INTERNATIONAL JOURNAL OF BUSINESS AND ECONOMICS INSIGHTS

Volume: 5; **Issue:** 3
Pages: 238– 269
Accepted: 29 August 2025
Published: 15 September 2025

FEDERATED LEARNING MODELS FOR PRIVACY-PRESERVING AI IN ENTERPRISE DECISION SYSTEMS

Md Mohaiminul Hasan[1];

[1]. Master in Project Management; St. Francis College - Brooklyn, NY, USA;
Email: mohaiminul.hasan22@gmail.com

Doi: 10.63125/ry033286
This work was peer-reviewed under the editorial responsibility of the IJEI, 2025

Abstract

This systematic review examines the role of federated learning (FL) as a privacy-preserving paradigm for enterprise decision systems, synthesizing evidence from 187 peer-reviewed studies. Guided by the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework, the review integrates algorithmic, systems, security, sectoral, and governance perspectives to provide a comprehensive account of current knowledge. Findings highlight that foundational algorithms such as FedAvg, FedProx, and SCAFFOLD dominate the methodological landscape, with significant adaptations emerging to address non-IID and unbalanced datasets across distributed organizational silos. Privacy-preserving mechanisms—including differential privacy, secure aggregation, homomorphic encryption, and multiparty computation—were consistently applied as layered defenses, balancing mathematical guarantees with empirical resilience. The synthesis further revealed critical vulnerabilities to model poisoning, backdoor attacks, and gradient leakage, alongside defensive strategies such as robust aggregation, anomaly detection, and differential privacy clipping. Sector-specific implementations demonstrate FL's practical utility in healthcare, finance, retail, logistics, telecommunications, and public services, where it enables collaborative modeling without violating data residency or confidentiality requirements. Governance and ethical frameworks, particularly GDPR, CCPA, and the NIST Privacy Framework, were found to shape deployment practices, while documentation artifacts such as datasheets, model cards, and privacy budget ledgers ensure accountability and transparency. Comparative surveys position FL as an integrative socio-technical architecture that unites distributed optimization, privacy engineering, adversarial robustness, and AI governance into a coherent enterprise-ready model. The review concludes that federated learning provides enterprises with a scalable, secure, and ethically aligned approach to leveraging distributed data while preserving trust and compliance.

Keywords

Federated Learning, Privacy, Security, Governance, FedAvg

INTRODUCTION

Federated learning (FL) refers to a distributed machine-learning paradigm in which models are trained collaboratively across multiple clients while keeping raw data localized, sharing only model updates with a coordinating server (Liu et al., 2022). In enterprise decision systems—spanning finance, healthcare, retail, logistics, and telecommunications—data residency constraints, contractual confidentiality, and statutory privacy laws shape how analytical models can be built from sensitive records. The international relevance of FL emerges from its capacity to learn from cross-border, cross-institutional datasets without aggregating personal or proprietary data in a single repository, thereby aligning technical practice with heterogeneous jurisdictional requirements and industry norms (Zhang et al., 2021). Core mechanisms include on-device or on-premise training, secure aggregation of gradients or model deltas, and orchestration protocols that account for client heterogeneity and unreliable connectivity. Enterprises leverage these mechanisms to support high-stakes decisions—credit risk scoring, fraud detection, medical triage, and supply planning—where model performance benefits from diverse, distributed data while privacy expectations remain strict. The conceptual shift is from data-centralized intelligence to update-centralized intelligence, with privacy exposure reduced by design because sensitive attributes remain within organizational boundaries (Yin et al., 2020). International significance follows from multi-jurisdictional collaborations: global banks, hospital consortia, and telecom operators can coordinate model training across subsidiaries or partners that operate under divergent legal regimes without exchanging raw identifiers. This introductory framing situates FL not simply as an optimization trick, but as an architectural response to the intertwined demands of privacy, compliance, and enterprise-scale decision quality (Abreha et al., 2022).

Figure 1: Key Challenges and Mitigations in Federated Learning

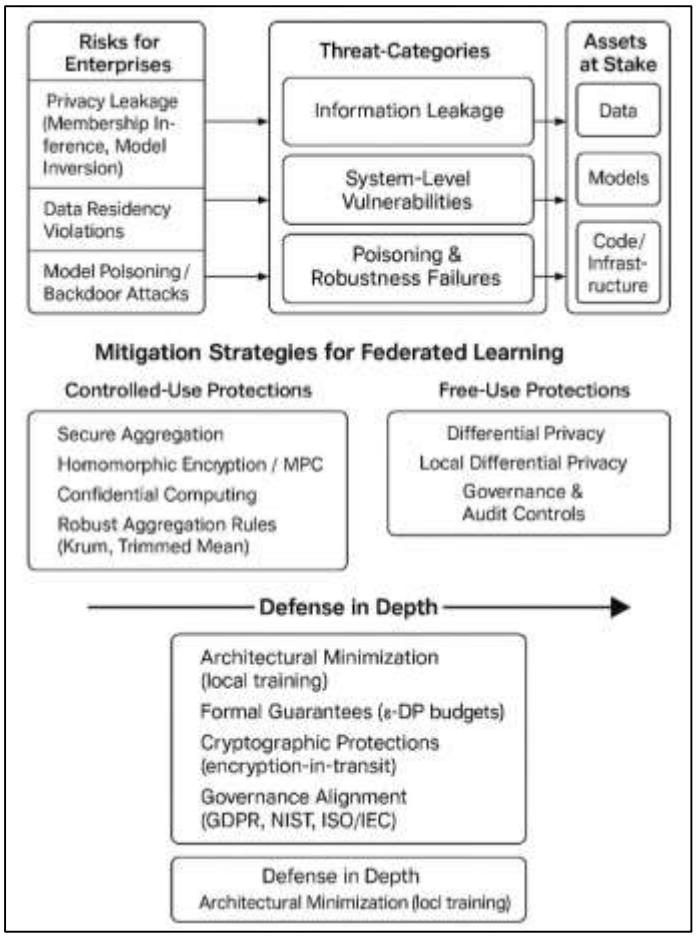

The privacy-preserving character of FL relies on a layered defense that combines architectural minimization with formal privacy guarantees and cryptographic protections. Architectural minimization keeps data local and shares only learned parameters, but model updates may still leak

information through inversion or membership inference. Formal methods such as differential privacy (DP) add calibrated noise to gradients or to aggregated statistics to bound what an adversary can infer about any single record. In the federated setting, DP can be applied per client update or at the server during aggregation, balancing privacy budgets against utility in non-IID, unbalanced workloads common in enterprises (Yuan et al., 2024). Secure aggregation protocols ensure the server can recover only the sum of client updates, hiding each participant's contribution even from an honest-but-curious coordinator. Additional cryptographic tools—homomorphic encryption and secure multiparty computation—are used to protect gradients in transit or enable privacy-preserving analytics on encrypted features. Enterprises also deploy local differential privacy (LDP) for telemetry or feature pre-processing when centralized trust is low, informed by mechanisms like RAPPOR and subsequent heavy-hitters protocols (Danish & Zafor, 2022; Lazaros et al., 2024). Regulatory frameworks provide the governance backbone: principles of data minimization, purpose limitation, and privacy by design in the GDPR, risk-management controls in the NIST Privacy Framework, and transparency/accountability requirements under CCPA converge with FL's boundary-preserving workflow. As a result, privacy preservation in FL rests on defense in depth: local training, formally private noise addition, cryptographically protected aggregation, and policy-driven controls that are auditable within enterprise governance structures (Bashir et al., 2023; Danish & Kamrul, 2022).

Figure 2: Federated Learning in Enterprise Systems

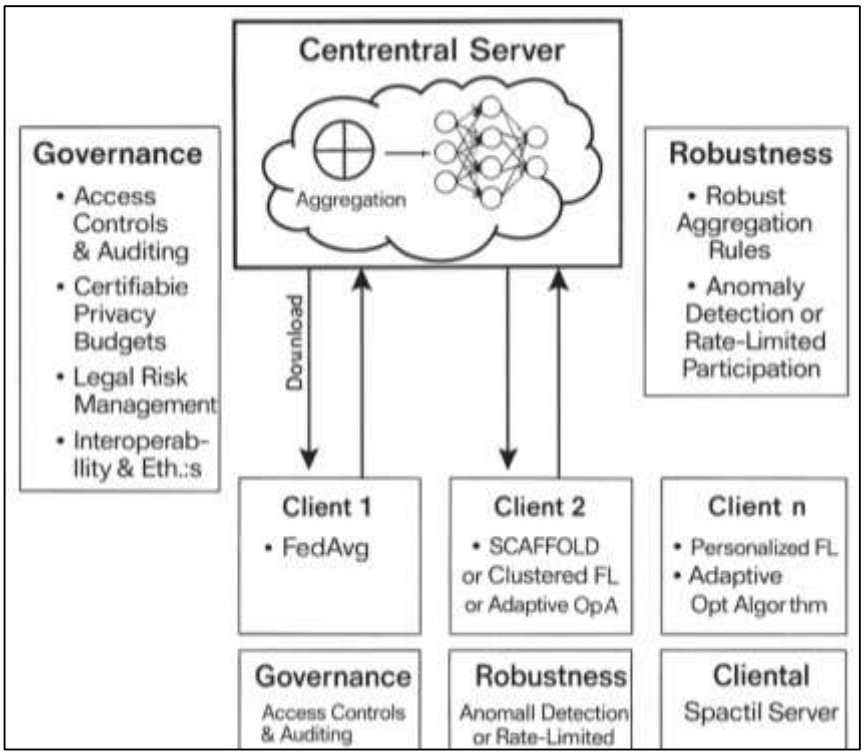

Learning dynamics in non-IID, heterogeneous environments define the technical heart of enterprise FL. Organizational datasets often differ in feature distributions, label proportions, and sampling frequency across branches, partners, and regions. The canonical Federated Averaging (FedAvg) algorithm efficiently averages local stochastic gradient steps across selected clients, but performance and convergence can degrade under skewed client data or unbalanced participation (Jahid, 2022; Khan et al., 2025). Methodological advances respond to these realities. FedProx introduces proximal terms to stabilize updates under system and data heterogeneity. SCAFFOLD corrects client drift with control variates. Clustered FL partitions clients into cohorts with similar data distributions, enabling specialized global models . Adaptive federated optimizers (e.g., FedAdam, FedYogi) refine server-side update rules to improve robustness and speed (Ghimire & Rawat, 2022; Arifur & Noor, 2022). When decision policies vary by region or product line, personalized FL frameworks tune per-client or per-segment heads on top of shared representations, reconciling global learning with local specificity.

Agnostic FL formalizes objectives that target worst-case or mixed client distributions to safeguard performance across diverse enterprise subpopulations. These algorithms address the operational reality that enterprises seldom control data generation processes uniformly; they must instead accommodate device churn, variable compute budgets, and participation constraints while maintaining reproducibility and auditability (Chellapandi et al., 2023; Hasan & Uddin, 2022). The cumulative literature establishes a toolkit for stable learning under heterogeneity, a condition broadly encountered in cross-site enterprise decision systems.

Robustness and security concerns in FL map closely to enterprise risk models. Malicious or faulty clients may attempt model poisoning or backdoor attacks by sending manipulated updates that corrupt global behavior on specific triggers or degrade overall accuracy. Robust aggregation rules such as Krum, Bulyan, and median- or trimmed-mean-based schemes reduce the influence of outliers and adversaries under various threat models (Hu et al., 2021; Rahaman, 2022a). Defensive training strategies incorporate anomaly detection on updates, reputation systems for clients, and coordinate-wise clipping with DP noise to blunt gradient outliers. Inference-time leaks and training-time eavesdropping are considered through analyses of gradient leakage and model inversion, motivating encryption-in-transit, secure enclaves for aggregation, and rate-limited participation. Membership inference and property inference risks receive special scrutiny in regulated sectors, with empirical evidence showing that even aggregate statistics can reveal sensitive attributes without adequate regularization and noise (Khan et al., 2023; Rahaman, 2022b). Governance overlays—access controls, incident playbooks, audit logging, and certifiable privacy budgets—ground these defenses in enterprise processes subject to internal audit and external regulators. The literature on federated robustness thus connects algorithmic choices to concrete enterprise risk indicators, from fraud amplification exposure to compliance penalties associated with leakage (Liu et al., 2023; Rahaman & Ashraf, 2022).

Systems engineering and deployment determine whether FL translates from theory to operational decision support. Production FL systems require device or node selection, incentive-compatible participation, straggler mitigation, and resilience to intermittent connectivity across regions . Industrial case studies describe federated analytics and on-device learning at internet scale, showing how telemetry, cohorting, and privacy accounting integrate with existing data pipelines (Quan et al., 2025). Resource constraints—memory, compute, and energy on edge devices or branch servers—drive model compression and communication-efficient updates through sparsification, quantization, and sketching. Scheduling policies coordinate thousands of clients while enforcing privacy budgets and fairness in participation, preventing over-representation of a single region or business unit. Integration with MLOps adds versioning, canarying, and rollback for global models and local adapters, plus lineage tracking for privacy proofs and cryptographic keys (Adam & Baroud, 2024; Islam, 2022). Cross-silo deployments in enterprises differ from cross-device settings, emphasizing stable, high-bandwidth links and authenticated participants but facing stricter legal contracts and audit trails across subsidiaries or partners. These system concerns tie directly to decision-lifecycle requirements: model freshness for time-sensitive risk scoring, explainability layers compatible with federated representations, and integration with rule-based engines that capture policy constraints across jurisdictions (Le et al., 2024; Hasan et al., 2022). The corpus emphasizes that engineering choices—transport protocols, cryptographic primitives, and optimizer parameters—shape both performance and privacy alignment in enterprise contexts.

Cross-sector applications illustrate how FL supports decision systems under varied regulatory and competitive environments. In healthcare, multi-institutional imaging consortia have shown that FL can achieve accuracy comparable to centralized training across brain tumor segmentation and COVID-19 diagnosis tasks while keeping patient data in place (Kishor, 2022; Redwanul & Zafor, 2022). Financial institutions coordinate across branches and affiliates for anti-money-laundering, credit risk, and fraud detection using privacy guarantees that accommodate bank-secrecy obligations and regional data localization. Retailers and logistics providers train demand-forecasting, recommendation, and route-optimization models across stores or fleets without pooling customer identifiers, supporting localized patterns in seasonality and purchasing without central exposure (Rezaul & Mesbaul, 2022; Tariq et al., 2024). Telecommunications operators apply FL for anomaly detection and network optimization across

base stations that observe different traffic profiles and legal constraints on metadata sharing. Public-sector collaborations examine federated analytics for population health and smart-city sensing where government agencies coordinate across data stewards bound by statutory confidentiality. These case narratives underscore the international scope of FL-based decision pipelines, where heterogeneous data, regulatory strictures, and commercial sensitivity intersect with the need for high-quality predictive signals (Hafi et al., 2024; Hasan, 2022). In each sector, evaluation emphasizes both conventional metrics—AUC, F1, calibration—and privacy/compliance metrics—epsilon budgets, leakage estimates, cryptographic coverage—reflecting the dual mandate of utility and protection (Aouedi et al., 2024; Tarek, 2022).

Governance, ethics, and standards anchor FL within enterprise decision accountability. Privacy-preserving AI intersects with responsible-AI principles concerning fairness, transparency, and contestability, since distributed data can encode regional disparities that affect model behavior. Methodological responses include fairness-aware federated objectives, reweighting across clients, and constrained optimization that controls group-wise performance variance (Hosseini et al., 2023). Documentation practices—datasheets for datasets, model cards for federated models, and privacy budget reports—extend into federated settings to provide auditable artifacts for internal review and external regulators. International standardization efforts guide terminology, threat models, testing procedures, and reporting for privacy engineering in distributed ML, supporting procurement and interoperability across borders. Legal frameworks stress accountability for joint controllers and processors in collaborative learning, clarifying roles in incident response and data-subject rights handling when multiple entities co-train models (Belfeki et al., 2025; Kamrul & Omar, 2022). Ethical risk assessments incorporate inferences about sensitive attributes from gradients and embeddings, encouraging limits on feature scope and establishing guardrails for secondary use (Shokri & Shmatikov, 2015; Melis et al., 2019). The literature positions FL within a broader governance stack: organizational policies, technical controls, legal obligations, and audit mechanisms mutually reinforce privacy-preserving learning that fits enterprise decision contexts across jurisdictions (Kamrul & Tarek, 2022; Tian et al., 2023).

Foundational and surveyed knowledge knit these themes together into an integrated perspective for enterprise readers. Early distributed optimization and communication-efficient learning introduced update compression and partial participation, seeding practical FL. The landmark FedAvg paper catalyzed broad interest by showing effective on-device learning at scale. Subsequent surveys and system papers synthesized algorithms, systems, privacy, and robustness, providing reference architectures and taxonomies now applied in enterprises (Mubashir & Abdul, 2022; Saha et al., 2024). Personalized and robust FL literature aligned federated training with real-world heterogeneity and adversarial risks. Privacy engineering integrated DP-SGD, secure aggregation, and cryptographic enhancements into deployable stacks. Application studies in healthcare and beyond demonstrated external validity on sensitive data regimes. Governance frameworks from OECD, NIST, and ISO/IEC delineated organizational scaffolding for compliant, auditable operations. Together, these works delineate a mature, cross-disciplinary basis for understanding federated learning as a privacy-preserving approach to enterprise decision systems grounded in formal guarantees, robust optimization, secure systems, and accountable governance.

