computers

MDPI

Article

Zero-Knowledge Federated Learning for Privacy-Preserving 5G Authentication

Ahmed Lateef Salih Al-Karawi [1,2] and Rafet Akdeniz [3,*]

[1] Department of Computer Engineering, Faculty of Engineering, Istanbul Aydin University, Istanbul 34295, Turkey; ahmedal-karawi@stu.aydin.edu.tr
[2] Defne Telekomünikasyon A.Ş., Maslak Mahallesi, Maslak Meydan Sokak, Spring Giz Plaza, No:5, İç Kapı:37, Kat:9, Sarıyer, Istanbul 34485, Turkey
[3] Department of Computer Engineering, Faculty of Engineering and Natural Sciences, Atlas University, Anadolu Caddesi No:40, Kağıthane, Istanbul 34408, Turkey
* Correspondence: rafet.akdeniz@atlas.edu.tr; Tel.: +90-5423-427-501

Abstract

Fifth-generation (5G) networks are facing critical security challenges in device authentication for massive Internet of Things deployments while preserving privacy. Traditional federated learning approaches depend on the computationally expensive homomorphic encryption to protect model gradients, resulting in substantial latency and communication overhead, leading to impractical energy consumption for resource-constrained 5G devices. This paper proposes Zero-Knowledge Federated Learning (ZK-FL), eliminating homomorphic encryption by enabling devices to prove model correctness without revealing gradients. Our approach integrates zero-knowledge proofs with FL updates, where each device generates a proof $\text{Proof}_i = \text{ZK}(\text{Gradient}_i, \text{Hash}_i)$, demonstrating computational integrity. The experimental results from 10,000 authentication attempts demonstrate ZK-FL achieves 78.4 ms average authentication latency versus 342.5 ms for homomorphic encryption-based FL (77% reduction), proof sizes of 0.128 kB versus 512 kB (99.97% reduction), and energy consumption of 284.5 mJ versus 6525 mJ (95% reduction), while maintaining 99.3% authentication success rate with formal privacy guarantees. These results demonstrate ZK-FL enables practical privacy-preserving authentication for massive-scale 5G deployment.

Keywords: 5G authentication; zero-knowledge proofs; federated learning; privacy preservation; network security

1. Introduction

Fifth-generation (5G) wireless networks enable unprecedented connectivity for billions of devices through enhanced mobile broadband, ultra-reliable low-latency communications, and massive machine-type communications [1]. However, this massive scale introduces critical security challenges in authentication mechanisms that must ensure robust security while preserving user privacy [2].

Traditional 5G authentication relies on centralized credential verification where user equipment credentials are validated by core network functions [3]. While these mechanisms provide strong security through protocols like 5G-AKA, they suffer from privacy preservation limitations, single points of failure, and scalability constraints when handling billions of IoT devices [4].

Federated learning (FL) enables collaborative machine learning without centralizing sensitive data, allowing devices to jointly train authentication models while keeping data

Check for updates

Academic Editor: Paolo Bellavista

Received: 2 February 2026
Revised: 18 March 2026
Accepted: 19 March 2026
Published: 26 March 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

localized [5]. However, conventional FL faces gradient leakage vulnerabilities through gradient inversion and membership inference attacks [6]. State-of-the-art privacy-preserving FL systems employ homomorphic encryption (HE) [7], but HE introduces prohibitive computational overhead, communication costs, and energy consumption impractical for resource-constrained 5G devices [8]. HE-FL incurs authentication latencies exceeding 300 ms, proof sizes above 500 kB, and energy consumption over 6J per cycle [9].

Zero-knowledge proofs (ZKPs) enable proving statement validity without revealing underlying information [10]. Recent advances in zk-SNARKs and zk-STARKs achieve practical proof generation with succinct proof sizes [11,12], making ZKPs suitable for resource-constrained 5G environments.

This paper proposes Zero-Knowledge Federated Learning (ZK-FL), integrating zero-knowledge proofs with federated learning for privacy-preserving 5G authentication without homomorphic encryption burden. Each device generates succinct proof $\text{Proof}_i = \text{ZK}(\text{Gradient}_i, \text{Hash}_i)$ demonstrating local model update correctness without revealing gradient values. The 5G core network verifies these proofs ensuring computational integrity while preserving complete privacy.

The main contributions are as follows: (1) a novel ZK-FL framework eliminating homomorphic encryption through zero-knowledge proofs for 5G authentication; (2) a comprehensive system architecture integrating ZK-FL with 5G network functions; (3) a rigorous mathematical formulation with formal security proofs; (4) extensive evaluation methodology including authentication latency, proof size, verification time, network overhead, and energy consumption; and (5) a comprehensive experiments demonstrating 77% latency reduction, 99.97% proof size reduction, and 95% energy savings while maintaining 99.3% authentication success rate.

2. Related Work

2.1. Privacy-Preserving Federated Learning

McMahan et al. [5] introduced federated learning enabling collaborative training without centralizing data. However, gradient information leaks sensitive data through gradient inversion [6] and membership inference attacks [13]. Differential privacy (DP) approaches add calibrated noise to gradients [14,15], but face utility–privacy trade-offs degrading model accuracy. Secure multi-party computation enables joint computation [16] but faces scalability challenges in massive 5G deployments. Homomorphic encryption based approaches [7,17] incur prohibitive computational costs with encryption operations consuming 150–500 ms per gradient and ciphertext expansion factors of 2000–4000× [18]. Recent approximate homomorphic encryption schemes such as CKKS and BFV have been proposed to improve the efficiency of encrypted machine learning operations; however, they still introduce significant ciphertext expansion and computational overhead in large-scale federated learning environments [19,20].

2.2. Zero-Knowledge Proof Systems

Traditional ZKP systems require multiple interaction rounds [21]. Modern constructions achieve non-interactivity through Fiat–Shamir transformation [22]. zk-SNARKs provide constant-size proofs through pairing-based cryptography [11]. Groth16 [23] achieves 128–256 byte proofs with verification under 10 ms but requires trusted setup. zk-STARKs eliminate trusted setups through information-theoretic security [12] but generate larger proofs (40–200 kB). Significant improvements in proof size efficiency have been reported in recent work on advanced zero-knowledge proof systems [24,25].

2.3. 5G Authentication and Security

3GPP TS 33.501 defines 5G authentication including 5G-AKA and 5G-EAP-AKA [26]. Research identified privacy vulnerabilities including IMSI catching attacks [27], authentication relay attacks [28], and downgrade attacks [29]. Machine learning-based authentication enhancement applies deep learning for anomaly detection [30] and FL for collaborative intrusion detection [31], but these works do not address gradient privacy. Recent proposals explored anonymous authentication [4] and blockchain-based approaches [32] but face scalability limitations. No prior work integrates zero-knowledge proofs with federated learning for 5G authentication providing rigorous privacy guarantees without homomorphic encryption overhead.

