXiv. A broad set of keywords related to federated learning was used, ranging from general terms such as "federated learning" and "federated training" to more targeted phrases addressing specific research dimensions, including heterogeneous FL, personalized FL, communication-efficient FL, privacy-preserving FL, federated optimization, federated aggregation, and client selection in FL.

To broaden coverage, we also performed backward and forward snowballing across the reference lists of seminal FL papers and existing surveys. This allowed us to capture works not easily discoverable by keyword search alone, such as papers that do not explicitly include "federated learning" in the title or metadata. Our primary scope focused on literature from top-tier ML conferences and journals, and papers 2020 onward, when modern FL research began to accelerate. Earlier foundational work in distributed optimization and privacy-preserving learning was included selectively for context. Finally, we prioritized high-quality and high-impact sources, including peer-reviewed journal articles, top-tier conference publications, and influential preprints spanning machine learning, distributed systems, networking, security, and application domains such as healthcare, IoT, and wireless communications. This multi-pronged approach yielded a curated corpus of FL-related publications that we organized using the six-challenge taxonomy described next.

4.0.3. Study Selection Criteria

To ensure both relevance and scholarly rigor, we applied explicit inclusion and exclusion criteria during the paper selection process. We included studies that explicitly consider federated learning (FL) or closely related decentralized learning paradigms in which model training is performed over distributed data while raw data remain local to participating clients. In addition, a paper was required to introduce, analyze, or empirically evaluate a technical mechanism addressing at least one fundamental FL challenge, including but not limited to data and system heterogeneity, computational constraints, communication efficiency, client selection, aggregation and optimization, privacy

and security, or the integration of FL with complementary learning paradigms. Only peer-reviewed publications or highly cited and widely recognized preprints were considered. Furthermore, each selected work was required to provide sufficient methodological depth, such as clear algorithmic formulations, theoretical analyses, system designs, or experimental evaluations, to enable accurate categorization within our proposed six-challenge taxonomy.

Conversely, we excluded studies that do not adhere to a genuine FL setting, such as distributed learning approaches that assume centralized data aggregation or unrestricted data sharing, which conflict with FL's core privacy and decentralization principles. We also removed papers that reference FL only superficially without making substantive methodological contributions. Works superseded by more comprehensive or updated versions were excluded unless they were of historical significance. Finally, we omitted purely application-driven or demonstration-oriented studies that merely apply existing FL techniques without offering new algorithmic, system-level, or analytical insights, except in cases where such studies exposed critical practical limitations or deployment challenges. Applying these criteria refined the literature corpus to a focused and high-impact set of works suitable for a rigorous, challenge-centric analysis.

 doi:10.20944/preprints202601.0271.v1

4.0.4. Taxonomy Construction

Based on the systematic analysis of the selected literature, we constructed a challenge-centric taxonomy that organizes federated learning (FL) research around six core challenges. These challenges collectively capture the principal algorithmic, system-level, and deployment obstacles encountered in the design and operation of practical FL systems. Rather than categorizing prior work solely by application domain or model architecture, our taxonomy emphasizes the fundamental constraints and trade-offs that arise throughout the FL lifecycle.

The first challenge concerns *heterogeneity*, encompassing statistical heterogeneity due to non-identically distributed client data, system heterogeneity stemming from diverse hardware and network capabilities, and model heterogeneity arising from variations in local architectures or computational budgets. The second challenge addresses *computation limitations* at the client side, where constraints on processing power, memory, energy consumption, and training time restrict the depth, frequency, or complexity of local updates. The third challenge focuses on *communication efficiency and protocol design*, including limited bandwidth, high latency, unreliable connectivity, and the need for communication-efficient mechanisms that reduce transmission overhead across training rounds. The fourth challenge centers on *client selection and participation management*, which involves determining which clients participate in each training round under practical constraints such as availability, reliability, fairness, incentive compatibility, and the presence of stragglers or dropouts. The fifth challenge pertains to *aggregation and optimization*, encompassing the development of robust global update rules capable of handling non-IID local updates, partial participation, and conflicting local objectives, while ensuring stable convergence and strong global model performance. The sixth challenge focuses on *privacy and security*, addressing threats such as inference attacks, data reconstruction, and model poisoning, and leveraging techniques including differential privacy, secure aggregation, cryptographic protocols, and robust optimization methods.

Each subsequent section of this survey is dedicated to one of these six challenge areas. For each challenge, we synthesize representative methods, analyze underlying assumptions and limitations, and highlight open research problems. We further discuss how advances in one challenge dimension may complement or trade off with progress in others, revealing important interdependencies across the FL pipeline. By adopting this challenge-centric taxonomy, the survey provides a structured and holistic view of the federated learning landscape and clarifies the interconnected nature of its core research challenges.

5. Challenge 1: Heterogeneity

Heterogeneity is one of the most fundamental and persistent challenges in federated learning (FL), arising from the inherent diversity of data distributions, system capabilities, and learning objectives across participating clients. Unlike centralized learning settings that typically assume independent and identically distributed (IID) data, FL operates over decentralized datasets that are frequently non-IID due to variations in user behavior, data-collection environments, sensing modalities, and domain characteristics. This statistical mismatch gives rise to the well-known *client drift* phenomenon, where locally trained models diverge from one another and from the global optimum, leading to slower convergence, degraded global accuracy, and unstable training dynamics [137–140]. In real-world deployments, statistical heterogeneity is often compounded by system heterogeneity, where clients differ significantly in computational power, memory capacity, energy constraints, and network reliability. These disparities further exacerbate optimization instability, introduce stragglers, and raise fairness concerns by disproportionately favoring resource-rich participants [141–143].

To mitigate heterogeneity, a wide range of approaches have been proposed across the learning objective, aggregation strategy, and information-sharing levels. Objective-level solutions, such as proximal and regularized federated optimization methods, explicitly constrain local training trajectories to remain close to the global model, thereby stabilizing updates and reducing client drift under non-IID data [143]. Complementarily, aggregation-level techniques design robust server-side update rules

 doi:10.20944/preprints202601.0271.v1

that account for biased gradients, statistical uncertainty, or unequal client contributions, improving convergence behavior in heterogeneous environments [144]. Another influential line of work focuses on constructing shareable intermediate information, including feature representations, prototypes, or latent embeddings, which enables collaborative distribution alignment across clients without revealing raw data [145,146]. By aggregating such intermediate representations, the server can obtain a richer approximation of the global data distribution while preserving privacy constraints.

Personalized federated learning (PFL) has emerged as a particularly effective paradigm for addressing heterogeneity by relaxing the assumption of a single universal global model. Instead of enforcing full model consensus, PFL methods aim to learn customized client-specific models that better reflect local data characteristics. Representative approaches include meta-learning, mixture-of-experts, representation sharing with local adaptation, fine-tuning strategies, and multi-task formulations [125]. By allowing controlled personalization, these methods significantly improve accuracy, robustness, and user satisfaction in highly heterogeneous environments, especially when client data distributions differ substantially. Recent advances further demonstrate the effectiveness of personalization in large-scale vision-centric and cross-domain settings [147].

Federated clustering methods offer an alternative perspective on heterogeneity by attempting to recover latent global structure from decentralized data without centralization. Typical approaches aggregate local proxies such as cluster centroids, basis matrices, or low-rank factors to approximate global distributions. For instance, federated k-means and fuzzy c-means methods collect local cluster centers [148–151], while federated spectral clustering reconstructs shared kernel representations from local sketches [152]. Federated non-negative matrix factorization (NMF) aggregates distributed basis matrices across clients [153]. Although these approaches preserve data locality, they often rely on biased local statistics computed under non-IID conditions, which can limit their ability to capture true global structure and result in degraded clustering performance. Recent methods explore synthetic data generation and cross-client regularization to alleviate these issues, but typically introduce trade-offs between privacy, fidelity, and computational overhead [154].

Knowledge distillation (KD) provides a flexible and architecture-agnostic mechanism for mitigating heterogeneity by enabling clients to exchange knowledge without sharing raw data or full model parameters. Soft-label or multi-path distillation allows clients to collaboratively transfer information through logits or intermediate representations, improving alignment across heterogeneous data distributions [155]. Extensions include asymmetric distillation, teacher–student coordination, and robust KD pipelines designed to tolerate corrupted or adversarial clients [156,157]. Frameworks that introduce shared feature spaces or lightweight coordination modules further reduce representation mismatch and client drift in the presence of skewed or imbalanced data [158,159]. KD-based methods are particularly effective in cross-device FL scenarios, where model architectures and resource constraints vary widely across participants.

In settings involving **graph-structured data**, heterogeneity manifests not only in node features and label distributions but also in graph topology and connectivity patterns. Recent work addresses this challenge by measuring inter-client similarity using gradient or embedding statistics [160], or by constructing global anchor graphs at the server to provide consistent structural references for local updates [161]. These structure-aware techniques reduce divergence across clients and improve convergence stability in federated graph learning tasks.

Overall, heterogeneity remains a central obstacle to scalable, fair, and reliable federated learning. While regularization, personalization, clustering, distillation, and structure-aware information sharing have led to substantial progress, non-IID data and system diversity continue to challenge convergence guarantees, robustness, and equitable participation. Future research is expected to explore multi-objective optimization, structured priors, privacy-preserving representation exchange, and more expressive global abstractions tailored to diverse client populations [162–166].

 doi:10.20944/preprints202601.0271.v1

6. Challenge 2: Computation Overhead

Federated Learning (FL) imposes substantial computation overhead on participating clients, particularly in cross-device environments where computational resources vary drastically, ranging from high-end servers to battery-limited mobile and IoT devices. In each communication round, clients are typically required to perform multiple epochs of local optimization on large, high-dimensional models, which can overwhelm devices with limited CPU capability, memory footprint, or energy budget. This imbalance gives rise to the well-known *straggler problem*, where slower or resource-constrained clients delay global aggregation and significantly reduce overall system throughput. Moreover, heterogeneous hardware capabilities introduce inconsistent training speeds and update frequencies across clients, increasing optimization variance and negatively affecting convergence stability and fairness across rounds. As modern deep learning architectures continue to grow in size and computational complexity, balancing model expressiveness with client-side feasibility remains a central bottleneck for scalable FL deployments.

A major line of research addresses computation overhead through **model lightweighting**, which reduces the cost of local training by compressing or simplifying neural network architectures. Pruning-based approaches remove redundant weights, filters, or channels to lower the number of floating-point operations required during training. For instance, FedMP dynamically prunes model parameters according to each client's computational capability, enabling heterogeneous devices to train capability-matched subnetworks while preserving global accuracy [167]. Similarly, sparse training and Lottery Ticket Hypothesis–based methods allow clients to train compact subnetworks tailored to their hardware constraints, thereby reducing both computation and memory usage [168]. Quantization-based FL methods further reduce computation by representing model parameters and gradients with low-precision arithmetic, such as 8-bit or mixed-precision formats, achieving significant speedups with minimal accuracy degradation [169]. Lightweight neural architectures, including MobileNet-style backbones or student models distilled from a server-side teacher, provide additional reductions in local FLOPs and energy consumption, making FL feasible for severely resource-limited devices [170,171].

Another class of techniques mitigates computation overhead by allowing **partial model training** or heterogeneous model sizes across clients. Instead of enforcing a single uniform model, these approaches adapt the depth, width, or parameter subset trained by each client. HeteroFL, for example, enables clients with different computational capabilities to train models of varying widths or depths, which are subsequently aligned and aggregated at the server [172]. Subnetwork-based FL methods similarly employ structured sparsity or layer dropping, allowing weaker clients to train only selected layers or blocks while still contributing useful updates to the global model [173]. By matching computational workload to device capability, these techniques prevent low-power clients from becoming bottlenecks and improve system-level efficiency without excluding them from participation.

Computation overhead can also be alleviated through **split learning and computation offloading**, which partition the neural network between clients and servers. In split learning, clients compute only the early layers of the model and transmit intermediate activations to the server, where deeper and more computationally expensive layers are trained. Frameworks such as SplitNN and SplitFed significantly reduce client-side computation by offloading heavy workloads to more powerful servers or edge nodes [174,175]. Hybrid variants that combine split learning with pruning or quantization further improve efficiency and robustness in heterogeneous environments [176]. These approaches are particularly effective when server-side resources are abundant and communication latency is manageable.

Beyond architectural adaptations, computation overhead is often mitigated through **adaptive client participation and scheduling**. Rather than involving all available clients in every round, resource-aware client selection strategies dynamically choose participants based on hardware availability, energy state, and expected utility. Oort, for instance, prioritizes clients that offer high training utility while avoiding consistently slow devices, thereby reducing time-to-accuracy and minimizing wasted computation [177]. Other methods adapt local training workloads by tuning the number of

 doi:10.20944/preprints202601.0271.v1

local epochs or batch sizes per client, allowing weaker devices to perform fewer updates while stronger devices contribute more [178]. Deadline-aware and energy-aware schedulers further ensure that clients participate only when they can complete training within round constraints, preventing stragglers from delaying aggregation.

Asynchronous federated learning provides an additional mechanism for reducing computation inefficiencies by relaxing strict round-level synchronization. In asynchronous FL, client updates are aggregated as they arrive, eliminating idle waiting time for faster devices and reducing the impact of slow or overloaded clients [179]. Staleness-aware aggregation rules and asynchronous optimization techniques improve convergence stability despite delayed updates [180,181]. Hybrid asynchronous approaches that combine low-precision updates or buffered aggregation further reduce both computation and communication overhead, enabling scalable FL under highly heterogeneous conditions.

Recent research also explores **resource-aware optimization and knowledge distillation** as computation-efficient alternatives. Resource-aware FL algorithms explicitly incorporate client computational budgets into the optimization process by dynamically adjusting learning rates, gradient sparsity, or update frequency based on device constraints [182]. Knowledge distillation enables clients to train lightweight student models while aligning them with a higher-capacity global or server-side teacher, significantly reducing local computation without sacrificing accuracy [170]. Such approaches are particularly attractive in cross-device settings, where maintaining a single large model across all participants is infeasible.

Overall, mitigating computation overhead remains a critical research direction in federated learning. While model lightweighting, partial training, split learning, adaptive scheduling, and asynchronous optimization have demonstrated substantial gains, preserving accuracy, fairness, and robustness across heterogeneous clients continues to present open challenges. Designing computation-efficient FL systems that scale gracefully with increasing model complexity and device diversity remains a key requirement for real-world deployment.

7. Challenge 3: Communication Bottlenecks

Communication remains one of the most significant bottlenecks in FL due to the repeated transmission of model updates between clients and the central server over bandwidth-limited, high-latency, or unreliable networks. In many real-world deployments, communication cost dominates the total training time, often exceeding local computation by orders of magnitude. Large deep learning models require transmitting millions of parameters per round, making FL expensive for mobile or IoT clients with constrained uplink capacities. Additionally, network instability and intermittent connectivity can cause client dropout, further degrading convergence. To mitigate these challenges, prior work investigates model update compression, gradient sparsification, quantization, sketching, communication-efficient optimizers, and periodic or asynchronous aggregation. Protocols such as secure aggregation, though essential for privacy, further increase communication load due to cryptographic overhead. Emerging approaches such as peer-to-peer communication, hierarchical aggregation, and device-to-edge offloading aim to reduce bandwidth consumption while preserving accuracy. Nevertheless, achieving scalable, reliable, and communication-efficient FL remains an open research problem, especially in cross-device scenarios where devices frequently join and leave the network.