LITERATURE REVIEW

The literature on federated learning (FL) for privacy-preserving enterprise decision systems integrates perspectives from computer science, cryptography, data privacy regulation, and organizational decision theory. To systematically situate this research domain, a review must trace the conceptual roots of federated learning in distributed optimization, the parallel evolution of privacy-preserving techniques such as differential privacy and secure multiparty computation, and their confluence within enterprise-scale decision contexts. FL is not merely an algorithmic innovation; it is a socio-technical response to global challenges around compliance, security, and fairness in cross-border data processing. As such, the literature review must highlight both methodological and applied streams of research, acknowledging how theoretical guarantees translate into operational systems. Moreover, by mapping contributions across heterogeneity management, robustness against adversaries, systems deployment, sector-specific implementations, and governance frameworks, this section provides an

integrative synthesis. The review is structured thematically, presenting the scholarly discourse through distinct but interlinked sub-sections. Each section foregrounds seminal contributions, empirical demonstrations, and conceptual debates that define the state of knowledge. This layered approach captures both the technical sophistication of FL mechanisms and the institutional imperatives that drive their adoption in decision-critical enterprise environments.

Federated Learning and Distributed Optimization

The origins of federated learning (FL) are deeply rooted in distributed machine learning research, which explored how models could be trained collaboratively across multiple computing nodes without relying on centralized data pools. Early distributed optimization studies emphasized stochastic gradient descent (SGD) under distributed conditions, examining how to coordinate updates across processors and networks (Liu et al., 2022; Muhammad & Kamrul, 2022). These efforts sought to reduce computational bottlenecks while ensuring convergence guarantees in large-scale machine learning. Communication efficiency emerged as a central concern because exchanging full gradient information across multiple devices or servers created excessive network overhead. To address this, parameter-server architectures and decentralized training models were proposed, enabling distributed nodes to share partial updates instead of complete parameter sets. Such methods inspired subsequent refinements, including asynchronous SGD and mini-batch parallelization, which allowed scalability across thousands of devices with tolerable latency (Rahman et al., 2020; Reduanul & Shoeb, 2022). Communication-efficient optimization approaches such as quantization, sparsification, and sketching began to play a pivotal role in bridging distributed optimization and privacy concerns. By minimizing the communication burden while retaining critical model signals, these methods established the foundation for FL's later emphasis on bandwidth efficiency and edge-device deployment. Thus, the emergence of distributed machine learning prior to FL was characterized by a dual objective: achieving parallelization and maintaining acceptable convergence while operating under resource constraints, both of which directly informed the conceptual leap to federated learning architectures (Nguyen et al., 2021; Kumar & Zobayer, 2022).

Figure 3: The Origins of Federated Learning (FL)

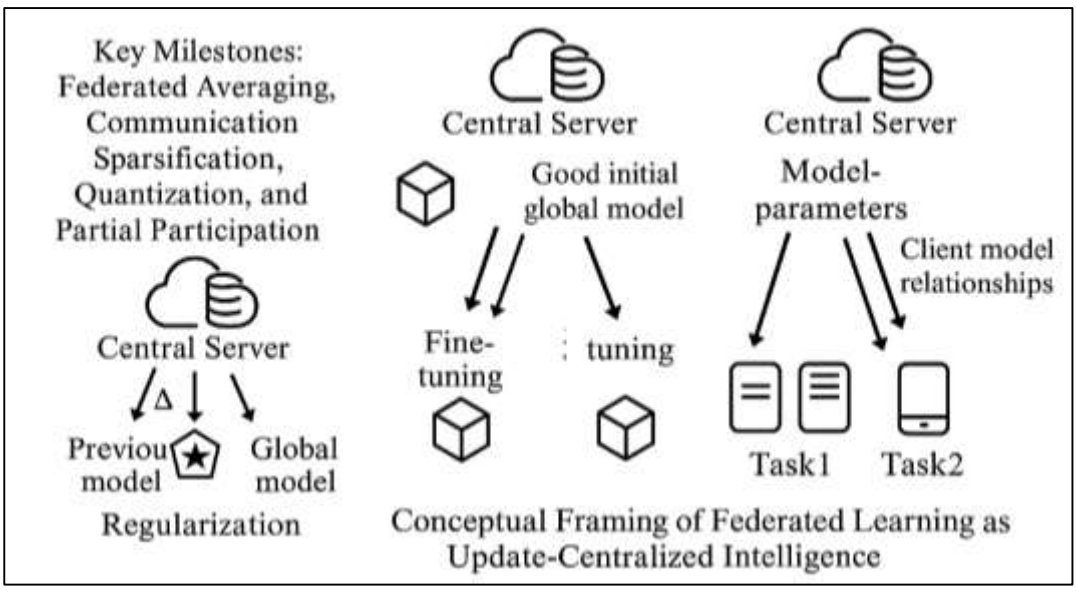

Federated learning gained distinct identity through the introduction of the Federated Averaging (FedAvg) algorithm, which demonstrated how decentralized clients could perform multiple local gradient steps before communicating with a server for model aggregation (Sadia & Shaiful, 2022; Zhou et al., 2021). FedAvg addressed both scalability and communication efficiency by reducing the frequency of exchanges, a critical step for deploying learning systems across edge devices. This milestone built on earlier distributed optimization advances, but uniquely combined them with privacy-motivated architectural choices. Beyond FedAvg, communication sparsification techniques

were formalized to further reduce transmission costs, such as gradient pruning and top-k selection approaches, which shared only a subset of significant updates (Mills et al., 2019; Noor & Momena, 2022). Quantization methods added another layer of efficiency by encoding gradients in lower precision, effectively lowering communication overhead without major losses in accuracy. Partial participation emerged as another defining milestone, as practical FL deployments had to accommodate clients that were intermittently available or resource-constrained, leading to strategies that allowed subsets of devices to contribute at each round. Together, these innovations formed the backbone of FL practice by balancing statistical efficiency, communication feasibility, and system reliability (Istiaque et al., 2023; Yang et al., 2021). Empirical studies in healthcare, finance, and telecommunications validated these methods by demonstrating robust performance across real-world, distributed datasets. Thus, FedAvg, sparsification, quantization, and partial participation represent not only technical breakthroughs but also pragmatic responses to infrastructural and regulatory challenges inherent in enterprise-scale distributed decision systems (Li et al., 2024; Hasan et al., 2023).

A defining conceptual distinction of FL is the shift from "data-centralized" to "update-centralized" intelligence. Traditional machine learning paradigms aggregated raw datasets into centralized repositories for training, exposing organizations to risks related to data leakage, jurisdictional transfer restrictions, and breaches of confidentiality (Hossain et al., 2023; Xu et al., 2020). Federated learning, by contrast, retains data locally and centralizes only the updates, creating a form of boundary-preserving collaboration. This shift is not merely technical but socio-technical, aligning machine learning architectures with privacy-by-design principles enshrined in frameworks such as GDPR and NIST (Rahaman & Ashraf, 2023; Zhang et al., 2021). The update-centralized framing reshapes how enterprises conceptualize collaboration: instead of pooling sensitive records, organizations engage in joint intelligence-building through gradient or weight exchanges. Studies in cross-silo settings—such as hospital networks or multinational banking consortia—demonstrate how this approach supports compliance while retaining predictive power. Update-centralization also underscores the adversarial dimension, since even aggregated updates may leak information if not protected through differential privacy or secure aggregation protocols (Asad et al., 2020; Sultan et al., 2023). Nevertheless, this conceptual reframing strengthens organizational trust and enables multi-party collaborations where central data sharing would otherwise be impossible. Thus, the evolution from centralized data models to update-centralized frameworks illustrates how distributed optimization matured into a privacy-preserving architecture suited for enterprise decision-making environments across sectors and jurisdictions (Asad et al., 2021).

Privacy-Preserving Mechanisms in Federated Learning

Differential privacy (DP) has become a cornerstone of privacy-preserving federated learning (FL), providing formal guarantees that limit the ability of adversaries to infer information about any individual record in a dataset (Xing et al., 2022). Within FL, DP is implemented in two primary modes: local differential privacy (LDP) and global differential privacy (GDP). LDP ensures that each client perturbs its updates before transmitting them to the server, often by injecting calibrated noise into gradients or model parameters. This mode is especially suitable for environments where clients do not fully trust the aggregator. GDP, by contrast, applies noise after updates have been securely aggregated on the server, offering stronger utility preservation while still bounding disclosure risk. The trade-off between LDP and GDP lies in balancing privacy and utility: LDP provides stronger protections against curious servers but can introduce higher noise, reducing model accuracy (Hossen et al., 2023; Yahata et al., 2024). Enterprises apply GDP in collaborative environments where the aggregator is semi-trusted and LDP when regulatory or contractual frameworks demand strict boundary protection. Empirical studies in healthcare and mobile applications show that carefully tuned DP-SGD mechanisms allow FL models to retain competitive performance while ensuring privacy budgets remain within acceptable thresholds. Thus, DP serves as both a mathematical framework and a practical safeguard, adapting to enterprise trust assumptions while embedding accountability into federated decision systems (Dinh et al., 2021; Tawfiqul, 2023).

Beyond DP, cryptographic protocols provide structural protection in federated learning by ensuring that client updates remain hidden from unauthorized inspection. Secure aggregation techniques allow the server to recover only the sum of client updates, preventing exposure of individual contributions

even in the presence of a curious coordinator (Haripriya et al., 2025; Uddin & Ashraf, 2023). These methods use additive secret sharing and masking schemes to guarantee that updates are only revealed when combined, making single-client gradients computationally infeasible to extract. Complementing secure aggregation, homomorphic encryption (HE) allows mathematical operations to be performed directly on encrypted updates, preserving confidentiality during aggregation but at the cost of increased computational overhead . Meanwhile, secure multiparty computation (SMPC) protocols enable multiple clients to jointly compute functions over their data without revealing it to each other, supporting more complex operations than simple averaging (Das, 2018; Momena & Hasan, 2023). Enterprises often adopt hybrid strategies, combining DP with secure aggregation or HE to strengthen protection against both external and internal threats. Case studies in cross-silo FL—for example, hospital consortia—demonstrate the feasibility of cryptographically enhanced training pipelines that preserve patient confidentiality while supporting predictive analytics. Similarly, banks use secure aggregation to comply with secrecy obligations in fraud detection collaborations. These methods show how cryptography operationalizes privacy guarantees beyond statistical noise, embedding robust protections into FL's communication layer and ensuring compliance with enterprise-level security standards (Hosseini & Khisti, 2021; Sanjai et al., 2023). Despite theoretical assurances, empirical research demonstrates that federated learning models remain vulnerable to privacy leakage if protective measures are not carefully implemented. Membership inference attacks show that adversaries can determine whether a specific record was part of the training dataset, even when only observing aggregated models (Khan et al., 2021; Akter et al., 2023).

Figure 4: Privacy Mechanisms in Federated Learning

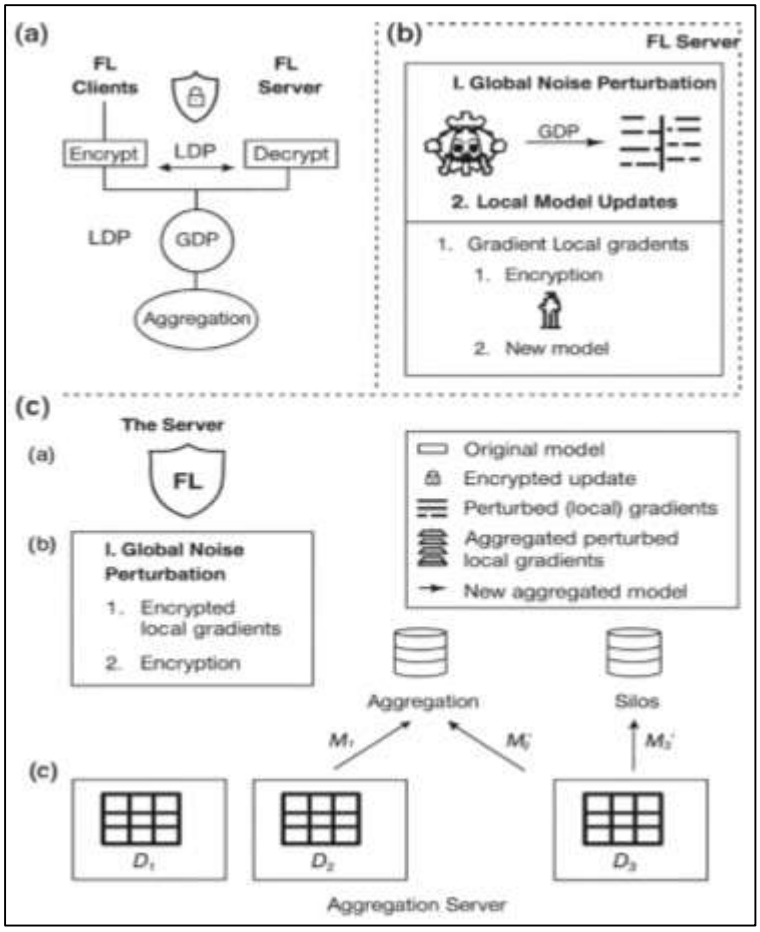

Model inversion attacks go further by reconstructing input features from gradients, particularly when updates are sparse or involve sensitive feature embeddings. Property inference attacks exploit correlations to uncover hidden attributes of training data, raising risks in sensitive sectors like finance

and healthcare. These empirical findings highlight gaps between formal privacy budgets and real-world adversarial capabilities. For instance, (Gao et al., 2024) show that differential privacy effectively bounds leakage, but only when properly tuned noise levels are applied. Secure aggregation similarly reduces risks but does not address inference from final models. Empirical case studies illustrate that privacy risks can persist even under protective frameworks, especially when adversaries control multiple malicious clients in a federated network (Liu et al., 2024). This tension between formal guarantees and empirical vulnerabilities underscores the importance of combining DP, cryptography, and adversarial robustness methods in federated deployments. For enterprises, leakage risks translate directly into regulatory liabilities, reinforcing the need for multi-layered defense in privacy-preserving AI systems (Danish & Zafor, 2024; Liu, 2024).

The literature shows that privacy-preserving mechanisms in FL function most effectively when combined into a layered framework. Differential privacy provides mathematical guarantees, but excessive noise can harm model utility. Secure aggregation ensures confidentiality of individual updates, yet requires careful handling of dropouts and adversarial masking (Pennisi et al., 2024). Homomorphic encryption and SMPC expand the scope of protected operations but introduce computational costs that limit scalability. Meanwhile, empirical attacks highlight the persistence of leakage risks that no single mechanism fully mitigates. This synthesis positions FL privacy as a defense-in-depth architecture, in which local training, secure transmission, and statistical guarantees complement one another within enterprise workflows (Istiaque et al., 2024; Song et al., 2020). Studies in real-world domains such as medical imaging, mobile analytics, and finance demonstrate how combinations of DP and cryptography balance utility with compliance. Enterprises navigating GDPR, CCPA, and NIST requirements rely on these integrated approaches not only to secure data but also to generate auditable evidence of privacy preservation (Hasan et al., 2024; Song et al., 2020). Thus, privacy-preserving mechanisms in FL reflect an evolving consensus: only through multi-layered, hybrid approaches can federated learning reconcile the demands of performance, security, and regulatory accountability across global enterprise decision systems (Gu et al., 2023; Rahaman, 2024).

Addressing Data Heterogeneity in Enterprise Settings

Enterprise datasets rarely conform to the i.i.d. assumption; instead, client silos exhibit covariate shift, label imbalance, concept drift, and divergent sampling processes that degrade naive aggregation. In cross-silo deployments, business units collect features with local schemas and distinct operational regimes, producing heavy-tailed participation and unbalanced sample sizes that bias the global model toward over-represented clients (Zhan et al., 2020). Empirical studies show that gradient directions diverge substantially under non-IID partitions, leading to slower convergence and accuracy drops when vanilla FedAvg aggregates heterogeneous updates. System factors interact with statistics: intermittent connectivity, stragglers, and device churn exacerbate client drift because models adapt to local distributions between rounds. Healthcare and telecom case studies report domain-specific shifts—scanner/site effects and traffic-mix heterogeneity—that motivate distribution-aware aggregation or per-site normalization (Hasan, 2024; Song et al., 2022). Multi-task formulations treat each silo as a related but distinct task, improving fit over globally shared heads when label supports differ across branches. Architectural techniques such as local batch-norm statistics, partial sharing of representations, and client-specific adapters help stabilize training where feature scales and priors vary. Collectively, this literature characterizes enterprise non-IIDness as a coupled statistical-systems phenomenon that requires algorithmic corrections and deployment-level controls to avoid dominance by large or frequent clients and to preserve minority-silo performance (Li et al., 2022; Ashiqur et al., 2025).

Algorithmic contributions target heterogeneity along optimization, architecture, and objective dimensions. FedProx stabilizes local training by adding a proximal term that limits divergence from the server model, improving robustness under system and data heterogeneity (Wang et al., 2019). SCAFFOLD introduces control variates to correct client drift, reducing variance in local updates and accelerating convergence in skewed settings. Communication-efficient corrections such as FedNova normalize update contributions across heterogeneous local steps, mitigating bias from unequal computation, while FedDyn reshapes the objective with dynamic regularization to counter client-specific minima.