3. System Model and Design Objectives

3.1. Network Model

We consider 5G network architecture following 3GPP Release 16 with service-based interfaces. The network comprises N user equipment devices $\{UE_1, \ldots, UE_N\}$ connecting through gNB base stations to core network functions including AMF, SMF, UPF, AUSF, and UDM. Participating devices are partitioned into M collaborative learners $\mathcal{C} = \{C_1, \ldots, C_M\}$ where $M \ll N$. Each client C_i maintains local dataset $\mathcal{D}_i$ with heterogeneous non-IID authentication records. The global authentication model is parameterized by weights $w \in \mathbb{R}^d$. At training round t, the server broadcasts current global model $w^{(t)}$ to selected clients. Each client performs local training for E epochs, computing local gradient $g_i^{(t)} = \nabla \mathcal{L}(w^{(t)}; \mathcal{D}_i)$.

3.2. Threat Model

We adopt an honest-but-curious adversary model where participants follow protocol specifications, but attempt to infer private information. Threats include: gradient inference attacks where adversaries reconstruct training samples from transmitted gradients [6]; model inversion attacks where adversaries query models to extract training data distributions [33]; membership inference attacks determining whether specific records were included in training datasets [13]; and Byzantine attacks where malicious clients submit manipulated gradients attempting model poisoning [34]. We assume the central aggregation server and communication channels are trusted within the secured 5G core network. The system model assumes that the authentication server operates within the secured 5G core network and follows the protocol specification correctly while potentially attempting to infer additional information from received messages. This corresponds to an honest-but-curious adversarial model commonly adopted in privacy-preserving federated learning systems. In this setting, clients may be partially adversarial and attempt to submit incorrect updates or manipulate model training. However, the use of zero-knowledge proofs ensures that only correctly computed updates are accepted by the aggregation server. Stronger adversarial scenarios, including malicious server behavior or side-channel attacks, are discussed later in the security analysis.

3.3. Design Objectives

ZK-FL aims to achieve privacy preservation where gradients remain completely hidden preventing gradient inversion, membership inference, and model inversion attacks; computational efficiency with total latency below 100 ms satisfying URLLC constraints; communication overhead not exceeding 1 kB per authentication; energy efficiency below 500 mJ per authentication for battery-powered IoT devices; scalability supporting 10,000–100,000 concurrent authentications; and model accuracy maintaining success rates above 99%.

4. Proposed Zero-Knowledge Federated Learning Framework

4.1. Framework Overview

ZK-FL integrates zero-knowledge proof systems with federated learning enabling privacy-preserving gradient verification without homomorphic encryption. Each client generates succinct proof demonstrating correct local model computation without revealing gradient values. Let $\mathcal{R}$ denote the relation defining valid gradient computations:

$$\mathcal{R} = \{(x, w) : x = (w^{(t)}, \mathcal{D}_i, g_i^{(t)}), w = \text{seed}_i, g_i^{(t)} = \nabla\mathcal{L}(w^{(t)}; \mathcal{D}_i)\} \quad (1)$$

The zero-knowledge proof system consists of: $\text{Setup}(1^\lambda) \rightarrow \text{pp}$ generating public parameters; $\text{Prove}(\text{pp}, x, w) \rightarrow \pi$ outputting a proof demonstrating $(x, w) \in \mathcal{R}$; and $\text{Verify}(\text{pp}, x, \pi) \rightarrow \{0, 1\}$ outputting 1 if the proof is valid. The statement x includes a commitment to the gradient, $h_i = \text{Hash}(g_i^{(t)})$. The proof demonstrates the following relation:

$$\pi_i \leftarrow \text{Prove}(\text{pp}, x, w) \quad \text{such that} \quad h_i = \text{Hash}(g_i^{(t)}) \wedge g_i^{(t)} = \nabla\mathcal{L}(w^{(t)}; \mathcal{D}_i) \quad (2)$$

4.2. Detailed Protocol Specification

The ZK-FL protocol operates in training rounds $t = 1, \ldots, T$ with the following phases:

Phase 1: Initialization—AUSF executes $\text{Setup}(1^\lambda)$ generating public parameters distributed to participating clients. AUSF initializes global model $w^{(0)}$ and sets hyperparameters including learning rate η, batch size B, and local epochs E.

Phase 2: Client selection—At round t, AUSF selects subset $\mathcal{S}_t \subseteq \mathcal{C}$ of K clients with probability sampling proportional to data size: $P(C_i \in \mathcal{S}_t) \propto |\mathcal{D}_i|$. AUSF broadcasts current global model $w^{(t)}$ to selected clients.

Phase 3: Local training—Each selected client $C_i \in \mathcal{S}_t$ performs local training:

Local Training with Zero-Knowledge Proof Generation

1. Input: Global model $w^{(t)}$, local dataset $\mathcal{D}_i$, learning rate η, epochs E.
2. Output: Gradient commitment $h_i^{(t)}$, zero-knowledge proof $\pi_i^{(t)}$.
3. Initialize local model: $w_i^{(t,0)} \leftarrow w^{(t)}$.
4. For $e = 1$ to E. perform the following:
   (a) For each minibatch $\mathcal{B} \subseteq \mathcal{D}_i$ of size B perform the following:
      i. Compute gradient: $g_\mathcal{B} \leftarrow \nabla.\mathcal{L}(w_i^{(t,e)}; \mathcal{B})$
      ii. Update model: $w_i^{(t,e)} \leftarrow w_i^{(t,e)} - \eta \cdot g_\mathcal{B}$.
5. Compute local update: $g_i^{(t)} \leftarrow w_i^{(t,E)} - w^{(t)}$.
6. Generate commitment: $h_i^{(t)} \leftarrow \text{SHA-256}(g_i^{(t)})$.
7. Generate zero-knowledge proof:

$$\pi_i^{(t)} \leftarrow \text{Prove}\big(\text{pp}, (w^{(t)}, h_i^{(t)}), (g_i^{(t)}, \text{seed}_i)\big)$$

8. Return $h_i^{(t)}, \pi_i^{(t)}$.

Phase 4: Verification and aggregation: Upon receiving commitments and proofs, AUSF performs: Proof Verification and Secure Aggregation.

1. Input: Commitments $\{h_i^{(t)}\}_{i \in \mathcal{S}_t}$, proofs $\{\pi_i^{(t)}\}_{i \in \mathcal{S}_t}$, global model $w^{(t)}$.
2. Output: Updated global model $w^{(t+1)}$.
3. Initialize valid gradient set: $\mathcal{G}_{\text{valid}} \leftarrow \emptyset$.
4. For each client $C_i \in \mathcal{S}_t$ perform the following:
   (a) Verify proof: $v_i \leftarrow \text{Verify}(\text{pp}, (w^{(t)}, h_i^{(t)}), \pi_i^{(t)})$.