Gradient sparsification and update compression: One of the most widely explored approaches to alleviate communication bottlenecks is compressing the model updates before transmission. Gradient sparsification techniques transmit only the top-k or most significant gradient elements, effectively reducing uplink traffic by orders of magnitude while maintaining convergence guarantees [183,184]. Complementary methods such as momentum correction and error feedback ensure that information lost due to sparsification is gradually incorporated over subsequent rounds [185]. Additional compression schemes, including random projection, sketching, and sign-based encoding, reduce message sizes

 doi:10.20944/preprints202601.0271.v1

further by communicating low-precision or binary gradient representations [186]. These methods are particularly effective in settings where clients possess severely limited bandwidth.

Quantization and low-precision communication: Quantization-based FL techniques reduce the number of bits needed to encode model parameters or gradients. Approaches such as QSGD quantize gradients to a small number of discrete levels, achieving significant bandwidth reductions with minimal accuracy loss [187]. Ternary or binary quantization schemes reduce update sizes even more aggressively and enable efficient hardware acceleration on edge devices [188]. Mixed-precision communication strategies allow different parts of the model to be quantized at different levels based on their sensitivity, balancing accuracy and communication cost [189]. Overall, quantization significantly reduces per-round communication overhead, enabling FL to scale across low-bandwidth clients.

Periodic, local, and partial model aggregation: Another family of techniques reduces communication frequency by allowing multiple local steps before global synchronization. Federated Averaging (FedAvg) is the canonical example, where clients perform several epochs locally before uploading updates [190]. More advanced methods adapt the communication interval based on model convergence rate, resource conditions, or gradient divergence [191]. Partial model communication, such as selective layer-sharing, transmitting only a subset of parameters, or layer-wise clustering, reduces communication size without requiring full-model transmission each round [192]. These approaches significantly lower the communication cost per round and mitigate bandwidth exhaustion in cross-device FL.

Asynchronous communication and decentralized aggregation: Asynchronous FL frameworks allow clients to transmit updates at different times without waiting for global synchronization, removing idle time and mitigating the impact of slow or intermittently connected devices [193]. Staleness-aware aggregation techniques weight updates based on their freshness to maintain convergence stability [194]. Decentralized or peer-to-peer FL frameworks eliminate the central server, enabling clients to exchange updates with neighbors in a communication graph, thereby reducing uplink congestion and server bandwidth requirements [195]. Such decentralized approaches are especially effective in large-scale networks with volatile connectivity.

Hierarchical and edge-assisted aggregation: Hierarchical FL introduces intermediate aggregation layers, such as edge servers, access points, or mobile base stations, to aggregate local updates before forwarding them to the cloud. This architecture reduces communication between resource-limited clients and the central server by leveraging high-speed intra-edge communication [196]. Approaches such as multi-tier FL, clustered FL, and device-to-edge offloading further distribute communication load and improve scalability in dense environments [197]. Hierarchical aggregation is particularly beneficial in mobile networks, where devices frequently join or leave and connectivity varies widely across locations.

Communication-efficient optimization algorithms: Recent work incorporates communication awareness directly into the optimization process. Methods such as adaptive gradient clipping, update skipping, and communication-triggering conditions reduce unnecessary transmissions by sending updates only when they meaningfully improve the global model [198]. Additionally, compressed or quantized variants of federated optimization algorithms (e.g., EF-SGD, SignSGD, and compressed FedAvg) explicitly balance local computation with communication efficiency [199]. These approaches represent a convergence of optimization theory and systems design, offering principled ways to minimize communication without sacrificing model quality.

In summary, communication bottlenecks remain a central challenge to scalable FL deployment. While compression, quantization, and hierarchical aggregation offer substantial improvements, ensuring reliable and efficient communication under real-world network constraints continues to require further innovation.

 doi:10.20944/preprints202601.0271.v1

8. Challenge 4: Client Selection

Client selection is a fundamental component of federated learning (FL) because it directly impacts convergence behavior, communication efficiency, fairness, and robustness of the global model. In each training round, the server must determine which subset of clients should participate, a task made difficult by the heterogeneous and unreliable nature of real-world clients. Uniform random sampling, although widely used in baseline FL systems such as FedAvg, often becomes inefficient under Non-IID data distributions or varied computational and communication capabilities. More advanced strategies incorporate system-awareness, selecting clients based on their resource availability, expected completion time, hardware capabilities, or network quality to avoid stragglers and reduce round latency [177]. These system-aware methods aim to maximize throughput and reduce failures by prioritizing clients that are more likely to complete local training efficiently.

Statistical and data-aware client selection: A second class of methods focuses on mitigating statistical heterogeneity by selecting clients whose data distributions are diverse or representative of the global population. Techniques such as clustering-based selection, gradient-similarity sampling, and distribution-aware scoring attempt to reduce Non-IID drift by ensuring that each round aggregates sufficiently balanced updates [200]. Representative sampling approximates global data coverage using proxies such as local label histograms or low-dimensional embeddings, enabling improved convergence and stability compared to uniform sampling [201]. These approaches help address data imbalance and reduce update bias in cross-device FL.

Incentive and economics-based selection: In environments where client participation is costly, uncertain, or voluntary, incentive-compatible mechanisms are required to ensure reliable engagement. Auction-based Federated Learning (AFL) introduces bidding and payment mechanisms that allow clients to "bid" their willingness or cost of participation; the server then selects clients that maximize utility under budget constraints [202]. Such mechanism-design–based strategies ensure truthful reporting from clients and improve resource usage by selecting participants who are both capable and motivated. Similarly, contract-theoretic and token-based incentive schemes encourage participation while maintaining energy fairness [203].

Learning-based and adaptive client selection: Recent works employ reinforcement learning, bi-level optimization, and meta-learning to learn dynamic client-selection policies that balance efficiency, fairness, and model accuracy [204]. These approaches treat selection as a sequential decision process where the server observes client performance, such as loss reduction, resource consumption, or reliability—and adapts selection over time. This enables more robust policies in environments with volatile connectivity, unpredictable dropouts, and changing data distributions. Learning-based methods have shown strong potential for optimizing multi-objective trade-offs that traditional heuristics struggle to manage.

Fairness and reliability-aware selection: Fairness-driven client selection aims to avoid over-selecting strong clients while neglecting weak or underrepresented groups. Such algorithms enforce fairness constraints, minimize client starvation, or promote long-term participation balance to ensure equitable model performance across demographic or device-level groups [205]. Reliability-aware methods incorporate client dropout prediction, connectivity modeling, or redundancy strategies to maintain robustness when clients fail mid-round or provide stale updates. These techniques improve fault tolerance and stabilize convergence in large-scale, cross-device FL deployments.

Despite these advances, several challenges remain unresolved: Non-IID drift can still degrade model quality if biased clients are repeatedly selected; device dropouts introduce instability; privacy constraints limit the extent of data-aware selection; and fairness–efficiency trade-offs remain difficult to optimize in massive FL systems. As highlighted in recent surveys, future research is increasingly moving toward multi-objective selection frameworks that jointly consider accuracy, fairness, resource constraints, privacy, and communication cost, as well as reliability-aware methods capable of operating effectively at scale under dynamic client behavior.

 doi:10.20944/preprints202601.0271.v1

9. Challenge 5: Aggregation and Optimization

Aggregation and optimization lie at the core of Federated Learning (FL), directly shaping convergence behavior, robustness, fairness, and global model quality. Due to statistical heterogeneity, partial participation, and inconsistent local training dynamics across clients, standard federated averaging (FedAvg) [32] often suffers from instability, client drift, and slow convergence. To overcome these issues, recent research focuses on enhancing aggregation rules, modifying optimization objectives, and designing principled mechanisms to stabilize cross-client learning.

Correction-term and regularization-based optimization: A major line of work improves FL optimization by incorporating correction terms that compensate for client drift under Non-IID data. FedProx [143] adds a proximal regularization term to each client's local objective, penalizing updates that deviate too far from the global model. SCAFFOLD [137] introduces control variates to estimate and correct for drift, enabling more stable convergence even under severe heterogeneity. Other approaches extend these ideas through normalized updates, adaptive learning rates, or server-side momentum (e.g., FedNova [206]), all aiming to harmonize client contributions during aggregation.

Knowledge Distillation (KD)-enhanced aggregation: Knowledge distillation provides an alternative path to robust aggregation by transferring knowledge through soft predictions, features, or synthetic proxy data instead of raw gradients. Methods such as FedDF [207], FedGen [208], and data-free distillation schemes [209,210] aggregate models via ensemble distillation on the server or by generating synthetic anchor samples to align local knowledge. KD-based aggregation improves robustness against client divergence, supports heterogeneous architectures, and facilitates communication-efficient optimization.

Model mutation and diversity-inducing optimization: Another stream of work promotes diversity during training by perturbing, mutating, or adaptively modifying global model parameters. FedMut [211] injects noise or mutations into the global model to encourage exploration in the optimization landscape, reducing overfitting to biased client data. FedQP [212] adopts quasi-Newton updates to introduce curvature-aware corrections, improving local convergence accuracy. These methods help FL escape poor local minima induced by statistically skewed client data.

Client clustering for structured aggregation: Client clustering aggregates clients based on similarity in data distribution, gradient direction, or representation space. ClusterSampling [213] selects statistically representative clients to improve aggregation stability. Other hierarchical and cluster-aware FL approaches [214–216] form groups of similar clients and perform cluster-wise aggregation before global fusion, mitigating Non-IID effects and reducing variance. Such structured aggregation significantly enhances FL robustness, particularly when client populations are large and diverse.

Multi-model search, architecture exploration, and ensemble optimization: Recent methods explore multiple model candidates or architectures during training to better capture heterogeneous data characteristics. Multi-model search techniques [217–219] jointly optimize several global models or submodels and dynamically select or merge them based on client feedback. This perspective treats aggregation as a multi-objective optimization problem, enabling FL systems to adaptively identify architectures or models that generalize well across diverse clients.

Unified optimization frameworks: Finally, many works design unified frameworks that integrate enhanced optimizers with improved aggregation rules, deploying them across multiple FL algorithms. For example, FedFed introduces a meta-aggregation layer compatible with FedAvg, FedProx, SCAFFOLD, and FedNova, offering consistent improvements across heterogeneous environments. These frameworks highlight the growing trend toward algorithm-agnostic optimization and aggregation pipelines in FL.

Overall, aggregation and optimization challenges in FL stem from the tension between local autonomy and global coordination. While correction terms, distillation, clustering, mutation, and multi-model exploration significantly improve robustness and convergence, developing universally stable, scalable, and fairness-aware optimization rules remains an open research frontier.

 doi:10.20944/preprints202601.0271.v1

10. Challenge 6: Privacy Preservation

Privacy preservation is a core requirement of Federated Learning (FL), where the central principle is that raw client data must remain local. Despite not sharing data directly, FL is still vulnerable to numerous privacy threats, including gradient inversion attacks, membership inference, and reconstruction of client-specific features [34,220,221]. These risks are amplified by statistical heterogeneity, where unique or skewed client data patterns leave strong signatures in local updates, making them easier to re-identify. Moreover, many recent advances in optimization, clustering, and model personalization provide improved convergence but inadvertently increase privacy leakage risk by exposing richer model updates or auxiliary client information.

Secure Aggregation and cryptographic protection: Secure aggregation (SecAgg) ensures that the server can only observe the aggregated model update, not individual client contributions. While essential in preventing direct reconstruction of local gradients, SecAgg introduces significant communication and computation overhead [222]. Furthermore, several advanced FL paradigms, such as client clustering, multi-model searching, or structured aggregation, are incompatible with SecAgg because they require access to individualized updates, thereby elevating privacy risks. This tension between algorithmic flexibility and cryptographic protection remains a key challenge.

Differential privacy and noise-based defenses: Differential Privacy (DP) provides a mathematically rigorous framework for limiting what can be inferred about individual clients by injecting calibrated noise into model updates [223]. DP-SGD and its federated variants reduce leakage risk but may degrade performance under Non-IID data, where noise disproportionately affects clients with smaller or highly skewed datasets. Achieving an optimal accuracy–privacy trade-off in FL thus remains a difficult and application-dependent problem.

Machine unlearning and erasure guarantees: Machine unlearning techniques enable the removal of a client's contribution from the global model, supporting the "right to be forgotten." Federated unlearning frameworks [224] revise or retrain global models to erase specific client updates without reconstructing the entire training process. These methods are essential for regulatory compliance but become complex under heterogeneous data distributions and partial client participation.

Privacy-preserving federated clustering and representation sharing: Federated clustering (FC) methods construct global clustering structures (e.g., centroids, basis matrices, kernel matrices) without accessing raw data. Secure FC adopts DP, secure aggregation, and cryptographic encoding techniques, including Lagrange coded computing [225], to enhance clustering privacy while improving global structure inference. Recent work achieves secure global k-means via homomorphic encryption or DP-protected centroid sharing [150,226–228] and secure spectral clustering via kernel reconstruction [152]. Despite progress, these methods inherit restrictive assumptions from centralized clustering, such as data compactness [229,230] or graph connectivity [231,232], limiting practical applicability under highly Non-IID settings.

Coded computing, masking, and redundancy-based privacy: Lagrange coded computing (LCC) partitions tasks into masked, coded sub-tasks that can be distributed across multiple servers, ensuring information-theoretic privacy even if some servers are compromised [225]. Although powerful, these techniques introduce substantial computational overhead and typically assume synchronous or structured participation patterns that do not align well with dynamic cross-device FL deployments.

Overall, privacy preservation in FL remains challenging due to inherent trade-offs between security, utility, scalability, and computation cost. While cryptographic protection, differential privacy, and coded computing provide robust defenses, they often struggle with the heterogeneity and unpredictability of real-world federated environments. Future research will likely explore adaptive privacy budgets, privacy–utility co-optimization, secure model compression, and hybrid approaches that combine DP, SecAgg, and learning-based privacy monitoring.

 doi:10.20944/preprints202601.0271.v1

11. Applications of Federated Learning

Federated learning (FL) has been adopted across a wide range of domains where data are inherently distributed, privacy-sensitive, or regulated. Based on an extensive review of the literature, we categorize FL applications into eight major domains, each characterized by distinct data modalities, system constraints, and deployment requirements.

Healthcare and Medical Research: Healthcare is one of the most prominent and impactful application domains for FL, driven by strict privacy regulations (e.g., HIPAA and GDPR) and the distributed nature of clinical data. Hospitals, diagnostic centers, and research institutions can collaboratively train models without sharing raw patient data. Representative applications include disease detection from medical imaging, such as tumor identification in CT/MRI scans and diabetic retinopathy classification from retinal images [233–235] as well as personalized treatment recommendations, clinical risk prediction, drug discovery, and workflow optimization [236–239]. FL also gained significant attention during the COVID-19 pandemic for CT-based diagnosis, epidemic forecasting, and hospital resource planning [240–245]. Beyond imaging, FL supports privacy-preserving electronic health record (EHR) analysis, hospitalization prediction, emotion and activity recognition, and multi-institutional patient modeling [246–250].