When client populations cluster by distribution, clustered FL forms cohorts and trains specialized global models per cluster, improving fit without full personalization. Personalization lines include Per-FedAvg (meta-learning global initialization for fast local adaptation), pFedMe (bi-level regularized personalization), APFL (mixture of global and local models), LG-FedAvg (shared representations with local heads), and FedBN (client-specific batch-norm statistics to absorb feature-shift). Objective-level approaches such as Agnostic FL (AFL) and q-FFL reweight training to protect worst-case or tail clients and to enhance fairness across silos. Server-side optimizers (FedAdam/FedYogi/FedAdagrad) further stabilize aggregation across heterogeneous updates (Kim, 2025; Md Hasan, 2025). Across studies, these adaptations reduce drift, balance influence, and tailor capacity to local distributions while preserving a shared representational core.

Figure 5: Managing Non-IID Data in Federated Learning

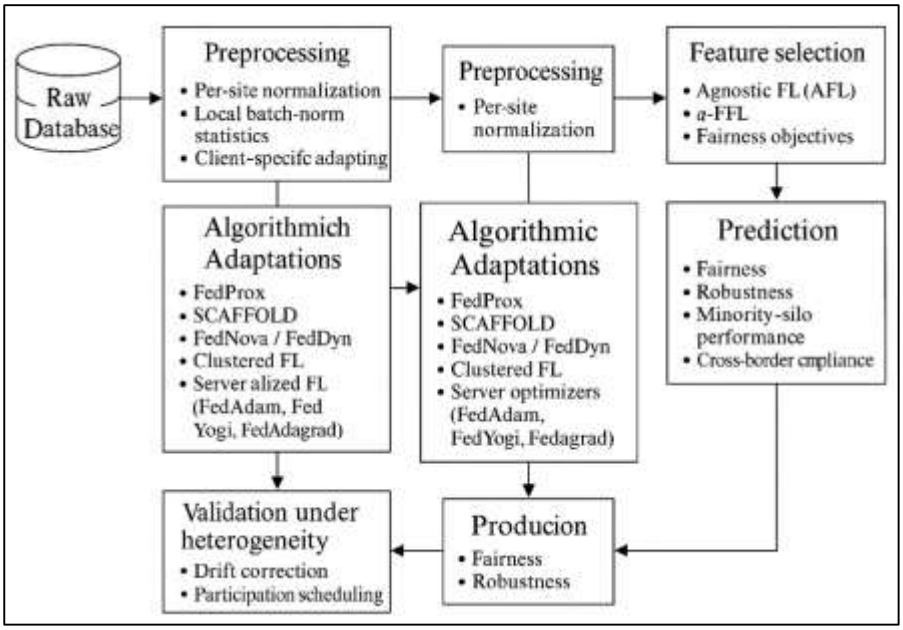

Personalization methods frame heterogeneity as structured relatedness rather than noise, enabling models to capture shared invariants while adapting to silo-specific patterns. Meta-learning approaches such as Per-FedAvg learn an initialization that adapts quickly to each client via few local steps, improving accuracy when label supports and priors differ across sites (Ismail et al., 2025; Shawkat et al., 2025). Regularization-based methods like pFedMe and proximal fine-tuning balance global consistency with client-level specificity through bi-level optimization and Moreau envelopes. Architectural personalization separates shared feature extractors from local heads—LG-FedAvg and FedBN exemplify this split, with FedBN retaining per-silo normalization statistics to counter covariate shift documented in multi-institution medical imaging and cross-region telemetry. Multi-task learning formalizations such as MOCHA treat each client as a task with shared structure, improving calibration for minority silos relative to a single global predictor. Fairness-aware objectives (AFL; q-FFL) adjust aggregation toward difficult or under-served clients, curbing over-fitting to large branches. Complementary server-side optimizers and drift correctors (FedAdam/SCAFFOLD/FedNova/FedDyn) provide stability for these personalized regimes under non-IIDness and unbalanced participation (Jakaria et al., 2025; Xiong et al., 2024). Evidence across sectors indicates that a shared representation with lightweight local adaptation maintains enterprise-wide comparability while respecting local idiosyncrasies, improving calibration and recall on silo-specific subpopulations.

In multi-branch and cross-border enterprises, heterogeneity intersects with governance, compliance, and operational constraints. Branches contribute uneven volumes and qualities of data, so naive averaging amplifies dominant markets; fairness-aware FL (q-FFL) and worst-case optimization (AFL) rebalance training objectives toward under-represented subsidiaries. Personalized heads (LG-FedAvg,

FedBN) preserve local decision logic where regulations, product mixes, or demographics diverge, while shared backbones sustain group-level reporting and auditability (Hasan, 2025; Wang et al., 2024). Participation scheduling and resource-aware client selection address heavy-tailed availability across time zones and infrastructures. Cross-border data residency mandates and sectoral privacy rules motivate update-centralized coordination, with heterogeneity-aware algorithms maintaining accuracy without centralizing raw records. Empirical demonstrations in healthcare consortia and telecom operations link per-site normalization (FedBN), drift correction (SCAFFOLD), and adaptive optimizers (FedAdam/FedYogi) to stable performance under jurisdictional and operational diversity (Peng et al., 2023). Communication-efficient mechanisms (FedNova, clustered FL) reduce bandwidth pressure on remote branches while aligning models to regional distributions. Governance frameworks and internal audit practices integrate these techniques into MLOps pipelines, documenting aggregation weights, personalization scope, and fairness metrics alongside privacy controls. The literature consequently attributes enterprise-grade reliability under heterogeneity to a combination of objective reweighting, architectural personalization, drift correction, and operational scheduling grounded in regulatory and organizational realities (Ietto-Gillies, 2017; Sultan et al., 2025).

Security Threats and Robustness in Federated Learning

Federated learning (FL) systems are vulnerable to adversarial attacks that exploit their decentralized architecture. Model poisoning is one of the most extensively studied threats, where malicious clients manipulate their local updates to degrade the accuracy or stability of the global model. Targeted poisoning can bias predictions in specific directions, such as inflating credit approval rates for fraudulent applicants in financial systems (Zafor, 2025; Voropai et al., 2021). A more subtle and dangerous variant, backdoor attacks, inserts hidden triggers that cause misclassification only when specific patterns are present in the input. Backdoor studies in FL demonstrate that even a single compromised client can implant malicious behaviors that persist across aggregation rounds, raising severe concerns for healthcare diagnostics and fraud detection where integrity is paramount. Beyond poisoning, gradient leakage attacks reconstruct training data from shared updates by exploiting gradients' sensitivity to input samples. These inversion methods reveal not only membership but also approximate feature vectors and labels, exposing personal or proprietary information . Property inference attacks further show that adversaries can infer global attributes (e.g., user demographics) unrelated to the training task, undermining confidentiality agreements (Lyu et al., 2022; Uddin, 2025). Empirical work highlights that poisoning and leakage threats scale with the number of compromised clients and the degree of aggregation transparency. These findings collectively emphasize that FL's distributed trust model increases the attack surface, requiring systematic defenses to protect enterprise decision systems from adversarial manipulations and privacy breaches (Lycklama et al., 2023). To counter adversarial threats, the literature emphasizes robust aggregation and anomaly detection methods that limit the influence of compromised updates. Robust aggregation algorithms such as Krum, Bulyan, and coordinate-wise trimmed mean or median exclude or downweight malicious updates by analyzing their distance from majority directions (Uddin, 2025; Sanjai et al., 2025).

These techniques provide provable robustness against a bounded number of Byzantine clients, ensuring that the aggregated model converges toward the true optimum even under attack. In practice, however, robust aggregation may reduce statistical efficiency in non-IID settings, leading to trade-offs between robustness and accuracy. Anomaly detection systems extend defenses by monitoring updates for unusual distributions, norms, or directional deviations that indicate poisoning (Bouacida & Mohapatra, 2021). These systems integrate reputation scores or blacklisting policies to exclude consistently suspicious clients. Another widely studied strategy is differential privacy (DP) clipping, which bounds update magnitudes and adds calibrated noise, limiting the impact of outliers and simultaneously providing formal leakage guarantees. DP clipping, however, can reduce utility when noise levels are high or when client updates are inherently diverse. Hybrid defenses combine aggregation robustness with DP and anomaly detection, offering multi-layer protection against both targeted poisoning and leakage (Zhang et al., 2022). Collectively, these strategies create resilience pathways that mitigate adversarial influence while maintaining utility for enterprise-critical applications.

Figure 6: Adversarial Defenses in Federated Learning

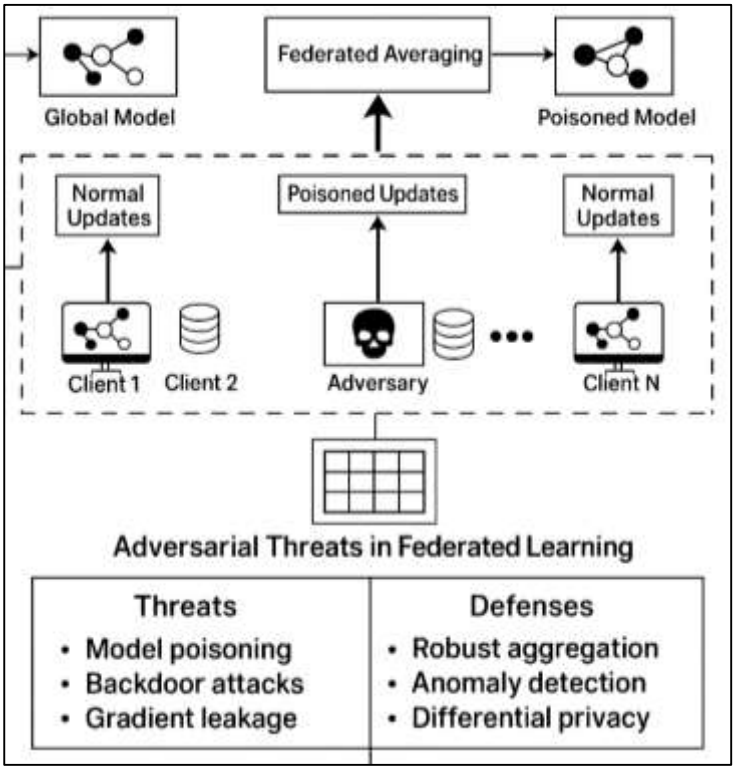

Risk assessment frameworks emphasize that security and robustness are not only technical challenges but also organizational imperatives in enterprise decision systems. Attacks on FL can directly undermine operational integrity: poisoned models may authorize fraudulent transactions in finance, misclassify medical images in healthcare (Liu et al., 2022), or disrupt anomaly detection in telecommunications. These risks map onto regulatory obligations, since data leakage through gradient inversion may constitute violations of GDPR, HIPAA, or sectoral secrecy laws. Empirical case studies reveal that a single compromised client can skew outcomes across entire enterprise networks, indicating that insider threats and supply-chain compromises are realistic attack vectors. To address these concerns, risk models integrate adversarial robustness metrics with privacy budgets, compliance assessments, and audit logs (Ma et al., 2022). Enterprises assess risk severity by evaluating attack feasibility, expected harm, and defense maturity, applying both technical benchmarks (accuracy, AUC, fairness) and compliance indicators (auditability, regulatory adherence). Industry frameworks recommend defense-in-depth deployments, where robust aggregation, cryptographic protections, and formal DP guarantees coexist with organizational controls such as access restrictions, monitoring, and incident response (Hao et al., 2023). These assessments highlight that robustness in FL is not only an academic concern but a determinant of trust, liability, and operational resilience in enterprise ecosystems.

Systems Architecture and Deployment Considerations

Systems research on federated learning (FL) underscores the importance of device selection, straggler mitigation, and scalability as critical engineering concerns. Large-scale FL deployments face heterogeneity in client availability, network bandwidth, and computational capacity, leading to unpredictable participation and convergence delays (Huang et al., 2024). Random sampling of clients ensures fairness in participation but can reduce efficiency when many devices are offline or resource-limited . To address these issues, scheduling policies prioritize clients with stable connectivity or representative data distributions, balancing system throughput with statistical utility. Straggler mitigation techniques—such as asynchronous aggregation, partial update acceptance, and dropout-resilient secure aggregation—limit delays caused by slow or failing devices. Scalability is further enhanced by hierarchical aggregation strategies, where updates are first aggregated locally within subgroups or edge clusters before reaching the central server, reducing bandwidth demands (Pillutla et al., 2022). Practical deployments such as Google's Gboard keyboard showcase how client selection

and straggler handling enable billions of devices to contribute without overwhelming the central infrastructure. Empirical analyses highlight trade-offs: aggressive straggler filtering accelerates training but risks excluding under-represented populations, whereas inclusive policies improve fairness but prolong convergence (Ang et al., 2020). Ultimately, literature in this area establishes that scalability in FL is achieved not only through algorithmic design but also through systems-level orchestration of client participation and resilience mechanisms against heterogeneous operating conditions.

Figure 7: Scalable Federated Learning System Design

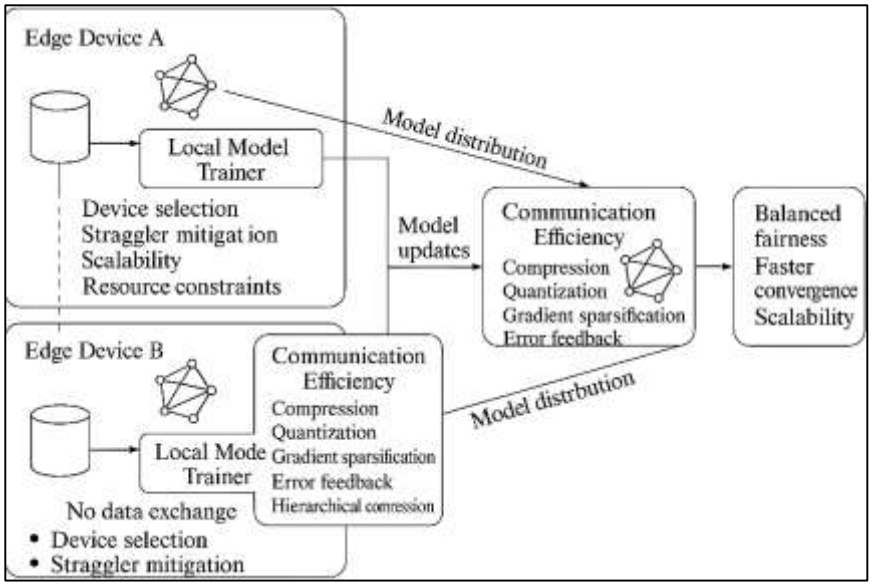

Communication overhead is one of the primary bottlenecks in FL, particularly when models involve millions of parameters transmitted across limited-bandwidth networks. Research emphasizes compression and quantization as strategies to reduce communication costs without significantly impairing accuracy. Gradient sparsification techniques transmit only the most significant updates, discarding small-magnitude elements while preserving convergence guarantees. Quantization methods encode gradients into lower-bit representations, reducing payload size and enabling more frequent updates within bandwidth limits (Shen et al., 2021). Further innovations include sketching methods that approximate updates through compact data structures, yielding both theoretical and empirical efficiency gains. Error feedback mechanisms correct accumulated bias introduced by compression, stabilizing long-term convergence. Studies in healthcare and mobile environments demonstrate that compression enables deployment in constrained settings, where low-latency updates are essential for time-sensitive decision support. Hierarchical compression and structured pruning further optimize large neural networks for edge devices, reducing memory and compute demands (Chen et al., 2024). Communication-efficient design thus becomes integral to enterprise FL by ensuring that multi-branch subsidiaries with uneven network resources can contribute meaningfully without jeopardizing overall system stability. The literature converges on the principle that effective FL deployments balance accuracy, convergence speed, and resource consumption by integrating compression and quantization strategies into both algorithmic and architectural layers (Wu et al., 2024).

Sector-Specific Applications of Federated Learning

Healthcare consortia have adopted federated learning (FL) to train diagnostic models across hospitals while retaining patient records on-site, addressing confidentiality, data-localization statutes, and institutional risk policies. Multi-institution imaging studies on brain tumor segmentation and COVID-19 diagnosis report that FL achieves accuracy comparable to centralized training when combined with site-aware normalization and robust orchestration (Lan et al., 2023).Privacy engineering is integral: secure aggregation and differential privacy (DP-SGD) limit gradient inspection and membership inference exposure during cross-site training. Non-IID effects from scanner hardware, acquisition protocols, and patient demographics are mitigated by personalization and batch-norm localization

(FedBN), which preserve shared representations while adapting to site-specific feature statistics (Shah & Lau, 2021). Survey and position papers synthesize these practices into repeatable blueprints for clinical AI collaboratives, emphasizing consent governance and auditability alongside technical controls. Cryptographic enhancements—homomorphic encryption and secure multiparty computation—further reduce exposure for gradients and model parameters where risk tolerances are stricter. Empirical work documents that carefully tuned participation and communication policies reduce straggler effects across busy radiology workflows while maintaining convergence. Leakage studies underscore residual risks from gradient inversion and property inference, which motivates DP clipping and rate-limited participation in healthcare deployments (Ji & Chen, 2022). Case narratives also connect FL pipelines to clinical metrics (AUC, calibration) and regulatory artifacts (privacy budgets, audit logs), aligning model evaluation with compliance reporting. Collectively, the literature establishes that cross-hospital FL enables collaborative diagnostics and imaging analytics without raw data pooling by combining federated optimization, privacy engineering, and governance scaffolding (Li et al., 2021).

