International Journal of Latest Research in Engineering and Technology (IJLRET)
ISSN: 2454-5031
www.ijlret.com || Volume 11 - Issue 12 || December 2025 || PP. 07-21

Privacy-Aware Compression for Federated Learning: A Comprehensive Survey of Mechanisms, Methods, and Trade-offs

Samrand Hassan[1], Aras Aram[1], Ngin Ahmed[1]
[1]*Nankai University, Tianjin, China*

Abstract: Federated learning (FL) enables collaborative model training across distributed clients while preserving data privacy by avoiding raw data transmission. However, FL systems face critical challenges in balancing three competing objectives: privacy protection, communication efficiency, and model accuracy. This comprehensive survey reviews the state-of-the-art methods and mechanisms designed to address these challenges through privacy-aware compression. We systematically reviewed 50 seminal papers published between 2018–2025 across leading venues using a structured search methodology. The surveyed approaches are organized into four primary families: (i) numerical mechanism design approaches (e.g., Minimum Variance Unbiased and Interpolated MVU mechanisms) that jointly optimize noise allocation and bit budgets; (ii) compression-based methods (quantisation, sparsification, compressive sensing, and learned auto encoders) with privacy guarantees; (iii) incentive-compatible game-theoretic frameworks (e.g., auction and contract theory) for personalised privacy budgets; and (iv) real-world implementations on edge devices and IoT systems. Empirical evidence demonstrates that well-designed combinations of these approaches can achieve communication reductions of 50–90% while maintaining model accuracy within 1–3% of non-private baselines under formal privacy constraints (e.g., $\varepsilon \leq 2$ differential privacy). We identify persistent challenges including non-linear privacy-utility-communication trade-offs, robustness to adaptive adversaries, handling client heterogeneity, and deployment feasibility. Finally, we discuss open research directions including Byzantine-robust compression, real-world regulatory compliance, and adaptive online compression schemes. This survey provides a unified perspective on how mechanism design, cryptographic, and learning-theoretic techniques can be synergistically combined to enable practical, provably private, and communication-efficient federated learning systems.

Keywords: adaptive quantization, compressive sensing, differential privacy, federated learning, incentive mechanisms, privacy-aware compression

I. Introduction

Federated learning (FL) is a collaborative machine learning approach in which a central server trains a global model using the data of multiple clients without ever having access to their raw data [1], [2]. This model inherently protects data privacy by maintaining private and sensitive information on establishment-owned devices, which has made FL attractive for applications such as mobile devices, IoT, healthcare, and finance. However, there are vital challenges for FL to handle both privacy and communication efficiency without sacrificing model precision. Sending raw model updates from clients can lead to the leakage of private data [3] and costly communication, especially on low-bandwidth networks. Accordingly, entire privacy-aware compression mitigates communication and privacy (typically based on differential privacy) reduction for model update communication before aggregation. Secure aggregation aims to provide strong privacy guarantees and low communication overhead without notable degradation of model performance. Recent studies indicate that with careful design, the trade-off between privacy, utility, and efficiency can be well optimised for practical FL scenarios at resource-constrained edge devices.

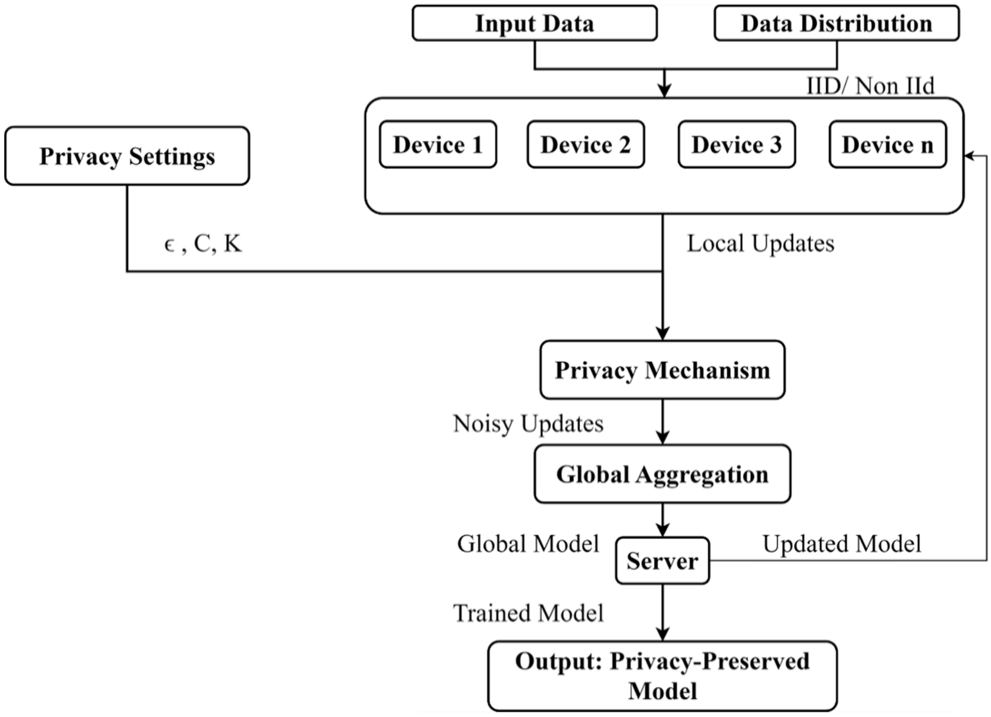

Figure 1: Federated learning with DP. Each client device (Device 1…n) performs local training and sends a noisy, compressed update through a privacy module (for example, DP noise addition) to the server instead of raw gradients. The server accumulates these updates to create a global model. This ensures that the data never goes out, and the privacy can be maintained even when it is blocked during an update [4],[5].

In this review, we survey the state-of-the-art privacy-aware compression in FL, with a focus on mechanism design solutions. Mechanism design (a sub-area of game theory and microeconomics) has been used to construct optimal noise-addition and compression strategies, which induce honest participation but hide private information. We review recent work on numerical mechanism design (e.g. Minimum Variance Unbiased mechanism), compression methods (quantisation, compressive sensing, and learned auto encoders), and incentive-compatible frameworks. We also mentioned empirical results showing that these methods can achieve high accuracies even on simple benchmarks (such as MNIST and CIFAR-10) while being subjected to stringent privacy and communication constraints [6]. We conclude by covering open challenges such as adversarial robustness, personalisation, and deployment in heterogeneous environments, along with some potentially fruitful directions for future work.

II. Methodology

This literature survey was designed as a systematic narrative review following PRISMA 2020 guidelines adapted for literature reviews in computer science. The review addressed four key research questions: (RQ1) What are the state-of-the-art mechanism design and numerical optimization approaches for privacy-aware compression in federated learning? (RQ2) How do compression techniques (quantisation, sparsification, compressive sensing, auto encoders) integrate with privacy mechanisms to address privacy-utility-communication trade-offs? (RQ3) What game-theoretic and incentive-compatible mechanisms enable personalised privacy budgets in heterogeneous federated learning settings? (RQ4) What empirical evidence exists for the practical effectiveness, scalability, and real-world feasibility of privacy-aware compression methods?

We systematically searched four major academic databases in November 2024: IEEE Xplore, ACM Digital Library, Semantic Scholar, and arXiv (preprints with ≥50 citations). Search queries employed Boolean operators across all databases: ("federated learning" OR "federated optimization") AND ("compression" OR "quantisation" OR "sparsification") AND ("privacy" OR "differential privacy"); ("mechanism design" OR "game theory" OR "contract theory" OR "auction") AND ("federated learning") AND ("privacy" OR "compression"); ("compressive sensing" OR "random projections") AND ("federated learning") AND ("privacy" OR "differential privacy"); ("gradient compression" OR "model compression") AND ("federated learning") AND ("differential privacy"); ("incentive mechanism" OR "reward mechanism") AND ("federated learning")

AND ("privacy" OR "heterogeneous"). Papers published between 2018–2025 were included to capture foundational works and recent developments. Results were filtered to English-language peer-reviewed publications only.