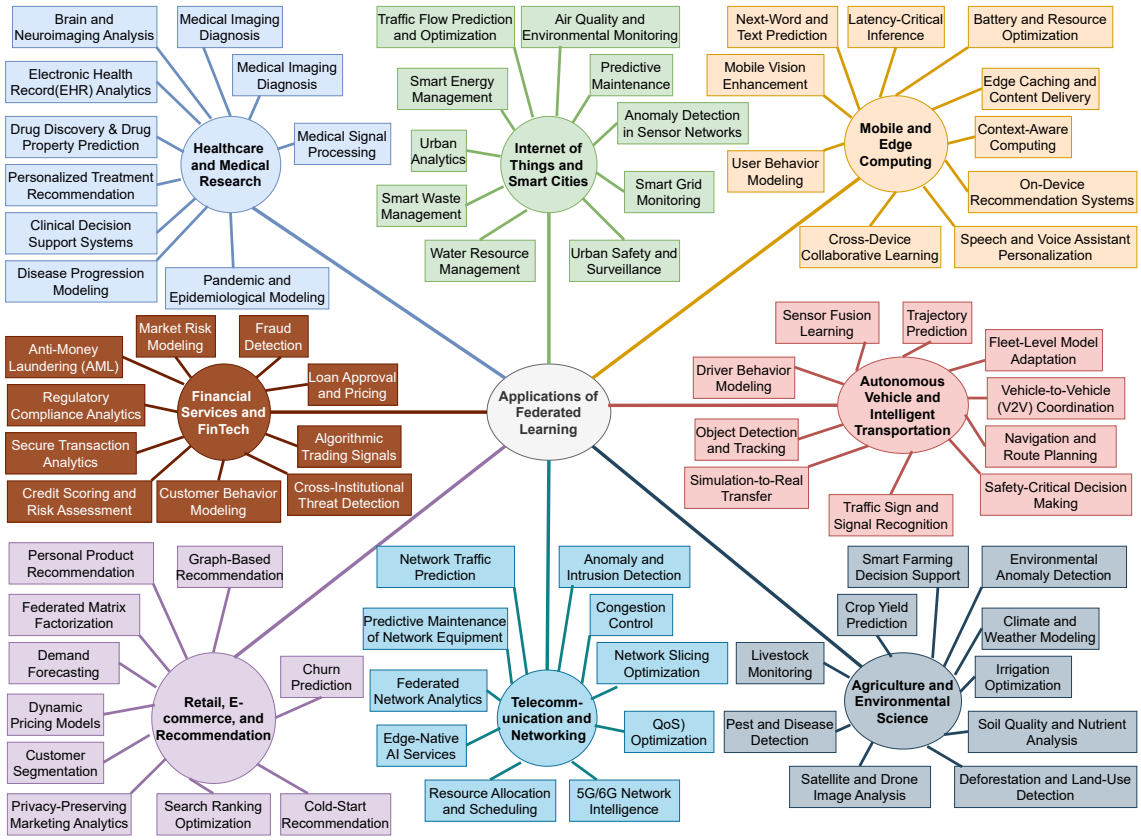

Figure 5. Application landscape of federated learning. The figure summarizes eight major application domains of federated learning, healthcare and medical research, internet of things and smart cities, mobile and edge computing, financial services and FinTech, autonomous vehicles and intelligent transportation, retail and e-commerce, telecommunications and networking, and agriculture and environmental science, along with representative subcategories within each domain. The diagram highlights the breadth of FL deployments across data modalities, system scales, and privacy-sensitive environments.

Internet of Things and Smart Cities: IoT ecosystems naturally align with FL due to their decentralized architecture, bandwidth limitations, and heterogeneous devices. FL enables collaborative intelligence for smart homes, industrial IoT, and smart-city infrastructures while minimizing communication overhead and preserving data locality. Applications include traffic flow prediction, air-quality

 doi:10.20944/preprints202601.0271.v1

monitoring, energy consumption forecasting, predictive maintenance, and anomaly detection in large-scale sensor networks [251–255].

Additional research explores cost-aware sensing [256], client-aware optimization [257], electric vehicle (EV) charging and energy management [258], and personalized IoT services [259,260]. These settings highlight the importance of communication efficiency and robustness under device churn.

Mobile and Edge Computing: Cross-device FL is widely deployed in mobile and edge environments, where data are generated continuously on user devices. A canonical example is Google's Gboard keyboard, which uses FL to improve next-word prediction without collecting raw keystrokes. Other applications include mobile recommendation engines, voice assistant personalization, app usage prediction, on-device image enhancement, and battery optimization. Edge computing platforms further benefit from FL by exploiting local computation while reducing reliance on centralized cloud infrastructure, particularly in latency-sensitive applications.

Financial Services and FinTech: Financial institutions employ FL to collaboratively build models for fraud detection, credit scoring, anti-money laundering, loan default prediction, and customer behavior analysis. FL enables multiple banks or financial entities to jointly learn risk patterns without exposing proprietary transaction records or violating regulatory constraints. These collaborative models improve detection of cross-institutional fraud and systemic risks while maintaining strict data confidentiality.

Autonomous Vehicles and Intelligent Transportation: Autonomous vehicles and intelligent transportation systems generate massive volumes of multimodal sensor data, including LiDAR, radar, GPS, and camera streams. FL enables fleet-level learning for object detection, trajectory prediction, motion planning, and vehicle-to-vehicle coordination without centralized data aggregation. By leveraging diverse driving conditions across regions and environments, FL improves model robustness while reducing privacy risks associated with raw data sharing.

Retail, E-commerce, and Recommendation Systems: Retail and e-commerce platforms increasingly adopt FL for privacy-preserving personalization, including product recommendation, demand forecasting, inventory management, customer churn prediction, and dynamic pricing. Federated recommendation systems allow collaborative model training across distributed retailers or user devices while protecting user preferences and business-sensitive data. Recent advances focus on federated matrix factorization, graph-based recommender models [261,262], and adaptive knowledge-fusion techniques for cold-start scenarios [263].

Telecommunications and Networking: Telecommunication networks leverage FL to optimize network operations across distributed base stations and edge servers. Applications include traffic prediction, anomaly detection, congestion control, resource allocation, and predictive maintenance of network infrastructure. By avoiding centralized aggregation of sensitive operational data, FL enables collaborative optimization of quality of service (QoS) and supports intelligent, network-native AI services [264].

Agriculture and Environmental Science FL supports precision agriculture and environmental monitoring by aggregating insights from geographically distributed data sources such as farm sensors, drones, and satellite imagery. Key applications include crop-yield prediction, soil and irrigation analytics, pest and disease detection, climate modeling, and large-scale environmental surveillance. These settings emphasize data sovereignty, spatial heterogeneity, and long-term sustainability.

Across these eight application domains, federated learning demonstrates its versatility in enabling collaborative intelligence under privacy, regulatory, and system constraints. Each domain presents unique challenges in terms of data distribution, communication efficiency, computation, and trust, reinforcing the need for domain-aware FL algorithm design and evaluation.

12. Open Source Systems

Flower (FLWR): Flower (FLWR) is a flexible, framework-agnostic federated learning platform that has gained significant traction in recent research due to its clean separation between federated orches-

 doi:10.20944/preprints202601.0271.v1

tration and local model training. By decoupling client-side learning logic from server-side coordination and aggregation, Flower enables rapid prototyping of novel federated algorithms across heterogeneous machine learning backends, including PyTorch, TensorFlow, JAX, and classical models. This modular design has made Flower particularly suitable for studying key federated learning challenges such as data heterogeneity, personalized aggregation, asynchronous training, and communication-efficient protocols. Recent works have adopted Flower as an experimental backbone for personalized and multimodal federated learning pipelines, as well as for benchmarking system-aware aggregation and client selection strategies under realistic deployment conditions. Owing to its extensibility and growing adoption in peer-reviewed literature, Flower serves as a representative research framework that bridges theoretical FL algorithm design and practical large-scale experimentation [265–267].

TensorFlow Federated (TFF): TensorFlow Federated (TFF) is a research-oriented federated learning framework developed to support principled algorithm design and theoretical analysis of federated optimization. TFF provides explicit abstractions for federated computations, enabling precise specification of client-side and server-side logic, aggregation operators, and communication rounds within a mathematically grounded programming model. This design facilitates rigorous experimentation with canonical and emerging federated learning algorithms, including variants addressing statistical heterogeneity, partial participation, and constrained optimization. Due to its close alignment with formal FL formulations and its tight integration with the TensorFlow ecosystem, TFF has been widely adopted in foundational studies on federated optimization, convergence behavior, and algorithmic robustness. As a result, TFF serves as a canonical research platform for validating theoretical advances in federated learning under controlled and reproducible experimental settings [32,197,268,269].

FedML: FedML is an open-source federated learning framework designed to support both algorithmic research and system-level experimentation at scale. It provides an end-to-end research pipeline that integrates federated optimization algorithms, benchmarking datasets, and deployment tools under a unified interface. FedML emphasizes reproducibility and scalability, enabling controlled evaluation of federated learning methods under statistical heterogeneity, partial client participation, communication constraints, and system dynamics. Its support for cross-device, cross-silo, and decentralized federated learning settings has made it a common platform for studying communication efficiency, client selection, and resource-aware training strategies. As a result, FedML has been widely adopted in recent research for benchmarking federated learning algorithms and analyzing the interaction between learning performance and system heterogeneity in realistic distributed environments [270,271].

PySyft (OpenMined): PySyft is an open-source federated learning and privacy-preserving machine learning framework developed by the OpenMined community, with a strong emphasis on data confidentiality and secure distributed computation. Unlike frameworks that primarily target optimization efficiency or system scalability, PySyft focuses on integrating federated learning with formal privacy-enhancing technologies, including differential privacy, secure multi-party computation, and encrypted tensor abstractions. This design enables researchers to study federated learning under explicit privacy constraints and adversarial threat models, making PySyft particularly suitable for investigating privacy–utility trade-offs, attack resilience, and secure aggregation protocols. PySyft has been adopted in research exploring private collaborative learning across sensitive domains such as healthcare and finance, and it serves as a representative platform for security- and privacy-centric federated learning studies [269,272].

FATE (Federated AI Technology Enabler): Federated AI Technology Enabler (FATE) is an open-source federated learning framework developed under the Linux Foundation with a primary focus on secure, cross-organizational collaboration. FATE is particularly well-suited for vertical federated learning, where Multiple parties share sample identities but hold complementary feature sets. The framework integrates privacy-preserving technologies such as secure multi-party computation and homomorphic encryption to enable a joint model training without revealing sensitive features or labels. Due to its emphasis on feature-partitioned learning and cryptographic security guarantees, FATE has been widely adopted in research on vertical federated learning, privacy-preserving data fusion, and

 doi:10.20944/preprints202601.0271.v1

collaborative modeling across regulated domains such as finance and healthcare. As a result, FATE serves as a representative research platform for studying secure and feature-distributed federated learning scenarios [273,274].

13. Future Directions

Federated learning continues to evolve beyond its original client-server paradigm, driven by emerging requirements in trust, security, scalability, and computational capability. This section highlights two promising research directions that are expected to significantly influence the next generation of federated learning systems: blockchain-enabled federated learning and quantum federated learning.

Blockchain-Enabled Federated Learning: Blockchain technology has recently attracted attention as a complementary mechanism for addressing trust, transparency, and decentralization challenges in federated learning [275]. In conventional FL systems, a central server coordinates client selection, aggregation, and model dissemination, which introduces a single point of trust and failure. Blockchain-based FL replaces or augments this centralized coordinator with a distributed ledger, enabling immutable recording of model updates, client participation, and aggregation outcomes. By leveraging smart contracts, blockchain-enabled FL systems can automate aggregation rules, incentive mechanisms, and access control without relying on a trusted third party [275]. This integration is particularly appealing for cross-silo and inter-organizational settings, where participants may not fully trust a central aggregator. Moreover, blockchain-based auditability provides verifiable guarantees on training integrity, which is critical for regulated domains such as finance and healthcare. Despite its promise, blockchain-enabled FL introduces new challenges, including increased communication latency, storage overhead, and limited transaction throughput. Future research is needed to design lightweight consensus mechanisms, hierarchical ledger architectures, and hybrid off-chain/on-chain solutions that balance security with efficiency. Exploring the interaction between blockchain protocols and federated optimization dynamics remains an open and important research direction.

Quantum Federated Learning: Quantum federated learning (QFL) represents an emerging interdisciplinary direction that combines federated learning with quantum computing and quantum communication technologies. The motivation behind QFL is twofold: to exploit the computational advantages of quantum machine learning models and to enhance security through quantum cryptographic primitives. Recent studies have explored federated training of quantum neural networks, variational quantum circuits, and hybrid quantum–classical models across distributed clients [276,277]. In this setting, clients may perform local optimization of quantum parameters while sharing only classical or quantum-derived updates with a coordinator. Other works investigate quantum-secure aggregation and encryption schemes to protect model updates against adversarial inference [278,279]. While QFL remains largely theoretical due to the current limitations of quantum hardware, it opens new research avenues at the intersection of distributed learning, quantum information, and secure communication. Key open challenges include handling quantum noise, limited qubit counts, communication constraints, and the integration of quantum learning protocols with classical federated optimization frameworks. As quantum technologies mature, QFL is expected to play an increasingly important role in privacy-preserving and high-assurance federated intelligence systems [280].

Large-Scale and Multi-Modal Federated Learning: Another major direction is federating large foundation models and multi-modal networks. The advent of billion-parameter models (LLMs, vision-language models, etc.) poses both challenges and opportunities for FL. Recent surveys on Federated LLMs note that FL can enable collaborative training of such models without centralizing sensitive data [281] . For instance, LLMs can be fine-tuned in federated settings (via prompt tuning or LoRA) across many clients, improving personalization while preserving privacy. Even more, LLMs themselves can augment FL: they can generate synthetic training data to enrich scarce or imbalanced local datasets [281] , or distill global knowledge back to clients via prompt engineering [281] . Future work will explore federated pre-training of large models, for example, splitting a base model and collaboratively training parts on different clients, as well as multimodal FL, where vision, text, and other data types

 doi:10.20944/preprints202601.0271.v1

are co-trained. One open question is how to co-optimize multiple models for different modalities so that the overall system is privacy-aware and efficient [281] . In short, federating foundation models and generative networks is a key frontier, likely yielding frameworks where a global transformer or diffusion model is incrementally improved across many devices, or where local models tap global generative priors for data augmentation [281].

Privacy, Security, and Trust Enhancements: While FL was born for privacy, there remain open directions in making it provably secure and trustworthy. Traditional tools (differential privacy, secure aggregation, homomorphic encryption) will continue to evolve, but new paradigms are emerging. For example, research is looking at hardware-rooted security: using Trusted Execution Environments (TEEs), Physical Unclonable Functions (PUFs), or even chaos-based and neuromorphic encryption to offload heavy cryptography [282]. One recent survey emphasizes the need for hardware-agnostic security techniques (so that FL protocols don't rely on specific chips) and standardized benchmarks for privacy methods [282]. Similarly, "hybrid architectures" that combine multiple approaches (e.g., enclaves+MPC+DP, or quantum encryption + classical schemes) are expected to emerge [282]. In practice, we foresee FL systems incorporating automated threat detection (e.g. federated poison-attack defenses), as well as protocols for mutual attestation among clients. Governance and trustworthiness will also be key: for instance, embedding identity- and blockchain-based verification in FL could ensure only vetted participants contribute. In summary, the frontier of FL security lies in blending cryptography, hardware, and system design to achieve scalable privacy without undermining efficiency [282].