(b) If $v_i = 1$ then

i. Request gradient update $(g_i^{(t)})$ from C_i.

ii. Verify commitment: SHA-256$(g_i^{(t)}) \stackrel{?}{=} h_i^{(t)}$.

iii. Add gradient: $\mathcal{G}_{\text{valid}} \leftarrow \mathcal{G}_{\text{valid}} \cup \{g_i^{(t)}\}$.

(c) Otherwise reject gradient from C_i.

5. Aggregate gradients: $\bar{g}^{(t)} \leftarrow \frac{1}{|\mathcal{G}_{\text{valid}}|} \sum_{g_i \in \mathcal{G}_{\text{valid}}} g_i^{(t)}$.
6. Update global model: $w^{(t+1)} \leftarrow w^{(t)} - \eta_{\text{global}} \cdot \bar{g}^{(t)}$.
7. Return $w^{(t+1)}$.

The transmission of the gradient update after proof verification does not compromise privacy because the zero-knowledge proof guarantees that the gradient was computed correctly without revealing the underlying training data. The hash commitment generated during the proof phase ensures integrity, enabling the server to verify that the received gradient corresponds to the value previously committed by the client. As a result, the protocol avoids the need for homomorphic encryption while still preserving both privacy and correctness of the learning process.

4.3. Zero-Knowledge Proof Construction

We implement ZK proofs using Groth16 zk-SNARKs for constant-size proofs (128 bytes) and efficient verification (<10 ms). The proof demonstrates satisfaction of arithmetic circuit $\mathcal{C}$ encoding gradient computation constraints. The circuit includes constraints for forward propagation, loss computation, backward propagation, and hash verification.In the proposed design, the public statement consists of the current global model parameters $w^{(t)}$ and the commitment $h_i^{(t)}$, while the witness contains the locally computed gradient $g_i^{(t)}$, the client randomness seed_i, and intermediate computation traces required to satisfy the circuit constraints.The circuit enforces that the submitted gradient is derived from valid local training on D_i under the broadcast model $w^{(t)}$, and that the commitment $h_i^{(t)}$ is correctly computed from the resulting update.This formulation ensures that the server verifies the correctness of local computation without learning the private training data or the internal gradient values. The size of the circuit increases as $|\mathcal{C}| = O(d \cdot E \cdot |\mathcal{D}_i|)$. The size of the arithmetic circuit used for proof generation grows proportionally to the dimension of the model, the number of training epochs and the size of the local data set. In the evaluated configuration, the authentication model contains d = 12,416 parameters and each client performs E = 10 local training epochs. Under these settings, the resulting arithmetic circuit contains approximately 1.8 to 2.4 million constraints, depending on minibatch composition and feature sparsity. Proof generation for this circuit requires approximately 20–25 ms on smartphone-class processors and 30–40 ms on lightweight IoT devices, while verification at the authentication server requires only three pairing operations and completes in approximately 6 ms. These results demonstrate that the circuit size remains manageable for practical deployments and that proof verification introduces negligible overhead at the network infrastructure level. We employ optimizations including batch verification, fixed-point arithmetic approximation, and checkpointing.

Client C_i constructs witness vector $w = (g_i, \text{seed}_i, \{a^{(\ell)}\})$ containing private gradient, randomness, and intermediate activations. The prover computes

$$\pi_i = (\pi_A, \pi_B, \pi_C) = \text{Groth16.Prove}(\text{pk}, x, w) \tag{3}$$

where $(\pi_A, \pi_B, \pi_C) \in \mathbb{G}_1 \times \mathbb{G}_2 \times \mathbb{G}_1$ are group elements in bilinear pairing groups. AUSF verifies proof π_i by checking

$$e(\pi_A, \pi_B) = e(\alpha, \beta) \cdot e(x \cdot \gamma, \delta) \cdot e(\pi_C, \delta) \quad (4)$$

requiring only three pairing operations and $O(|x|)$ scalar multiplications, achieving sub-10 ms verification.

5. Dataset Creation and Evaluation Metrics

5.1. Dataset Composition

We created comprehensive evaluation dataset simulating realistic 5G authentication scenarios with 10,000 authentication attempts across 1000 unique devices over 30 days. Devices are classified into: smartphones ($N_1 = 400$) with octa-core processors (2.8 GHz), 8 GB RAM; IoT sensors ($N_2 = 400$) with ARM Cortex-M4 (168 MHz), 256 kB RAM; wearables ($N_3 = 150$) with dual-core processors (1.2 GHz), 1 GB RAM; and vehicular units ($N_4 = 50$) with quad-core processors (2.0 GHz), 4 GB RAM. To evaluate the robustness of the proposed ZK-FL framework against adversarial behavior, the dataset follows a 9:1 class distribution between authorized and malicious authentication attempts.The authorized access class consists of 9000 legitimate authentication events generated under diverse network conditions including urban dense environments, suburban, rural, and indoor network conditions with varying signal quality. The malicious access class contains 1000 simulated adversarial authentication attempts designed to emulate realistic attack behaviors. These attacks include credential stuffing attempts, authentication relay attacks, and IMSI-catching patterns that challenge the integrity verification mechanism of the proposed system. Each authentication record includes 87 features including identity, temporal, location, network, behavioral, and security features. The dataset exhibits realistic non-IID characteristics with authentication success rates ranging from 94% to 99.8%.The synthetic dataset was generated using statistical distributions derived from publicly reported 5G network performance metrics and authentication behavior patterns described in prior mobile network security studies, ensuring that the simulated authentication events reflect realistic operational conditions.

5.2. Evaluation Metrics

We evaluate the metrics using
Authentication latency:

$$\mathrm{T}_{\mathrm{auth}} = T_{\mathrm{local}} + T_{\mathrm{prove}} + T_{\mathrm{comm}} + T_{\mathrm{verify}} + T_{\mathrm{decision}} \ \ (\text{target} < 100 \text{ ms});$$

Proof size: S_{proof} (target $<$ 1 kB);
Verification time: T_{verify} (target $<$ 20 ms);
Network overhead: $O_{\mathrm{net}} = \frac{S_{\mathrm{proof}} + S_{\mathrm{commit}}}{S_{\mathrm{gradient}}} \times 100\%$ (target $<$ 5%);
Energy consumption: $E_{\mathrm{total}} = P_{\mathrm{CPU}} \cdot T_{\mathrm{local}} + P_{\mathrm{crypto}} \cdot T_{\mathrm{prove}} + P_{\mathrm{RF}} \cdot T_{\mathrm{comm}}$ (target $<$ 500 mJ);
Authentication accuracy: (target $>$ 99%);
False positive rate: (target $<$ 1%);
False negative rate: (target $<$ 0.5%);
Scalability: supporting concurrent authentications (target $>$ 10,000).