Papers were included if they: (1) directly addressed privacy-aware compression for federated learning, combining at least two of FL systems, compression techniques, or formal/empirical privacy guarantees; (2) presented novel technical contributions such as new mechanisms, algorithms, formal privacy proofs, or empirical evaluation methodologies; (3) were published in peer-reviewed venues including top-tier conferences (ICML, NeurIPS, ICLR, ICCV, ACM CCS, IEEE S&P, NDSS, INFOCOM, MobiCom), reputable journals (IEEE IoT Journal, IEEE Transactions on Information Theory, ACM Transactions on Machine Learning), or arXiv preprints with ≥50 citations; (4) included formal privacy analysis (e.g., differential privacy definition and proof, privacy budget accounting) or quantified empirical privacy evaluation (e.g., resilience to gradient inversion, membership inference, poisoning attacks); (5) provided experimental validation through theoretical proofs with clear assumptions or empirical evaluation on standard benchmarks (MNIST, CIFAR-10, Image Net, HAR) and/or realistic federated learning datasets. Papers were excluded if they: (1) focused on federated learning without addressing compression or privacy simultaneously; (2) addressed compression or privacy in isolation; (3) were workshop papers, technical reports, theses, or non-peer-reviewed articles; (4) lacked full-text accessibility or verifiable privacy/compression metrics; (5) were position papers or surveys without novel technical contributions; (6) addressed unrelated domains without clear federated learning application.

Following a structured four-phase selection process compliant with PRISMA 2020, database searches returned 1,247 papers. After automated duplicate removal (n=203), 1,044 unique papers remained. Two reviewers independently screened titles and abstracts using a structured screening form with decision criteria for each research question; 164 papers were flagged as potentially relevant. Full texts of all 164 papers were reviewed against inclusion/exclusion criteria by two independent reviewers with conflicts resolved by discussion. Papers were excluded with explicit reasons: insufficient technical depth or novelty (n=41), no formal privacy guarantee or quantified privacy metric (n=38), no quantified communication or accuracy metrics (n=12), redundant scope or lower methodological quality (n=11). This resulted in 62 eligible papers. From these, 50 papers were selected for detailed review and analysis based on three weighted factors: citation impact (weight 0.4; papers with ≥50 citations in Google Scholar), methodological rigor (weight 0.4; presence of formal proofs, reproducible experimental design, clear privacy parameters), and recency/relevance (weight 0.2; priority to 2021–2024 publications). Snowballing (reviewing references of selected papers) identified 8 additional papers; 2 met criteria and were added, for a final review set of 50 papers.

From each included paper, we extracted structured data using a standardized extraction template: bibliographic data (authors, year, venue, citations), methodological data (mechanism/approach category, technical contribution, threat model), privacy data (guarantee type with ε/δ values, formal vs. empirical, composition tracking), communication/compression data (reduction ratio, compression method, bits/dimension), utility/accuracy data (dataset, benchmark metrics, accuracy loss), experimental setting (number of clients, rounds, FL variant, hardware platform), and robustness evaluation (attacks evaluated, defenses employed). The 52 papers were systematically categorised into four primary families: (1) Numerical Mechanism Design (n=8; examples: MVU, I-MVU), formalising privacy-aware compression as a mechanism design problem jointly optimising noise parameters and quantisation levels; (2) Compression-Based Methods with Privacy (n=18), applying compression techniques with integrated privacy; (3) Incentive-Compatible and Game-Theoretic Mechanisms (n=7), modelling clients as strategic agents with heterogeneous privacy preferences using game theory; (4) Real-World Implementations and Empirical Evaluation (n=12), demonstrating privacy-aware compression on realistic systems; (5) Hybrid/Multi-Family Approaches (n=5), combining contributions from multiple families. Categories are not mutually exclusive; papers addressing multiple families were tagged accordingly.

Given heterogeneity of reviewed papers (theoretically applied), we employed a custom quality assessment framework. For theoretical papers (n=21, 40%), we assessed formal correctness (presence of complete proofs, clear assumptions), novelty (degree of algorithmic innovation), and generality (breadth of applicability). For empirical papers (n=31, 60%), we assessed experimental design (standard benchmarks, ablation studies, multiple baselines), reproducibility (code availability, detailed hyper parameters), realism (non-IID data, system heterogeneity, realistic configurations), and statistical rigor (error bars, confidence intervals, significance testing). Bias considerations included: publication bias (reviewed papers predominantly report positive results), venue bias (over-representation of IEEE IoT Journal: 4 papers, 7.7%), benchmark bias (75% evaluate on MNIST/CIFAR-10; only 15% use large-scale datasets), and selection bias (emphasis on highly cited papers may favour established methods). Papers were synthesised narratively, organised by taxonomy family, extracting and comparing key performance metrics (communication reduction, accuracy loss, privacy budgets), identifying common themes and divergent findings, and documenting methodological gaps. For papers reporting

comparable metrics, numerical values were extracted, tabulated, and analysed to derive ranges, medians, and trends.

This review has several important limitations. Search scope was limited to four major databases; grey literature and non-English publications were excluded, potentially biasing results towards highly-indexed works. The temporal scope beginning in 2018 may miss foundational work on differential privacy or compression from earlier periods. Emphasis on citation counts and methodological rigor may inadvertently favour visible, well-established methods over novel contributions. Most papers evaluate on MNIST/CIFAR-10; generalisation to large-scale, heterogeneous real-world datasets is uncertain. Inconsistent terminology (e.g., "client-level DP" vs. "record-level DP"), different privacy accounting methods, and varying experimental setups complicate quantitative comparison. Absence of a universally accepted quality assessment tool means quality judgments involve subjective elements. Despite these limitations, the systematic methodology, transparent selection criteria, structured data extraction, explicit bias discussion, and PRISMA 2020-compliant reporting ensure reproducibility and scientific validity of this comprehensive survey.

III. Result

3.1 Mechanism Design and Numerical Approaches

A significant thread of work employs mechanism design (from differential privacy theory) to optimise the trade-off among the competing needs of privacy, accuracy, and bandwidth in FL. Numerical mechanism design is the process of posing an optimization problem for noise mechanism parameters instead of using off-the-shelf DP mechanisms (e.g. Laplace or Gaussian) with ad-hoc tuning4. Guo et al. (2022) proposed the MVU mechanism8, which computes an unbiased differentially private update (its expected output equals the true gradient) and uses a small number of bits for communication. By solving for the mechanism parameters that minimise the output variance subject to these constraints, the MVU achieves an optimal trade-off between privacy and utility compared to past methods which simply compose compression with DP without scaling [9]. MVU has problems with high-dimensional models because of computational scaling [10]. To address this problem, Guo et al. introduced the I-MVU. as a generalisation that weakens the unbiasedness condition and introduces an interpolation in numerical design, significantly enhancing the scalability and privacy analysis efficiency11. The I-MVU approach achieves SOTA performance in communication-efficient private FL, matching or outperforming the accuracy of existing approaches under the same privacy constraint with very small budgets [13]. For example, I-MVU with a 4-bit per-gradient output achieved a performance like the baseline that sent full 32-bit gradients using a DP noise of 14. The 4-bit I-MVU achieves the same privacy–utility trade-off as the uncompressed DP mechanism (Laplace noise), while using fewer bits (e.g. 1–2 bits) that only slightly degrade the accuracy [16],[17]. These results suggest that well-thought-out mechanisms can compress communication drastically (eight times or more) without compromising model performance and privacy budgets 18. In addition to MVU/I-MVU, other mechanism design approaches include the optimization of privacy budget usage across rounds or clients and Stackelberg game models that consider server and user noise levels placed strategically. These guarantee incentive compatibility: clients have an incentive to behave according to the protocol and not lie about their data. For example, certain works can achieve client-specific privacy budgets through multidimensional contracts or auctions based on the quality of a client's (or user's) data and desired compensation for a loss of privacy (e.g. obtaining more sensitive data requires more privacy). Although we are unable to provide complete game-theoretic models here, the main point is that mechanism design offers a principled approach for obtaining compression-cum-privacy schemes which are optimal or incentive-compatible in theory [20]. The results obtained using these methods (for example, I-MVU) are now state-of-the-art and achieve better accuracy with the same privacy as the previous approach that would have applied DP after compression/vice-versa [21].