Personalization and Adaptation: Real-world FL networks are highly heterogeneous, so personalized federated learning is a major trend. Instead of learning one global model, future FL will tailor models or objectives for individual clients. Techniques like model interpolation, meta-learning, or local fine-tuning are being developed so each user gets a model adapted to their data distribution. Such personalization can also address fairness across diverse participants: as one review notes, future work will explore "fairness, robustness, and personalized federated learning mechanisms" to improve deployment [275]. Alongside, new privacy notions e.g. sample-level or user-level privacy budgets, will allow more granular trade-offs per client. We also expect adaptive FL algorithms that handle concept drift: for example, as data patterns change over time or new devices join, the system will dynamically re-weight updates or re-train sub-models. In short, embracing heterogeneity will drive new FL frameworks where multiple local and global models coexist, each honed to client needs while still benefiting from collective training [275].

Beyond Supervised Learning: So far most FL work focuses on supervised tasks, but the future will expand FL to other learning paradigms. One direction is unsupervised and self-supervised FL: developing federated algorithms for representation learning, clustering or density estimation when labels are scarce. Another is federated reinforcement learning (FedRL): here multiple agents (e.g. robots, vehicles) collaboratively learn policies without sharing raw trajectories. Tackling these requires new methods for non-i.i.d. sequential data and multi-agent coordination. Indeed, surveys highlight that moving "beyond supervised learning in federated networks" e.g. exploring reinforcement learning or federated analytics is a key open problem [283] . Relatedly, one can imagine federated multi-task or transfer learning, where different clients solve related but not identical problems, exchanging model knowledge. As FL reaches maturity, we will likely see frameworks that seamlessly integrate these broader ML tasks into the federated infrastructure [283].

Benchmarks, Frameworks, and Deployment: Finally, for FL to advance, the community must build common benchmarks, platforms, and case studies. Early efforts like the LEAF benchmark and TensorFlow Federated have helped, but the field needs more realistic datasets, standardized evaluation suites, and open-source implementations. Future directions include large-scale "federation-as-a-service" testbeds, cross-industry consortia for FL standards, and federated MLaaS platforms that support many algorithms. For instance, one review urges researchers to "build upon existing implementations and benchmarking tools, to facilitate reproducibility and dissemination of new

 doi:10.20944/preprints202601.0271.v1

solutions for FL" [283] . Work will also focus on real-world deployments, for example in healthcare or smart grids, to understand FL at scale (e.g. handling device churn, network variability, regulatory constraints). In sum, establishing common frameworks and best practices will be crucial for federated learning's next phase of growth [283].

14. Conclusion

Federated learning (FL) has emerged as a foundational paradigm for privacy-preserving and distributed machine learning, motivated by the growing need to train models across large-scale, heterogeneous, and sensitive data sources without centralized data aggregation. This survey presented a comprehensive, challenge-centric overview of the FL landscape, systematically examining six core challenges that fundamentally shape the design, deployment, and performance of FL systems. These challenges include data and system heterogeneity, communication efficiency, computational constraints, client selection, aggregation and optimization, privacy and security, and the integration of FL with other machine learning paradigms. Across each challenge area, we reviewed representative state-of-the-art approaches, identified recurring design patterns and trade-offs, and highlighted persistent limitations that continue to impede real-world deployment. Despite significant progress, FL remains far from a universally mature or turnkey solution. Practical systems must contend with highly non-IID data distributions, unreliable and resource-constrained clients, communication bottlenecks, and the inherent tension between model utility, privacy guarantees, and robustness. Moreover, as FL extends beyond its original scope into emerging paradigms such as continual learning, graph learning, and multimodal learning, additional complexities arise, demanding algorithms that can adapt to dynamic data distributions, evolving client populations, and heterogeneous model architectures.

Looking forward, advancing FL will require a shift toward holistic and unified frameworks that jointly address efficiency, robustness, fairness, privacy, and adaptability, rather than optimizing these objectives in isolation. Promising research directions include privacy–utility co-optimization, adaptive and event-driven communication protocols, personalized and hierarchical FL architectures, secure and scalable aggregation mechanisms, and deeper integration with foundation models and large pre-trained architectures. Equally important are efforts toward standardized benchmarks, reproducible evaluation pipelines, and system-level validation in real-world settings, supported by closer collaboration among academia, industry, and regulatory stakeholders.

In conclusion, federated learning stands at a critical juncture. While its foundational principles are now well established, realizing scalable, trustworthy, and general-purpose federated intelligence remains an open and interdisciplinary challenge. By synthesizing the challenges, methods, and open problems outlined in this survey, we aim to provide a unified reference point for the community and to inspire future research that bridges the gap between theoretical advances and practical, large-scale FL deployments across diverse application domains.

References

1. Kuutti, S.; Bowden, R.; Jin, Y.; Barber, P.; Fallah, S. A Survey of Deep Learning Applications to Autonomous Vehicle Control. *IEEE Transactions on Intelligent Transportation Systems* **2021**, *22*, 712–733. https://doi.org/10.1109/TITS.2019.2962338.
2. Mozaffari, S.; Al-Jarrah, O.Y.; Dianati, M.; Jennings, P.; Mouzakitis, A. Deep Learning-Based Vehicle Behavior Prediction for Autonomous Driving Applications: A Review. *IEEE Transactions on Intelligent Transportation Systems* **2022**, *23*, 33–47. https://doi.org/10.1109/TITS.2020.3012034.
3. Liu, L.; Lu, S.; Zhong, R.; Wu, B.; Yao, Y.; Zhang, Q.; Shi, W. Computing Systems for Autonomous Driving: State of the Art and Challenges. *IEEE Internet of Things Journal* **2021**, *8*, 6469–6486. https://doi.org/10.1109/JIOT.2020.3043716.
4. Wang, B.; Zheng, Y.; Han, X.; et al. A Systematic Literature Review on Integrating AI-Powered Smart Glasses into Digital Health Management for Proactive Healthcare Solutions. *npj Digital Medicine* **2025**, *8*, 410. https://doi.org/10.1038/s41746-025-01715-x.

 doi:10.20944/preprints202601.0271.v1

5. Yalcin, N.; Alisawi, M. Enhancing Social Interaction for the Visually Impaired: A Systematic Review of Real-Time Emotion Recognition Using Smart Glasses and Deep Learning. *IEEE Access* **2025**, *13*, 102092–102108. https://doi.org/10.1109/ACCESS.2025.3577106.
6. Hoang, M.L. A Review of Developments and Metrology in Machine Learning and Deep Learning for Wearable IoT Devices. *IEEE Access* **2025**, *13*, 106035–106054. https://doi.org/10.1109/ACCESS.2025.3573937.
7. Xiong, J.; Hsiang, E.L.; He, Z.; et al. Augmented Reality and Virtual Reality Displays: Emerging Technologies and Future Perspectives. *Light: Science & Applications* **2021**, *10*, 216. https://doi.org/10.1038/s41377-021-00658-8.
8. Liberatore, M.J.; Wagner, W.P. Virtual, Mixed, and Augmented Reality: A Systematic Review for Immersive Systems Research. *Virtual Reality* **2021**, *25*, 773–799. https://doi.org/10.1007/s10055-020-00492-0.
9. Tong, Y.; Liu, H.; Zhang, Z. Advancements in Humanoid Robots: A Comprehensive Review and Future Prospects. *IEEE/CAA Journal of Automatica Sinica* **2024**, *11*, 301–328. https://doi.org/10.1109/JAS.2023.124140.
10. Carpentier, J.; Wieber, P.B. Recent Progress in Legged Robots Locomotion Control. *Current Robotics Reports* **2021**, *2*, 231–238. https://doi.org/10.1007/s43154-021-00059-0.
11. Kotha, S.S.; Akter, N.; Abhi, S.H.; Das, S.K.; Islam, M.R.; Ali, M.F.; Ahamed, M.H.; Islam, M.M.; Sarker, S.K.; Badal, M.F.R.; et al. Next Generation Legged Robot Locomotion: A Review on Control Techniques. *Heliyon* **2024**, *10*. https://doi.org/10.1016/j.heliyon.2024.e37237.
12. Ahmed, F.; Mohanta, J.C.; Keshari, A.; et al. Recent Advances in Unmanned Aerial Vehicles: A Review. *Arabian Journal for Science and Engineering* **2022**, *47*, 7963–7984. https://doi.org/10.1007/s13369-022-06738-0.
13. Küçükerdem, H.; Yilmaz, C.; Kahraman, H.T.; et al. Autonomous Control of Unmanned Aerial Vehicles: Applications, Requirements, Challenges. *Cluster Computing* **2025**, *28*, 734. https://doi.org/10.1007/s10586-025-05418-6.
14. Mohsan, S.A.H.; Othman, N.Q.H.; Li, Y.; et al. Unmanned Aerial Vehicles (UAVs): Practical Aspects, Applications, Open Challenges, Security Issues, and Future Trends. *Intelligent Service Robotics* **2023**, *16*, 109–137. https://doi.org/10.1007/s11370-022-00452-4.
15. IoT Analytics. State of IoT 2025: Number of connected IoT devices growing 1421.1 billion globally and projected to 39 billion by 2030. https://iot-analytics.com/number-connected-iot-devices/, 2025. Accessed: 2025-12-16.
16. Fayyad, U.; Piatetsky-Shapiro, G.; Smyth, P. The KDD process for extracting useful knowledge from volumes of data. *Commun. ACM* **1996**, *39*, 27–34. https://doi.org/10.1145/240455.240464.
17. Shearer, C. The CRISP-DM Model: The New Blueprint for Data Mining, 2000.
18. Baylor, D.; Breck, E.; Cheng, H.T.; et al. TFX: A TensorFlow-Based Production-Scale Machine Learning Platform. In Proceedings of the Proceedings of the 23rd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining, 2017.
19. Mohd Noor, M.H.; Ige, A.O. A survey on state-of-the-art deep learning applications and challenges. *Engineering Applications of Artificial Intelligence* **2025**, *159*, 111225. https://doi.org/10.1016/j.engappai.2025.111225.
20. Baduwal, M. Hybrid(Transformer+CNN)-based Polyp Segmentation, 2025, [arXiv:eess.IV/2508.09189].
21. General Data Protection Regulation (GDPR). https://gdpr.eu/, 2016.
22. California Consumer Privacy Act (CCPA). https://oag.ca.gov/privacy/ccpa, 2018.
23. Challenges of Centralized Machine Learning in Modern Data Ecosystems. *arXiv* **2020**.
24. Silver, D.; Huang, A.; Maddison, C.J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature* **2016**, *529*, 484–489. https://doi.org/10.1038/nature16961.
25. NVIDIA. Nemotron 3 Nano: Open, Efficient Mixture-of-Experts Hybrid Mamba-Transformer Model for Agentic Reasoning, 2025. Technical report.
26. Team, Q. Qwen3 Technical Report, 2025, [arXiv:cs.CL/2505.09388].
27. ChatGPT: Proprietary AI Agent and Conversational Assistant. OpenAI Product, 2025. Closed-source large language model agent for conversational AI and automated workflows.
28. Microsoft Copilot: AI-Driven Agentic Assistance. Microsoft Product, 2025. Enterprise-grade proprietary agent integrated with Microsoft 365 and developer tools.
29. Google Antigravity: Proprietary Agent-First AI IDE. Google Product, 2025. Agent-centric proprietary coding environment powered by Gemini 3 Pro and integrated agents.

 doi:10.20944/preprints202601.0271.v1

30. Anthropic Claude with Opus Agentic Capabilities. Anthropic AI Model, 2025. Proprietary agentic reasoning and task automation enhancements in Claude powered by Opus 4.5.
31. IBM watsonx: Enterprise-Grade Proprietary AI Agents. IBM Product Suite, 2025. Proprietary AI agents and orchestration within IBM's watsonx platform for businesses.
32. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; y Arcas, B.A. Communication-efficient learning of deep networks from decentralized data. In Proceedings of the Artificial intelligence and statistics. PMLR, 2017, pp. 1273–1282.
33. European Data Protection Supervisor. Opinion on privacy and federated learning. https://edps.europa.eu, 2020.
34. Mothukuri, V.; Parizi, R.M.; Pouriyeh, S.; Huang, Y.; Dehghantanha, A.; Srivastava, G. A survey on security and privacy of federated learning. *Future Generation Computer Systems* **2021**, *115*, 619–640. https://doi.org/https://doi.org/10.1016/j.future.2020.10.007.
35. Federated Learning: Applications and Opportunities. *Frontiers in Artificial Intelligence* **2021**.
36. Federated Learning for Medical Applications. *Frontiers in Medicine* **2021**.
37. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S. Communication-Efficient Learning of Deep Networks from Decentralized Data. *Proceedings of AISTATS* **2017**.
38. Federated Intelligence in Smart Cities. *Frontiers in Sustainable Cities* **2021**.
39. Federated Learning for COVID-19 Diagnosis and Analysis. *Nature Communications* **2020**.
40. Privacy-preserving Analytics during the COVID-19 Pandemic with Federated Learning. *PMC* **2020**.
41. Stripelis, D.; Ambite, J.L.; Lam, P.; Thompson, P. Scaling neuroscience research using federated learning. In Proceedings of the 2021 IEEE 18th international symposium on biomedical imaging (ISBI). Ieee, 2021, pp. 1191–1195.
42. Stripelis, D.; Gupta, U.; Saleem, H.; Dhinagar, N.; Ghai, T.; Anastasiou, C.; Sánchez, R.; Ver Steeg, G.; Ravi, S.; Naveed, M.; et al. A federated learning architecture for secure and private neuroimaging analysis. *Patterns* **2024**, *5*.
43. Li, X.; Gu, Y.; Dvornek, N.; Staib, L.H.; Ventola, P.; Duncan, J.S. Multi-site fMRI analysis using privacy-preserving federated learning and domain adaptation: ABIDE results. *Medical image analysis* **2020**, *65*, 101765.
44. Thapaliya, B.; et al. Efficient Federated Learning for Distributed Neuroimaging Data. *Frontiers in Neuroinformatics* **2024**, *18*.
45. Sadilek, A.; Liu, L.; Nguyen, D.; Kamruzzaman, M.; Serghiou, S.; Rader, B.; Ingerman, A.; Mellem, S.; Kairouz, P.; Nsoesie, E.O.; et al. Privacy-first health research with federated learning. *NPJ digital medicine* **2021**, *4*, 132.
46. Xu, J.; Glicksberg, B.S.; Su, C.; Walker, P.; Bian, J.; Wang, F. Federated learning for healthcare informatics. *Journal of healthcare informatics research* **2021**, *5*, 1–19.
47. Li, S.; Miao, D.; Wu, Q.; Hong, C.; D'Agostino, D.; Li, X.; Ning, Y.; Shang, Y.; Wang, Z.; Liu, M.; et al. Federated learning in healthcare: A benchmark comparison of engineering and statistical approaches for structured data analysis. *Health Data Science* **2024**, *4*, 0196.
48. Diniz, J.M. The Missing Subject in Health Federated Learning: Preventive and Personalized Care. In *Federated Learning Systems: Towards Privacy-Preserving Distributed AI*; Springer, 2025; pp. 107–127.
49. Zhu, W.; Luo, J.; White, A.D. Federated learning of molecular properties with graph neural networks in a heterogeneous setting. *Patterns* **2022**, *3*.
50. Guo, Y.; Gao, Y.; Song, J. Molcfl: A personalized and privacy-preserving drug discovery framework based on generative clustered federated learning. *Journal of biomedical informatics* **2024**, *157*, 104712.
51. Smajić, A.; Grandits, M.; Ecker, G.F. Privacy-preserving techniques for decentralized and secure machine learning in drug discovery. *Drug Discovery Today* **2023**, *28*, 103820.
52. Chen, S.; Xue, D.; Chuai, G.; Yang, Q.; Liu, Q. FL-QSAR: a federated learning-based QSAR prototype for collaborative drug discovery. *Bioinformatics* **2020**, *36*, 5492–5498.
53. Li, C. Breaking data silos in drug discovery with federated learning. *Nature Chemical Engineering* **2025**, *2*, 288–289.
54. Huang, D.; Ye, X.; Sakurai, T. Multi-party collaborative drug discovery via federated learning. *Computers in Biology and Medicine* **2024**, *171*, 108181.
55. Hanser, T.; Ahlberg, E.; Amberg, A.; Anger, L.T.; Barber, C.; Brennan, R.J.; Brigo, A.; Delaunois, A.; Glowienke, S.; Greene, N.; et al. Data-driven federated learning in drug discovery with knowledge distillation. *Nature Machine Intelligence* **2025**, pp. 1–14.