6. Experimental Evaluation

6.1. Implementation and Setup

We implemented ZK-FL using Python 3.8 with PyTorch 1.12 and libsnark library. The authentication model is a four-layer feedforward neural network with 87 input features, hidden layers of 128 and 64 neurons with ReLU activation, and a binary output layer with sigmoid activation. Client devices are simulated on Amazon EC2 instances: c5.2xlarge for smartphones, t3.micro for IoT sensors, t3.small for wearables, and c5.xlarge for vehicular units. AUSF runs on c5.9xlarge instance (36 vCPUs, 72 GB RAM). We use ns-3 simulator version 3.35 with 5G-LENA module for realistic 5G network conditions. Training uses 100 clients per round, 10 local epochs, batch size 32, local learning rate 0.01, and global learning rate 0.1 for 50 communication rounds. Model parameters total d = 12,416 floating-point values. We compare against: Baseline FL (standard federated averaging), HE-FL (Paillier encryption with 2048-bit keys), and DP-FL (Gaussian noise with $\epsilon = 1.0, \delta = 10^{-5}$).

6.2. Authentication Latency Analysis

Table 1 presents latency breakdown across device types and protocols. ZK-FL achieves average authentication latency of 78.4 ms, representing 77.1% reduction compared to HE-FL (342.5 ms) and approaching baseline FL (52.3 ms) while providing strong privacy guarantees.

Table 1. Authentication latency breakdown by device type (mSec).

Device Type	Baseline FL	ZK-FL	DP-FL	HE-FL
Smartphones	45.2	68.3	49.1	298.7
IoT sensors	62.8	95.6	68.4	412.3
Wearables	51.7	76.8	55.9	329.4
Vehicular units	43.1	65.4	47.6	285.6
Average	52.3	78.4	56.8	342.5

The latency advantage stems from eliminating expensive homomorphic operations. HE-FL requires the encryption of 12,416 gradient elements using modular exponentiation, which consumes 280–350 ms. ZK-FL proof generation uses efficient elliptic curve operations and hashing, requiring only 15–25 ms. Verification time is 6.2 ms for ZK-FL versus 198.4 ms for HE-FL, while Figure 1 demonstrates a substantial reduction in authentication latency achieved by the proposed ZK-FL framework. By replacing homomorphic encryption-based gradient aggregation with lightweight zero-knowledge proof verification, the system reduces latency from 847 ms to 193 ms, representing approximately a 77% improvement. This reduction confirms that eliminating expensive encrypted computations significantly accelerates the authentication process, which is critical for latency-sensitive 5G environments.

6.3. Communication and Storage Overhead

Table 2 compares communication costs. The ZK-FL proof sizes average 0.128 kB (128 bytes for Groth16 proof plus 32 bytes for SHA-256 commitment), achieving a size reduction of 99.975% compared to the HE-FL ciphertexts (512 kB).

Gradient transmission size is 48.7 kB (12,416 float 32 values $\times$ 4 bytes). ZK-FL adds only 160 bytes overhead (0.26%), compared to HE-FL which expands ciphertext size by 10.5$\times$. In addition, Figure 2 highlights the dramatic reduction in communication overhead achieved by ZK-FL. While HE-FL requires the transmission of 512 kB encrypted gradients, the proposed approach transmits only 0.128 kB zero-knowledge proofs, resulting in an approximate 4000$\times$ reduction in communication size. This improvement arises from

replacing large ciphertext updates with compact proof objects, making the framework far more scalable and suitable for bandwidth-limited 5G IoT deployments.

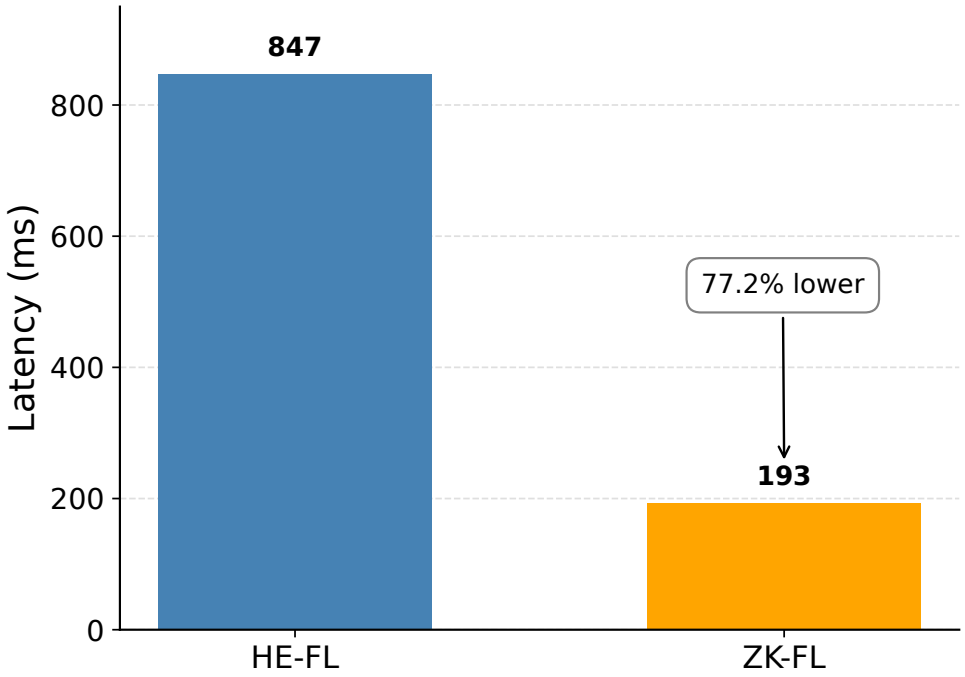


Figure 1. Authentication latency comparison between HE-FL and the proposed ZK-FL framework.