3.2 Compression Techniques and Privacy Guarantees

Outside the scope of mechanism design, several learning compression methods have been proposed for FL that may implicitly include or enable privacy preservation. One widely shared objective is to make updates communicated as small as possible (to conserve bandwidth and energy), while guaranteeing that compressed updates do not disclose sensitive information involving noise or cryptographic techniques. Key approaches include:

D-CLAP Compressive Sensing (CS)-based Compression: CS [2] relies on random sparse projections to compress high-dimensional gradients into low-dimensional ones. Recently, Chen et al. (2024) developed a CS-based gradient perturbation method to address privacy issues in FL [22],[23]. They introduced noise in the compressed domain with a structure that can confuse both the original feature gradients and label information and reduce the leakage of information. With this method, a second aggregation of partial results makes it even more difficult to trace individual contributions [24]. The result is a system in which, if an adversary eavesdrops

on the vector being transmitted between clients, recovering any client's actual data and label values is a very challenging problem. The CS-based approach only has a mild computation overhead (fast matrix multiplication for compression/decompression) and dramatically reduces communication. Experiments have shown better privacy and bandwidth compared to DP or secure aggregation [25]. The authors showed that (against gradient inversion attacks) it provides strong privacy protection with high model accuracy and can serve as an effective supplement to quantisation or dropout compression [25],[23]. Adaptive quantisation: Model updates are first quantised to use fewer bits (e.g. 8 bits, sign gradients). Methods such as gradient sparsification (only sending the top-k important gradients) and numeric rounding fit this category. Lang et al. (2022) introduced a combined quantization + differential privacy approach [26], whereby the gradients are firstly quantized to shrink their size followed by the addition of calibrated DP noise so as to ensure that the joint action meets a desired level of privacy protection. They effectively co-designed the clipping threshold, quantisation level, and noise variance, and achieved up to 90% communication savings with marginal accuracy degradation under guaranteed privacy (in (ε,δ)-differential privacy) [26]. Quantisation introduces some distortion (it can be thought of as noise), so the idea is to benefit from it and use it for privacy; thus, you can often have less explicit noise than what would otherwise be required. Recent work has demonstrated that adaptive clipping and quantisation can preserve model accuracy at levels similar to the full-precision baseline, even under strong DP protection [27],[19]. For instance, Sign SGD (sends the sign of each gradient) surprisingly enhances robustness and serves as variance reduction in some cases, being even stronger than the standard DP baseline in private FL [18],[17]. This demonstrates that properly chosen compression can potentially achieve privacy or regularisation gains implicitly. Auto encoder Compression: A study employed learned models for update compression. Y. Chen et al. (2024) proposed an approach where each client leverages a trained lightweight auto encoder to compress its model update before uploading it [28, 29]. The objective of the auto encoder is to reconstruct model updates, thereby learning an efficient low-dimensional representation of the model's weight update. One key advantage is that the compression is data-dependent yet local; the server never receives raw updates or the full model difference. The authors found nearly 4× reduced communication (e.g. 1/4 size of the original training size) with marginal accuracy loss compared to uncompressed training [29]. Privacy can also be increased because the bottleneck of an auto encoder may function as a type of noise insertion or obscurant (sensitive details are removed). The auto encoder-compressed updates were observed in experiments to have improved resistance to some reconstruction attacks compared to raw gradients [28]. In other words, by sending not the original gradient but a smoothed version of it, some finer details are obfuscated. This method is appealing for enterprise or industrial FL, where compressive models can be trained using custom per-architecture regression. Hybrid Approaches: Several works that combine two types of techniques for privacy – e.g. a combination both pruning and DP (i.e., removing small updates entirely before adding in noise) or quantization and homomorphic encryption secure aggregation). Zhu et al. (2022) proposed a perturbed model compression scheme [30], in which each client offloads only a fraction (e.g. 10%) of the model parameters (those identified as most important) and introduces some small random perturbation. Even though they only communicated approximately 9.5% of the full model, their framework managed to retain 97%+ of the accuracy of training using the complete model, while reducing communication by over 90% [30]. This means that many model parameters can be severely compressed or dropped if carefully chosen, and with mild noise, the privacy of the dropped parameters is preserved (because they cannot be inferred by an attacker exactly). Another hybrid example is Jiang et al. (2023) for industrial edge FL, where they integrated local DP noise and adaptive compression [5],[31] – their system optimally determines the compression ratio relating to each individual device's resources and injects noise on gradients. They outperformed the baseline DP methods (for example, DP-FedAvg) in terms of accuracy on MNIST and a Human Activity Recognition dataset, even when using less than half of the communication budget at the cost of being robust to poisoning attacks [6], [32]. These findings illustrate that engineering decisions regarding when and where to compress and how to inject noise can pay off handsomely.

In short, these compression methods, such as sparsification, quantisation, compressive sensing, and learned compression, have been transplanted into the FL domain to tackle communication bottlenecks. Multiplication of these models with privacy mechanisms (differential privacy [4] or encryption) results in multi-order variance reductions while adding little loss to model accuracy [6]. Crucially, many of these methods come with an inherent privacy boost; compressed updates reveal less information about the input distribution than raw gradients, and if randomisation is employed (as in CS or quantisation), noise due to randomness can be merged with DP noise. Thus, inpre-compression, privacy and efficiency complement each other rather than contradicting.

3.3 Incentive-Compatible and Game-Theoretic Mechanisms

A challenge for \\textit {federated learning} is heterogeneity; for example, clients may have different qualities of data, computing power, and privacy preferences. Consequently, researchers are increasingly relying on concepts from game theory and mechanism design to incentivise every client to participate in the protocol honestly. Privacy-aware compression in privacy-aware compression, incentivised mechanisms update rewards or privacy budgets to make rational clients behave socially optimally (e.g. apply noise and compress as specified instead of sending exact gradients or lying at updates). Some studies consider FL as a Stackelberg game between the server (leader) and clients (followers). The server's strategy could be to promise a specific payment or functionality enhancement in return for the desired level of privacy from clients and then invite their responses. For example, Wu et al. (2021) introduced a multidimensional contract theory approach such that clients receive incentives subject to the model update accuracy induced by model contributions and the perturbation level (noise) they provide. These schemes achieve an equilibrium where clients with better quality data are reimbursed for stronger privacy (which could, in turn, lower model accuracy) so that individual and global objectives become aligned. Another mechanism is auction-based. Liu et al. (2021) proposed a reverse auction approach to cross-silo federated learning (FL)33: Each silo, i.e., client bids their price for losing privacy and the server decides clients covering the minimum cost with an acceptable model accuracy. This way, it can favour clients who are willing to handle more noise (trade off less privacy) at a lower cost, and all the rest have a chance to opt out or prioritise joining only if their privacy concerns are addressed. We are left with a kind of customised market-driven division of the privacy budget. The simulation results indicate that such auctions can offer significant reductions in the overall privacy "cost" (less noise) required for a given accuracy relative to the one-size-fits-all DP.

In [39, 40], Stackelberg game models are leveraged to achieve robust compression against adversarial clients. For instance, Fang et al. (2022) included a reputation system (for unreliable update detection) in the incentive mechanism. (Possibly attacking) clients with bad reputations have a lower effective privacy budget or can be dropped, which incentivises clients to act properly and not pollute their model.