 doi:10.20944/preprints202601.0271.v1

56. Oldenhof, M.; et al. Industry-Scale Orchestrated Federated Learning for Drug Discovery. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2023, Vol. 37, pp. 497–505.
57. Dritsas, E.; Trigka, M.; Kavallieros, D. Federated Learning for IoT: A Survey of Techniques, Challenges, and Applications. *Sensors* **2025**, *25*.
58. War, M.R.; Singh, Y.; Sheikh, Z.A.; Singh, P.K. Review on the Use of Federated Learning Models for the Security of Cyber-Physical Systems. *Scalable Computing: Practice and Experience* **2025**, *26*, 16–33.
59. Lu, Y.; Huang, X.; Dai, Y.; Maharjan, S.; Zhang, Y. Federated Learning for Data Privacy Preservation in Vehicular Cyber-Physical Systems. *IEEE Network* **2020**, *34*, 50–56. https://doi.org/10.1109/MNET.011.1900317.
60. Long, G.; Tan, Y.; Jiang, J.; Zhang, C. Federated Learning for Open Banking, 2021, [arXiv:cs.DC/2108.10749].
61. Awosika, T.; Shukla, R.M.; Pranggono, B. Transparency and Privacy: The Role of Explainable AI and Federated Learning in Financial Fraud Detection. *IEEE Access* **2024**, *12*, 64551–64560. https://doi.org/10.1109/ACCESS.2024.3394528.
62. Wang, Z.; Xiao, J.; Wang, L.; Yao, J. A novel federated learning approach with knowledge transfer for credit scoring. *Decision Support Systems* **2024**, *177*, 114084. https://doi.org/https://doi.org/10.1016/j.dss.2023.114084.
63. Zhao, L.; Cai, L.; Lu, W.S. Federated Learning for Data Trading Portfolio Allocation With Autonomous Economic Agents. *IEEE Transactions on Neural Networks and Learning Systems* **2025**, *36*, 1467–1481. https://doi.org/10.1109/TNNLS.2023.3332315.
64. Abadi, A.; Doyle, B.; Gini, F.; Guinamard, K.; Murakonda, S.K.; Liddell, J.; Mellor, P.; Murdoch, S.J.; Naseri, M.; Page, H.; et al. Starlit: Privacy-Preserving Federated Learning to Enhance Financial Fraud Detection, 2024, [arXiv:cs.LG/2401.10765].
65. McMahan, H.B.; Moore, E.; Ramage, D.; Hampson, S.; y Arcas, B.A. Communication-Efficient Learning of Deep Networks from Decentralized Data, 2023, [arXiv:cs.LG/1602.05629].
66. Chalamala, S.R.; et al. Federated Learning to Comply with Data Protection Regulations. *CSI Transactions on ICT* **2022**, *10*, 47–60.
67. Blika, A.; Palmos, S.; Doukas, G.; Lamprou, V.; Pelekis, S.; Kontoulis, M.; Ntanos, C.; Askounis, D. Federated Learning for Enhanced Cybersecurity and Trustworthiness in 5G and 6G Networks: A Comprehensive Survey. *IEEE Open Journal of the Communications Society* **2025**, *6*, 3094–3130. https://doi.org/10.1109/OJCOMS.2024.3449563.
68. Yang, Z.; Chen, M.; Wong, K.K.; Poor, H.V.; Cui, S. Federated Learning for 6G: Applications, Challenges, and Opportunities. *Engineering* **2022**, *8*, 33–41. https://doi.org/https://doi.org/10.1016/j.eng.2021.12.002.
69. Lee, J.; Solat, F.; Kim, T.Y.; Poor, H.V. Federated Learning-Empowered Mobile Network Management for 5G and Beyond Networks: From Access to Core. *IEEE Communications Surveys & Tutorials* **2024**, *26*, 2176–2212. https://doi.org/10.1109/COMST.2024.3352910.
70. Challenges and Future Directions in Federated Learning. *arXiv preprint* **2022**.
71. A Comprehensive Survey of Challenges in Federated Learning. *Frontiers in AI* **2023**.
72. Survey of Federated Learning: Taxonomies and Challenges. *IEEE Communications Surveys & Tutorials* **2022**.
73. Finn, C.; Abbeel, P.; Levine, S. Model-Agnostic Meta-Learning for Fast Adaptation. In Proceedings of the International Conference on Machine Learning, 2017.
74. Contrastive Federated Learning: Mitigating Non-IID with Representation Learning. In Proceedings of the CVPR, 2021.
75. Konečný, J.; McMahan, H.B.; Ramage, D.; Richtárik, P. Federated Optimization: Distributed Machine Learning for On-Device Intelligence, 2016, [arXiv:cs.LG/1610.02527].
76. Konečný, J.; McMahan, H.B.; Yu, F.X.; Richtárik, P.; Suresh, A.T.; Bacon, D. Federated Learning: Strategies for Improving Communication Efficiency, 2017, [arXiv:cs.LG/1610.05492].
77. Li, T.; Sahu, A.K.; Talwalkar, A.; Smith, V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine* **2020**.
78. Konečný, J.; et al. Federated Optimization: Distributed Machine Learning for On-Device Intelligence. In Proceedings of the arXiv preprint arXiv:1610.02527, 2016.
79. Sheller, M.J.e.a. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific Reports* **2020**.
80. Shi, W.; Dustdar, S. Edge Computing: Vision and Challenges. *IEEE Internet of Things Journal* **2016**, *3*, 637–646.
81. Satyanarayanan, M. The Emergence of Edge Computing. *Computer* **2017**, *50*, 30–39.

 doi:10.20944/preprints202601.0271.v1

82. Li, T.; Sahu, A.K.; Talwalkar, A.; Smith, V. Federated Learning: Challenges, Methods, and Future Directions. *IEEE Signal Processing Magazine* **2020**, *37*, 50–60.
83. Lim, W.Y.B.; et al. Federated Learning in Mobile Edge Networks: A Comprehensive Survey. *IEEE Communications Surveys & Tutorials* **2020**, *22*, 2031–2063.
84. Yang, Q.; Liu, Y.; Chen, T.; Tong, Y. Federated Machine Learning: Concept and Applications. *ACM Trans. Intell. Syst. Technol.* **2019**, *10*. https://doi.org/10.1145/3298981.
85. Li, T.; Sahu, A.K.; Talwalkar, A.; Smith, V. Federated Learning: Challenges, Methods, and Future Directions. *IEEE Signal Processing Magazine* **2020**, *37*, 50–60. https://doi.org/10.1109/MSP.2020.2975749.
86. Kairouz, P.; McMahan, H.B.; et al. Advances and Open Problems in Federated Learning. *Foundations and Trends in Machine Learning* **2021**.
87. Li, Q.; Wen, Z.; Wu, Z.; Hu, S.; Wang, N.; Li, Y.; Liu, X.; He, B. A Survey on Federated Learning Systems: Vision, Hype and Reality for Data Privacy and Protection. *IEEE Transactions on Knowledge and Data Engineering* **2023**, *35*, 3347–3366. https://doi.org/10.1109/TKDE.2021.3124599.
88. Liu, J.; Huang, J.; Zhou, Y.; Liu, X.; Liu, S.; Gu, Q. From distributed machine learning to federated learning: a survey. *Knowledge and Information Systems* **2022**, *64*, 885–917. https://doi.org/10.1007/s10115-022-01664-x.
89. Liu, B.; Lv, N.; Guo, Y.; Li, Y. Recent advances on federated learning: A systematic survey. *Neurocomputing* **2024**, *597*, 128019. https://doi.org/https://doi.org/10.1016/j.neucom.2024.128019.
90. Wen, J.; Zhang, Z.; Lan, Y.; Cui, Z.; Cai, J.; Zhang, W. A survey on federated learning: challenges and applications. *International Journal of Machine Learning and Cybernetics* **2023**, *14*, 513–535. https://doi.org/10.1007/s13042-022-01647-y.
91. Yurdem, B.; Kuzlu, M.; Gullu, M.K.; Catak, F.O.; Tabassum, M. Federated learning: Overview, strategies, applications, tools and future directions. *Heliyon* **2024**, *10*, e38137. https://doi.org/10.1016/j.heliyon.2024.e38137.
92. Chaudhary, R.K.; Kumar, R.; Saxena, N. A systematic review on federated learning system: a new paradigm to machine learning. *Knowledge and Information Systems* **2025**, *67*, 1811–1914. https://doi.org/10.1007/s10115-024-02257-6.
93. Nasim, M.A.A.; Soshi, F.T.J.; Biswas, P.; Ferdous, A.S.M.A.; Rashid, A.; Biswas, A.; Gupta, K.D. Principles and Components of Federated Learning Architectures, 2025, [arXiv:cs.LG/2502.05273].
94. Aledhari, M.; Razzak, R.; Parizi, R.M.; Saeed, F. Federated Learning: A Survey on Enabling Technologies, Protocols, and Applications. *IEEE Access* **2020**, *8*, 140699–140725. https://doi.org/10.1109/access.2020.3013541.
95. Lo, S.K.; Lu, Q.; Wang, C.; Paik, H.Y.; Zhu, L. A Systematic Literature Review on Federated Machine Learning: From a Software Engineering Perspective. *ACM Comput. Surv.* **2021**, *54*. https://doi.org/10.1145/3450288.
96. Gupta, R.; Alam, T. Survey on Federated-Learning Approaches in Distributed Environment. *Wireless Personal Communications* **2022**, *125*, 1631–1652. https://doi.org/10.1007/s11277-022-09624-y.
97. Pouriyeh, S.; Shahid, O.; Parizi, R.M.; Sheng, Q.Z.; Srivastava, G.; Zhao, L.; Nasajpour, M. Secure Smart Communication Efficiency in Federated Learning: Achievements and Challenges. *Applied Sciences* **2022**, *12*. https://doi.org/10.3390/app12188980.
98. Mahlool, D.H.; Alsalihi, M.H. A Comprehensive Survey on Federated Learning: Concept and Applications. In Proceedings of the Mobile Computing and Sustainable Informatics; Shakya, S.; Ntalianis, K.; Kamel, K.A., Eds., Singapore, 2022; pp. 539–553.
99. Nguyen, D.C.; Ding, M.; Pathirana, P.N.; Seneviratne, A.; Li, J.; Vincent Poor, H. Federated Learning for Internet of Things: A Comprehensive Survey. *IEEE Communications Surveys & Tutorials* **2021**, *23*, 1622–1658. https://doi.org/10.1109/COMST.2021.3075439.
100. Lim, W.Y.B.; Luong, N.C.; Hoang, D.T.; Jiao, Y.; Liang, Y.C.; Yang, Q.; Niyato, D.; Miao, C. Federated Learning in Mobile Edge Networks: A Comprehensive Survey. *IEEE Communications Surveys & Tutorials* **2020**, *22*, 2031–2063. https://doi.org/10.1109/COMST.2020.2986024.
101. Wahab, O.A.; Mourad, A.; Otrok, H.; Taleb, T. Federated Machine Learning: Survey, Multi-Level Classification, Desirable Criteria and Future Directions in Communication and Networking Systems. *IEEE Communications Surveys & Tutorials* **2021**, *23*, 1342–1397. https://doi.org/10.1109/COMST.2021.3058573.
102. Mammen, P.M. Federated Learning: Opportunities and Challenges, 2021, [arXiv:cs.LG/2101.05428].
103. Abreha, H.G.; Hayajneh, M.; Serhani, M.A. Federated Learning in Edge Computing: A Systematic Survey. *Sensors* **2022**, *22*. https://doi.org/10.3390/s22020450.