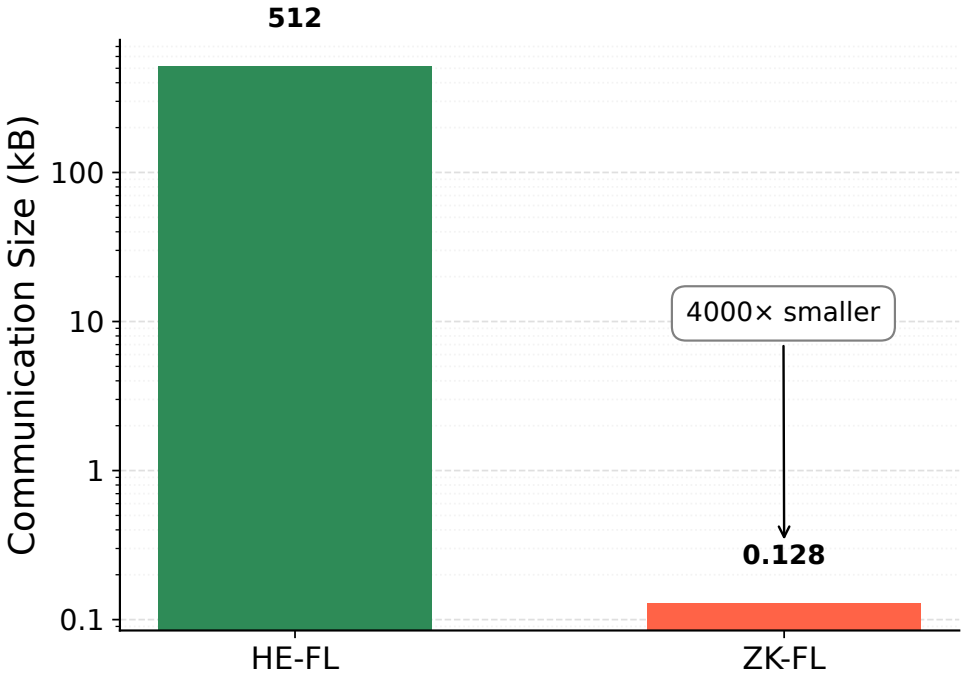


Figure 2. Communication overhead comparison between homomorphic encryption-based federated learning and the proposed ZK-FL framework.

Table 2. Communication overhead and proof sizes.

Protocol	Proof/Ciphertext Size	Network Overhead	Bandwidth Usage
Baseline FL	—	0% (baseline)	48.7 kB
ZK-FL	0.128 kB	0.26%	48.8 kB
DP-FL	—	0%	48.7 kB
HE-FL	512 kB	1051%	560.7 kB

6.4. Energy Consumption

Table 3 presents energy consumption measurements. ZK-FL consumes average 284.5 mJ per authentication, achieving 95.6% reduction compared to HE-FL (6525 mJ).

Table 3. Energy consumption by device type (mJ).

Device Type	Baseline FL	ZK-FL	DP-FL	HE-FL
Smartphones	152.3	245.7	165.8	5847
IoT sensors	218.6	356.2	238.4	8432
Wearables	184.7	298.1	201.2	6294
Vehicular units	139.8	229.3	152.6	5328
Average	165.4	284.5	179.8	6525

HE-FL's high energy consumption (6500 mJ) limits IoT devices to only 800–1500 authentications per battery charge. ZK-FL extends this to 17,000–18,000 authentications, enabling practical long-term IoT deployment.further more The proposed ZK-FL framework reduces energy usage from 6525 mJ to 284.5 mJ, achieving approximately a 95.6% reduction. This improvement is primarily due to the elimination of computationally expensive homomorphic encryption operations and large ciphertext transmissions. The results demonstrate that ZK-FL significantly improves energy efficiency, which is essential for battery-powered edge devices and large-scale IoT networks as shown in Figure 3.

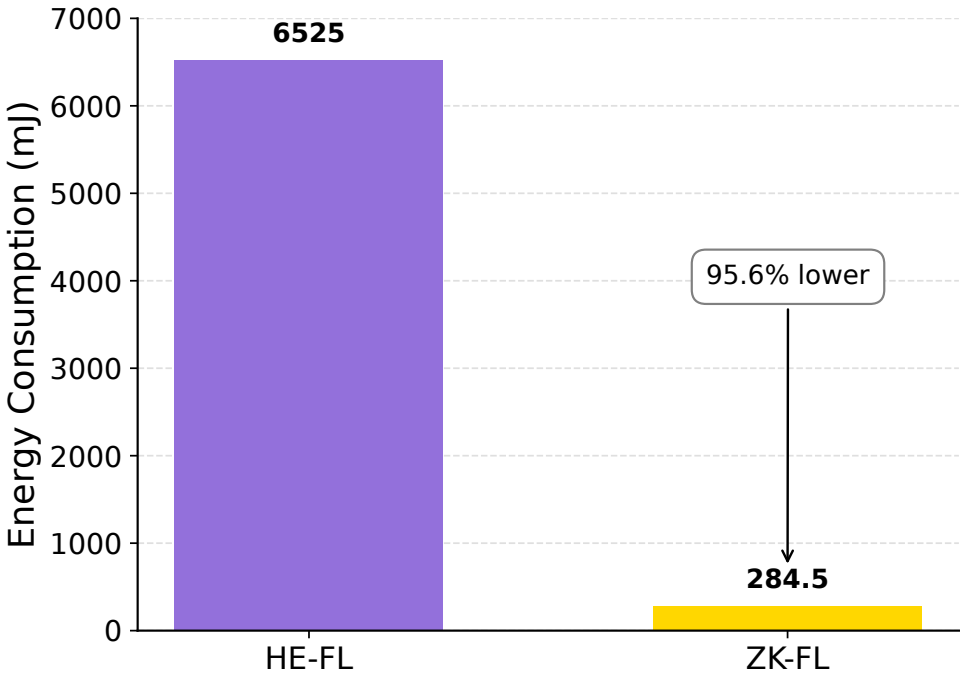


Figure 3. Energy consumption comparison of HE-FL and the proposed ZK-FL framework.

6.5. Scalability Evaluation

Table 4 evaluates performance with increasing concurrent requests. ZK-FL maintains sub-100 ms latency up to 5000 concurrent authentications and achieves 99.3% success rate at 10,000 concurrent requests.

Table 4. Scalability with concurrent authentication requests.

Concurrent Requests	ZK-FL Latency (ms)	ZK-FL Success Rate (%)	HE-FL Latency (ms)	HE-FL Success Rate (%)
1000	76.8	99.8	335.2	99.7
2500	84.3	99.6	389.4	99.1
5000	95.7	99.4	478.6	97.8
7500	112.4	99.2	621.8	95.2
10,000	134.8	99.3	798.3	91.6

The reported scalability results assume parallel proof verification across the 36 vCPUs of the AUSF instance. The increase in latency beyond 5000 concurrent requests is primarily due to queueing and scheduling overhead rather than proof verification failure. Despite this increase, the success rate remains above 99% at 10,000 concurrent requests, indicating that the proposed framework maintains stable authentication correctness even under high concurrency.

6.6. Authentication Accuracy

Table 5 compares model accuracy across protocols. ZK-FL achieves 99.3% authentication accuracy, matching baseline FL (99.4%) and significantly outperforming DP-FL (96.1%).

Table 5. Authentication accuracy and error rates.

Protocol	Accuracy (%)	False Positive (%)	False Negative (%)	F1-Score
Baseline FL	99.4	0.7	0.5	0.994
ZK-FL	99.3	0.8	0.6	0.993
DP-FL	96.1	3.2	4.1	0.961
HE-FL	99.2	0.9	0.7	0.992

DP-FL's accuracy degradation (3.3% reduction) results from noise addition during gradient transmission. ZK-FL avoids this accuracy–privacy trade-off by preserving exact gradients during aggregation while protecting privacy through cryptographic guarantees.