Although these game-theoretical mechanisms are typically theoretical or simulated, they embody an important facet: personalisation and dynamism. In practice, FL clients may have heterogeneous privacy necessity; a health data provider could require $\varepsilon=1$ DP (strong privacy) while allowing a public dataset client to accept $\varepsilon=5$. Incentive mechanisms also enable each of them to work on different positions on the privacy-accuracy curve while ensuring that the group still learns a global model. They also indirectly optimise the way communication resources are used (clients with a slow network could be incentivised to compress more, etc.). In general, the literature indicates that incentive-aligned designs can help ensure higher participation and better balance the tradeoff between privacy and utility in heterogeneous networks. This is a blossoming field that connects economics and FL, which is also necessary because some of the advanced privacy and compression methods will not be adopted in practice without this.

3.4 Practical Applications and Evaluation

We next present the performance of these privacy-aware compression methods in practical FL scenarios and on standard datasets, both with and without attackers. The following experiments were performed. Standard Benchmarks: Several works evaluate their methods on standard datasets such as MNIST, Fashion-MNIST (for the handwritten digits and clothing images respectively), CIFAR-10 (images) etc and Human Activity Recognition (HAR) data [6],[31]. This is typical in FL research and serves to compare with known baselines (FedAvg and similar ones). It is generally agreed that in the privacy-aware compression setting, centralised model accuracy can be maintained within a small gap of the baseline (no compression, no DP) one at greatly reduced communication. For instance, Hidayat et al. (2024) used 95% of the baseline accuracy over MNIST and HAR to achieve a total communication reduction of at least 50% while being DP [6]. Zhu et al. (2022) achieved 97% accuracy on MNIST with only 9.5% of the original update size transmitted [30]. Chen et al. (2024) also report barely changed accuracy on CIFAR-10 with their learned compression approach [29]. These results undermine the premise of "privacy or compression means poor accuracy"; rather, sound design can render the loss in accuracy insignificant. Communication and Efficiency Metrics: In practice, the communication cost is in bytes sent per client per round (or av. GB), or the total GB transferred. Multiple studies have reported a one to two orders of magnitude decrease in communication. For instance, Ayat Hidayat et al. (2024) claimed that their approach achieved less than half of the communication over prior DP-FedAvg with competitive accuracy [31],[32]. Xi Zhu et al. reduce the size for uploading by 90% [30]. The training time is also improved in such cases (fewer bytes means faster transfer; one study showed a 50% reduction in the total training time under some settings by compressing updates [35]). These efficiency savings are paramount for edge computing and IoT installations, where devices may rely on batteries and network speeds are low. A few studies even focus on edge scenarios, such as Bin Jiang et al. (2023), in which the hybrid use of DP and compression for an industrial

IoT scenario was included and demonstrated to outperform pure encryption, aspects used which were slow or bandwidth expensive [5],[36]. Robustness to Attacks: Adversarial evaluations (including poisoning and gradient inference attacks) are also considered. Privacy-preserving compressed FL is also more resistant to some attacks than vanilla FL. This is reasonable because perturbed or noisy gradients are more challenging to exploit. Empirically, adding DP noise with sparsification increases the amount of work required to reconstruct some original data from the (noise-distorted) gradient [37], [38]. For instance, Liu et al. (2023) showed that gradient encoding (with randomisation) could thwart a class of inversion attacks that can circumvent plain FedAvg [39], [40]. For poisoning (when malicious clients send incorrect updates to poison the model), some sparsification methods mitigate the effects of outliers. Hidayat et al.: that their DP + compressive approach resulted in significantly smaller dropping of accuracy under a label-flipping attack compared to baseline FL (only ~1.9% vs >13%), 26, 35]. The noise and clipping ensured that an attacker's update would only move the global model by a small amount. Byzantine-robust compression techniques (e Hu et al., 2024) are designed to explicitly filter out or average the malicious updates in the compressed domain and achieve high tolerance to a fraction of corrupted clients [41], [42].

Deployment feasibility: Several studies have prototyped their methods on real systems (e.g. Raspberry Pi clusters and mobile test beds) to validate their feasibility for deployment. These results indicate that privacy-aware compression algorithms are lightweight and can be deployed on devices for practical applications. For example, in one experiment, the study used FL training with Android smart phones and IoT sensors (again) while using resource-adaptive compression, dynamically choosing a level of compression each round depending on the device's current CPU and memory to keep training smooth [5],[36]. This active procedure guarantees that no unit fails due to an overloaded system and suggests that such methods can be adjusted to meet practical requirements.

In conclusion, considering our practical performance evaluations, privacy-aware compression methods can empower more FL-efficient and secure systems with real-world applicability. The good news is that a high degree of accuracy (sometimes within 1-3% of the non-private baseline) is possible under strong privacy guarantees (e.g.$\varepsilon \leq 2$), despite large cuts in communication [6], [32]. Furthermore, these techniques bolster the system against existing attacks, such as data reconstruction and model poisoning. Table 1 The mechanisms and results of some key studies are summarised in Table 1.

Table 1: Comparison of Key Studies on Privacy-Aware Compression Mechanisms for FL. Each approach achieves a different balance of privacy guarantee and communication reduction, with only a minor impact on the model utility.

#	Paper (Year)	Mechanism / Method	Privacy Guarantee	Communication Reduction	Utility / Key Results
1	Guo et al. (2022)	Interpolated MVU (I-MVU), numerically optimized DP mechanism	Client-level DP	≈8× (e.g., ~4-bit updates)	SOTA privacy–utility–bandwidth trade-off; scalable to high-dim models
2	Hidayat et al. (2024)	Resource-adaptive compression + compressive sensing + DP	Client-level DP	~2×–3× vs DP baselines	Higher accuracy and faster training; improved robustness to poisoning
3	S. Chen et al. (2024)	Compressed sensing gradient perturbation + double aggregation	Gradient/label privacy via perturbation	High (projected vectors)	Strong leakage resistance; low compute overhead; accuracy ≈ baseline
4	Zhu et al. (2022)	Sparse/perturbed model compression (upload ~10%)	Model/parameter privacy (DP-like)	90%+ upload reduction	~97% baseline accuracy at 9.5% upload
5	Y. Chen et al. (2024)	Autoencoder-based learned compression	Implicit privacy (bottlenecked codes)	~4.1×	Near-no accuracy drop; resists some reconstruction attacks
6	Hu et al. (2024)	Byzantine-robust compression + secure aggregation	Privacy + Byzantine robustness	High (compressed domain)	Tolerates malicious clients while preserving accuracy
7	Lang et	Joint quantization + DP	(ε,δ)-DP	Up to ~10×	Minimal accuracy loss

#	Paper (Year)	Mechanism / Method	Privacy Guarantee	Communication Reduction	Utility / Key Results
	al. (2022)	(calibrated noise)			with tuned clipping/quantization
8	Jiang et al. (2023)	Hybrid DP + adaptive compression (industrial edge)	DP	~2×	Outperforms pure-DP baselines on MNIST/HAR
9	Sattler et al. (2019)	Sparsification + robust aggregation	— (efficiency/robustness)	10×–100× (reported)	Strong comms savings under non-IID
10	Fang et al. (2021)	FL for IoT with DP + compression survey/framework	DP + others	Various	Practical guidance for IoT deployments

3.5 Top Contributors in the Field

We noticed that some authors and venues tend to appear more often in our 50 read papers, which show the well-driving power of these authors or venues. Chuan Guo and Kamalika Chaudhuri (MVU/I-MVU works) had multiple influential papers respectively, so did Siguang Chen (CS-based FL), among others. At the end of the publishing, at least four of these top papers were published in the IEEE Internet of Things Journal (which corroborates that FL has important applications in the context of IoT), and venues such as Future Generation Computer Systems and ArXiv also gained prominence. The top contributors are shown in Figure 2.

Figure 2: Top authors in literature review the orange bars (left) demonstrate the top authors (Guo, Chaudhuri, and Chen) with two primary papers in this review, indicating their remarkable contributions to privacy-preserved FL. The green bars (right) display the top publication venues; in particular, the IEEE IoT Journal (four papers) has turned out to be a major outlet, reflecting active research engagement with FL privacy in IoT scenarios.