 doi:10.20944/preprints202601.0271.v1

104. Sirohi, D.; Kumar, N.; Rana, P.S.; Verma, A.; Singh, M.; Kumar, G. Federated learning for 6G-enabled secure communication systems: a comprehensive survey. *Artificial Intelligence Review* **2023**, *56*, 11297–11389. https://doi.org/10.1007/s10462-023-10417-3.
105. Albshaier, L.; Almarri, S.; Albuali, A. Federated Learning for Cloud and Edge Security: A Systematic Review of Challenges and AI Opportunities. *Electronics* **2025**, *14*. https://doi.org/10.3390/electronics14051019.
106. Jia, N.; Qu, Z.; Ye, B.; Wang, Y.; Hu, S.; Guo, S. A Comprehensive Survey on Communication-Efficient Federated Learning in Mobile Edge Environments. *IEEE Communications Surveys & Tutorials* **2025**, pp. 1–1. https://doi.org/10.1109/COMST.2025.3535957.
107. Li, L.; Fan, Y.; Tse, M.; Lin, K.Y. A review of applications in federated learning. *Computers & Industrial Engineering* **2020**, *149*, 106854. https://doi.org/https://doi.org/10.1016/j.cie.2020.106854.
108. Bharati, S.; Mondal, M.R.H.; Podder, P.; Prasath, V.S. Federated learning: Applications, challenges and future directions. *International Journal of Hybrid Intelligent Systems* **2022**, *18*, 19–35. https://doi.org/10.3233/his-220006.
109. Moshawrab, M.; Adda, M.; Bouzouane, A.; Ibrahim, H.; Raad, A. Reviewing Federated Machine Learning and Its Use in Diseases Prediction. *Sensors* **2023**, *23*. https://doi.org/10.3390/s23042112.
110. Moshawrab, M.; Adda, M.; Bouzouane, A.; Ibrahim, H.; Raad, A. Reviewing Federated Learning Aggregation Algorithms; Strategies, Contributions, Limitations and Future Perspectives. *Electronics* **2023**, *12*. https://doi.org/10.3390/electronics12102287.
111. Martínez Beltrán, E.T.; Pérez, M.Q.; Sánchez, P.M.S.; Bernal, S.L.; Bovet, G.; Pérez, M.G.; Pérez, G.M.; Celdrán, A.H. Decentralized Federated Learning: Fundamentals, State of the Art, Frameworks, Trends, and Challenges. *IEEE Communications Surveys & Tutorials* **2023**, *25*, 2983–3013. https://doi.org/10.1109/COMST.2023.3315746.
112. Almanifi, O.R.A.; Chow, C.O.; Tham, M.L.; Chuah, J.H.; Kanesan, J. Communication and computation efficiency in Federated Learning: A survey. *Internet of Things* **2023**, *22*, 100742. https://doi.org/https://doi.org/10.1016/j.iot.2023.100742.
113. Blanco-Justicia, A.; Domingo-Ferrer, J.; Martínez, S.; Sánchez, D.; Flanagan, A.; Tan, K.E. Achieving security and privacy in federated learning systems: Survey, research challenges and future directions. *Engineering Applications of Artificial Intelligence* **2021**, *106*, 104468. https://doi.org/https://doi.org/10.1016/j.engappai.2021.104468.
114. Zhang, J.; Zhu, H.; Wang, F.; Zhao, J.; Xu, Q.; Li, H. Security and Privacy Threats to Federated Learning: Issues, Methods, and Challenges. *Security and Communication Networks* **2022**, *2022*, 2886795, [https://onlinelibrary.wiley.com/doi/pdf/10.1155/2022/2886795]. https://doi.org/https://doi.org/10.1155/2022/2886795.
115. Saha, S.; Hota, A.; Chattopadhyay, A.K.; Banerjee, S.; De, D. A multifaceted survey on privacy preservation of federated learning: progress, challenges, and opportunities. *Artificial Intelligence Review* **2024**, *57*, 184. https://doi.org/10.1007/s10462-024-10766-7.
116. Hu, K.; Gong, S.; Zhang, Q.; Cheng, S.; Li, Q.; Xu, Y. An overview of implementing security and privacy in federated learning. *Artificial Intelligence Review* **2024**, *57*, 204. https://doi.org/10.1007/s10462-024-10846-8.
117. Gupta, R.; Gupta, J. Federated learning using game strategies: State-of-the-art and future trends. *Computer Networks* **2023**, *225*, 109650. https://doi.org/https://doi.org/10.1016/j.comnet.2023.109650.
118. Neto, H.N.C.; Hribar, J.; Dusparic, I.; Mattos, D.M.F.; Fernandes, N.C. A Survey on Securing Federated Learning: Analysis of Applications, Attacks, Challenges, and Trends. *IEEE Access* **2023**, *11*, 41928–41953. https://doi.org/10.1109/ACCESS.2023.3269980.
119. Qammar, A.; Karim, A.; Ning, H.; Li, D.; Sajid, A.; Liu, X. Securing federated learning with blockchain: a systematic literature review. *Artificial Intelligence Review* **2023**, *56*, 3951–3985. https://doi.org/10.1007/s10462-022-10271-9.
120. Zhu, J.; Cao, J.; Saxena, D.; Jiang, S.; Ferradi, H. Blockchain-empowered Federated Learning: Challenges, Solutions, and Future Directions. *ACM Comput. Surv.* **2023**, *55*. https://doi.org/10.1145/3570953.
121. Rahman, R. Federated Learning: A Survey on Privacy-Preserving Collaborative Intelligence, 2025, [arXiv:cs.LG/2504.17703].
122. Tariq, A.; Serhani, M.A.; Sallabi, F.M.; Barka, E.S.; Qayyum, T.; Khater, H.M.; Shuaib, K.A. Trustworthy Federated Learning: A Comprehensive Review, Architecture, Key Challenges, and Future Research Prospects. *IEEE Open Journal of the Communications Society* **2024**, *5*, 4920–4998. https://doi.org/10.1109/OJCOMS.2024.3438264.

 doi:10.20944/preprints202601.0271.v1

123. Ye, M.; Fang, X.; Du, B.; Yuen, P.C.; Tao, D. Heterogeneous Federated Learning: State-of-the-art and Research Challenges. *ACM Comput. Surv.* **2023**, *56*. https://doi.org/10.1145/3625558.
124. Wu, Q.; He, K.; Chen, X. Personalized Federated Learning for Intelligent IoT Applications: A Cloud-Edge Based Framework. *IEEE Open Journal of the Computer Society* **2020**, *1*, 35–44. https://doi.org/10.1109/OJCS.2020.2993259.
125. Tan, A.Z.; Yu, H.; Cui, L.; Yang, Q. Towards Personalized Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems* **2023**, *34*, 9587–9603. https://doi.org/10.1109/TNNLS.2022.3160699.
126. Zhu, H.; Zhang, H.; Jin, Y. From federated learning to federated neural architecture search: a survey. *Complex & Intelligent Systems* **2021**, *7*, 639–657. https://doi.org/10.1007/s40747-020-00247-z.
127. Che, L.; Wang, J.; Zhou, Y.; Ma, F. Multimodal Federated Learning: A Survey. *Sensors* **2023**, *23*. https://doi.org/10.3390/s23156986.
128. Niknam, S.; Dhillon, H.S.; Reed, J.H. Federated Learning for Wireless Communications: Motivation, Opportunities and Challenges, 2020, [arXiv:eess.SP/1908.06847].
129. Kulkarni, V.; Kulkarni, M.; Pant, A. Survey of Personalization Techniques for Federated Learning. In Proceedings of the 2020 Fourth World Conference on Smart Trends in Systems, Security and Sustainability (WorldS4), 2020, pp. 794–797. https://doi.org/10.1109/WorldS450073.2020.9210355.
130. Yin, X.; Zhu, Y.; Hu, J. A Comprehensive Survey of Privacy-preserving Federated Learning: A Taxonomy, Review, and Future Directions. *ACM Comput. Surv.* **2021**, *54*. https://doi.org/10.1145/3460427.
131. Khan, L.U.; Saad, W.; Han, Z.; Hossain, E.; Hong, C.S. Federated Learning for Internet of Things: Recent Advances, Taxonomy, and Open Challenges. *IEEE Communications Surveys & Tutorials* **2021**, *23*, 1759–1799. https://doi.org/10.1109/COMST.2021.3090430.
132. Gao, D.; Yao, X.; Yang, Q. A Survey on Heterogeneous Federated Learning, 2022, [arXiv:cs.LG/2210.04505].
133. Asad, M.; Shaukat, S.; Hu, D.; Wang, Z.; Javanmardi, E.; Nakazato, J.; Tsukada, M. Limitations and Future Aspects of Communication Costs in Federated Learning: A Survey. *Sensors* **2023**, *23*. https://doi.org/10.3390/s23177358.
134. Alotaibi, B.; Khan, F.A.; Mahmood, S. Communication Efficiency and Non-Independent and Identically Distributed Data Challenge in Federated Learning: A Systematic Mapping Study. *Applied Sciences* **2024**, *14*. https://doi.org/10.3390/app14072720.
135. Xie, Q.; Jiang, S.; Jiang, L.; Huang, Y.; Zhao, Z.; Khan, S.; Dai, W.; Liu, Z.; Wu, K. Efficiency Optimization Techniques in Privacy-Preserving Federated Learning With Homomorphic Encryption: A Brief Survey. *IEEE Internet of Things Journal* **2024**, *11*, 24569–24580. https://doi.org/10.1109/JIOT.2024.3382875.
136. Kaur, H.; Rani, V.; Kumar, M.; Singh, A.; Gupta, S. Federated learning: a comprehensive review of recent advances and applications. *Multimedia Tools and Applications* **2024**, *83*, 54165–54188. https://doi.org/10.1007/s11042-023-17737-0.
137. Karimireddy, S.P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; Suresh, A.T. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In Proceedings of the Proceedings of the 37th International Conference on Machine Learning; III, H.D.; Singh, A., Eds. PMLR, 13–18 Jul 2020, Vol. 119, *Proceedings of Machine Learning Research*, pp. 5132–5143.
138. Qi, Z.; Meng, L.; Chen, Z.; Hu, H.; Lin, H.; Meng, X. Cross-Silo Prototypical Calibration for Federated Learning with Non-IID Data. In Proceedings of the Proceedings of the 31st ACM International Conference on Multimedia, New York, NY, USA, 2023; MM '23, p. 3099–3107. https://doi.org/10.1145/3581783.3612481.
139. Qi, Z.; Zhou, S.; Meng, L.; Hu, H.; Yu, H.; Meng, X. Federated Deconfounding and Debiasing Learning for Out-of-Distribution Generalization, 2025, [arXiv:cs.CV/2505.04979].
140. Ma, Y.; Dai, W.; Huang, W.; Chen, J. Geometric Knowledge-Guided Localized Global Distribution Alignment for Federated Learning. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2025, pp. 20958–20968.
141. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; Arcas, B.A.y. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Proceedings of the Proceedings of the 20th International Conference on Artificial Intelligence and Statistics; Singh, A.; Zhu, J., Eds. PMLR, 20–22 Apr 2017, Vol. 54, *Proceedings of Machine Learning Research*, pp. 1273–1282.
142. Li, Q.; Diao, Y.; Chen, Q.; He, B. Federated Learning on Non-IID Data Silos: An Experimental Study. In Proceedings of the 2022 IEEE 38th International Conference on Data Engineering (ICDE), 2022, pp. 965–978. https://doi.org/10.1109/ICDE53745.2022.00077.

 doi:10.20944/preprints202601.0271.v1

143. Li, T.; Sahu, A.K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; Smith, V. Federated Optimization in Heterogeneous Networks. In Proceedings of the Proceedings of Machine Learning and Systems; Dhillon, I.; Papailiopoulos, D.; Sze, V., Eds., 2020, Vol. 2, pp. 429–450.
144. Yurochkin, M.; Agarwal, M.; Ghosh, S.; Greenewald, K.; Hoang, N.; Khazaeni, Y. Bayesian Nonparametric Federated Learning of Neural Networks. In Proceedings of the Proceedings of the 36th International Conference on Machine Learning; Chaudhuri, K.; Salakhutdinov, R., Eds. PMLR, 09–15 Jun 2019, Vol. 97, *Proceedings of Machine Learning Research*, pp. 7252–7261.
145. Zhao, Y.; Li, M.; Lai, L.; Suda, N.; Civin, D.; Chandra, V. Federated Learning with Non-IID Data **2018**. https://doi.org/10.48550/ARXIV.1806.00582.
146. Luo, M.; Chen, F.; Hu, D.; Zhang, Y.; Liang, J.; Feng, J. No Fear of Heterogeneity: Classifier Calibration for Federated Learning with Non-IID Data. In Proceedings of the Advances in Neural Information Processing Systems; Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; Vaughan, J.W., Eds. Curran Associates, Inc., 2021, Vol. 34, pp. 5972–5984.
147. Zheng, H.; Hu, Z.; Yang, L.; Zheng, M.; Xu, A.; Wang, B. FedCALM: Conflict-aware Layer-wise Mitigation for Selective Aggregation in Deeper Personalized Federated Learning. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2025, pp. 15444–15453.
148. Dennis, D.K.; Li, T.; Smith, V. Heterogeneity for the Win: One-Shot Federated Clustering. In Proceedings of the Proceedings of the 38th International Conference on Machine Learning; Meila, M.; Zhang, T., Eds. PMLR, 18–24 Jul 2021, Vol. 139, *Proceedings of Machine Learning Research*, pp. 2611–2620.
149. Stallmann, M.; Wilbik, A. Towards Federated Clustering: A Federated Fuzzy c-Means Algorithm (FFCM), 2022, [arXiv:cs.LG/2201.07316].
150. Pan, C.; Sima, J.; Prakash, S.; Rana, V.; Milenkovic, O. Machine Unlearning of Federated Clusters, 2023, [arXiv:cs.LG/2210.16424].
151. Xu, J.; Chen, H.Y.; Chao, W.L.; Zhang, Y. Jigsaw Game: Federated Clustering, 2024, [arXiv:cs.LG/2407.12764].
152. Qiao, D.; Ding, C.; Fan, J. Federated Spectral Clustering via Secure Similarity Reconstruction. In Proceedings of the Advances in Neural Information Processing Systems; Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; Levine, S., Eds. Curran Associates, Inc., 2023, Vol. 36, pp. 58520–58555.
153. Wang, S.; Chang, T.H. Federated Matrix Factorization: Algorithm Design and Application to Data Clustering. *IEEE Transactions on Signal Processing* **2022**, *70*, 1625–1640. https://doi.org/10.1109/TSP.2022.3151505.
154. Yan, J.; Liu, J.; Ning, Y.Z.; Zhang, Z.Y. SDA-FC: Bridging federated clustering and deep generative model. *Information Sciences* **2024**, *681*, 121203. https://doi.org/https://doi.org/10.1016/j.ins.2024.121203.
155. Huang, S.; Fu, L.; Ye, F.; Liao, T.; Deng, B.; Zhang, C.; Chen, C. Soft-consensual Federated Learning for Data Heterogeneity via Multiple Paths.
156. Fang, X.; Ye, M.; Du, B. Robust Asymmetric Heterogeneous Federated Learning With Corrupted Clients. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **2025**, *47*, 2693–2705. https://doi.org/10.1109/TPAMI.2025.3527137.
157. Fang, X.; Ye, M. Noise-Robust Federated Learning With Model Heterogeneous Clients. *IEEE Transactions on Mobile Computing* **2025**, *24*, 4053–4071. https://doi.org/10.1109/TMC.2024.3522573.
158. Huang, S.; Fu, L.; Li, Y.; Chen, C.; Zheng, Z.; Dai, H.N. A Cross-Client Coordinator in Federated Learning Framework for Conquering Heterogeneity. *IEEE Transactions on Neural Networks and Learning Systems* **2025**, *36*, 8828–8842. https://doi.org/10.1109/TNNLS.2024.3439878.
159. Qi, Z.; Meng, L.; Li, Z.; Hu, H.; Meng, X. Cross-Silo Feature Space Alignment for Federated Learning on Clients with Imbalanced Data. *Proceedings of the AAAI Conference on Artificial Intelligence* **2025**, *39*, 19986–19994. https://doi.org/10.1609/aaai.v39i19.34201.
160. Xie, H.; Ma, J.; Xiong, L.; Yang, C. Federated Graph Classification over Non-IID Graphs. In Proceedings of the Advances in Neural Information Processing Systems; Ranzato, M.; Beygelzimer, A.; Dauphin, Y.; Liang, P.; Vaughan, J.W., Eds. Curran Associates, Inc., 2021, Vol. 34, pp. 18839–18852.
161. Baek, J.; Jeong, W.; Jin, J.; Yoon, J.; Hwang, S.J. Personalized Subgraph Federated Learning. In Proceedings of the Proceedings of the 40th International Conference on Machine Learning; Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; Scarlett, J., Eds. PMLR, 23–29 Jul 2023, Vol. 202, *Proceedings of Machine Learning Research*, pp. 1396–1415.
162. Yang, Z.; Zhang, Y.; Zheng, Y.; Tian, X.; Peng, H.; Liu, T.; Han, B. Fedfed: Feature distillation against data heterogeneity in federated learning. *Advances in neural information processing systems* **2023**, *36*, 60397–60428.
163. Yan, Y.; Fu, H.; Li, Y.; Xie, J.; Ma, J.; Yang, G.; Zhu, L. A Simple Data Augmentation for Feature Distribution Skewed Federated Learning, 2024, [arXiv:cs.LG/2306.09363].