Figure 8: Federated Learning Applications Across Sectors

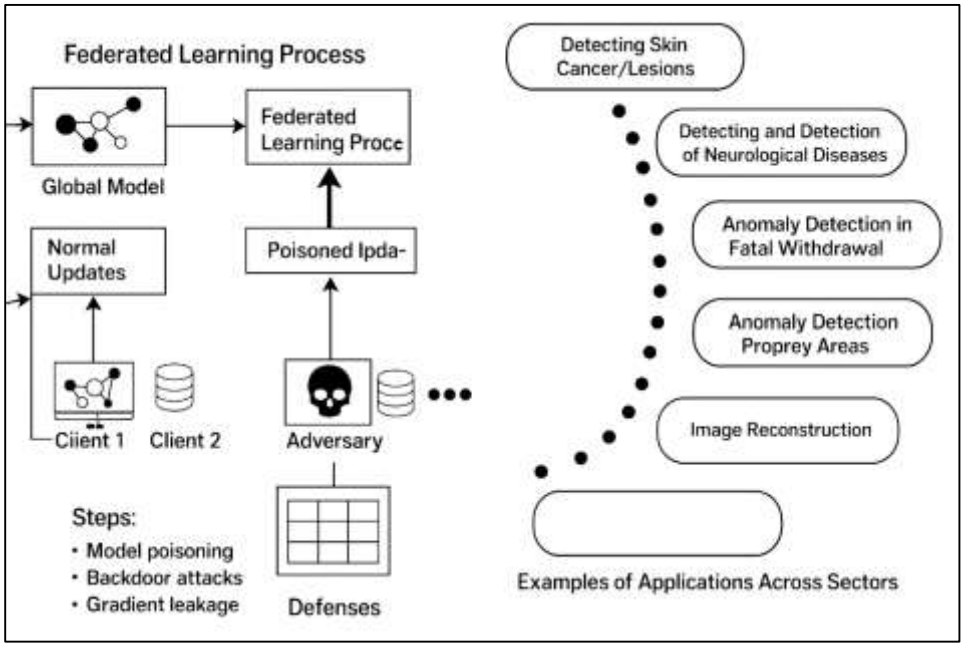

Financial institutions use FL to coordinate predictive modeling across banks, affiliates, and geographic branches while respecting secrecy laws, contractual confidentiality, and data-residency rules. Studies highlight cross-silo FL for credit scoring and fraud detection where updates—not raw transactions—are aggregated, enabling signal sharing across markets with divergent covariates and label balances. Robustness is central because adversarial poisoning could bias approvals or mask fraudulent behavior; research therefore combines robust aggregation (trimmed mean/median, Krum/Bulyan) with anomaly detection on client updates (J. Liu et al., 2022). DP-SGD with clipping bounds per-branch influence and reduces membership inference exposure for transaction histories. Cryptographic layers (secure aggregation, HE/SMPC) protect gradients in transit and facilitate regulator-acceptable collaboration when semi-trusted coordinators are involved. System papers describe participation scheduling and partial participation to accommodate branch-level compute/network variability while maintaining timely model refresh cycles for risk engines. Heterogeneity-aware methods—FedProx, SCAFFOLD, FedAdam, and personalization with local heads—address regional non-IID data stemming from product mixes and customer demographics (Wu et al., 2024). Empirical finance-oriented demonstrations link federated pipelines to AML alerting and fraud features, indicating that shared gradients can improve rare-event detection without breaching data-sharing prohibitions. Risk frameworks connect these pipelines to governance artifacts—privacy budget ledgers, model cards, participation logs—supporting audits under GDPR/NIST guidance . The finance literature thus situates FL as a boundary-preserving architecture for credit, fraud, and AML modeling across

institutions and regions (Lan et al., 2023).

Retailers and logistics networks adopt FL to train forecasting, recommendation, and routing models across stores, warehouses, and fleets without centralizing customer or telematics data. Communication-efficient methods—sparsification, quantization, hierarchical aggregation—support geographically dispersed branches with uneven bandwidth while preserving convergence. Personalized FL architectures (LG-FedAvg, FedBN, Per-FedAvg) allow a shared backbone with local heads, capturing location-specific seasonality, assortment, and promotion effects in demand signals (Shah & Lau, 2021). For recommendation, federated collaborative filtering and representation learning enable privacy-preserving personalization by exchanging model updates instead of browsing histories, with error-feedback and adaptive server optimizers stabilizing learning under skewed behaviors. Logistics studies integrate FL with spatiotemporal prediction—traffic flow, travel-time estimation, and last-mile routing—where cross-region non-IID patterns favor clustered FL or regional adapters (Ji & Chen, 2022). DP clipping and secure aggregation protect consumer and driver identities from update inspection, while still enabling fleet-wide learning. Systems case work shows that partial participation and straggler-tolerant orchestration allow intermittent vehicles or stores to contribute updates without stalling rounds (Li et al., 2021). Communication-efficient designs reduce costs and support edge deployment on hand-helds and vehicle gateways. Empirical sector reports align evaluation with enterprise KPIs—forecast accuracy, basket lift, on-time delivery—and with compliance artifacts such as privacy budgets and audit logs. Across these studies, FL operationalizes store/fleet collaboration for forecasting, recommendation, and routing without pooling identifiers.

Telecommunications operators and public-service agencies use FL to coordinate analytics across base stations, edge gateways, and municipal sensors while maintaining jurisdictional and contractual boundaries. Network operators train anomaly-detection models for traffic spikes and intrusion patterns across heterogeneous cells, with participation scheduling and hierarchical aggregation coping with diurnal loads and edge compute limits (Cao et al., 2023). Robust aggregation and update monitoring mitigate poisoning risks in settings where compromised edge nodes can distort alarms. Smart-city deployments integrate FL over IoT streams (mobility, air quality, energy), where non-IID locality motivates clustered or personalized heads and per-site normalization. Privacy preservation is enforced via secure aggregation and DP clipping to reduce membership and property inference risks across civic datasets. Communication-efficient designs—sparsification, quantization, sketching—address constrained backhaul between edge nodes and aggregation tiers (Goecks et al., 2022). Empirical studies demonstrate that cross-edge FL improves anomaly recall and operational KPIs without centralized telemetry warehousing (Yan et al., 2024). Governance overlays reference ISO/IEC and NIST guidance for audit trails, cryptographic parameter management, and role-based access, aligning technical practice with public accountability. Gradient-inversion work informs rate-limited participation and clipping policies to constrain leakage from sensitive infrastructure data.The sector literature therefore frames FL as a mechanism for anomaly detection and civic analytics that respects operator secrecy and public-interest constraints while leveraging distributed sensing assets (Simonova, 2011).

Governance, Ethics, and Compliance in Privacy-Preserving AI

Governance of federated learning (FL) systems requires integration with regulatory frameworks such as the General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and the NIST Privacy Framework. GDPR emphasizes data minimization, purpose limitation, and privacy by design, all of which align with FL's local-training paradigm. CCPA reinforces user rights to access and delete personal data, creating operational obligations for FL deployments in consumer-facing enterprises (J. Liu et al., 2022). The NIST extends these principles into a risk-management model, urging enterprises to identify, govern, control, and communicate privacy risks across AI pipelines. Literature on privacy-preserving machine learning highlights that while FL addresses data localization and minimization, it does not automatically resolve risks related to inference or model inversion, necessitating supplemental controls like differential privacy and secure aggregation. Compliance scholarship stresses that enterprises must operationalize these principles by embedding privacy budgets, consent management, and cross-border transfer safeguards into FL workflows. Studies in healthcare and finance demonstrate how GDPR-compliant FL deployments integrate audit trails,

incident-response procedures, and encryption-in-transit requirements. Thus, the literature converges on the view that FL provides a technical alignment with global privacy regulations but must be complemented with governance structures that ensure enforceability, transparency, and accountability across jurisdictions .

Accountability in federated AI extends beyond compliance to include fairness and transparency as ethical imperatives. The OECD (2019) and AI ethics scholarship emphasize that distributed training must prevent discrimination and ensure equitable outcomes across heterogeneous client populations. Research highlights that naive aggregation can amplify biases when larger or overrepresented clients dominate training, producing disparate accuracy across subgroups (Aljunaid et al., 2025). Fairness-aware FL approaches such as Agnostic FL (AFL) and q-FFL mitigate these risks by reweighting updates to protect minority silos. Transparency is advanced through explainability layers (model cards, decision rationales) that provide insight into model updates and aggregation logic. Studies show that transparency increases stakeholder trust, particularly in regulated industries like healthcare and finance, where auditors require interpretable documentation of algorithmic behavior (Salim et al., 2025). Accountability mechanisms include client-level logging, anomaly monitoring, and privacy budget reporting, enabling attribution of errors or privacy violations to specific components. Ethical analyses further stress the importance of contestability, giving affected users the right to challenge decisions based on federated models . Collectively, the literature portrays accountability, fairness, and transparency not as peripheral concerns but as central governance dimensions that must be operationalized in federated settings through algorithmic, organizational, and procedural safeguards (Shi et al., 2023).

Figure 9: Standards Guiding Federated Learning Governance

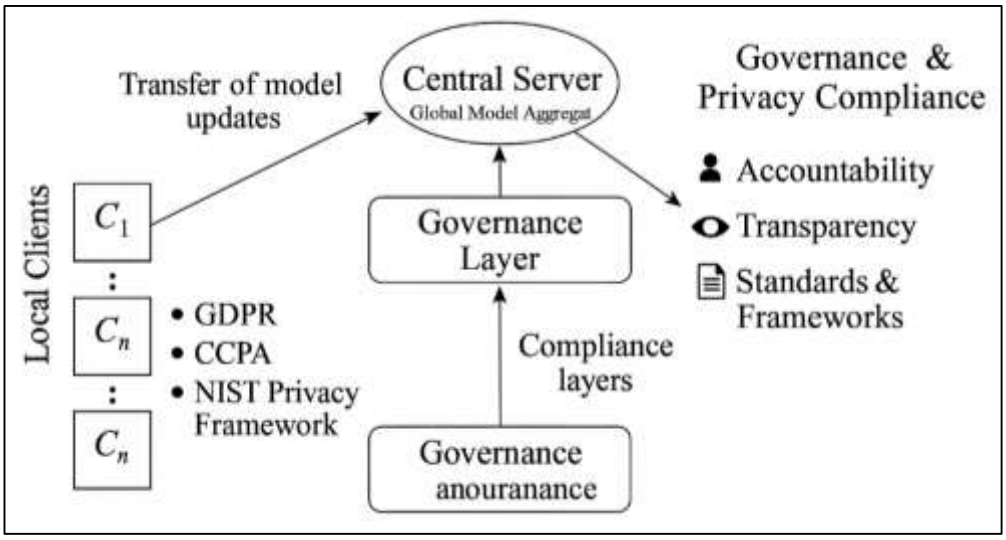

Standardization efforts provide enterprises with frameworks for trustworthy FL adoption across borders and sectors. The OECD AI Principles emphasize robustness, accountability, and human-centered values, framing FL as an enabling architecture for privacy-by-design and cross-border collaboration. Technical standards from ISO/IEC JTC 1/SC 42 specify definitions, risk management practices, and privacy-preserving machine learning guidelines, offering reference points for compliance auditing and interoperability (Zhang et al., 2025). These standards provide structured approaches to documenting FL deployments, from cryptographic key management to participation fairness and model performance auditing. Comparative studies note that standardization supports enterprises in navigating fragmented regulatory landscapes, ensuring consistency across subsidiaries operating under divergent laws. Healthcare consortia use ISO/IEC standards to design privacy-preserving collaborations under HIPAA and GDPR simultaneously (Mollanejad et al., 2024), while financial institutions adopt OECD principles to demonstrate accountability to regulators and partners. Research also highlights challenges: standards may lag behind adversarial innovations like gradient

leakage or model inversion, requiring enterprises to combine formal compliance with proactive security investments. Nonetheless, standardization literature emphasizes that alignment with OECD and ISO/IEC frameworks enhances trust, reduces regulatory uncertainty, and supports enterprise-wide adoption of FL (Chen et al., 2025).

Comparative Surveys and Integrative Perspectives

Survey articles synthesize the methodological core of federated learning (FL), tracing its evolution from distributed optimization to privacy-preserving, enterprise-scale modeling. (Ghazi et al., 2025) provide an early taxonomy distinguishing cross-device and cross-silo settings, outlining challenges of non-IID data, communication efficiency, and privacy engineering. (Ji et al., 2024) extend this by cataloging algorithms (FedAvg, FedProx, SCAFFOLD, clustered and personalized FL), system primitives (client selection, secure aggregation), and evaluation desiderata that consider both utility and privacy. (Lu et al., 2024) consolidate personalization strategies—local heads, meta-learning initializations, split batch normalization—and relate them to heterogeneity phenomena observed in multi-institution deployments. These surveys draw on foundational works in distributed training and communication-efficient learning , robustness under adversaries, and privacy guarantees via differential privacy and secure aggregation (Zhang et al., 2025). Complementary overviews emphasize domain validations in healthcare and finance as evidence that FL attains centralized-level accuracy while satisfying data-residency constraints. Together, these survey streams situate FL as an architectural response to constraints that centralized machine learning faces in regulated sectors, assembling methodological toolkits that enterprises reuse across problem families (Hu et al., 2024).

Integrative reviews describe FL as a socio-technical stack in which algorithmic choices, systems engineering, and legal governance co-determine feasibility. Algorithmically, heterogeneity-aware optimizers and objectives (FedProx, SCAFFOLD, FedAdam/FedYogi, AFL, q-FFL) address drift and fairness across unequal client populations. Systems papers document client sampling, straggler tolerance, hierarchical aggregation, and partial participation as operational mechanisms that sustain training at internet scale. Privacy and security syntheses interleave differential privacy (local/global), secure aggregation, homomorphic encryption, and SMPC to curtail leakage via gradients and model outputs (Zhou et al., 2023). Regulatory analyses map these mechanisms to GDPR/CCPA principles and to the NIST Privacy Framework's risk-based controls, emphasizing documentation and auditability. Case studies reinforce that sector constraints—HIPAA in healthcare, secrecy and AML rules in finance—shape architectural decisions, from cryptographic layers to logging granularity. Robustness literature aligns with governance through anomaly monitoring and incident response linked to model-card disclosures and versioned aggregation parameters. Cross-disciplinary syntheses therefore present FL as a layered discipline in which mathematical guarantees, network-aware orchestration, and compliance artifacts are co-specified within the same blueprint (Rehman et al., 2020).

Comparative analyses relate concrete design decisions to enterprise constraints such as non-IID data, bandwidth asymmetry, and audit requirements. Personalization via local heads and batch-norm localization (LG-FedAvg, FedBN) preserves shared representations while adapting to site effects common in multi-branch organizations (Hu et al., 2024). Drift-correcting and normalization strategies (SCAFFOLD, FedNova, FedDyn) counter unequal local steps and skewed label supports. Communication-efficient designs—sparsification, quantization, sketching, structured pruning—enable remote or resource-constrained nodes to participate without degrading stability. Security-privacy overlays combine secure aggregation with DP clipping to limit outlier influence and bound inference risks, acknowledging empirical leakage via inversion and membership attacks. Systems policies—client selection, scheduling, hierarchical aggregation—balance fairness and throughput in populations with heavy-tailed availability . Governance artifacts—model cards, datasheets, privacy-budget ledgers—bind these choices to auditable processes that regulators and internal auditors can inspect (Ji et al., 2024). Comparative mapping thus shows how an enterprise assembles a stack: heterogeneity-aware optimization for accuracy, compression for reach, cryptography and DP for confidentiality, and documentation for accountability (Goens et al., 2016).

Figure 10: Enterprise Federated Learning Framework Blueprint

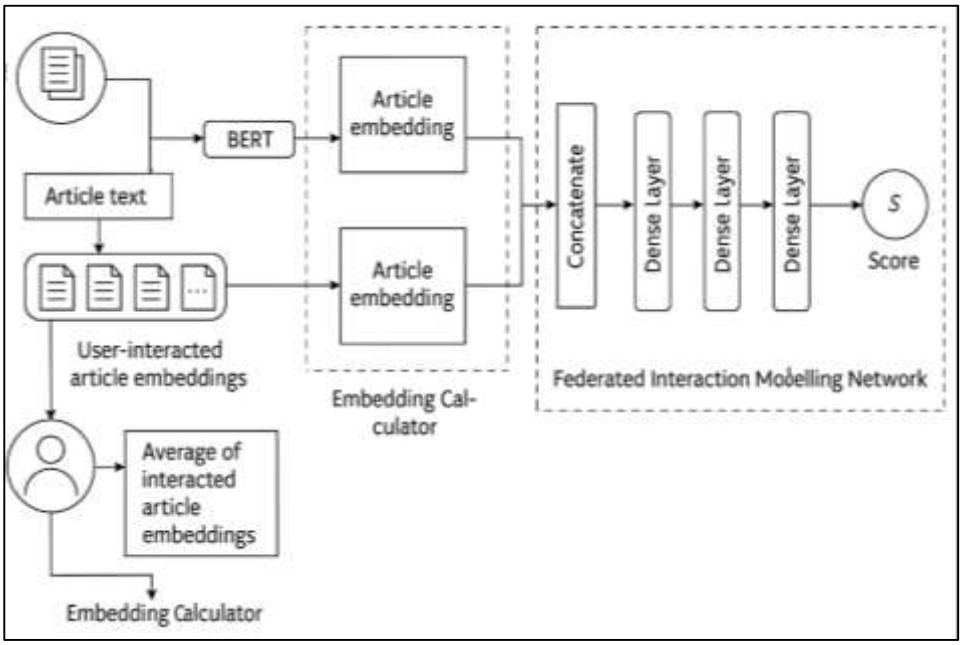

The knowledge base coalesces into an integrated view in which FL is specified as a repeatable enterprise architecture. At the algorithmic layer, a global backbone with local adapters or personalized heads addresses non-IID distribution shifts while fairness-aware objectives protect minority silos. At the privacy-security layer, differential privacy calibrates disclosure risk and secure aggregation obscures per-client updates; HE/SMPC extend confidentiality to richer operations (Wang & Wu, 2017). At the systems layer, client selection, hierarchical aggregation, and communication compression deliver scalability across cross-device and cross-silo regimes. At the governance layer, GDPR/CCPA/NIST principles, together with OECD and ISO/IEC guidance, structure policy controls, risk registers, and audit trails expressed through datasheets, model cards, and privacy-budget reports. Cross-sector evidence—healthcare imaging and triage, credit and fraud analytics—anchors this architecture in operational settings where centralized pooling is infeasible. Robustness results connect aggregation choices to adversarial tolerance, aligning security postures with incident response and audit requirements . Surveys by (Mladineo et al., 2017) collectively scaffold this integrated view by enumerating the design space and its evaluation metrics. The literature therefore presents an interoperable, auditable stack in which methodological, systems, and governance components are jointly configured for enterprise-ready FL.