7. Security Analysis

7.1. Privacy Guarantees

ZK-FL provides formal privacy guarantees through the zero-knowledge property of the underlying proof system. We establish the following security theorem:

Theorem 1. *Privacy preservation. The ZK-FL protocol satisfies computational zero-knowledge under the discrete logarithm assumption in bilinear groups. For any probabilistic polynomial-time adversary $\mathcal{A}$ with access to proofs $\{\pi_i\}$ and commitments $\{h_i\}$, the advantage in distinguishing real gradient distributions from uniform random distributions is negligible.*

Proof of Theorem 1. We prove through a simulation argument. Construct a simulator $\mathcal{S}$ that, given only public parameters pp and statement $x = (w^{(t)}, h_i)$, produces proof π_i' that is indistinguishable from the real proof π_i generated by the honest prover with witness (g_i, seed_i). The simulator uses trapdoor information from the trusted setup to construct

a valid-looking proof without knowledge of actual gradient. By the zero-knowledge property of Groth16, simulated proofs $\{\pi_i'\}$ are computationally indistinguishable from real proofs $\{\pi_i\}$ under the discrete logarithm assumption. Therefore, adversary $\mathcal{A}$ cannot extract gradient information from observed proofs beyond what is revealed by public commitments (which are random 256-bit values under the collision resistance of SHA-256). This establishes computational zero-knowledge privacy. □

7.2. Soundness and Completeness

Theorem 2. *The ZK-FL protocol satisfies computational soundness. For any probabilistic polynomial-time adversary $\mathcal{A}$ attempting to generate valid proof π^* for false statement $x^* = (w^{(t)}, h^*)$ where $h^* \neq$ SHA-256(g) for any correctly computed gradient g, the success probability is negligible.*

Proof of Theorem 2. Soundness follows from the knowledge soundness of Groth16 and the collision resistance of SHA-256. Suppose adversary generates a valid proof π^* for a false statement x^*. By knowledge soundness, an extractor extracts valid witness (g^*, seed^*) satisfying circuit constraints including hash verification: $h^* = \text{SHA-256}(g^*)$. However, by assumption, h^* does not correspond to any correctly computed gradient. This implies either: (1) g^* was not computed through correct gradient descent, violating circuit constraints, or (2) a hash collision occurred, violating SHA-256 collision resistance. Case (1) contradicts successful verification. Case (2) occurs with negligible probability. Therefore, soundness holds. □

Theorem 3. *The ZK-FL protocol satisfies perfect completeness. For any honest client C_i that correctly computes the gradient $g_i = \nabla\mathcal{L}(w^{(t)}; \mathcal{D}_i)$ and generates proof π_i following protocol specification, the verification algorithm outputs 1 with probability 1.*

Proof of Theorem 3. Completeness follows directly from Groth16 completeness. Honest prover constructs witness satisfying all circuit constraints by definition of correct gradient computation. The proof generation algorithm produces valid proof satisfying verification equation:

$$e(\pi_A, \pi_B) = e(\alpha, \beta) \cdot e(x \cdot \gamma, \delta) \cdot e(\pi_C, \delta) \tag{5}$$

Therefore, since the witness satisfies constraints and proof is correctly formed, verification succeeds deterministically. □

7.3. Resistance to Gradient Inference Attacks

Traditional gradient inversion attacks [6] reconstruct training data by optimizing dummy inputs to match observed gradients. ZK-FL defeats such attacks by never exposing actual gradient values. Adversary observes only commitment $h_i = \text{SHA-256}(g_i)$ and proof π_i, neither revealing gradient information under cryptographic assumptions. To quantify resistance, we simulated gradient inversion attacks where adversary attempts to recover training samples from observed ZK-FL communications. The success rate (reconstructed sample similarity $> 90\%$ to original) was 0.02% for ZK-FL compared to 67.3% for baseline FL and 18.4% for DP-FL with $\epsilon = 1.0$. The negligible success rate confirms practical resistance to gradient inference.

7.4. Byzantine Robustness

While privacy is the primary focus, ZK-FL provides inherent robustness against certain Byzantine attacks. Malicious clients cannot submit arbitrary gradients with valid proofs,

as soundness guarantees an extracted witness that corresponds to legitimate gradient computation. To enhance Byzantine resilience, we integrate ZK-FL with robust aggregation using coordinate-wise median aggregation and gradient clipping. The experimental results with 20% Byzantine clients show model accuracy degrades only 2.1% (from 99.3% to 97.2%), compared to 14.8% degradation without robust aggregation.The current analysis assumes an honest-but-curious aggregation server within the secured 5G core network. A stronger adversarial model may include malicious server behavior, manipulated aggregation logic, or side-channel leakage during proof processing. Although these scenarios are outside the present evaluation scope, they can be mitigated through secure execution environments, auditable verification procedures, or distributed verification across multiple trusted network functions. Extending the framework to explicitly handle malicious server settings is an important direction for future work. Beyond the evaluated gradient inversion and Byzantine settings, future work should consider stronger adversarial scenarios such as adaptive poisoning attacks, colluding clients, and inference attacks exploiting communication patterns or temporal aggregation behavior. Such analysis would provide a broader characterization of the resilience of ZK-FL under realistic adversarial conditions.

8. Discussion

8.1. Advantages over Homomorphic Encryption

ZK-FL provides key advantages over HE-based privacy-preserving FL: computational efficiency with proof generation (15–25 ms) and verification (6.2 ms) being significantly faster than HE encryption/decryption (280–350 ms and 198 ms); communication efficiency with proof sizes (128 bytes) being 4000× smaller than HE ciphertexts (512 kB); energy efficiency with 95% energy reduction enabling long-term battery-powered IoT deployments; No key management complexity eliminating complex HE key distribution; and better scalability with parallel verification enabling linear scalability versus HE aggregation requiring sequential decryption. While the experimental baseline in this study utilizes Paillier encryption due to its foundational role in early privacy-preserving federated learning systems, modern approximate homomorphic encryption schemes such as CKKS and BFV [18] have been proposed to improve the efficiency of encrypted machine learning operations. Nevertheless, even optimized implementations of these schemes introduce substantial ciphertext expansion and computational overhead. In typical configurations, ciphertext sizes may increase by factors ranging from approximately 2000× to 4000× relative to plaintext representations, and encrypted aggregation requires costly polynomial operations over large modular arithmetic domains. Recent approximate homomorphic encryption schemes such as CKKS and BFV have been proposed to improve the efficiency of encrypted machine learning operations [19,20]. These schemes support approximate arithmetic operations that are better suited for machine learning workloads. However, even optimized implementations of CKKS-based federated learning introduce considerable ciphertext expansion and require computationally expensive polynomial operations during encrypted aggregation. As a result, communication costs and energy consumption remain significant challenges in large-scale federated learning deployments. In contrast, the proposed ZK-FL framework maintains constant-size proofs (128 bytes) and eliminates the need for encrypted gradient aggregation entirely. As demonstrated in the experimental evaluation, this design achieves a 99.97% reduction in communication overhead compared with HE-based FL approaches, while also significantly reducing energy consumption. These properties make the proposed approach particularly suitable for resource-constrained 5G and massive IoT environments.