 doi:10.20944/preprints202601.0271.v1

164. Wang, Z.; Wang, Z.; Wang, Z.; Fan, X.; Wang, C. Federated Learning with Domain Shift Eraser, 2025, [arXiv:cs.CV/2503.13063].
165. Zhang, X.; Li, S.; Li, A.; Liu, Y.; Zhang, F.; Zhu, C.; Zhang, L. Subspace Constraint and Contribution Estimation for Heterogeneous Federated Learning. In Proceedings of the 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2025, pp. 20632–20642. https://doi.org/10.1109/CVPR52734.2025.01921.
166. Raswa, F.H.; Lu, C.S.; Wang, J.C. HistoFS: Non-IID Histopathologic Whole Slide Image Classification via Federated Style Transfer with RoI-Preserving. In Proceedings of the Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR), June 2025, pp. 30251–30260.
167. Yin, H.; et al. FedMP: Federated Learning with Model Pruning for Efficient Communication and Computation. *IEEE Transactions on Neural Networks and Learning Systems* **2022**.
168. Li, Z.; Wen, Y.; et al. LotteryFL: Personalized and Communication-Efficient Federated Learning with Lottery Ticket Hypothesis. In Proceedings of the ICLR, 2020.
169. Xu, J.; colleagues. Quantized Federated Learning: Optimizing Training Efficiency for Edge Devices. *IEEE Internet of Things Journal* **2022**.
170. Seo, H.; et al. Federated Knowledge Distillation for Resource-Constrained Edge Devices. In Proceedings of the NeurIPS Workshop on Federated Learning, 2021.
171. Deng, Y.; Lyu, L.; Chen, J. ScaleFL: Resource-Adaptive Federated Learning with Scalable Neural Networks. *IEEE Transactions on Mobile Computing* **2022**.
172. Diao, E.; Ding, J.; Tarokh, V. HeteroFL: Computation and Communication Efficient Federated Learning for Heterogeneous Clients. In Proceedings of the ICML, 2021.
173. Gao, R.; et al. Sub-FL: Training Subnetworks for Efficient Federated Learning on Heterogeneous Devices. *Pattern Recognition* **2022**.
174. Gupta, S.; Raskar, R. Split Learning for Distributed Deep Learning. *arXiv:1905.08821* **2019**.
175. Thapa, C.; Arachchige, P.; et al. SplitFed: When Federated Learning Meets Split Learning. In Proceedings of the ACL Workshop on Federated Learning, 2020.
176. Chen, W.; et al. Hybrid Split-Pruned Federated Learning for Resource-Constrained Edge Devices. *IEEE Transactions on Mobile Computing* **2023**.
177. Lai, F.; et al. Oort: Efficient Federated Learning via Intelligent Client Selection. In Proceedings of the OSDI, 2021.
178. Zhou, F.; et al. Adaptive Federated Optimization for Heterogeneous Devices. *Neurocomputing* **2022**.
179. Xie, C.; et al. Asynchronous Federated Optimization. In Proceedings of the AISTATS, 2020.
180. Wang, S.; et al. Staleness-Aware Asynchronous Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems* **2021**.
181. Nguyen, T.D.; Pham, M.; Tran, T. Federated Learning with Buffered Asynchronous Aggregation. *Proceedings of the MLSys Conference* **2021**.
182. He, C.; et al. Resource-Aware Federated Learning: Optimizing Performance under System Constraints. *IEEE Transactions on Mobile Computing* **2021**.
183. Aji, A.F.; Heafield, K. Sparse communication for distributed gradient descent. In Proceedings of the EMNLP, 2017.
184. Lin, Y.; Han, S. Deep gradient compression: Reducing the communication bandwidth for distributed training. In Proceedings of the ICLR, 2018.
185. Stich, S. Sparsified SGD with Memory. *NeurIPS* **2018**.
186. Bernstein, J.; et al. signSGD: Compressed Optimization for Non-Convex Problems. In Proceedings of the ICML, 2018.
187. Alistarh, D.; et al. QSGD: Communication-Efficient SGD via Gradient Quantization and Encoding. In Proceedings of the NeurIPS, 2017.
188. Wen, W.; et al. TernGrad: Ternary Gradients to Reduce Communication in Distributed Deep Learning. In Proceedings of the NeurIPS, 2017.
189. Horvath, S.; et al. Natural Compression for Distributed Deep Learning. *NeurIPS* **2019**.
190. McMahan, B.; et al. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Proceedings of the AISTATS, 2017.
191. Wang, S.; et al. Adaptive Federated Learning in Resource-Constrained Edge Computing Systems. In Proceedings of the IEEE INFOCOM, 2019.

 doi:10.20944/preprints202601.0271.v1

192. Li, T.; et al. Federated Optimization in Heterogeneous Networks. In Proceedings of the Proceedings of MLSys, 2020.
193. Xie, C.; et al. Asynchronous Federated Optimization. In Proceedings of the AISTATS, 2020.
194. Chen, J.; et al. Asynchronous Distributed Learning via Stale Gradient Methods. *IEEE Journal on Selected Areas in Communications* **2019**.
195. Lalitha, A.; et al. Peer-to-Peer Federated Learning on Graphs. *arXiv:1901.11173* **2019**.
196. Shi, S.; et al. Communication-Efficient Edge-Assisted Federated Learning. In Proceedings of the IEEE GLOBECOM, 2020.
197. Bonawitz, K.; et al. Towards Federated Learning at Scale: System Design. In Proceedings of the SysML, 2019.
198. Karimireddy, S.P.; et al. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In Proceedings of the ICML, 2020.
199. Basu, D.; et al. Qsparse-local-SGD: Distributed SGD with Quantization, Sparsification, and Local Computation. In Proceedings of the NeurIPS, 2019.
200. Sattler, F.; et al. Clustered Federated Learning: Model-Agnostic Distributed Multi-Task Optimization. In Proceedings of the NeurIPS, 2020.
201. Ji, Y.; et al. Towards Statistical-Quality-Aware Client Selection for Federated Learning. *IEEE Transactions on Mobile Computing* **2022**.
202. Tang, X.; Yu, Y. Auction-Based Federated Learning via Truthful Mechanism Design. *IEEE Transactions on Mobile Computing* **2025**.
203. Tang, X.; Yu, H. A Cost-Aware Utility-Maximizing Bidding Strategy for Auction-Based Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems* **2025**, *36*, 12866–12879. https://doi.org/10.1109/TNNLS.2024.3474102.
204. Cho, Y.; et al. Reinforcement Learning-Based Client Selection for Efficient Federated Learning. In Proceedings of the AAAI, 2021.
205. Li, Q.; et al. Fair Resource Allocation in Federated Learning. In Proceedings of the ICML Workshop on Federated Learning, 2020.
206. Wang, J.; Liu, Q.; Liang, H.; Joshi, G.; Poor, H.V. Tackling the Objective Inconsistency Problem in Heterogeneous Federated Optimization. In Proceedings of the Advances in Neural Information Processing Systems; Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; Lin, H., Eds. Curran Associates, Inc., 2020, Vol. 33, pp. 7611–7623.
207. Lin, T.; Kong, L.; Stich, S.U.; Jaggi, M. Ensemble Distillation for Robust Model Fusion in Federated Learning. In Proceedings of the Advances in Neural Information Processing Systems; Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; Lin, H., Eds. Curran Associates, Inc., 2020, Vol. 33, pp. 2351–2363.
208. Zhu, Z.; Hong, J.; Zhou, J. Data-Free Knowledge Distillation for Heterogeneous Federated Learning. In Proceedings of the Proceedings of the 38th International Conference on Machine Learning; Meila, M.; Zhang, T., Eds. PMLR, 18–24 Jul 2021, Vol. 139, *Proceedings of Machine Learning Research*, pp. 12878–12889.
209. Wang, H.; Li, Y.; Xu, W.; Li, R.; Zhan, Y.; Zeng, Z. DaFKD: Domain-Aware Federated Knowledge Distillation. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2023, pp. 20412–20421.
210. Zhao, S.; Liao, T.; Fu, L.; Chen, C.; Bian, J.; Zheng, Z. Data-free knowledge distillation via generator-free data generation for Non-IID federated learning. *Neural Networks* **2024**, *179*, 106627. https://doi.org/https://doi.org/10.1016/j.neunet.2024.106627.
211. Hu, M.; Cao, Y.; Li, A.; Li, Z.; Liu, C.; Li, T.; Chen, M.; Liu, Y. FedMut: Generalized Federated Learning via Stochastic Mutation. *Proceedings of the AAAI Conference on Artificial Intelligence* **2024**, *38*, 12528–12537. https://doi.org/10.1609/aaai.v38i11.29146.
212. Weng, J.; Xia, Z.; Li, R.; Hu, M.; Chen, M. FedQP: Towards Accurate Federated Learning using Quadratic Programming Guided Mutation, 2024, [arXiv:cs.LG/2411.15847].
213. Fraboni, Y.; Vidal, R.; Kameni, L.; Lorenzi, M. Clustered Sampling: Low-Variance and Improved Representativity for Clients Selection in Federated Learning. In Proceedings of the Proceedings of the 38th International Conference on Machine Learning; Meila, M.; Zhang, T., Eds. PMLR, 18–24 Jul 2021, Vol. 139, *Proceedings of Machine Learning Research*, pp. 3407–3416.
214. Li, A.; Wang, G.; Hu, M.; Sun, J.; Zhang, L.; Tuan, L.A.; Yu, H. Joint Client-and-Sample Selection for Federated Learning via Bi-Level Optimization. *IEEE Transactions on Mobile Computing* **2024**, *23*, 15196–15209. https://doi.org/10.1109/TMC.2024.3455331.

 doi:10.20944/preprints202601.0271.v1

215. Chen, C.; Chen, Z.; Zhou, Y.; Kailkhura, B. FedCluster: Boosting the Convergence of Federated Learning via Cluster-Cycling. In Proceedings of the 2020 IEEE International Conference on Big Data (Big Data), 2020, pp. 5017–5026. https://doi.org/10.1109/BigData50022.2020.9377960.
216. Qi, Z.; Wang, Y.; Chen, Z.; Wang, R.; Meng, X.; Meng, L. Clustering-based Curriculum Construction for Sample-Balanced Federated Learning. In Proceedings of the Artificial Intelligence; Fang, L.; Povey, D.; Zhai, G.; Mei, T.; Wang, R., Eds., Cham, 2022; pp. 155–166.
217. Hu, M.; Yue, Z.; Xie, X.; Chen, C.; Huang, Y.; Wei, X.; Lian, X.; Liu, Y.; Chen, M. Is Aggregation the Only Choice? Federated Learning via Layer-wise Model Recombination. In Proceedings of the Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, New York, NY, USA, 2024; KDD '24, p. 1096–1107. https://doi.org/10.1145/3637528.3671722.
218. Hu, M.; Zhou, P.; Yue, Z.; Ling, Z.; Huang, Y.; Li, A.; Liu, Y.; Lian, X.; Chen, M. FedCross: Towards Accurate Federated Learning via Multi-Model Cross-Aggregation. In Proceedings of the 2024 IEEE 40th International Conference on Data Engineering (ICDE), 2024, pp. 2137–2150. https://doi.org/10.1109/ICDE60146.2024.00170.
219. Xia, Z.; Hu, M.; Yan, D.; Liu, R.; Li, A.; Xie, X.; Chen, M. MultiSFL: Towards Accurate Split Federated Learning via Multi-Model Aggregation and Knowledge Replay. *Proceedings of the AAAI Conference on Artificial Intelligence* **2025**, *39*, 914–922. https://doi.org/10.1609/aaai.v39i1.32076.
220. Nasr, M.; Shokri, R.; Houmansadr, A. Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-box Inference Attacks against Centralized and Federated Learning. In Proceedings of the 2019 IEEE Symposium on Security and Privacy (SP), 2019, pp. 739–753. https://doi.org/10.1109/SP.2019.00065.
221. Yan, H.; Li, S.; Wang, Y.; Zhang, Y.; Sharif, K.; Hu, H.; Li, Y. Membership Inference Attacks Against Deep Learning Models via Logits Distribution. *IEEE Transactions on Dependable and Secure Computing* **2023**, *20*, 3799–3808. https://doi.org/10.1109/TDSC.2022.3222880.
222. Bonawitz, K.; Ivanov, V.; Kreuter, B.; Marcedone, A.; McMahan, H.B.; Patel, S.; Ramage, D.; Segal, A.; Seth, K. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In Proceedings of the Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, New York, NY, USA, 2017; CCS '17, p. 1175–1191. https://doi.org/10.1145/3133956.3133982.
223. Dwork, C.; Roth, A. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science* **2014**, *9*, 211–407. https://doi.org/10.1561/0400000042.
224. Cao, Y.; Yang, J. Towards Making Systems Forget with Machine Unlearning. In Proceedings of the 2015 IEEE Symposium on Security and Privacy, 2015, pp. 463–480. https://doi.org/10.1109/SP.2015.35.
225. Yu, Q.; Li, S.; Raviv, N.; Kalan, S.M.M.; Soltanolkotabi, M.; Avestimehr, S.A. Lagrange Coded Computing: Optimal Design for Resiliency, Security, and Privacy. In Proceedings of the Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics; Chaudhuri, K.; Sugiyama, M., Eds. PMLR, 16–18 Apr 2019, Vol. 89, *Proceedings of Machine Learning Research*, pp. 1215–1225.
226. Li, S.; Hou, S.; Buyukates, B.; Avestimehr, S. Secure Federated Clustering, 2022, [arXiv:cs.LG/2205.15564].
227. Wang, Y.; Pang, W.; Pedrycz, W. One-Shot Federated Clustering Based on Stable Distance Relationships. *IEEE Transactions on Industrial Informatics* **2024**, *20*, 13262–13272. https://doi.org/10.1109/TII.2024.3435420.
228. Scott, J.; Lampert, C.H.; Saulpic, D. Differentially Private Federated k-Means Clustering with Server-Side Data, 2025, [arXiv:cs.CR/2506.05408].
229. MCQUEEN, J. Some methods of classification and analysis of multivariate observations. *Proc. of 5th Berkeley Symposium on Math. Stat. and Prob.* **1967**, pp. 281–297.
230. Ikotun, A.M.; Ezugwu, A.E.; Abualigah, L.; Abuhaija, B.; Heming, J. K-means clustering algorithms: A comprehensive review, variants analysis, and advances in the era of big data. *Information Sciences* **2023**, *622*, 178–210. https://doi.org/https://doi.org/10.1016/j.ins.2022.11.139.
231. Ng, A.; Jordan, M.; Weiss, Y. On Spectral Clustering: Analysis and an algorithm. In Proceedings of the Advances in Neural Information Processing Systems; Dietterich, T.; Becker, S.; Ghahramani, Z., Eds. MIT Press, 2001, Vol. 14.
232. Ding, L.; Li, C.; Jin, D.; Ding, S. Survey of spectral clustering based on graph theory. *Pattern Recognition* **2024**, *151*, 110366. https://doi.org/https://doi.org/10.1016/j.patcog.2024.110366.
233. Naeem, A.; Anees, T.; Naqvi, R.A.; Loh, W.K. A comprehensive analysis of recent deep and federated-learning-based methodologies for brain tumor diagnosis. *Journal of Personalized Medicine* **2022**, *12*, 275.
234. Lo, J.; Timothy, T.Y.; Ma, D.; Zang, P.; Owen, J.P.; Zhang, Q.; Wang, R.K.; Beg, M.F.; Lee, A.Y.; Jia, Y.; et al. Federated learning for microvasculature segmentation and diabetic retinopathy classification of OCT data. *Ophthalmology Science* **2021**, *1*, 100069.