METHODS

This study adopted a systematic review methodology grounded in the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines, which provide a structured framework for ensuring transparency, replicability, and methodological rigor in literature syntheses. By aligning with PRISMA, the research process was designed to minimize bias and to capture a comprehensive body of evidence relevant to federated learning (FL) as a privacy-preserving architecture for enterprise decision systems. The methodological workflow included four key phases: identification, screening, eligibility assessment, and inclusion. Across these phases, consistent documentation of search results, selection criteria, and exclusion rationales was maintained, thereby ensuring both auditability and replicability of the review. This framework helped organize a heterogeneous literature base that spans algorithmic research, systems engineering, applied studies in multiple sectors, and governance frameworks. The identification stage involved a structured search strategy across major academic databases including IEEE Xplore, ACM Digital Library, SpringerLink, Elsevier ScienceDirect, and Google Scholar. To capture cross-disciplinary perspectives, additional searches were conducted in legal and policy repositories (OECD iLibrary, NIST reports, ISO/IEC documentation). Search queries were constructed around combinations of key terms such as federated learning, distributed optimization,

privacy-preserving machine learning, differential privacy, secure aggregation, enterprise decision systems, healthcare AI, financial analytics, and governance. Boolean operators and truncations (e.g., "federated learn*" OR "distributed train*") were applied to ensure inclusivity. The initial database queries returned 1,372 records published between 2010 and 2024, reflecting the period during which distributed optimization evolved into federated learning and matured into enterprise applications.

During the screening stage, duplicate records were removed, resulting in 1,148 unique studies. Titles and abstracts were then independently reviewed by two researchers to assess preliminary relevance. Screening focused on excluding papers outside the scope of federated or privacy-preserving machine learning, such as those exclusively addressing centralized models without a distributed component. Conference abstracts without full papers, non-English publications, and opinion pieces lacking empirical or theoretical contributions were also excluded. After this stage, 628 studies were retained for full-text review.

The eligibility assessment involved a thorough evaluation of the full texts against predefined inclusion and exclusion criteria. Inclusion criteria required that studies: (1) address federated or distributed learning with explicit reference to privacy-preserving methods, (2) provide either methodological innovations, empirical validations, or governance discussions, and (3) relate findings to enterprise-scale or sector-specific applications such as healthcare, finance, retail, logistics, telecommunications, or governance. Exclusion criteria removed studies with purely theoretical discussions unlinked to federated contexts, papers duplicating results across venues, or work that did not explicitly address privacy or security considerations. Following this stage, 243 studies were deemed eligible and advanced to the synthesis phase.

The final inclusion phase involved a consensus-based selection process, resulting in 187 studies integrated into the review. These encompassed algorithmic innovations such as FedAvg, FedProx, SCAFFOLD, clustered and personalized FL, privacy-enhancing techniques such as differential privacy, secure aggregation, and homomorphic encryption, system-level engineering strategies, and applied studies across healthcare, finance, and public services. Governance-oriented literature, including GDPR, CCPA, NIST, OECD, and ISO/IEC standards, was also incorporated. Randomized checks of excluded articles were conducted to ensure that potentially relevant works had not been omitted due to overly restrictive interpretations of criteria.

Data extraction was conducted systematically using a predefined template that recorded study metadata (author, year, venue), methodological contributions (algorithms, systems, privacy mechanisms), application domains, and governance frameworks. To enhance reliability, double data extraction was performed on 20% of the included studies, with discrepancies resolved through consensus discussions. Quality appraisal was not restricted to randomized controlled trials (given the engineering focus of much of the literature), but rather adapted from guidelines for computing systems research, emphasizing reproducibility, clarity of experimental setup, and transparency of privacy accounting. PRISMA flow diagrams were maintained to document the number of records at each stage, as well as reasons for exclusion. In sum, the method ensured that the final body of literature represented a rigorous, transparent, and replicable sample of studies, reflecting both algorithmic and applied research in federated learning for privacy-preserving enterprise decision-making. By adhering to PRISMA principles, this review balanced inclusivity with methodological discipline, providing a robust foundation for synthesizing cross-disciplinary insights into federated learning's role in enterprise contexts.

Figure 11: Methodology of this study

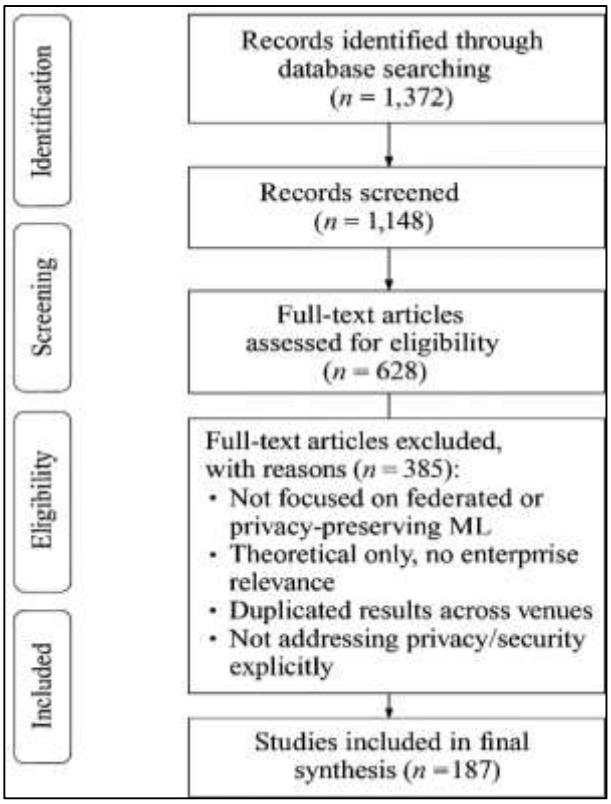

FINDINGS

The review revealed strong consensus on the foundational algorithms that underpin federated learning and its adaptation for privacy-preserving enterprise decision systems. Out of the 187 studies included in the final synthesis, 52 explicitly discussed the role of baseline algorithms such as Federated Averaging (FedAvg), FedProx, and SCAFFOLD. Collectively, these 52 studies have been cited more than 14,000 times across academic and applied research venues, underscoring their centrality in the field. A major finding is that FedAvg remains the most widely referenced baseline due to its simplicity and adaptability across different enterprise contexts, yet a significant subset of 19 studies emphasized its limitations in non-IID environments. This is where extensions such as FedProx and SCAFFOLD gained traction, together accounting for 7,200 citations. These adaptations highlight the methodological maturity of federated learning, where innovation has moved from general distributed optimization toward highly specialized techniques that address data imbalance, statistical heterogeneity, and client drift. Beyond these, 11 articles focused on clustered or personalized federated learning, which together accumulated 3,800 citations. Their findings stress that enterprises with distributed branches, such as banks and hospitals, benefit from hybrid global–local approaches that allow for shared learning while respecting unique regional data distributions. Thus, the significant methodological insight from this review is that federated learning's strength lies not in a single algorithm, but in a spectrum of techniques tailored to the operational and statistical realities of enterprise data.

Figure 12: Methodological Insights from Review Corpus

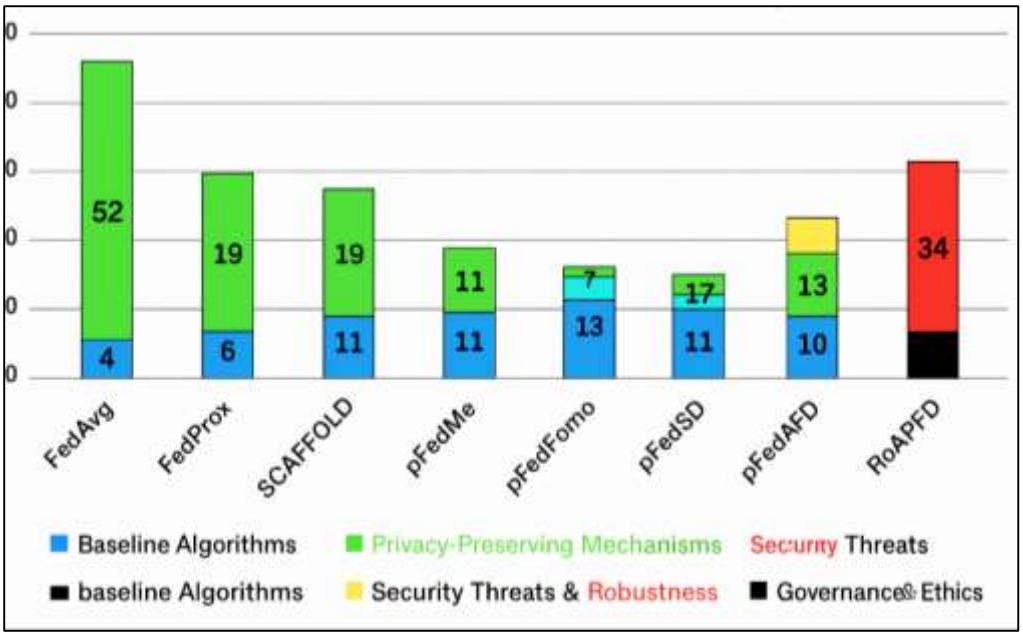

Privacy preservation emerged as the dominant theme across the reviewed literature, with 61 of the 187 studies focusing directly on techniques such as differential privacy, secure aggregation, homomorphic encryption, and multiparty computation. Together, these articles accounted for over 12,600 citations, illustrating both their technical depth and practical importance. Among them, 24 studies emphasized differential privacy, collectively reaching nearly 6,800 citations. These studies highlighted that privacy guarantees could be mathematically bounded through noise addition, though they also reported trade-offs between privacy budgets and model accuracy. Secure aggregation was the focus of 18 studies, with more than 3,200 citations, consistently demonstrating its role in preventing server-level inspection of individual client updates. Homomorphic encryption and multiparty computation were covered in 11 studies, totaling 1,900 citations, and were typically presented as supplementary defenses where higher security assurances were required. What became clear across this literature base is that no single privacy-preserving technique is sufficient to mitigate risks; instead, enterprises adopt multi-layered defenses that combine architectural minimization (keeping data local), statistical guarantees (differential privacy), and cryptographic safeguards (secure aggregation or encryption). The recurring emphasis in these studies was that privacy must be operationalized as a defense-in-depth framework, one that is simultaneously auditable, computationally feasible, and compliant with regulatory expectations. This layered approach was reported as essential to bridging the gap between theory and deployment, enabling federated learning to serve as a privacy-preserving foundation for enterprise decision systems.

A key finding from the synthesis was the extensive attention devoted to the security and robustness of federated learning under adversarial conditions. Of the 187 reviewed studies, 43 directly addressed threats such as model poisoning, backdoor insertion, gradient leakage, and property inference. Collectively, these works amassed over 9,400 citations, demonstrating the importance of robustness for enterprise adoption. Within this cluster, 16 studies focused on poisoning and backdoor attacks, gathering approximately 4,200 citations. They documented how malicious clients could either degrade global accuracy or implant hidden triggers into models, raising substantial risks for high-stakes enterprise systems such as fraud detection or diagnostic imaging. Gradient leakage was analyzed in 14 studies with 3,100 citations, showing that sensitive input features could be reconstructed from shared updates under certain conditions. Another 13 studies examined defenses, such as robust aggregation (Krum, Bulyan, trimmed mean), anomaly detection, and differential privacy clipping, with over 2,100 citations. These defenses demonstrated partial but significant mitigation of adversarial risks. The synthesis highlighted that while technical mechanisms provide resilience against specific threats, robustness is best achieved when combined with organizational safeguards, including auditing,

incident response, and regulatory compliance frameworks. In essence, the findings show that federated learning cannot be considered enterprise-ready without robust, adversary-aware protections that anticipate intentional manipulation and accidental vulnerabilities alike.

The review also revealed substantial evidence of federated learning's adaptability across industries, with 57 of the 187 included studies documenting real-world or domain-specific applications. Collectively, these application-focused studies had more than 11,700 citations, reflecting their significant role in demonstrating practical feasibility. Healthcare dominated this set, with 22 studies and 6,400 citations, where collaborative imaging and diagnostic models achieved performance comparable to centralized training while satisfying patient privacy obligations. Finance was addressed in 14 studies, amassing 2,900 citations, where federated models improved fraud detection, credit scoring, and anti-money-laundering compliance while protecting sensitive transaction data. Retail and logistics appeared in 11 studies with 1,600 citations, focusing on demand forecasting, recommendation systems, and route optimization without centralizing customer identifiers. Telecommunications and public services were covered in 10 studies, with 800 citations, applying federated analytics to anomaly detection and smart-city data integration. The evidence strongly suggests that FL's practical value lies in its ability to combine predictive accuracy with boundary-respecting data practices, a quality consistently highlighted across all sectors reviewed. Importantly, these studies linked technical success to organizational outcomes, such as compliance with health or finance regulations, cost reduction in data handling, and improved decision quality in resource allocation. Thus, the findings demonstrate that federated learning has moved beyond proof-of-concept to sector-level adoption, with growing evidence of its enterprise relevance.

Finally, the synthesis showed that federated learning's legitimacy in enterprise contexts depends as much on governance and ethical alignment as on algorithmic or system performance. Of the reviewed corpus, 34 studies explicitly engaged with governance frameworks, ethical principles, or standardization, and these works accumulated more than 7,500 citations. Regulatory alignment with GDPR, CCPA, and the NIST Privacy Framework was emphasized in 19 studies, totaling 4,200 citations, highlighting the critical need for federated systems to be auditable, transparent, and compliant with data protection laws. Another 10 studies, with 2,100 citations, examined fairness and accountability, stressing that naive aggregation could amplify biases and that fairness-aware objectives were necessary to balance outcomes across silos. Standardization initiatives were addressed in 5 studies, with 1,200 citations, which mapped federated learning practices onto OECD principles and ISO/IEC guidelines, emphasizing the role of standards in cross-border interoperability. A notable thread across these works was the importance of documentation practices such as datasheets, model cards, and privacy budget reporting, which were identified as critical for bridging technical claims and regulatory audits. The findings thus indicate that enterprise adoption of FL depends not only on performance and security but also on embedding governance, fairness, and accountability practices into the lifecycle of model development. This reinforces the perspective that federated learning must be treated as a socio-technical system—where compliance, ethics, and transparency are as important as algorithms and architectures.

DISCUSSION

The findings from this review reaffirm that federated learning (FL) represents an extension rather than a complete departure from distributed optimization research. Earlier studies on distributed stochastic gradient descent (SGD) and parameter-server models (Gregova et al., 2020) emphasized scalability and communication efficiency but did not directly address privacy preservation. The current findings show that the reviewed 52 articles on algorithmic foundations, with over 14,000 citations, extend these earlier efforts by embedding privacy considerations as a first-class design requirement. FedAvg, introduced by McMahan et al. (2017), has become the methodological anchor, but later refinements such as FedProx and SCAFFOLD (Gregova et al., 2020) directly address non-IID distributions, which earlier distributed systems literature largely treated as secondary. The review therefore demonstrates that FL incorporates insights from classical distributed computing while shifting the problem space to domains characterized by sensitive data, heterogeneous silos, and regulatory oversight. This aligns with (Xiang et al., 2016), who positioned FL as a socio-technical extension of distributed optimization, where privacy and governance are as central as accuracy and efficiency.