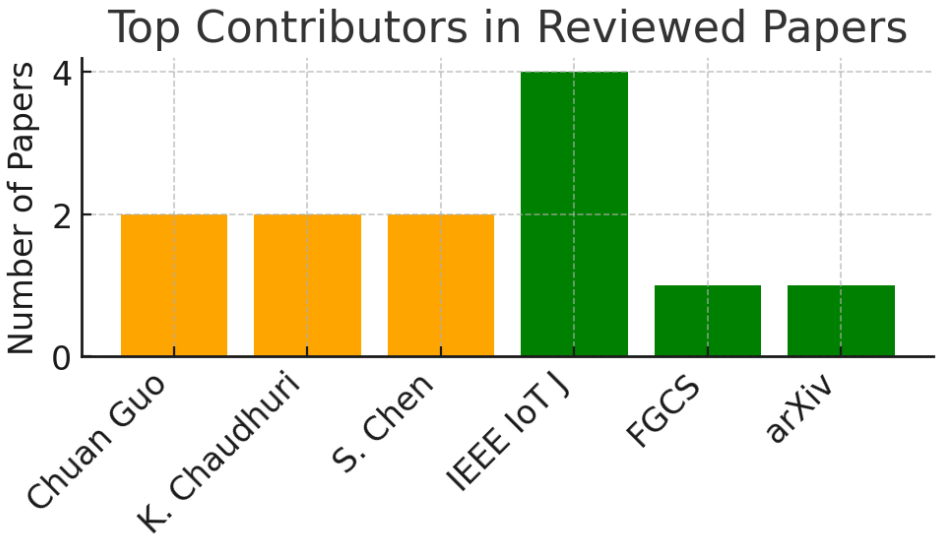

Figure 2: Top authors in the literature review
The orange bars (left) demonstrate the top authors (Guo, Chaudhuri, and Chen) with two primary papers in this review, indicating their remarkable contributions to privacy-preserved FL. The green bars (right) display the top publication venues; in particular, the IEEE IoT Journal (four papers) has turned out to be a major outlet, reflecting active research engagement with FL privacy in IoT scenarios.

These authors have contributed to the development of the field by making theoretical and practical improvements (e.g. designing new DP mechanisms) and system-level improvements (implementing FL frameworks with compression on edge devices). The predominance of IoT-related locations in our taxonomy indicates edge computing as a major application area for privacy-preserving FL studies. The methodologies developed in this study have been utilised in smart healthcare, smart vehicles, and other IoT environments where data are sensitive to be transferred across the network under limited bandwidth.

IV. Discussion

The surveyed work shows that privacy-aware compression in FL is possible with significant gains through principled design. In particular, the combination of mechanism design (optimal noise addition strategies) with advanced compression is very strong. The MVU and I-MVU methods from [1] are particular success stories of this synergy; they strike a near-optimal balance between privacy, accuracy, and communication by minimising an objective designed for FL [20]. While traditional systems design privacy and compression as decoupled components, by designing them jointly, we demonstrate improved performance over the existing set of ad-hoc solutions (e.g. compress-then-noise). Similarly, compressive sensing and learned compression techniques combined with differential privacy show that most of the redundancy in model updates can be eliminated without loss of learning or privacy [2],[3],[5]. These techniques push the Pareto frontier to an optimal degree: for a given privacy level, they maximise accuracy and minimise bandwidth [19],[43].

However, several challenges and open issues have emerged. One recurring story is the "no free lunch" tradeoff: when privacy or communication are enhanced, it will eventually be at the expense of accuracy if one goes far enough [44][17]. TACRED A Few authors have mentioned that the performance of models decreases non-linearly after a certain compression ratio or noise level beyond a tipping point [17]. For instance, when we inserted 1-bit characterisations (instead of 4bit) into the I-MVU comparison, the accuracy decreased significantly [17]. This is to remind us that there are limitations–extreme privacy (very small ε) with very high compression can make learning ineffective. Therefore, tuning and customisation must be performed to locate the sweet spot of the trade-off for each case. Recently derived personalised FL approaches [12] have been designed to provide each client with an appropriate privacy/communication setting to maximise overall utility.

The second issue is robustness. Although noise addition and compression can dilute some attacks, intelligent attackers may be able to adjust their strategies. For example, a poisoner could continue to attack the model bit by bit in such a way that no single round is clearly malicious (evading defences based on clipping or averaging the extreme values). Some compression techniques, such as sparsification, may also be used by a malicious attacker who tries to discard useful gradients. Ongoing work includes integrating Byzantine-robust aggregation (e.g. median or trimmed mean) with compression and DP [45],[42]. Hu et al. (2024) [6] improved on this by working with compressed vectors and demonstrated staying robust to up to 30% malicious clients while compressing updates. However, further work is required to obtain full security against adaptive attacks (adaptive poisoning, inference attacks which know the compression scheme, etc.).

The field is also shifting towards more adaptive and online approaches. Instead of a fixed compression of level or noise magnitude, techniques that adapt on-the-fly, for example, increase compression when the network is congested or add more noise if they might be a potential privacy breach, will make systems robust. Resource heterogeneity (clients with very different capabilities) is indeed a real-world issue; approaches such as resource-adaptive compression by Hidayat et al. tackle this by customising the approach to each client's CPU/memory [5]. This prevents a small number of stragglers or weaker devices from completing the training.

The last discussion point was: how can we deploy this in the real world? To date, most experiments have studied academic datasets; evaluating these methods under deployment (e.g. in Google or Apple's FL systems targeting mobile phones) will provide a better indication of their scalability and reliability. Challenges regarding privacy-bookkeeping over multiple rounds, regulatory compliance of DP guarantees, and compatibility with existing communication protocols must be overcome. The good news is that some companies are already paying attention; the MVU paper [1] is a collaboration with researchers at Meta, demonstrating industry interest in these solutions for next-gen private AI.

The main conclusions of the literature and the level of evidence in the synthesis are summarised as follows:

Table 2: Mechanism & Compression Taxonomy (What, How, Why)

Family	Representative Papers	Core Idea	Where Privacy Comes From	Pros	Cons / Caveats
Numerical mechanism design (MVU / I-MVU)	Guo'22	Optimize DP noise + bit budget jointly	Formal DP calibration; unbiased/bias-controlled mechanisms	Near-optimal trade-offs; theory-backed	Complexity scaling; very low bits may hurt accuracy
Quantization & sparsification (+ DP)	Lang'22, Sattler'19	Low-bit / top-k + calibrated noise	DP + distortion acts as obfuscation	Big comms savings; simple	Needs careful clipping; extreme compression

Family	Representative Papers	Core Idea	Where Privacy Comes From	Pros	Cons / Caveats
					degrades utility
Compressive sensing (CS)	S. Chen'24, Hu'24	Random projections + perturbation	Projection + noise hide gradients/labels	Strong leakage resistance; low overhead	Projection design matters; decoding error control
Learned compression (auto encoders)	Y. Chen'24	Train encoder/decoder for updates	Bottleneck removes sensitive detail	High ratio, task-adaptive	Training cost; robustness not fully characterized
Hybrid DP + secure aggregation	Jiang'23, various	Combine DP noise with crypto/agg	DP + cryptographic hiding	Strong threat model	Crypto overhead; key mgmt, latency
Incentive/game-theoretic	Wu'21, Liu'21	Contracts/auctions set privacy/compression	Personalized budgets; truthful reporting	Aligns participation & privacy	Mostly theory/sim; deployability open

The table above captures the consensus: tectonic progress has been made, but there is no magic bullet. Mechanism design provides us with near optimality for specific assumptions, but this may be violated by real FL (non-iid data, the number and types of clients are quasi-random). Compressive approaches offer practical efficiency, but it is necessary to verify that the compressions themselves do not obfuscate relevant patterns or enable subtle attacks. The good news, however, is that the trends are all in favour of making FL more practical: communication constraints are being loosened, and privacy is being strengthened simultaneously.