 doi:10.20944/preprints202601.0271.v1

235. Zhou, S.; Landman, B.A.; Huo, Y.; Gokhale, A. Communication-efficient federated learning for multi-institutional medical image classification. In Proceedings of the Medical Imaging 2022: Imaging Informatics for Healthcare, Research, and Applications. SPIE, 2022, Vol. 12037, pp. 6–12.
236. Antunes, R.S.; André da Costa, C.; Küderle, A.; Yari, I.A.; Eskofier, B. Federated Learning for Healthcare: Systematic Review and Architecture Proposal **2022**. *13*. https://doi.org/10.1145/3501813.
237. Liu, Z.; Chen, Y.; Zhao, Y.; Yu, H.; Liu, Y.; Bao, R.; Jiang, J.; Nie, Z.; Xu, Q.; Yang, Q. Contribution-Aware Federated Learning for Smart Healthcare. *Proceedings of the AAAI Conference on Artificial Intelligence* **2022**, *36*, 12396–12404. https://doi.org/10.1609/aaai.v36i11.21505.
238. Yang, Q.; Zhang, J.; Hao, W.; Spell, G.P.; Carin, L. Flop: Federated learning on medical datasets using partial networks. In Proceedings of the Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, 2021, pp. 3845–3853.
239. Ogier du Terrail, J.; Ayed, S.S.; Cyffers, E.; Grimberg, F.; He, C.; Loeb, R.; Mangold, P.; Marchand, T.; Marfoq, O.; Mushtaq, E.; et al. Flamby: Datasets and benchmarks for cross-silo federated learning in realistic healthcare settings. *Advances in Neural Information Processing Systems* **2022**, *35*, 5315–5334.
240. Sarma, K.V.; Harmon, S.; Sanford, T.; Roth, H.R.; Xu, Z.; Tetreault, J.; Xu, D.; Flores, M.G.; Raman, A.G.; Kulkarni, R.; et al. Federated learning improves site performance in multicenter deep learning without data sharing. *Journal of the American Medical Informatics Association* **2021**, *28*, 1259–1264.
241. Kumar, R.; Khan, A.A.; Kumar, J.; Golilarz, N.A.; Zhang, S.; Ting, Y.; Zheng, C.; Wang, W.; et al. Blockchain-federated-learning and deep learning models for covid-19 detection using ct imaging. *IEEE Sensors Journal* **2021**, *21*, 16301–16314.
242. Yan, B.; Wang, J.; Cheng, J.; Zhou, Y.; Zhang, Y.; Yang, Y.; Liu, L.; Zhao, H.; Wang, C.; Liu, B. Experiments of federated learning for COVID-19 chest X-ray images. In Proceedings of the International Conference on Artificial Intelligence and Security. Springer, 2021, pp. 41–53.
243. Yang, D.; Xu, Z.; Li, W.; Myronenko, A.; Roth, H.R.; Harmon, S.; Xu, S.; Turkbey, B.; Turkbey, E.; Wang, X.; et al. Federated semi-supervised learning for COVID region segmentation in chest CT using multi-national data from China, Italy, Japan. *Medical image analysis* **2021**, *70*, 101992.
244. Qayyum, A.; Ahmad, K.; Ahsan, M.A.; Al-Fuqaha, A.; Qadir, J. Collaborative federated learning for healthcare: Multi-modal covid-19 diagnosis at the edge. *IEEE Open Journal of the Computer Society* **2022**, *3*, 172–184.
245. Abdul Salam, M.; Taha, S.; Ramadan, M. COVID-19 detection using federated machine learning. *PloS one* **2021**, *16*, e0252573.
246. Duan, R.; Boland, M.R.; Liu, Z.; Liu, Y.; Chang, H.H.; Xu, H.; Chu, H.; Schmid, C.H.; Forrest, C.B.; Holmes, J.H.; et al. Learning from electronic health records across multiple sites: A communication-efficient and privacy-preserving distributed algorithm. *Journal of the American Medical Informatics Association* **2020**, *27*, 376–385.
247. Huang, L.; Shea, A.L.; Qian, H.; Masurkar, A.; Deng, H.; Liu, D. Patient clustering improves efficiency of federated machine learning to predict mortality and hospital stay time using distributed electronic medical records. *Journal of biomedical informatics* **2019**, *99*, 103291.
248. Li, Z.; Roberts, K.; Jiang, X.; Long, Q. Distributed learning from multiple EHR databases: contextual embedding models for medical events. *Journal of biomedical informatics* **2019**, *92*, 103138.
249. Vaid, A.; Jaladanki, S.K.; Xu, J.; Teng, S.; Kumar, A.; Lee, S.; Somani, S.; Paranjpe, I.; De Freitas, J.K.; Wanyan, T.; et al. Federated learning of electronic health records to improve mortality prediction in hospitalized patients with COVID-19: machine learning approach. *JMIR medical informatics* **2021**, *9*, e24207.
250. Brisimi, T.S.; Chen, R.; Mela, T.; Olshevsky, A.; Paschalidis, I.C.; Shi, W. Federated learning of predictive models from federated electronic health records. *International journal of medical informatics* **2018**, *112*, 59–67.
251. Jia, C.; Hu, M.; Chen, Z.; Yang, Y.; Xie, X.; Liu, Y.; Chen, M. AdaptiveFL: Adaptive heterogeneous federated learning for resource-constrained AIoT systems. In Proceedings of the Proceedings of the 61st ACM/IEEE Design Automation Conference, 2024, pp. 1–6.
252. Xia, Z.; Hu, M.; Yan, D.; Xie, X.; Li, T.; Li, A.; Zhou, J.; Chen, M. CaBaFL: Asynchronous Federated Learning via Hierarchical Cache and Feature Balance. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **2024**, *43*, 4057–4068. https://doi.org/10.1109/TCAD.2024.3446881.
253. Chen, Z.; Jia, C.; Hu, M.; Xie, X.; Li, A.; Chen, M. FlexFL: Heterogeneous Federated Learning via APoZ-Guided Flexible Pruning in Uncertain Scenarios. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **2024**, *43*, 4069–4080. https://doi.org/10.1109/TCAD.2024.3444695.

 doi:10.20944/preprints202601.0271.v1

254. Dai, C.; Wei, S.; Dai, S.; Garg, S.; Kaddoum, G.; Shamim Hossain, M. Federated Self-Supervised Learning Based on Prototypes Clustering Contrastive Learning for Internet of Vehicles Applications. *IEEE Internet of Things Journal* **2025**, *12*, 4692–4700. https://doi.org/10.1109/JIOT.2024.3453336.
255. Yan, D.; Yang, Y.; Hu, M.; Fu, X.; Chen, M. MMDFL: Multi-Model-based Decentralized Federated Learning for Resource-Constrained AIoT Systems. In Proceedings of the 2025 62nd ACM/IEEE Design Automation Conference (DAC), 2025, pp. 1–7. https://doi.org/10.1109/DAC63849.2025.11133116.
256. Moulik, S.; Misra, S.; Gaurav, A. Cost-effective mapping between wireless body area networks and cloud service providers based on multi-stage bargaining. *IEEE Transactions on Mobile Computing* **2016**, *16*, 1573–1586.
257. Cui, Y.; Cao, K.; Cao, G.; Qiu, M.; Wei, T. Client scheduling and resource management for efficient training in heterogeneous IoT-edge federated learning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **2021**, *41*, 2407–2420.
258. Yu, H.; Xu, R.; Zhang, H.; Yang, Z.; Liu, H. EV-FL: Efficient verifiable federated learning with weighted aggregation for industrial IoT networks. *IEEE/ACM Transactions on Networking* **2023**, *32*, 1723–1737.
259. Jia, Z.; Zhou, T.; Yan, Z.; Hu, J.; Shi, Y. Personalized meta-federated learning for IoT-enabled health monitoring. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **2024**, *43*, 3157–3170.
260. Li, X.; Li, S.; Li, Y.; Zhou, Y.; Chen, C.; Zheng, Z. A Personalized Federated Tensor Factorization Framework for Distributed IoT Services QoS Prediction From Heterogeneous Data. *IEEE Internet of Things Journal* **2022**, *9*, 25460–25473. https://doi.org/10.1109/JIOT.2022.3197172.
261. Tan, B.; Liu, B.; Zheng, V.; Yang, Q. A Federated Recommender System for Online Services. In Proceedings of the Proceedings of the 14th ACM Conference on Recommender Systems, New York, NY, USA, 2020; RecSys '20, p. 579–581. https://doi.org/10.1145/3383313.3411528.
262. Muhammad, K.; Wang, Q.; O'Reilly-Morgan, D.; Tragos, E.; Smyth, B.; Hurley, N.; Geraci, J.; Lawlor, A. FedFast: Going Beyond Average for Faster Training of Federated Recommender Systems. In Proceedings of the Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, New York, NY, USA, 2020; KDD '20, p. 1234–1242. https://doi.org/10.1145/3394486.3403176.
263. Li, Y.; Shan, Y.; Liu, Y.; Wang, H.; Wang, W.; Wang, Y.; Li, R. Personalized Federated Recommendation for Cold-Start Users via Adaptive Knowledge Fusion. In Proceedings of the Proceedings of the ACM on Web Conference 2025, New York, NY, USA, 2025; WWW '25, p. 2700–2709. https://doi.org/10.1145/3696410.3714635.
264. Chen, M.; Yang, Z.; Saad, W.; Yin, C.; Poor, H.V.; Cui, S. A Joint Learning and Communications Framework for Federated Learning Over Wireless Networks. *IEEE Transactions on Wireless Communications* **2021**, *20*, 269–283. https://doi.org/10.1109/TWC.2020.3024629.
265. Beutel, D.; Topal, T.; Mathur, A.; Qiu, X.; Parcollet, T.; Lane, N.D. Flower: A Friendly Federated Learning Research Framework. *arXiv preprint arXiv:2007.14390* **2022**.
266. Chen, M.; et al. OpenFed: A Comprehensive and Versatile Open-Source Federated Learning Framework. In Proceedings of the CVPR Workshops (FedVision), 2023.
267. He, C.; et al. Benchmarking Federated Learning Algorithms under System and Statistical Heterogeneity. *IEEE Transactions on Neural Networks and Learning Systems* **2023**.
268. Kairouz, P.; McMahan, H.B.; Avent, B.; et al. Advances and Open Problems in Federated Learning. *Foundations and Trends® in Machine Learning* **2021**, *14*, 1–210.
269. Abadi, M.; Chu, A.; Goodfellow, I.; et al. Deep Learning with Differential Privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* **2016**.
270. He, C.; Annavaram, M.; Avestimehr, S. FedML: A Research Library and Benchmark for Federated Machine Learning. *arXiv preprint arXiv:2007.13518* **2020**.
271. He, C.; Avestimehr, S. Towards Scalable and Robust Federated Learning Systems. *IEEE Transactions on Parallel and Distributed Systems* **2021**.
272. Ryffel, T.; Trask, A.; Dahl, M.; Wagner, B.; Mancuso, J.; Rueckert, D. A Generic Framework for Privacy Preserving Deep Learning. *arXiv preprint arXiv:1811.04017* **2018**.
273. Yang, Q.; Liu, Y.; Chen, T.; Tong, Y. Federated Machine Learning: Concept and Applications. *ACM Transactions on Intelligent Systems and Technology* **2019**.
274. Liu, Y.; Chen, T.; Yang, Q. Secure Federated Learning for Vertical Partitioned Data. *IEEE Intelligent Systems* **2020**, *35*, 90–97.
275. Ning, W.; Zhu, Y.; Song, C.; Li, H.; Zhu, L.; Xie, J.; Chen, T.; Xu, T.; Xu, X.; Gao, J. Blockchain-Based Federated Learning: A Survey and New Perspectives. *Applied Sciences* **2024**, *14*. https://doi.org/10.3390/app14209459.

 doi:10.20944/preprints202601.0271.v1

276. Larasati, H.T.; Firdaus, M.; Kim, H. Quantum federated learning: Remarks and challenges. In Proceedings of the 2022 IEEE 9th International Conference on Cyber Security and Cloud Computing (CSCloud)/2022 IEEE 8th International Conference on Edge Computing and Scalable Cloud (EdgeCom). IEEE, 2022, pp. 1–5.
277. Innan, N.; Marchisio, A.; Bennai, M.; Shafique, M. Qfnn-ffd: Quantum federated neural network for financial fraud detection. In Proceedings of the 2025 IEEE International Conference on Quantum Software (QSW). IEEE, 2025, pp. 41–47.
278. Chu, C.; Jiang, L.; Chen, F. Cryptoqfl: quantum federated learning on encrypted data. In Proceedings of the 2023 IEEE International Conference on Quantum Computing and Engineering (QCE). IEEE, 2023, Vol. 1, pp. 1231–1237.
279. Zhang, Y.; Zhang, C.; Zhang, C.; Fan, L.; Zeng, B.; Yang, Q. Federated learning with quantum secure aggregation. *arXiv preprint arXiv:2207.07444* **2022**.
280. Ren, C.; Yan, R.; Zhu, H.; Yu, H.; Xu, M.; Shen, Y.; Xu, Y.; Xiao, M.; Dong, Z.Y.; Skoglund, M.; et al. Toward Quantum Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems* **2025**.
281. Yao, Y.; Zhang, J.; Wu, J.; Huang, C.; Xia, Y.; Yu, T.; Zhang, R.; Kim, S.; Rossi, R.; Li, A.; et al. Federated Large Language Models: Current Progress and Future Directions, 2025, [arXiv:cs.LG/2409.15723].
282. Abouelmagd, A.A.; Hilal, A. Emerging Paradigms for Securing Federated Learning Systems, 2025, [arXiv:cs.CR/2509.21147].
283. Li, T.; Sahu, A.K.; Talwalkar, A.; Smith, V. Federated Learning: Challenges, Methods, and Future Directions. *IEEE Signal Processing Magazine* **2020**, *37*, 50–60. https://doi.org/10.1109/msp.2020.2975749.