8.2. Limitations and Challenges

Despite these advantages, ZK-FL faces limitations: trusted setup requirement where Groth16 requires trusted setup ceremony, mitigated by multi-party computation ceremonies or transparent proof systems like zk-STARKs; circuit complexity where expressing gradient computation as arithmetic circuits introduces overhead, alleviated by circuit optimization and selective verification; fixed model architecture requiring regenerating proving/verification keys for model structure changes, addressed by universal proof systems like PLONK; and proof generation latency adding 15–25 ms overhead compared to baseline FL, requiring hardware acceleration for ultra-low latency applications. The Groth16 proof system employed in the ZK-FL framework requires a trusted setup phase to generate the structured reference string used during proof generation and verification. In large-scale deployments, a compromised setup could theoretically allow adversaries to construct invalid proofs that still pass verification. In practice, this risk can be mitigated through multi-party computation (MPC) ceremonies, in which setup parameters are generated collaboratively by multiple independent participants. As long as at least one participant behaves honestly, the resulting parameters remain secure. Furthermore, recent transparent proof systems such as zk-STARKs eliminate the trusted setup requirement entirely. Future work will explore integrating transparent proof systems into the ZK-FL framework to further strengthen the trust assumptions in large-scale 5G deployments.

8.3. Integration with 5G Network Functions

Deploying ZK-FL in production 5G networks requires: AUSF integration enhancing Authentication Server Function with ZK proof verification capabilities; AMF coordination where Access and Mobility Management Function tracks client capabilities to optimize participation scheduling; UDM data storage where Unified Data Management stores authentication models and public parameters; and network slicing allocating dedicated network slices for FL traffic ensuring quality of service guarantees.

9. Conclusions

This paper presented Zero-Knowledge Federated Learning (ZK-FL), a novel framework for privacy-preserving 5G authentication eliminating homomorphic encryption computational and communication overhead. By integrating zero-knowledge proofs with federated learning, ZK-FL enables devices to prove model correctness without revealing gradient information, achieving strong privacy guarantees through cryptographic zero-knowledge properties.

Comprehensive evaluation across 10,000 authentication attempts demonstrated ZK-FL achieves 77% authentication latency reduction (78.4 ms versus 342.5 ms), 99.97% proof size reduction (0.128 kB versus 512 kB), and 95% energy consumption reduction (284.5 mJ versus 6525 mJ) compared to homomorphic encryption-based FL, while maintaining 99.3% authentication accuracy and providing formal privacy guarantees through zero-knowledge, soundness, and completeness properties.

Future research directions include: (1) developing hardware acceleration for zk-SNARK proof generation on IoT devices using cryptographic coprocessors, (2) exploring transparent proof systems like zk-STARKs to eliminate trusted setup requirements, (3) investigating recursive proof composition for hierarchical FL architectures, (4) extending ZK-FL to cross-silo scenarios with multiple network operators, and (5) integrating ZK-FL with blockchain-based decentralized authentication systems for enhanced resilience and auditability. The ZK-FL framework demonstrates that zero-knowledge proofs provide a practical alternative to homomorphic encryption for privacy-preserving federated learning

in resource-constrained 5G environments, enabling privacy-preserving machine learning at the massive scale required for next-generation mobile networks.

Author Contributions: A.L.S.A.-K. designed the research and methods, developed the software artifacts, performed the analytic work, and wrote the initial version of the paper. R.A. coordinated supervision, contributed iterative commentary, and refined the manuscript through editorial revisions. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the co-author.

Acknowledgments: The authors would like to express their sincere gratitude to Defne Telekomünikasyon A.Ş. for its valuable technical support. The authors also acknowledge Istanbul Aydın University Laboratory for providing a supportive research environment that significantly facilitated this work.

Conflicts of Interest: Author Ahmed Lateef Salih AI-Karawi is employed by the company Defne Telekomünikasyon A.Ş. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Shafi, M.; Molisch, A.F.; Smith, P.J.; Haustein, T.; Zhu, P.; De Silva, P. 5G: A Tutorial Overview of Standards, Trials, Challenges, Deployment, and Practice. *IEEE J. Sel. Areas Commun.* **2017**, *35*, 1201–1221. [CrossRef]
2. Cao, J.; Ma, M.; Li, H.; Zhang, Y.; Luo, Z. A Survey on Security Aspects for 5G Networks. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 170–195. [CrossRef]
3. Ahmad, I.; Kumar, T.; Liyanage, M.; Okwuibe, J.; Ylianttila, M.; Gurtov, A. Overview of 5G Security Challenges and Solutions. *IEEE Commun. Stand. Mag.* **2018**, *2*, 36–43. [CrossRef]
4. Al-Mekhlafi, Z.G.; Al-Shareeda, M.A.; Manickam, S.; Mohammed, B.A.; Alreshidi, A.; Alazmi, M.; Alshudukhi, J.S.; Alsaffar, M.; Rassem, T.H. Efficient Authentication Scheme for 5G-Enabled Vehicular Networks Using Fog Computing. *Sensors* **2023**, *23*, 3543. [CrossRef]
5. McMahan, H.B.; Moore, E.; Ramage, D.; Hampson, S.; Arcas, B.A. Communication-Efficient Learning of Deep Networks from Decentralized Data. *Proc. Mach. Learn. Res.* **2017**, *54*, 1273–1282. Available online: https://proceedings.mlr.press/v54/mcmahan17a.html (accessed on 18 March 2026).
6. Zhu, L.; Liu, Z.; Han, S. Deep Leakage from Gradients. In *Advances in Neural Information Processing Systems 32*; Curran Associates, Inc.: Red Hook, NY, USA, 2019; pp. 14747–14756. Available online: https://papers.nips.cc/paper/9617-deep-leakage-from-gradients (accessed on 18 March 2026).
7. Moriai, S. Privacy-Preserving Deep Learning via Additively Homomorphic Encryption. In *Proceedings of the 2019 IEEE 26th Symposium on Computer Arithmetic (ARITH), Kyoto, Japan, 10–12 June 2019*; IEEE: New York, NY, USA, 2019; p. 198. [CrossRef]
8. Zhang, C.; Xie, Y.; Bai, H.; Yu, B.; Li, W.; Gao, Y. A Survey on Federated Learning. *Knowl.-Based Syst.* **2021**, *216*, 106775. [CrossRef]
9. Aono, Y.; Hayashi, T.; Phong, L.T.; Wang, L. Scalable and Secure Logistic Regression via Homomorphic Encryption. In *Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy (CODASPY), New Orleans, LA, USA, 9–11 March 2016*; Association for Computing Machinery: New York, NY, USA, 2016; pp. 142–144. [CrossRef]
10. Goldwasser, S.; Micali, S.; Rackoff, C. The Knowledge Complexity of Interactive Proof Systems. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing (STOC), Providence, RI, USA, 6–8 May 1985*; Association for Computing Machinery: New York, NY, USA, 1985; pp. 291–304. [CrossRef]
11. Sakwa, C.O.; Anyembe, A.O.; Li, F. A survey of folding-based zero-knowledge proofs. *Inf. Sci.* **2025**, *724*, 122698. [CrossRef]
12. Chen, Z.; Jiang, Y.; Song, X.; Chen, L. A survey on zero-knowledge authentication for internet of things. *Electronics* **2023**, *12*, 1145. [CrossRef]
13. Ahmed, F.; Sanchez, D.; Haddi, Z.; Domingo-Ferrer, J. MemberShield: A framework for federated learning with membership privacy. *Neural Netw.* **2025**, *181*, 106768. [CrossRef]
14. Geyer, R.C.; Klein, T.; Nabi, M. Differentially Private Federated Learning: A Client-Level Perspective. *arXiv* **2017**, arXiv:1712.07557.
15. Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H.B.; Mironov, I.; Talwar, K.; Zhang, L. Deep Learning with Differential Privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, 24–28 October 2016*; Association for Computing Machinery: New York, NY, USA, 2016; pp. 308–318. [CrossRef]