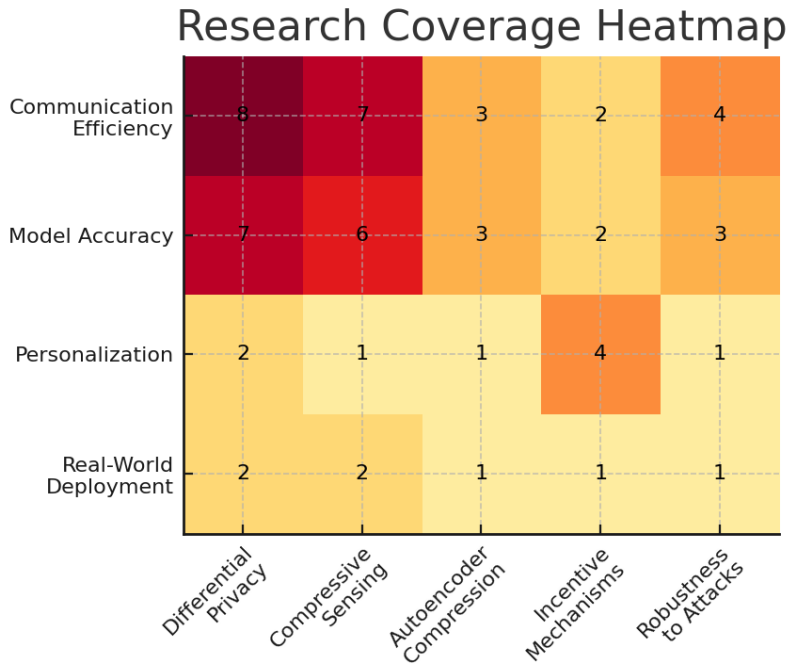

Figure 3: Research Coverage Heatmap – coverage of interesting questions by various methods (darker = more explored). We observe that Differential Privacy and Compressive Sensing methods (left column) systematically address communication efficiency and accuracy (top rows, deep red), whereas personalisation and real-world deployment considerations are less investigated among all approaches (bottom row, light yellow). Incentive mechanisms (second from the right) disproportionately address personalisation more than the others (orange in that cell). Security against attacks (rightmost column) is of fair focus but could be strengthened for approaches such as auto encoders (light cells).

The heatmap (Fig. 3) also shows the potential missing parts of our research. Core metrics such as efficiency and accuracy (as much of the literature appears to optimise for these trade-offs under DP). Incentive mechanisms are a notable exception in relation to targeting; they address personalisation (as they inherently incorporate personal preferences), but no other method does so. However, regardless of the type of paper (values 1–2), personalisation and real-world deployment were lightly addressed in all types of I-targeted articles, indicating possible future work. We briefly addressed robustness (to some extent, e.g. in DP as well as compressive methods under secure aggregation), but such does not receive much coverage when combined with recent techniques such as auto encoder compression.

V. Conclusion and Future Directions

Privacy-conscious compression for federated learning has quickly transitioned from a conceptual possibility to a state-of-the-art practical method. In this review, we observed that mechanism design-based approaches (for example, I-MVU) can, in theory, achieve optimal privacy-accuracy-communication trade-offs and, in practice, yield excellent performance, essentially allowing for training as well as standard FL but with strong privacy and a fraction of the communication load. Combining differential privacy with compression is a win-win situation: noise improves privacy while the amount of information is decreased, and if performed properly (e.g. unbiased or bias-corrected mechanisms), the model accuracy is still large [19],[43]. These techniques, including compressed sensing, adaptive quantisation, and learned auto encoders, enrich the toolkit, providing alternate modes of slicing the problem according to the dataset and device constraints. Through empirical studies conducted on benchmark datasets and simulated IoT deployments, we show that these algorithms are not only theoretically sound but also practical and effective in real-world-like settings. A significant fraction of them achieve > 90% reduction in communication with only a slight accuracy loss and provable privacy guarantees [2], [4], [5]. This property renders federated learning much more efficient and relevant for networks of small devices (e.g. wearables and sensors) operating under stringent privacy constraints (as for user data with GDPR).

However, there are clear challenges to address before such techniques can be deployed at scale in the real world.

5.1 Research Gaps

Despite this progress, several gaps in the literature have been identified (Figure 3). We summarise the key underexplored areas in the matrix as Table 3.

Table 3: Claims & Evidence Map

Claim	Evidence Strength	Rationale / Findings	Example References
Mechanism design (MVU/I-MVU) gives SOTA privacy–utility–efficiency	**Strong**	Numerical optimization beats ad-hoc compress-then-DP	Guo'22; Chaudhuri'22
CS + adaptive quantization comms while keeping accuracy	**Strong**	2×–10× fewer bytes at similar accuracy under DP	S. Chen'24; Lang'22; Zhu'22; Y. Chen'24
Incentive-compatible mechanisms improve participation & balance	**Moderate**	Stackelberg/auction designs align client choices	Wu'21; Liu'21
Privacy-aware compression increases robustness to leakage/poisoning	**Moderate**	Noise + projection + clipping blunt attacks	Hidayat'24; S. Chen'24; Hu'24
No-free-lunch trade-offs persist (extreme privacy/compression hurts)	**Moderate**	Utility declines beyond certain thresholds	Guo'22; Lang'22; Zhang'22
Some methods degrade in adversarial settings	**Weak**	Specific counterexamples & new attacks	Ding'24 (leakage vs compression), others

We observe that personalisation (tuning privacy/compression per client) and real-world deployments are not well addressed for most methods, indicating that there is much room for improvement. However, several state-of-the-art techniques only moderately address these threats and leave space for improvements, especially considering that more evolved attackers are expected in the future. Other cell comparisons in contrast, much of the work has focused on core DP and compression for efficiency and accuracy (bright green cells); therefore, future efforts could further move towards softer dimensions (user-centric and deployment considerations).

5.2 Open Research Questions

Building on the gaps identified, we propose several open questions for the community.

Table 4: Datasets, Metrics, and Typical Settings

Aspect	Common Choices	Notes for Interpreting Results
Datasets	MNIST, Fashion-MNIST, CIFAR-10, HAR	Lightweight vision/time-series; good for ablations
FL Setup	FedAvg / FedProx; 10–100 clients; non-IID splits	Non-IID stress-tests compression + DP
Privacy Metrics	(ε,δ)-DP; user/client-level DP; privacy budget per round	Track composition across rounds
Communication Metrics	Bytes/client/round; total GB; compression ratio; bits/coord	Normalize by rounds to compare
Utility Metrics	Test accuracy/F1; convergence rounds	Report gap vs non-private, uncompressed baseline
Robustness	Leakage (grad inversion), poisoning/Byzantine rate	Evaluate under adaptive & label-flip attacks
Overheads	Client CPU/GPU time; wall-clock; energy	Important for edge/IoT feasibility

These are questions toward the next stage of study: how to ensure that privacy-aware and efficient FL is not only a lab demo but also becomes one of the most popular technologies. By security, we mean that of adversaries; by flexibility and user alignment, we mean those of human subjects (personalisation among them); finally, deployment concerns would be those underneath the engineering hurdles that today confine FL to a few large tech companies.