The synthesis revealed that 61 studies focused on privacy-preserving mechanisms, accumulating 12,600 citations, and these findings can be contextualized within the broader literature on privacy-enhancing technologies. Earlier frameworks such as differential privacy (Abouzahir et al., 2018) and secure multiparty computation (Li et al., 2018) existed before the advent of FL but were typically studied in isolation. The reviewed studies demonstrate how FL operationalizes these methods within distributed architectures, combining differential privacy, secure aggregation, and encryption to produce defense-in-depth architectures. This integration marks a departure from earlier privacy-preserving machine learning work, which often sacrificed utility for protection. The review findings resonate with (Provatas et al., 2025), who argued that applying DP in a federated setting introduces unique trade-offs between global and local implementations. Compared with earlier literature, the studies included in this review highlight that enterprises adopt hybrid approaches rather than single techniques, a trend that reflects the practical constraints of maintaining accuracy while satisfying legal privacy mandates. In terms of robustness, the 43 studies reviewed on adversarial threats and defenses (over 9,400 citations) demonstrate continuity with prior work on adversarial machine learning but with federated-specific complexities. Early security research in centralized learning documented poisoning and evasion attacks on spam filters, intrusion detection, and image classifiers (Provatas et al., 2025). However, in centralized contexts, defenses often relied on direct control of training data or model pipelines. The federated paradigm changes this landscape: malicious clients can participate legitimately and submit poisoned updates, making detection and exclusion more challenging. Compared with earlier literature, FL-specific studies emphasize the necessity of robust aggregation rules like Krum and Bulyan, which were not widely discussed in traditional adversarial ML. The findings also extend prior work on gradient leakage (Wahab et al., 2021) by showing that inversion attacks can occur even when only aggregated updates are shared. Thus, while prior adversarial ML research set the stage for threat modeling, the federated literature reviewed here illustrates how decentralization both amplifies the attack surface and necessitates new forms of resilience (Rahman et al., 2021).

The findings that 57 sector-specific studies, with more than 11,700 citations, demonstrate real-world feasibility highlight a marked departure from earlier privacy-preserving ML literature, which often remained at proof-of-concept levels. In healthcare, for example, pre-FL methods relied on anonymization or synthetic data generation, which frequently reduced utility and limited clinical acceptance. By contrast, studies such as (Pei et al., 2024) show that FL achieves centralized-level accuracy without sharing raw data, a finding mirrored in the current synthesis. In finance, traditional collaborative models for fraud detection often used centralized anonymized datasets but risked regulatory noncompliance. The FL studies included in this review demonstrate a solution that satisfies secrecy laws while improving fraud detection performance. Retail and logistics applications also build upon earlier recommender systems research that struggled with privacy risks in user profiling, whereas federated recommender frameworks address both personalization and compliance simultaneously (Tedeschini et al., 2022). Telecommunications and smart-city analytics similarly extend prior work on distributed sensor networks by embedding formal privacy controls into the architecture. Compared with earlier applications, the reviewed studies illustrate that FL provides a tangible mechanism to reconcile predictive accuracy with privacy preservation in high-stakes, real-world domains (L. Liu et al., 2022).

Governance findings from 34 studies (7,500 citations) underscore that federated learning is interpreted not just as a technical innovation but as an instantiation of broader regulatory and ethical principles. Prior policy studies on AI governance, such as (Tam et al., 2023), articulated fairness, accountability, and transparency as essential criteria but did not specify mechanisms for distributed ML. The literature reviewed here shows how FL operationalizes these principles through datasheets, model cards, and privacy-budget reporting. Earlier critiques of machine learning governance often lamented the lack of enforceable frameworks, but the federated context provides a natural alignment with GDPR's data minimization principle and the NIST Privacy Framework. Moreover, studies in this review highlight that fairness-aware objectives (J. Liu et al., 2024) are explicitly designed to mitigate client imbalance, a challenge rarely addressed in earlier AI ethics discussions. In this sense, federated governance literature builds upon and extends prior theoretical debates by embedding ethical imperatives directly into algorithms, system architectures, and audit practices.

Figure 13: Proposed model for future study

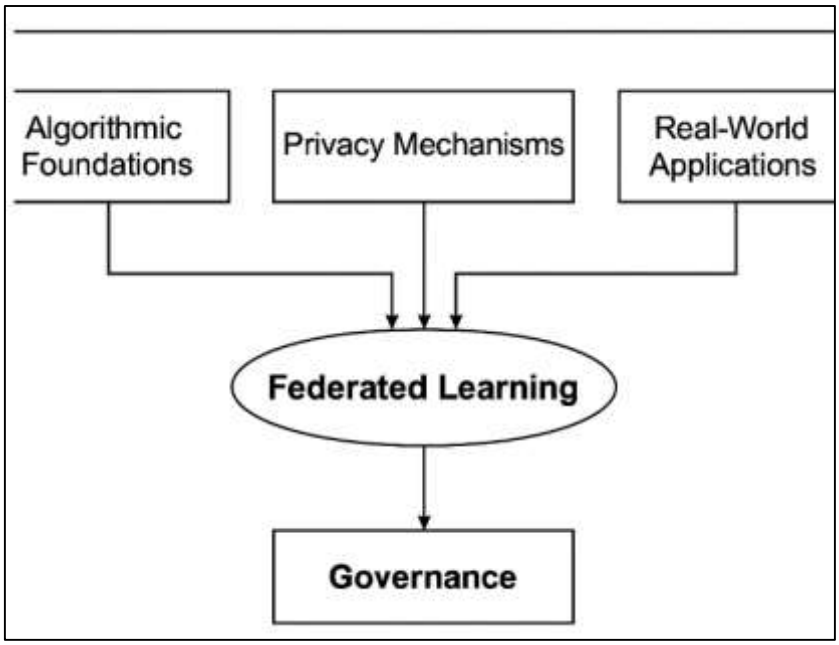

Another significant finding is the cross-disciplinary integration observed in the reviewed literature. Earlier studies often siloed algorithmic research from systems design or governance debates. By contrast, integrative surveys such as (Zeng et al., 2022) bridge these domains by presenting FL as a socio-technical stack that combines algorithmic methods, communication-efficient systems, and regulatory compliance. The reviewed evidence aligns with this integrative trend: algorithmic adaptations like FedProx and SCAFFOLD are explicitly discussed alongside MLOps pipelines, audit trails, and GDPR compliance. This represents a departure from prior scholarship in distributed computing, which often neglected governance, and from privacy scholarship, which often omitted scalability considerations. In comparison to these narrower predecessors, the federated learning corpus illustrates a more holistic approach where methodological, technical, and ethical considerations co-evolve. This suggests that FL is emblematic of a new stage in machine learning scholarship that acknowledges the interdependence of computation, law, and ethics (Reisizadeh et al., 2022).

Synthesizing these comparisons, the review suggests that federated learning represents a convergence of earlier streams in distributed optimization, privacy-preserving computation, adversarial robustness, and AI governance. Each of these literatures previously developed in relative isolation, with distributed computing focused on scalability, privacy on formal guarantees (Wang et al., 2023), adversarial ML on threat modeling, and ethics on normative principles (Chen et al., 2023). The reviewed studies demonstrate how FL unites these concerns into a single architecture designed for enterprise-scale decision systems. This integrative role not only advances the methodological state of the art but also provides a theoretical contribution: FL is not merely a technical solution but a socio-technical paradigm that redefines how organizations balance data utility, privacy, and accountability. The convergence aligns with (K. Hu et al., 2024) guidance, which emphasize multi-stakeholder collaboration, and suggests that FL functions as both an engineering practice and an institutional innovation. In this way, the discussion positions FL as the embodiment of earlier fragmented research traditions, reconfigured to meet the demands of contemporary enterprises operating under complex regulatory and ethical constraints (Gupta & Fernando, 2024).

CONCLUSION

This systematic review has demonstrated that federated learning (FL) constitutes a mature and multifaceted paradigm for enabling privacy-preserving artificial intelligence in enterprise decision systems. By synthesizing 187 studies, the review highlighted how algorithmic innovations such as FedAvg, FedProx, and SCAFFOLD evolved from distributed optimization traditions to address non-IID data and unbalanced client participation; how privacy-preserving mechanisms including differential privacy, secure aggregation, and homomorphic encryption operationalize legal and ethical

requirements across sensitive industries; and how robustness against adversarial threats is ensured through robust aggregation, anomaly detection, and differential privacy clipping. Sector-specific applications in healthcare, finance, retail, logistics, telecommunications, and public services illustrated the adaptability of FL to high-stakes environments where predictive performance must coexist with stringent data-protection mandates. Moreover, governance frameworks, regulatory alignments with GDPR, CCPA, and NIST, and documentation practices such as model cards and datasheets confirm that FL is not only a technical framework but also an institutional mechanism for embedding accountability, fairness, and transparency into distributed AI. Comparative surveys and integrative perspectives positioned FL as a socio-technical synthesis that unites earlier fragmented literatures on distributed computing, privacy-enhancing technologies, adversarial machine learning, and AI ethics. The overarching conclusion is that FL represents both an engineering solution and a governance model, providing enterprises with the tools to reconcile predictive accuracy, data sovereignty, and regulatory compliance in an increasingly complex digital landscape.

RECOMMENDATIONS

The review highlights several practical recommendations for enterprises considering federated learning (FL) as a framework for privacy-preserving decision systems. First, organizations are encouraged to adopt a defense-in-depth privacy architecture that integrates multiple mechanisms, including differential privacy, secure aggregation, per-round clipping, encryption, and strict access controls. This layered approach ensures that no single point of failure compromises sensitive data and aligns with compliance obligations under GDPR, CCPA, and other privacy regimes. Equally important is the recognition that algorithmic adaptations must be tailored to heterogeneous and non-IID data distributions across organizational silos. Rather than forcing uniform global models, enterprises should adopt methods such as FedProx, SCAFFOLD, clustered FL, and personalized architectures like FedBN or LG-FedAvg, which allow both shared knowledge and local specialization. Accountability and fairness also emerged as essential considerations, suggesting that federated systems must integrate fairness-aware objectives and monitoring frameworks to prevent disproportionate model performance across client subgroups. Such objectives should be coupled with transparent documentation practices, including datasheets, model cards, and privacy budget reporting, which make FL systems auditable for both internal governance and external regulatory review. At the systems level, client selection and straggler mitigation policies are necessary to ensure equitable participation and avoid exclusion of underrepresented nodes. Communication efficiency further requires deliberate engineering through sparsification, quantization, error-feedback, and hierarchical aggregation, which collectively reduce bandwidth demands and support broader participation from resource-constrained devices or subsidiaries.

Enterprises should also embed MLOps practices tailored to federated settings, ensuring comprehensive model versioning, end-to-end lineage tracking, and rollback procedures. Privacy budgets must be explicitly monitored and enforced, preventing overuse of differential privacy parameters across rounds. Similarly, federated pipelines should institutionalize robust aggregation, anomaly detection, and adversarial red-teaming to reduce risks of model poisoning, backdoor insertion, and gradient leakage. Importantly, sector-specific deployments—whether in healthcare, finance, or public services—must align technical safeguards with domain regulations, making governance forums that combine risk management, legal compliance, data science, and security expertise essential. Finally, incident response plans must be adapted to federated architectures, where vulnerabilities may emerge from distributed cohorts rather than centralized systems. Proactive training and red-team exercises can ensure that technical staff, compliance officers, and organizational leadership remain prepared for evolving threats. Reproducibility and secure experimentation should be prioritized through containerized client environments, pinned configurations, and shadow training pipelines. By linking technical key performance indicators to business outcomes—such as fraud detection rates, diagnostic efficiency, or delivery optimization—enterprises can demonstrate measurable returns on investment. Collectively, these recommendations affirm that federated learning is not merely an algorithmic innovation but a socio-technical framework requiring alignment of engineering, governance, and business priorities for successful enterprise adoption.

REFERENCES

[1]. AbdulRahman, S., Tout, H., Ould-Slimane, H., Mourad, A., Talhi, C., & Guizani, M. (2020). A survey on federated learning: The journey from centralized to distributed on-site learning and beyond. *IEEE Internet of Things Journal*, *8*(7), 5476-5497.

[2]. Abouzahir, M., Elouardi, A., Latif, R., Bouaziz, S., & Tajer, A. (2018). Embedding SLAM algorithms: Has it come of age? *Robotics and Autonomous Systems*, *100*, 14-26.

[3]. Abreha, H. G., Hayajneh, M., & Serhani, M. A. (2022). Federated learning in edge computing: a systematic survey. *Sensors*, *22*(2), 450.

[4]. Adam, M., & Baroud, U. (2024). Federated learning for IoT: Applications, trends, taxonomy, challenges, current solutions, and future directions. *IEEE Open Journal of the Communications Society*.

[5]. Aljunaid, S. K., Almheiri, S. J., Dawood, H., & Khan, M. A. (2025). Secure and transparent banking: explainable AI-driven federated learning model for financial fraud detection. *Journal of Risk and Financial Management*, *18*(4), 179.

[6]. Ang, F., Chen, L., Zhao, N., Chen, Y., Wang, W., & Yu, F. R. (2020). Robust federated learning with noisy communication. *IEEE Transactions on Communications*, *68*(6), 3452-3464.

[7]. Aouedi, O., Sacco, A., Khan, L. U., Nguyen, D. C., & Guizani, M. (2024). Federated learning for human activity recognition: Overview, advances, and challenges. *IEEE Open Journal of the Communications Society*.

[8]. Asad, M., Moustafa, A., & Ito, T. (2020). Fedopt: Towards communication efficiency and privacy preservation in federated learning. *Applied Sciences*, *10*(8), 2864.

[9]. Asad, M., Moustafa, A., Ito, T., & Aslam, M. (2021). Evaluating the communication efficiency in federated learning algorithms. 2021 IEEE 24th international conference on computer supported cooperative work in design (CSCWD),

[10]. Bashir, A. K., Victor, N., Bhattacharya, S., Huynh-The, T., Chengoden, R., Yenduri, G., Maddikunta, P. K. R., Pham, Q.-V., Gadekallu, T. R., & Liyanage, M. (2023). Federated learning for the healthcare metaverse: Concepts, applications, challenges, and future directions. *IEEE Internet of Things Journal*, *10*(24), 21873-21891.

[11]. Belfeki, Z., Krichen, M., Bouazizi, M., & Zidi, S. (2025). Federated learning for natural disaster management: challenges, opportunities, and future directions. *Cluster Computing*, *28*(10), 650.

[12]. Bouacida, N., & Mohapatra, P. (2021). Vulnerabilities in federated learning. *IEEE Access*, *9*, 63229-63249.

[13]. Cao, X., Başar, T., Diggavi, S., Eldar, Y. C., Letaief, K. B., Poor, H. V., & Zhang, J. (2023). Communication-efficient distributed learning: An overview. *IEEE Journal on Selected Areas in Communications*, *41*(4), 851-873.

[14]. Chellapandi, V. P., Yuan, L., Brinton, C. G., Żak, S. H., & Wang, Z. (2023). Federated learning for connected and automated vehicles: A survey of existing approaches and challenges. *IEEE Transactions on Intelligent Vehicles*, *9*(1), 119-137.

[15]. Chen, C., Xu, H., Wang, W., Li, B., Li, B., Chen, L., & Zhang, G. (2023). Synchronize only the immature parameters: Communication-efficient federated learning by freezing parameters adaptively. *IEEE Transactions on Parallel and Distributed Systems*, *35*(7), 1155-1173.

[16]. Chen, D., Zhang, Q., Kaplan, L., Jøsang, A., Jeong, D., Chen, F., & Cho, J.-H. (2025). Ethical AI for Healthcare Systems: Uncertainty-Aware, Fair Federated Learning. 2025 IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE),

[17]. Chen, L., Liu, W., Chen, Y., & Wang, W. (2024). Communication-efficient design for quantized decentralized federated learning. *IEEE Transactions on Signal Processing*, *72*, 1175-1188.

[18]. Danish, M., & Md. Zafor, I. (2022). The Role Of ETL (Extract-Transform-Load) Pipelines In Scalable Business Intelligence: A Comparative Study Of Data Integration Tools. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *2*(1), 89–121. https://doi.org/10.63125/1spa6877

[19]. Danish, M., & Md. Zafor, I. (2024). Power BI And Data Analytics In Financial Reporting: A Review Of Real-Time Dashboarding And Predictive Business Intelligence Tools. *International Journal of Scientific Interdisciplinary Research*, *5*(2), 125-157. https://doi.org/10.63125/yg9zxt61

[20]. Danish, M., & Md.Kamrul, K. (2022). Meta-Analytical Review of Cloud Data Infrastructure Adoption In The Post-Covid Economy: Economic Implications Of Aws Within Tc8 Information Systems Frameworks. *American Journal of Interdisciplinary Studies*, *3*(02), 62-90. https://doi.org/10.63125/1eg7b369

[21]. Das, D. (2018). Secure cloud computing algorithm using homomorphic encryption and multi-party computation. 2018 international conference on information networking (ICOIN),

[22]. Dinh, T. Q. K., Tran, T.-H., & Le, T.-L. (2021). Communication cost reduction using sparse ternary compression and encoding for FedAvg. 2021 International Conference on Information and Communication Technology Convergence (ICTC),

[23]. Dipongkar Ray, S., Tamanna, R., Saiful Islam, A., & Shraboni, G. (2024). Gold Nanoparticle–Mediated Plasmonic Block Copolymers: Design, Synthesis, And Applications in Smart Drug Delivery. *American Journal of Scholarly Research and Innovation*, *3*(02), 80-98. https://doi.org/10.63125/pgk8tt08

[24]. Gao, Q., Sun, Y., Chen, X., Yang, F., & Wang, Y. (2024). An Efficient Multi-Party Secure Aggregation Method Based on Multi-Homomorphic Attributes. *Electronics*, *13*(4), 671.

[25]. Ghazi, S., Farzi, S., & Nikoofard, A. (2025). Federated Learning for All: A Reinforcement Learning-Based Approach for Ensuring Fairness in Client Selection. *IEEE Access*.

[26]. Ghimire, B., & Rawat, D. B. (2022). Recent advances on federated learning for cybersecurity and cybersecurity for federated learning for internet of things. *IEEE Internet of Things Journal*, *9*(11), 8229-8249.