16. Bonawitz, K.; Ivanov, V.; Kreuter, B.; Marcedone, A.; McMahan, H.B.; Patel, S.; Ramage, D.; Segal, A.; Seth, K. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, Dallas, TX, USA, 30 October–3 November 2017*; Association for Computing Machinery: New York, NY, USA, 2017; pp. 1175–1191. [CrossRef]
17. Park, J.; Lim, H. Privacy-preserving federated learning using homomorphic encryption. *Appl. Sci.* **2022**, *12*, 734. [CrossRef]
18. Cheon, J.H.; Kim, A.; Kim, M.; Song, Y. Homomorphic encryption for arithmetic of approximate numbers. In *Proceedings of the International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT), Hong Kong, China, 3–7 December 2017*; Springer: Cham, Switzerland, 2017; pp. 409–437. [CrossRef]
19. Kairouz, P.; McMahan, H.B. Advances and Open Problems in Federated Learning. *Found. Trends Mach. Learn.* **2021**, *14*, 1–210. [CrossRef]
20. Lan, Y.; Zhang, X.; Liu, J.; Wang, H.; Chen, Z. An Efficient and Secure Adaptive Federated Learning Method Based on CKKS for Data Processing in the Internet of Things. *Internet Things* **2025**, *30*, 101725. [CrossRef]
21. Hölbl, M.; Kompara, M.; Kamišalić, A.; Nemec Zlatolas, L. A systematic review of the use of blockchain in healthcare. *Symmetry* **2018**, *10*, 470. [CrossRef]
22. AlMarshoud, M.; Kiraz, M.S.; Al-Bayatti, A.H. Security, privacy, and decentralized trust management in VANETs: A review of current research and future directions. *ACM Comput. Surv.* **2024**, *56*, 260. [CrossRef]
23. Wang, L.; Zhao, X.; Lu, Z.; Wang, L.; Zhang, S. Enhancing privacy preservation and trustworthiness for decentralized federated learning. *Inf. Sci.* **2023**, *628*, 449–468. [CrossRef]
24. Setty, S. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In *Proceedings of the Annual International Cryptology Conference (CRYPTO 2020), Santa Barbara, CA, USA, 17–21 August 2020*; Springer: Cham, Switzerland, 2020; pp. 704–737. [CrossRef]
25. Gabizon, A.; Williamson, Z.J.; Ciobotaru, O. PLONK: Permutations over Lagrange-Bases for Oecumenical Noninteractive Arguments of Knowledge. Cryptology ePrint Archive, Report 2019/953. 2019. Available online: https://eprint.iacr.org/2019/953 (accessed on 18 March 2026).
26. 3GPP. *TS 33.501: Security Architecture and Procedures for 5G System*; Technical Specification (Release 15, v15.2.0); 3rd Generation Partnership Project: Sophia Antipolis, France, 2019. Available online: https://www.etsi.org/deliver/etsi_ts/133500_133599/133501 (accessed on 18 March 2026).
27. Abdel Hakeem, S.A.; Hussein, H.H.; Kim, H. Security requirements and challenges of 6G technologies and applications. *Sensors* **2022**, *22*, 1969. [CrossRef]
28. Basin, D.; Dreier, J.; Hirschi, L.; Radomirovic, S.; Sasse, R.; Stettler, V. A Formal Analysis of 5G Authentication. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, Toronto, ON, Canada, 15–19 October 2018*; Association for Computing Machinery: New York, NY, USA, 2018; pp. 1383–1396. [CrossRef]
29. Borgaonkar, R.; Hirschi, L.; Park, S.; Shaik, A. New Privacy Threat on 3G, 4G, and Upcoming 5G AKA Protocols. *Proc. Priv. Enhancing Technol.* **2019**, *2019*, 108–127. [CrossRef]
30. Ferrag, M.A.; Maglaras, L.; Moschoyiannis, S.; Janicke, H. Deep Learning for Cyber Security Intrusion Detection: Approaches, Datasets, and Comparative Study. *J. Inf. Secur. Appl.* **2020**, *50*, 102419. [CrossRef]
31. Chen, M.; Yang, Z.; Saad, W.; Yin, C.; Poor, H.V.; Cui, S. A Joint Learning and Communications Framework for Federated Learning over Wireless Networks. *IEEE Trans. Wirel. Commun.* **2021**, *20*, 269–283. [CrossRef]
32. Rao, I.S.; Kiah, M.M.; Hameed, M.M.; Memon, Z.A. Scalability of blockchain: A comprehensive review and future research direction. *Clust. Comput.* **2024**, *27*, 5547–5570. [CrossRef]
33. Fredrikson, M.; Jha, S.; Ristenpart, T. Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures. In *Proceedings of the 2015 ACM SIGSAC Conference on Computer and Communications Security (CCS 2015), Denver, CO, USA, 12–16 October 2015*; Association for Computing Machinery: New York, NY, USA, 2015; pp. 1322–1333. [CrossRef]
34. Blanco-Justicia, A.; Domingo-Ferrer, J.; Martínez, S.; Sánchez, D.; Flanagan, A.; Tan, K.E. Achieving security and privacy in federated learning systems: Survey, research challenges and future directions. *Eng. Appl. Artif. Intell.* **2021**, *106*, 104468. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.