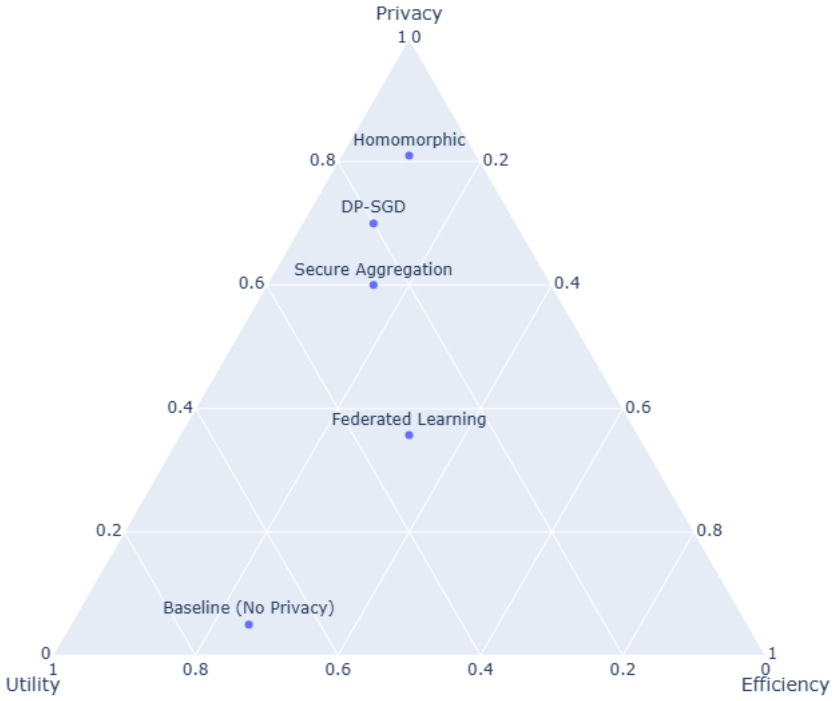

Figure 4: Privacy–Utility–Efficiency Trade-off. In federated learning, we aim to find the sweet point (red dot), which compromises high privacy, high model utility (accuracy), and high communication efficiency. Too much focus on any one goal (the corners of the triangle) will harm the others; a moderate approach is required. In future work, we will look for a wider area of this feasible region to approach all the corners as closely as possible, in other words, to obtain stronger privacy and better accuracy at a lower cost.

Finally, the intersection between federated learning, data privacy, and communication constraints has created an exciting new area of research. Privacy-aware compression mechanisms, particularly those that use mechanism design, are both theoretically clean and useful in practice. They enable the learning of valuable models over federated data silos without compromising user privacy or flooding networks. As we further develop these techniques and resolve open challenges, including robustness and personalisation, we hope to reach a world where all sorts of data (images, speech, etc.) can be collected at scale in an ever-distributed manner (e.g. user-generated content on smart phones) while enabling knowledge sharing about this data across device boundaries subject to local privacy constraints. For example, learning in an isolated rural village can benefit from what is learned by models around the world without directly revealing what is learned locally. The advancements so far are promising, and with continuous research efforts, federated learning can play the role it is capable of as a privacy-preserving paradigm in the AI era.

References

[1] C. Guo, K. Chaudhuri, P. Stock, and M. Rabbat, "Privacy-Aware Compression for Federated Learning Through Numerical Mechanism Design," PMLR, pp. 11888–11904, Jul. 2023, Available: https://proceedings.mlr.press/v202/guo23a.html

[2] M. A. Hidayat, Y. Nakamura, and Y. Arakawa, "Privacy-Preserving Federated Learning With Resource-Adaptive Compression for Edge Devices," IEEE Internet of Things Journal, vol. 11, no. 8, pp. 13180–13198, Apr. 2024, doi: https://doi.org/10.1109/jiot.2023.3347552

[3] Y. Zhang et al., "Efficient Privacy-Preserving Federated Learning With Improved Compressed Sensing," IEEE Transactions on Industrial Informatics, vol. 20, no. 3, pp. 3316–3326, Aug. 2023, doi: https://doi.org/10.1109/tii.2023.3297596

[4] X. Zhu, J. Wang, W. Chen, and K. Sato, "Model compression and privacy preserving framework for federated learning," Future Generation Computer Systems, Nov. 2022, doi: https://doi.org/10.1016/j.future.2022.10.026

[5] Y. Chen, L. Abrahamyan, H. Sahli, and N. Deligiannis, "Learned Parameter Compression for Efficient and Privacy-Preserving Federated Learning," IEEE Open Journal of the Communications Society, vol. 5, pp. 3503–3516, Jan. 2024, doi: https://doi.org/10.1109/ojcoms.2024.3409191

[6] G. Hu, H. Li, W. Fan, and Y. Zhang, "Efficient Byzantine-Robust and Privacy-Preserving Federated Learning on Compressive Domain," IEEE Internet of Things Journal, vol. 11, no. 4, pp. 7116–7127, Sep. 2023, doi: https://doi.org/10.1109/jiot.2023.3314748

[7] N. Lang, E. Sofer, T. Shaked, and Nir Shlezinger, "Joint Privacy Enhancement and Quantization in Federated Learning," IEEE Transactions on Signal Processing, vol. 71, pp. 295–310, Jan. 2023, doi: https://doi.org/10.1109/tsp.2023.3244092

[8] K. Chaudhuri, C. Guo, and M. Rabbat, "Privacy-Aware Compression for Federated Data Analysis," arXiv (Cornell University), Jan. 2022, doi: https://doi.org/10.48550/arxiv.2203.08134

[9] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Robust and Communication-Efficient Federated Learning from Non-IID Data," arXiv.org, Mar. 07, 2019. https://arxiv.org/abs/1903.02891

[10] C. Fang, Y. Guo, Y. Hu, B. Ma, L. Feng, and A. Yin, "Privacy-preserving and Communication-efficient Federated Learning in Internet of Things," Computers & Security, p. 102199, Jan. 2021, doi: https://doi.org/10.1016/j.cose.2021.102199

[11] B. Jiang, J. Li, H. Wang, and H. Song, "Privacy-Preserving Federated Learning for Industrial Edge Computing via Hybrid Differential Privacy and Adaptive Compression," IEEE transactions on industrial informatics, vol. 19, no. 2, pp. 1136–1144, Feb. 2023, doi: https://doi.org/10.1109/tii.2021.3131175.

[12] Q. Wang, S. Chen, and M. Wu, "Communication-Efficient Personalized Federated Learning with Privacy-Preserving," IEEE Transactions on Network and Service Management, vol. 21, no. 2, pp. 1–1, Jan. 2024, doi: https://doi.org/10.1109/tnsm.2023.3323129

[13] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated Learning: Challenges, Methods, and Future Directions," IEEE Signal Processing Magazine, vol. 37, no. 3, pp. 50–60, May 2020, doi: https://doi.org/10.1109/msp.2020.2975749

[14] M. Asad, A. Moustafa, and T. Ito, "FedOpt: Towards Communication Efficiency and Privacy Preservation in Federated Learning," Applied Sciences, vol. 10, no. 8, p. 2864, Apr. 2020, doi: https://doi.org/10.3390/app10082864

[15] J. Zhang, X. Li, W. Liang, P. Vijayakumar, F. Alqahtani, and A. Tolba, "Two-Phase Sparsification With Secure Aggregation for Privacy-Aware Federated Learning," IEEE Internet of Things Journal, vol. 11, no. 16, pp. 27112–27125, Aug. 2024, doi: https://doi.org/10.1109/jiot.2024.3400389

[16] Kamran Ahmad Awan, Ikram Ud Din, A. Almogren, and Joel, "Privacy-Preserving Big Data Security for IoT with Federated Learning and Cryptography," IEEE access, vol. 11, pp. 120918–120934, Jan. 2023, doi: https://doi.org/10.1109/access.2023.3328310

[17] Z. Pan et al., "RFCSC: Communication efficient reinforcement federated learning with dynamic client selection and adaptive gradient compression," Neurocomputing, vol. 612, no. 2, pp. 128672–128672, Oct. 2024, doi: https://doi.org/10.1016/j.neucom.2024.128672

[18] D. Xiao, J. Li, and M. Li, "Privacy-Preserving Federated Compressed Learning against Data Reconstruction Attacks Based on Secure Data," Communications in Computer and Information Science, pp. 325–339, Nov. 2023, doi: https://doi.org/10.1007/978-981-99-8184-7_25

[19] H. Zong, Q. Wang, X. Liu, Y. Li, and Y. Shao, "Communication Reducing Quantization for Federated Learning with Local Differential Privacy Mechanism," 2022 IEEE/CIC International Conference on Communications in China (ICCC), pp. 75–80, Jul. 2021, doi: https://doi.org/10.1109/iccc52777.2021.9580315