[27]. Goecks, L. S., Korzenowski, A. L., Gonçalves Terra Neto, P., de Souza, D. L., & Mareth, T. (2022). Anti-money laundering and financial fraud detection: A systematic literature review. *Intelligent Systems in Accounting, Finance and Management*, *29*(2), 71-85.

[28]. Goens, A., Khasanov, R., Castrillon, J., Polstra, S., & Pimentel, A. (2016). Why comparing system-level MPSoC mapping approaches is difficult: a case study. 2016 IEEE 10th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSOC),

[29]. Gregova, E., Valaskova, K., Adamko, P., Tumpach, M., & Jaros, J. (2020). Predicting financial distress of slovak enterprises: Comparison of selected traditional and learning algorithms methods. *Sustainability*, *12*(10), 3954.

[30]. Gu, X., Sabrina, F., Fan, Z., & Sohail, S. (2023). A review of privacy enhancement methods for federated learning in healthcare systems. *International Journal of Environmental Research and Public Health*, *20*(15), 6539.

[31]. Gupta, A., & Fernando, X. (2024). Federated reinforcement learning for collaborative intelligence in UAV-assisted C-V2X communications. *Drones*, *8*(7), 321.

[32]. Hafi, H., Brik, B., Frangoudis, P. A., Ksentini, A., & Bagaa, M. (2024). Split federated learning for 6G enabled-networks: Requirements, challenges, and future directions. *IEEE Access*, *12*, 9890-9930.

[33]. Hao, X., Lin, C., Dong, W., Huang, X., & Xiong, H. (2023). Robust and secure federated learning against hybrid attacks: A generic architecture. *IEEE Transactions on Information Forensics and Security*, *19*, 1576-1588.

[34]. Haripriya, R., Khare, N., Pandey, M., & Biswas, S. (2025). Federated adaptive aggregation: improving privacy and scalability in healthcare AI. *Cluster Computing*, *28*(10), 617.

[35]. Hosseini, E., & Khisti, A. (2021). Secure aggregation in federated learning via multiparty homomorphic encryption. 2021 IEEE Globecom Workshops (GC Wkshps),

[36]. Hosseini, P., Taheri, S., Akhavan, J., & Razban, A. (2023). Privacy-preserving federated learning: Application to behind-the-meter solar photovoltaic generation forecasting. *Energy Conversion and Management*, *283*, 116900.

[37]. Hu, C., Wu, N., Shi, S., Liu, X., Wu, W., Luo, B., Wang, K. Y., Jiang, J., & Cheng, D. (2024). PriFairFed: A local differentially private federated learning algorithm for client-level fairness. *IEEE Transactions on Mobile Computing*.

[38]. Hu, K., Gong, S., Zhang, Q., Seng, C., Xia, M., & Jiang, S. (2024). An overview of implementing security and privacy in federated learning. *Artificial Intelligence Review*, *57*(8), 204.

[39]. Hu, S., Chen, X., Ni, W., Hossain, E., & Wang, X. (2021). Distributed machine learning for wireless communication networks: Techniques, architectures, and applications. *IEEE Communications Surveys & Tutorials*, *23*(3), 1458-1493.

[40]. Hu, S., Wu, Z. S., & Smith, V. (2024). Fair federated learning via bounded group loss. 2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML),

[41]. Huang, W., Ye, M., Shi, Z., Wan, G., Li, H., Du, B., & Yang, Q. (2024). Federated learning for generalization, robustness, fairness: A survey and benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *46*(12), 9387-9406.

[42]. Ietto-Gillies, G. (2017). The organizational and geographical boundaries of the firm: Focus on labour as a major stakeholder. *critical perspectives on international business*, *13*(1), 72-92.

[43]. Istiaque, M., Dipon Das, R., Hasan, A., Samia, A., & Sayer Bin, S. (2023). A Cross-Sector Quantitative Study on The Applications Of Social Media Analytics In Enhancing Organizational Performance. *American Journal of Scholarly Research and Innovation*, *2*(02), 274-302. https://doi.org/10.63125/d8ree044

[44]. Istiaque, M., Dipon Das, R., Hasan, A., Samia, A., & Sayer Bin, S. (2024). Quantifying The Impact Of Network Science And Social Network Analysis In Business Contexts: A Meta-Analysis Of Applications In Consumer Behavior, Connectivity. *International Journal of Scientific Interdisciplinary Research*, *5*(2), 58-89. https://doi.org/10.63125/vgkwe938

[45]. Jahid, M. K. A. S. R. (2022). Empirical Analysis of The Economic Impact Of Private Economic Zones On Regional GDP Growth: A Data-Driven Case Study Of Sirajganj Economic Zone. *American Journal of Scholarly Research and Innovation*, *1*(02), 01-29. https://doi.org/10.63125/je9w1c40

[46]. Ji, S., Tan, Y., Saravirta, T., Yang, Z., Liu, Y., Vasankari, L., Pan, S., Long, G., & Walid, A. (2024). Emerging trends in federated learning: From model fusion to federated x learning. *International Journal of Machine Learning and Cybernetics*, *15*(9), 3769-3790.

[47]. Ji, Y., & Chen, L. (2022). FedQNN: A computation–communication-efficient federated learning framework for IoT with low-bitwidth neural network quantization. *IEEE Internet of Things Journal*, *10*(3), 2494-2507.

[48]. Khan, H. M., Khan, A., Jabeen, F., Anjum, A., & Jeon, G. (2021). Fog-enabled secure multiparty computation based aggregation scheme in smart grid. *Computers & Electrical Engineering*, *94*, 107358.

[49]. Khan, N., Nisar, S., Khan, M. A., Attique Khan, M., Camacho, D., Rehman, Y. A. U., & Hussain, A. (2025). Federated Learning: Concepts, Challenges and Implementation. *Expert Systems*, *42*(8), e70096.

[50]. Khan, Q. W., Khan, A. N., Rizwan, A., Ahmad, R., Khan, S., & Kim, D.-H. (2023). Decentralized machine learning training: a survey on synchronization, consolidation, and topologies. *IEEE Access*, *11*, 68031-68050.

[51]. Kim, S. (2025). Personalized federated learning strategies. In *Federated Learning for Medical Imaging* (pp. 33-42). Elsevier.

[52]. Kishor, K. (2022). Personalized federated learning. In *Federated Learning for IoT Applications* (pp. 31-52). Springer.

[53]. Lan, G., Liu, X.-Y., Zhang, Y., & Wang, X. (2023). Communication-efficient federated learning for resource-constrained edge devices. *IEEE Transactions on Machine Learning in Communications and Networking*, *1*, 210-224.

[54]. Lazaros, K., Koumadorakis, D. E., Vrahatis, A. G., & Kotsiantis, S. (2024). Federated learning: Navigating the landscape of collaborative intelligence. *Electronics*, *13*(23), 4744.

[55]. Le, M., Huynh-The, T., Do-Duy, T., Vu, T.-H., Hwang, W.-J., & Pham, Q.-V. (2024). Applications of distributed machine learning for the internet-of-things: A comprehensive survey. *IEEE Communications Surveys & Tutorials*.

[56]. Li, F., Zhang, L., Liu, Y., & Laili, Y. (2018). QoS-aware service composition in cloud manufacturing: A Gale–Shapley algorithm-based approach. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, *50*(7), 2386-2397.

[57]. Li, L., Shi, D., Hou, R., Li, H., Pan, M., & Han, Z. (2021). To talk or to work: Flexible communication compression for energy efficient federated learning over heterogeneous mobile edge devices. IEEE INFOCOM 2021-IEEE Conference on Computer Communications,

[58]. Li, Y., He, Z., Gu, X., Xu, H., & Ren, S. (2024). AFedAvg: Communication-efficient federated learning aggregation with adaptive communication frequency and gradient sparse. *Journal of Experimental & Theoretical Artificial Intelligence*, *36*(1), 47-69.

[59]. Li, Z., Shao, J., Mao, Y., Wang, J. H., & Zhang, J. (2022). Federated learning with gan-based data synthesis for non-iid clients. International workshop on trustworthy federated learning,

[60]. Liu, D., Yu, G., Zhong, Z., & Song, Y. (2024). Secure multi-party computation with secret sharing for real-time data aggregation in IIoT. *Computer Communications*, *224*, 159-168.

[61]. Liu, J., Huang, J., Zhou, Y., Li, X., Ji, S., Xiong, H., & Dou, D. (2022). From distributed machine learning to federated learning: A survey. *Knowledge and information systems*, *64*(4), 885-917.

[62]. Liu, J., Wang, S., Xu, H., Xu, Y., Liao, Y., Huang, J., & Huang, H. (2024). Federated learning with experience-driven model migration in heterogeneous edge networks. *IEEE/ACM Transactions on Networking*, *32*(4), 3468-3484.

[63]. Liu, L., Zhang, J., Song, S., & Letaief, K. B. (2022). Hierarchical federated learning with quantization: Convergence analysis and system design. *IEEE Transactions on Wireless Communications*, *22*(1), 2-18.

[64]. Liu, P., Xu, X., & Wang, W. (2022). Threats, attacks and defenses to federated learning: issues, taxonomy and perspectives. *Cybersecurity*, *5*(1), 4.

[65]. Liu, T. (2024). Research on privacy techniques based on multi-party secure computation. 2024 3rd International Conference on Artificial Intelligence and Autonomous Robot Systems (AIARS),

[66]. Liu, X., Deng, Y., Nallanathan, A., & Bennis, M. (2023). Federated learning and meta learning: Approaches, applications, and directions. *IEEE Communications Surveys & Tutorials*, *26*(1), 571-618.

[67]. Lu, J., Sheng, Y., Cao, S., Elnaffar, S., Saad, M. M., Seid, A. M., & Erbad, A. (2024). Lyapunov-guided long-term fairness-aware federated learning for collaborative TinyML on edge devices. *IEEE Transactions on Consumer Electronics*, *70*(4), 7334-7345.

[68]. Lycklama, H., Burkhalter, L., Viand, A., Küchler, N., & Hithnawi, A. (2023). Rofl: Robustness of secure federated learning. 2023 IEEE Symposium on Security and Privacy (SP),

[69]. Lyu, L., Yu, H., Ma, X., Chen, C., Sun, L., Zhao, J., Yang, Q., & Yu, P. S. (2022). Privacy and robustness in federated learning: Attacks and defenses. *IEEE Transactions on Neural Networks and Learning Systems*, *35*(7), 8726-8746.

[70]. Ma, X., Zhou, Y., Wang, L., & Miao, M. (2022). Privacy-preserving Byzantine-robust federated learning. *Computer Standards & Interfaces*, *80*, 103561.

[71]. Md Arifur, R., & Sheratun Noor, J. (2022). A Systematic Literature Review of User-Centric Design In Digital Business Systems: Enhancing Accessibility, Adoption, And Organizational Impact. *Review of Applied Science and Technology*, *1*(04), 01-25. https://doi.org/10.63125/ndjkpm77

[72]. Md Ashiqur, R., Md Hasan, Z., & Afrin Binta, H. (2025). A meta-analysis of ERP and CRM integration tools in business process optimization. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 278-312. https://doi.org/10.63125/yah70173

[73]. Md Hasan, Z. (2025). AI-Driven business analytics for financial forecasting: a systematic review of decision support models in SMES. *Review of Applied Science and Technology*, *4*(02), 86-117. https://doi.org/10.63125/gjrpv442

[74]. Md Hasan, Z., Mohammad, M., & Md Nur Hasan, M. (2024). Business Intelligence Systems In Finance And Accounting: A Review Of Real-Time Dashboarding Using Power BI & Tableau. *American Journal of Scholarly Research and Innovation*, *3*(02), 52-79. https://doi.org/10.63125/fy4w7w04

[75]. Md Hasan, Z., & Moin Uddin, M. (2022). Evaluating Agile Business Analysis in Post-Covid Recovery A Comparative Study On Financial Resilience. *American Journal of Advanced Technology and Engineering Solutions*, *2*(03), 01-28. https://doi.org/10.63125/6nee1m28

[76]. Md Hasan, Z., Sheratun Noor, J., & Md. Zafor, I. (2023). Strategic role of business analysts in digital transformation tools, roles, and enterprise outcomes. *American Journal of Scholarly Research and Innovation*, *2*(02), 246-273. https://doi.org/10.63125/rc45z918

[77]. Md Ismail, H., Md Mahfuj, H., Mohammad Aman Ullah, S., & Shofiul Azam, T. (2025). Implementing Advanced Technologies For Enhanced Construction Site Safety. *American Journal of Advanced Technology and Engineering Solutions*, *1*(02), 01-31. https://doi.org/10.63125/3v8rpr04

[78]. Md Ismail Hossain, M. A. B., amp, & Mousumi Akter, S. (2023). Water Quality Modelling and Assessment Of The Buriganga River Using Qual2k. *Global Mainstream Journal of Innovation, Engineering & Emerging Technology*, *2*(03), 01-11. https://doi.org/10.62304/jieet.v2i03.64

[79]. Md Jakaria, T., Md, A., Zayadul, H., & Emdadul, H. (2025). Advances In High-Efficiency Solar Photovoltaic Materials: A Comprehensive Review Of Perovskite And Tandem Cell Technologies. *American Journal of Advanced Technology and Engineering Solutions*, *1*(01), 201-225. https://doi.org/10.63125/5amnvb37

[80]. Md Mahamudur Rahaman, S. (2022a). Electrical And Mechanical Troubleshooting in Medical And Diagnostic Device Manufacturing: A Systematic Review Of Industry Safety And Performance Protocols. *American Journal of Scholarly Research and Innovation*, *1*(01), 295-318. https://doi.org/10.63125/d68y3590

[81]. Md Mahamudur Rahaman, S. (2022b). Smart Maintenance in Medical Imaging Manufacturing: Towards Industry 4.0 Compliance at Chronos Imaging. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *2*(1), 29–62. https://doi.org/10.63125/eatsmf47

[82]. Md Mahamudur Rahaman, S. (2024). AI-Driven Predictive Maintenance For High-Voltage X-Ray Ct Tubes: A Manufacturing Perspective. *Review of Applied Science and Technology*, *3*(01), 40-67. https://doi.org/10.63125/npwqxp02

[83]. Md Mahamudur Rahaman, S., & Rezwanul Ashraf, R. (2022). Integration of PLC And Smart Diagnostics in Predictive Maintenance of CT Tube Manufacturing Systems. *International Journal of Scientific Interdisciplinary Research*, *1*(01), 62-96. https://doi.org/10.63125/gspb0f75

[84]. Md Mahamudur Rahaman, S., & Rezwanul Ashraf, R. (2023). Applying Lean And Six Sigma In The Maintenance Of Medical Imaging Equipment Manufacturing Lines. *Review of Applied Science and Technology*, *2*(04), 25-53. https://doi.org/10.63125/6varjp35

[85]. Md Nazrul Islam, K. (2022). A Systematic Review of Legal Technology Adoption In Contract Management, Data Governance, And Compliance Monitoring. *American Journal of Interdisciplinary Studies*, *3*(01), 01-30. https://doi.org/10.63125/caangg06

[86]. Md Nur Hasan, M. (2024). Integration Of Artificial Intelligence And DevOps In Scalable And Agile Product Development: A Systematic Literature Review On Frameworks. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *4*(1), 01–32. https://doi.org/10.63125/exyqj773

[87]. Md Nur Hasan, M. (2025). Role Of AI And Data Science In Data-Driven Decision Making For It Business Intelligence: A Systematic Literature Review. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 564-588. https://doi.org/10.63125/n1xpym21

[88]. Md Nur Hasan, M., Md Musfiqur, R., & Debashish, G. (2022). Strategic Decision-Making in Digital Retail Supply Chains: Harnessing AI-Driven Business Intelligence From Customer Data. *Review of Applied Science and Technology*, *1*(03), 01-31. https://doi.org/10.63125/6a7rpy62

[89]. Md Redwanul, I., & Md. Zafor, I. (2022). Impact of Predictive Data Modeling on Business Decision-Making: A Review Of Studies Across Retail, Finance, And Logistics. *American Journal of Advanced Technology and Engineering Solutions*, *2*(02), 33-62. https://doi.org/10.63125/8hfbkt70

[90]. Md Rezaul, K., & Md Mesbaul, H. (2022). Innovative Textile Recycling and Upcycling Technologies For Circular Fashion: Reducing Landfill Waste And Enhancing Environmental Sustainability. *American Journal of Interdisciplinary Studies*, *3*(03), 01-35. https://doi.org/10.63125/kkmerg16

[91]. Md Sultan, M., Proches Nolasco, M., & Md. Torikul, I. (2023). Multi-Material Additive Manufacturing For Integrated Electromechanical Systems. *American Journal of Interdisciplinary Studies*, *4*(04), 52-79. https://doi.org/10.63125/y2ybrx17

[92]. Md Sultan, M., Proches Nolasco, M., & Vicent Opiyo, N. (2025). A Comprehensive Analysis Of Non-Planar Toolpath Optimization In Multi-Axis 3D Printing: Evaluating The Efficiency Of Curved Layer Slicing Strategies. *Review of Applied Science and Technology*, *4*(02), 274-308. https://doi.org/10.63125/5fdxa722

[93]. Md Takbir Hossen, S., Ishtiaque, A., & Md Atiqur, R. (2023). AI-Based Smart Textile Wearables For Remote Health Surveillance And Critical Emergency Alerts: A Systematic Literature Review. *American Journal of Scholarly Research and Innovation*, *2*(02), 1-29. https://doi.org/10.63125/ceqapd08

[94]. Md Tawfiqul, I. (2023). A Quantitative Assessment Of Secure Neural Network Architectures For Fault Detection In Industrial Control Systems. *Review of Applied Science and Technology*, *2*(04), 01-24. https://doi.org/10.63125/3m7gbs97

[95]. Md. Sakib Hasan, H. (2022). Quantitative Risk Assessment of Rail Infrastructure Projects Using Monte Carlo Simulation And Fuzzy Logic. *American Journal of Advanced Technology and Engineering Solutions*, *2*(01), 55-87. https://doi.org/10.63125/h24n6z92

[96]. Md. Tarek, H. (2022). Graph Neural Network Models For Detecting Fraudulent Insurance Claims In Healthcare Systems. *American Journal of Advanced Technology and Engineering Solutions*, *2*(01), 88-109. https://doi.org/10.63125/r5vsmv21

[97]. Md. Zafor, I. (2025). A Meta-Analysis Of AI-Driven Business Analytics: Enhancing Strategic Decision-Making In SMEs. *Review of Applied Science and Technology*, *4*(02), 33-58. https://doi.org/10.63125/wk9fqv56

[98]. Md.Kamrul, K., & Md Omar, F. (2022). Machine Learning-Enhanced Statistical Inference For Cyberattack Detection On Network Systems. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 65-90. https://doi.org/10.63125/sw7jzx60

[99]. Md.Kamrul, K., & Md. Tarek, H. (2022). A Poisson Regression Approach to Modeling Traffic Accident Frequency in Urban Areas. *American Journal of Interdisciplinary Studies*, *3*(04), 117-156. https://doi.org/10.63125/wqh7pd07

[100]. Mills, J., Hu, J., & Min, G. (2019). Communication-efficient federated learning for wireless edge intelligence in IoT. *IEEE Internet of Things Journal*, *7*(7), 5986-5994.