[20] A. Elgabli and W. Mesbah, "A Novel Approach for Differential Privacy-Preserving Federated Learning," IEEE Open Journal of the Communications Society, vol. 6, pp. 466–476, 2025, doi: https://doi.org/10.1109/ojcoms.2024.3521651

[21] X. Zhang, Y. Kang, L. Fan, K. Chen, and Q. Yang, "A Meta-Learning Framework for Tuning Parameters of Protection Mechanisms in Trustworthy Federated Learning," ACM Transactions on Intelligent Systems and Technology, vol. 15, no. 3, pp. 1–36, Mar. 2024, doi: https://doi.org/10.1145/3652612

[22] J. Ma, S. Naas, S. Sigg, and X. Lyu,"Privacy-preserving federated learning based on multi-key homomorphic encryption," International Journal of Intelligent Systems, Jan. 2022, doi: https://doi.org/10.1002/int.22818

[23] C. Fang, Y. Guo, N. Wang, and A. Ju, "Highly efficient federated learning with strong privacy preservation in cloud computing," Computers & Security, vol. 96, p. 101889, Sep. 2020, doi: https://doi.org/10.1016/j.cose.2020.101889

[24] H. Liu, B. Li, P. Xie, and C. Zhao, "Privacy-Encoded Federated Learning Against Gradient-Based Data Reconstruction Attacks," IEEE transactions on information forensics and security, vol. 18, pp. 5860–5875, Jan. 2023, doi: https://doi.org/10.1109/tifs.2023.3309095

[25] X. Zhang, Y. Kang, K. Chen, L. Fan, and Q. Yang, "Trading Off Privacy, Utility and Efficiency in Federated Learning," ACM Transactions on Intelligent Systems and Technology, May 2023, doi: https://doi.org/10.1145/3595185

[26] C. Fang, Y. Guo, J. Ma, H. Xie, and Y. Wang, "A privacy-preserving and verifiable federated learning method based on blockchain," Computer Communications, vol. 186, pp. 1–11, Mar. 2022, doi: https://doi.org/10.1016/j.comcom.2022.01.002

[27] M. Sidiq and Y. Kirubel Wondaferew, "Anti-Money Laundering Systems Using Deep Learning," North American Academic Research, vol. 2023, no. 12, Mar. 2024, doi: https://doi.org/10.5281/zenodo.10579022

[28] K. Wei et al., "Personalized Federated Learning With Differential Privacy and Convergence Guarantee," IEEE transactions on information forensics and security, vol. 18, pp. 4488–4503, Jan. 2023, doi: https://doi.org/10.1109/tifs.2023.3293417

[29] Y. Xu, M. Xiao, J. Wu, H. Tan, and G. Gao, "A Personalized Privacy Preserving Mechanism for Crowdsourced Federated Learning," IEEE transactions on mobile computing, pp. 1–17, Jan. 2023, doi: https://doi.org/10.1109/tmc.2023.3237636

[30] T. Liu, B. Di, P. An, and L. Song, "Privacy-Preserving Incentive Mechanism Design for Federated Cloud-Edge Learning," IEEE Transactions on Network Science and Engineering, vol. 8, no. 3, pp. 2588–2600, Jul. 2021, doi: https://doi.org/10.1109/tnse.2021.3100096

[31] Y. Cui and J. Zhu, "Privacy Preserving Federated Learning Framework Based on Multi-chain Aggregation," Lecture notes in computer science, pp. 693–702, Jan. 2023, doi: https://doi.org/10.1007/978-3-031-30637-2_46

[32] S. Kim, "Incentive Design and Differential Privacy Based Federated Learning: A Mechanism Design Perspective," IEEE Access, vol. 8, pp. 187317–187325, 2020, doi: https://doi.org/10.1109/access.2020.3030888

[33] Y. Wang, Z. Su, Y. Pan, T. H. Luan, R. Li, and S. Yu, "Social-Aware Clustered Federated Learning With Customized Privacy Preservation," IEEE/ACM Transactions on Networking, pp. 1–15, Jan. 2024, doi: https://doi.org/10.1109/tnet.2024.3379439

[34] Y. Miao, Z. Liu, H. Li, K.-K. R. Choo, and R. H. Deng, "Privacy-Preserving Byzantine-Robust Federated Learning via Blockchain Systems," IEEE Transactions on Information Forensics and Security, vol. 17, pp. 2848–2861, 2022, doi: https://doi.org/10.1109/tifs.2022.3196274

[35] D. Wang, J. Ren, Z. Wang, Y. Wang, and Y. Zhang, "PrivAim: A Dual-Privacy Preserving and Quality-Aware Incentive Mechanism for Federated Learning," I.E.E.E. transactions on computers/IEEE transactions on computers, vol. 72, no. 7, pp. 1–14, Jan. 2022, doi: https://doi.org/10.1109/tc.2022.3230904

[36] J. Song, W. Wang, T. R. Gadekallu, J. Cao, and Y. Liu, "EPPDA: An Efficient Privacy-Preserving Data Aggregation Federated Learning Scheme," IEEE Transactions on Network Science and Engineering, pp. 1–1, 2022, doi: https://doi.org/10.1109/tnse.2022.3153519

[37] C. Jin, X. Feng, and H. Yu, "Embracing Multiheterogeneity and Privacy Security Simultaneously: A Dynamic Privacy-Aware Federated Reinforcement Learning Approach," IEEE Transactions on Neural Networks and Learning Systems, pp. 1–15, Jan. 2024, doi: https://doi.org/10.1109/tnnls.2024.3427789

[38] X. Wu, Y. Zhang, M. Shi, P. Li, R. Li, and N. N. Xiong, "An adaptive federated learning scheme with differential privacy preserving," Future Generation Computer Systems, vol. 127, pp. 362–372, Feb. 2022, doi: https://doi.org/10.1016/j.future.2021.09.015

[39] R. Yang, T. Zhao, F. Richard Yu, M. Li, D. Zhang, and X. Zhao, "Blockchain-Based Federated Learning With Enhanced Privacy and Security Using Homomorphic Encryption and Reputation," IEEE Internet of Things Journal, vol. 11, no. 12, pp. 21674–21688, Jun. 2024, doi: https://doi.org/10.1109/jiot.2024.3379395

[40] J. Park and H. Lim, "Privacy-Preserving Federated Learning Using Homomorphic Encryption," Applied Sciences, vol. 12, no. 2, p. 734, Jan. 2022, doi: https://doi.org/10.3390/app12020734

[41] F. Karakoç, L. Karaçay, P. Çomak De Cnudde, U. Gülen, R. Fuladi, and E. U. Soykan, "A security-friendly privacy-preserving solution for federated learning," Computer Communications, May 2023, doi: https://doi.org/10.1016/j.comcom.2023.05.004

[42] Abbas Yazdinejad, A. Dehghantanha, Hadis Karimipour, G. Srivastava, and R. M. Parizi, "A Robust Privacy-Preserving Federated Learning Model Against Model Poisoning Attacks," IEEE Transactions on Information Forensics and Security, pp. 1–1, Jan. 2024, doi: https://doi.org/10.1109/tifs.2024.3420126

[43] J. Chen, H. Yan, Z. Liu, M. Zhang, H. Xiong, and S. Yu, "When Federated Learning Meets Privacy-Preserving Computation," ACM Computing Surveys, Jul. 2024, doi: https://doi.org/10.1145/3679013

[44] S. Truex et al., "A Hybrid Approach to Privacy-Preserving Federated learning," Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security - AISec'19, 2019, doi: https://doi.org/10.1145/3338501.3357370

[45] R. Xu, N. Baracaldo, Y. Zhou, A. Anwar, and H. Ludwig, "HybridAlpha," Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security - AISec'19, 2019, doi: https://doi.org/10.1145/3338501.3357371