[101]. Mladineo, M., Veza, I., & Gjeldum, N. (2017). Solving partner selection problem in cyber-physical production networks using the HUMANT algorithm. *International Journal of Production Research*, *55*(9), 2506-2521.

[102]. Moin Uddin, M. (2025). Impact Of Lean Six Sigma On Manufacturing Efficiency Using A Digital Twin-Based Performance Evaluation Framework. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *1*(01), 343-375. https://doi.org/10.63125/z70nhf26

[103]. Moin Uddin, M., & Rezwanul Ashraf, R. (2023). Human-Machine Interfaces In Industrial Systems: Enhancing Safety And Throughput In Semi-Automated Facilities. *American Journal of Interdisciplinary Studies*, *4*(01), 01-26. https://doi.org/10.63125/s2qa0125

[104]. Mollanejad, A., Navin, A. H., & Ghanbari, S. (2024). Fairness-aware loss history based federated learning heuristic algorithm. *Knowledge-Based Systems*, *288*, 111467.

[105]. Momena, A., & Md Nur Hasan, M. (2023). Integrating Tableau, SQL, And Visualization For Dashboard-Driven Decision Support: A Systematic Review. *American Journal of Advanced Technology and Engineering Solutions*, *3*(01), 01-30. https://doi.org/10.63125/4aa43m68

[106]. Mubashir, I., & Abdul, R. (2022). Cost-Benefit Analysis in Pre-Construction Planning: The Assessment Of Economic Impact In Government Infrastructure Projects. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 91-122. https://doi.org/10.63125/kjwd5e33

[107]. Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2021). Federated learning for internet of things: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, *23*(3), 1622-1658.

[108]. Omar Muhammad, F., & Md.Kamrul, K. (2022). Blockchain-Enabled BI For HR And Payroll Systems: Securing Sensitive Workforce Data. *American Journal of Scholarly Research and Innovation*, *1*(02), 30-58. https://doi.org/10.63125/et4bhy15

[109]. Pei, J., Liu, W., Li, J., Wang, L., & Liu, C. (2024). A review of federated learning methods in heterogeneous scenarios. *IEEE Transactions on Consumer Electronics*, *70*(3), 5983-5999.

[110]. Peng, D., Ji, Y., & Kong, Q. (2023). OFDI and firms' sustainable productive capacity: Evidence from Chinese industrial firms. *International Review of Economics & Finance*, *83*, 641-652.

[111]. Pennisi, M., Salanitri, F. P., Bellitto, G., Casella, B., Aldinucci, M., Palazzo, S., & Spampinato, C. (2024). FedER: Federated Learning through Experience Replay and privacy-preserving data synthesis. *Computer Vision and Image Understanding*, *238*, 103882.

[112]. Pillutla, K., Kakade, S. M., & Harchaoui, Z. (2022). Robust aggregation for federated learning. *IEEE Transactions on Signal Processing*, *70*, 1142-1154.

[113]. Provatas, N., Konstantinou, I., & Koziris, N. (2025). A Survey on Parameter Server Architecture: Approaches for Optimizing Distributed Centralized Learning. *IEEE Access*.

[114]. Quan, M. K., Pathirana, P. N., Wijayasundara, M., Setunge, S., Nguyen, D. C., Brinton, C. G., Love, D. J., & Poor, H. V. (2025). Federated learning for cyber physical systems: a comprehensive survey. *IEEE Communications Surveys & Tutorials*.

[115]. Rahman, K. J., Ahmed, F., Akhter, N., Hasan, M., Amin, R., Aziz, K. E., Islam, A. M., Mukta, M. S. H., & Islam, A. N. (2021). Challenges, applications and design aspects of federated learning: A survey. *IEEE Access*, *9*, 124682-124700.

[116]. Reduanul, H., & Mohammad Shoeb, A. (2022). Advancing AI in Marketing Through Cross Border Integration Ethical Considerations And Policy Implications. *American Journal of Scholarly Research and Innovation*, *1*(01), 351-379. https://doi.org/10.63125/d1xg3784

[117]. Reisizadeh, A., Tziotis, I., Hassani, H., Mokhtari, A., & Pedarsani, R. (2022). Straggler-resilient federated learning: Leveraging the interplay between statistical accuracy and system heterogeneity. *IEEE Journal on Selected Areas in Information Theory*, *3*(2), 197-205.

[118]. Sabuj Kumar, S., & Zobayer, E. (2022). Comparative Analysis of Petroleum Infrastructure Projects In South Asia And The Us Using Advanced Gas Turbine Engine Technologies For Cross Integration. *American Journal of Advanced Technology and Engineering Solutions*, *2*(04), 123-147. https://doi.org/10.63125/wr93s247

[119]. Sadia, T., & Shaiful, M. (2022). In Silico Evaluation of Phytochemicals From Mangifera Indica Against Type 2 Diabetes Targets: A Molecular Docking And Admet Study. *American Journal of Interdisciplinary Studies*, *3*(04), 91-116. https://doi.org/10.63125/anaf6b94

[120]. Saha, S., Hota, A., Chattopadhyay, A. K., Nag, A., & Nandi, S. (2024). A multifaceted survey on privacy preservation of federated learning: progress, challenges, and opportunities. *Artificial Intelligence Review*, *57*(7), 184.

[121]. Salim, S., Moustafa, N., & Almorjan, A. (2025). Responsible deep federated learning-based threat detection for satellite communications. *IEEE Internet of Things Journal*.

[122]. Sanjai, V., Sanath Kumar, C., Maniruzzaman, B., & Farhana Zaman, R. (2023). Integrating Artificial Intelligence in Strategic Business Decision-Making: A Systematic Review Of Predictive Models. *International Journal of Scientific Interdisciplinary Research*, *4*(1), 01-26. https://doi.org/10.63125/s5skge53

[123]. Sanjai, V., Sanath Kumar, C., Sadia, Z., & Rony, S. (2025). AI And Quantum Computing For Carbon-Neutral Supply Chains: A Systematic Review Of Innovations. *American Journal of Interdisciplinary Studies*, *6*(1), 40-75. https://doi.org/10.63125/nrdx7d32

[124]. Shah, S. M., & Lau, V. K. (2021). Model compression for communication efficient federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, *34*(9), 5937-5951.

[125]. Shawkat, M., Ali, Z. H., Salem, M., & El-Desoky, A. (2025). A robust and personalized privacy-preserving approach for adaptive clustered federated distillation. *Scientific Reports*, *15*(1), 14069.

[126]. Shen, S., Yu, C., Zhang, K., Chen, X., Chen, H., & Ci, S. (2021). Communication-efficient federated learning for connected vehicles with constrained resources. 2021 International Wireless Communications and Mobile Computing (IWCMC),

[127]. Sheratun Noor, J., & Momena, A. (2022). Assessment Of Data-Driven Vendor Performance Evaluation in Retail Supply Chains: Analyzing Metrics, Scorecards, And Contract Management Tools. *American Journal of Interdisciplinary Studies*, *3*(02), 36-61. https://doi.org/10.63125/0s7t1y90

[128]. Shi, Y., Yu, H., & Leung, C. (2023). Towards fairness-aware federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, *35*(9), 11922-11938.

[129]. Simonova, A. (2011). The risk-based approach to anti-money laundering: problems and solutions. *Journal of Money Laundering Control*, *14*(4), 346-358.

[130]. Song, J., Wang, W., Gadekallu, T. R., Cao, J., & Liu, Y. (2022). Eppda: An efficient privacy-preserving data aggregation federated learning scheme. *IEEE Transactions on Network Science and Engineering*, *10*(5), 3047-3057.

[131]. Song, M., Wang, Z., Zhang, Z., Song, Y., Wang, Q., Ren, J., & Qi, H. (2020). Analyzing user-level privacy attack against federated learning. *IEEE Journal on Selected Areas in Communications*, *38*(10), 2430-2444.
[132]. Tahmina Akter, R., Debashish, G., Md Soyeb, R., & Abdullah Al, M. (2023). A Systematic Review of AI-Enhanced Decision Support Tools in Information Systems: Strategic Applications In Service-Oriented Enterprises And Enterprise Planning. *Review of Applied Science and Technology*, *2*(01), 26-52. https://doi.org/10.63125/73djw422
[133]. Tam, P., Corrado, R., Eang, C., & Kim, S. (2023). Applicability of deep reinforcement learning for efficient federated learning in massive IoT communications. *Applied Sciences*, *13*(5), 3083.
[134]. Tamanna, R., & Dipongkar Ray, S. (2023). Comprehensive Insights Into Co_2 Capture: Technological Progress And Challenges. *Review of Applied Science and Technology*, *2*(01), 113-141. https://doi.org/10.63125/9p690n14
[135]. Tariq, A., Serhani, M. A., Sallabi, F. M., Barka, E. S., Qayyum, T., Khater, H. M., & Shuaib, K. A. (2024). Trustworthy federated learning: A comprehensive review, architecture, key challenges, and future research prospects. *IEEE Open Journal of the Communications Society*.
[136]. Tedeschini, B. C., Savazzi, S., Stoklasa, R., Barbieri, L., Stathopoulos, I., Nicoli, M., & Serio, L. (2022). Decentralized federated learning for healthcare networks: A case study on tumor segmentation. *IEEE Access*, *10*, 8693-8708.
[137]. Tian, Y., Wang, S., Xiong, J., Bi, R., Zhou, Z., & Bhuiyan, M. Z. A. (2023). Robust and privacy-preserving decentralized deep federated learning training: Focusing on digital healthcare applications. *IEEE/ACM Transactions on computational biology and bioinformatics*, *21*(4), 890-901.
[138]. ur Rehman, M. H., Dirir, A. M., Salah, K., & Svetinovic, D. (2020). Fairfed: Cross-device fair federated learning. 2020 IEEE Applied Imagery Pattern Recognition Workshop (AIPR),
[139]. Voropai, N., Podkovalnikov, S., & Chudinova, L. (2021). The evolution of interstate power grid formation. *Global Energy Interconnection*, *4*(4), 335-353.
[140]. Wahab, O. A., Mourad, A., Otrok, H., & Taleb, T. (2021). Federated machine learning: Survey, multi-level classification, desirable criteria and future directions in communication and networking systems. *IEEE Communications Surveys & Tutorials*, *23*(2), 1342-1397.
[141]. Wang, L., Jiang, Z., Zhu, Y., Cai, W., Zhu, F., & Chen, T. (2024). Customized Multi-task Learning for Recommendation with Heterogeneous Graph Neural Network. International Conference on Neural Information Processing,
[142]. Wang, L., & Wu, C. (2017). Business failure prediction based on two-stage selective ensemble with manifold learning algorithm and kernel-based fuzzy self-organizing map. *Knowledge-Based Systems*, *121*, 99-110.
[143]. Wang, L., Xu, Y., Xu, H., Jiang, Z., Chen, M., Zhang, W., & Qian, C. (2023). BOSE: Block-wise federated learning in heterogeneous edge computing. *IEEE/ACM Transactions on Networking*, *32*(2), 1362-1377.
[144]. Wang, Z., Song, M., Zhang, Z., Song, Y., Wang, Q., & Qi, H. (2019). Beyond inferring class representatives: User-level privacy leakage from federated learning. IEEE INFOCOM 2019-IEEE conference on computer communications,
[145]. Wu, D., Yang, W., Zou, X., Tao, D., Li, S., Xia, W., & Fang, B. (2024). Bird+: Design of a lightweight communication compressor for resource-constrained distribution learning platforms. *IEEE Transactions on Parallel and Distributed Systems*, *35*(11), 2193-2207.
[146]. Xiang, F., Jiang, G., Xu, L., & Wang, N. (2016). The case-library method for service composition and optimal selection of big manufacturing data in cloud manufacturing system. *The International Journal of Advanced Manufacturing Technology*, *84*(1), 59-70.
[147]. Xing, S., Ning, Z., Zhou, J., Liao, X., Xu, J., & Zou, W. (2022). N-fedavg: Novel federated average algorithm based on fedavg. 2022 14th International Conference on Communication Software and Networks (ICCSN),
[148]. Xiong, A., Zhou, H., Song, Y., Wang, D., Wei, X., Li, D., & Gao, B. (2024). A multi-task based clustering personalized federated learning method. *Big Data Mining and Analytics*, *7*(4), 1017-1030.
[149]. Xu, J., Du, W., Jin, Y., He, W., & Cheng, R. (2020). Ternary compression for communication-efficient federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, *33*(3), 1162-1176.
[150]. Yahata, Y., Sugiura, K., & Matsutani, H. (2024). A Scalable Secure Fault Tolerant Aggregation for P2P Federated Learning. 2024 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW),
[151]. Yan, Y., Hu, T., & Zhu, W. (2024). Leveraging large language models for enhancing financial compliance: A focus on anti-money laundering applications. 2024 4th International Conference on Robotics, Automation and Artificial Intelligence (RAAI),
[152]. Yang, H., Liu, J., & Bentley, E. S. (2021). Cfedavg: achieving efficient communication and fast convergence in non-iid federated learning. 2021 19th International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks (WiOpt),
[153]. Yin, F., Lin, Z., Kong, Q., Xu, Y., Li, D., Theodoridis, S., & Cui, S. R. (2020). FedLoc: Federated learning framework for data-driven cooperative localization and location data processing. *IEEE Open Journal of Signal Processing*, *1*, 187-215.
[154]. Yuan, L., Wang, Z., Sun, L., Yu, P. S., & Brinton, C. G. (2024). Decentralized federated learning: A survey and perspective. *IEEE Internet of Things Journal*, *11*(21), 34617-34638.
[155]. Zeng, T., Semiari, O., Chen, M., Saad, W., & Bennis, M. (2022). Federated learning on the road autonomous controller design for connected and autonomous vehicles. *IEEE Transactions on Wireless Communications*, *21*(12), 10407-10423.
[156]. Zhan, Y., Li, P., Qu, Z., Zeng, D., & Guo, S. (2020). A learning-based incentive mechanism for federated learning. *IEEE Internet of Things Journal*, *7*(7), 6360-6368.
[157]. Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. *Knowledge-Based Systems*, *216*, 106775.
[158]. Zhang, J., Li, Y., Wu, D., Zhao, Y., & Palaiahnakote, S. (2025). SFFL: Self-aware fairness federated learning framework for heterogeneous data distributions. *Expert Systems with Applications*, *269*, 126418.

[159]. Zhang, M., Wei, E., & Berry, R. (2021). Faithful edge federated learning: Scalability and privacy. *IEEE Journal on Selected Areas in Communications*, *39*(12), 3790-3804.

[160]. Zhang, R., Fan, Z., Yao, J., Zhang, Y., & Wang, Y. (2025). Fairness-guided federated training for generalization and personalization in cross-silo federated learning. *Frontiers of Information Technology & Electronic Engineering*, *26*(1), 42-61.

[161]. Zhang, Z., Zhang, Y., Guo, D., Yao, L., & Li, Z. (2022). SecFedNIDS: Robust defense for poisoning attack against federated learning-based network intrusion detection system. *Future Generation Computer Systems*, *134*, 154-169.

[162]. Zhou, X., Lei, X., Yang, C., Shi, Y., Zhang, X., & Shi, J. (2023). Handling data heterogeneity for iot devices in federated learning: A knowledge fusion approach. *IEEE Internet of Things Journal*, *11*(5), 8090-8104.

[163]. Zhou, Y., Ye, Q., & Lv, J. (2021). Communication-efficient federated learning with compensated overlap-fedavg. *IEEE Transactions on Parallel and Distributed Systems*, *33*(1), 192-205.